ANL/DIS-06-5

Argonne NATIONAL LABORATORY

Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms

Co-hosted by

Argonne National Laboratory
The University of Chicago

In association with

North American Association for Computational Social and Organizational Science

Gleacher Center

450 North Cityfront Plaza Drive
Chicago, Illinois

October 13–15, 2005

THE UNIVERSITY OF CHICAGO

A U.S. Department of Energy laboratory
managed by The University of Chicago

Agent 2005

ISBN 0-9679168-6-0
9 780967 916866 90000

Suggested Citations:

Proceedings: Macal, C.M., M.J. North, and D. Sallach (editors), 2005, *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms,* ANL/DIS-06-5, ISBN 0-9679168-6-0, co-sponsored by Argonne National Laboratory and The University of Chicago, October 13–15.

Paper: Jones, A., and B.C. Smith, 2005, "Title of Paper," in *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms,* ANL/DIS-06-5, ISBN 0-9679168-6-0, C.M. Macal, M.J. North, and D. Sallach (editors), co-sponsored by Argonne National Laboratory and The University of Chicago, October 13–15.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by The University of Chicago under contract W-31-109-Eng-38. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne, see www.anl.gov.

Availability of This Report

Decision and Information Sciences Division Office
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439-4832
Telephone: (630) 252-5464
Home page: http://www.dis.anl.gov/
Agent home page: http://www.agent2005.anl.gov

Reproduced from best available originals.

ANL/DIS-06-5

Argonne NATIONAL LABORATORY

Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms

Co-hosted by

Argonne National Laboratory
The University of Chicago

In association with

North American Association for Computational Social and Organizational Science

Gleacher Center

450 North Cityfront Plaza Drive
Chicago, Illinois

October 13–15, 2005

THE UNIVERSITY OF CHICAGO

A U.S. Department of Energy laboratory managed by The University of Chicago

Agent 2005

Edited by:
Charles Macal, Michael North, and David Sallach

CONTENTS

FRIDAY, OCTOBER 14, 2005

Social Simulation Applications

Invited Speaker: J. Epstein

Public Policy

Parallel Applications Session I – Economics and Environmental Policy

SATURDAY, OCTOBER 15, 2005

Computational Social Theory

Invited Speaker: L.-E. Cederman

Social and Conflict Dynamics

Patterns and Actions

Evolutionary Models

FOREWORD

Welcome to Agent 2005, co-sponsored by Argonne National Laboratory and the University of Chicago. Since this Chicago conference series on agent simulation began in 1999, it has focused on three priorities: (1) social simulation models, (2) tools and toolkits, and (3) how these first two priorities are applied in diverse application areas. These priorities, which are also Special Interest Groups in the North American Association for Computational Social and Organizational Science (NAACSOS), have helped us to attract quality papers and, thus, to keep the conference fresh and stimulating. As supporters of NAACSOS, we encourage you to become active as well and, thereby, advance the field of computational social science.

The theme of this conference is *Generative Social Processes, Models, and Mechanisms*. Social agent simulation is increasingly recognized as an effective methodology within the social sciences, particularly applicable to a growing range of policy issues. At the same, the deep and dynamic complexity of the domain continues to challenge social modelers. Several approaches have emerged with the potential to address social complexities. One strategy is to design and employ multi-mechanism models. This strategy provides a way for specialized mechanisms to progress at their own pace, while participating in a broader model, and also facilitates the cooperation of diverse subject-matter experts. A second approach is more embryonic. Social complexity is generated by the interaction of social agents. Accordingly, much of the richness of social institutions and processes is emergent. Using generative software to generate social dynamics thus appears to provide a natural strategy. Software can be generative in two senses: it generates a customized simulation application from a common ontology or set of principles and, also, situated agents, as individuals and collectivities, dynamically generate and evolve preferences, plans, communication, action, and consequences. These two approaches by no means exhaust research creativity in social simulation, and the Agent 2005 conference is pleased to highlight all such innovations.

As in previous years, our goal is to share our models and results, stimulate and learn from each other, and identify areas in which progress is both necessary and possible. We believe you will find the regular sessions to be rich and substantial, and the invited speakers to be stimulating and insightful. We also value (and record) the discussions in each section and, thereby, make them available for future reference in the conference *Proceedings*.

We hope you enjoy Agent 2005 and become increasingly committed to the kinds of social science progress that computational modeling makes possible. Once again, welcome.

The Center for Complex Adaptive Agent Systems Simulation
Argonne National Laboratory and The University of Chicago

Charles Macal
Michael North
David Sallach
Thomas Wolsko

ACKNOWLEDGMENTS

We acknowledge the support of many people in organizing *Agent 2005*. Kathy Ruffatto handled registration, administration, and logistics, along with assistance from Dee Albarado. Guy Pandola managed the conference Web site. Argonne's Technical Services Division prepared the program book and the proceedings for publication. Margaret Clemmons served as project manager and editor, with support from Michele Nelson in graphic design and Argonne's Document Processing Center.

ORGANIZING COMMITTEE

Charles Macal, Argonne National Laboratory and The University of Chicago
Michael North, Argonne National Laboratory and The University of Chicago
David Sallach, Argonne National Laboratory and The University of Chicago
Thomas Wolsko, Argonne National Laboratory

Thursday, October 13, 2005

Methods, Toolkits, and Techniques

Integrating Agent Modeling Toolkits and Critical Applications

LINKING REPAST AND COMPUTATIONAL MATHEMATICS SYSTEMS: *MATHEMATICA* AND MATLAB

C.M. MACAL* and T.R. HOWE,
Argonne National Laboratory, Argonne, IL

ABSTRACT

This paper describes investigations into linking computational mathematics systems (CMSs) with the Java-based Repast agent-based modeling toolkit. The goal is to build an interactive and seamless agent simulation environment that benefits from the strongest points of each component. In general, CMSs such as *Mathematica* and MATLAB are fully integrated development environments. Their interpretative nature and seamless integration of graphical and statistical analysis capabilities provide immediate feedback to users during the model prototyping and development phases. These features make CMSs particularly useful as rapid prototype development tools or as part of large-scale model development efforts that use agent-based modeling toolkits. Large-scale agent-based modeling environments, such as Repast, support features specific to agent modeling, including the availability of sophisticated time schedulers, agent communications mechanisms, flexible interaction topologies, and facilities for storing and displaying agent states. Typically, Repast users build models by incorporating Repast library components into their own programs or by using a visual scripting environment that creates program code automatically.

Of the two CMSs investigated, *Mathematica* was found to be more suitable for linking to Repast because of its sophisticated capabilities for linking to Java, as provided by the *J/Link* environment. In the system described here, *Mathematica* is used to implement the models of agent behavior and interaction, while Repast is used for its discrete-event simulation controls and features. In this configuration, Repast simulation classes are extended to call user-written *Mathematica* programs that model agent behaviors. Interactive simulation controls provided by Repast allow the user to control and interact with the simulation as it progresses through time. *Mathematica's* full mathematical modeling libraries, visualization capabilities, statistical analysis routines, and database capabilities are available to the agent model during and after the simulation. The linked *Mathematica*-Repast system is an interpreted environment that requires no compilation or linking steps, which are needed by other general programming languages or by agent-based modeling toolkits alone. The linked system also allows users to access the full range of Repast-provided Java classes and *Mathematica*-provided Java classes through the *J/Link* libraries. One feature of the linked system is that it naturally provides a real-time, interactive animation system for agent-based simulations without accumulating the overhead entailed in storing a long stream of graphics images. Two examples of using the linked system for social agent simulation experimentation are presented.

Keywords: Agent-based modeling, computational mathematics systems, *Mathematica*, Repast, toolkit, rapid prototype

---

* *Corresponding author address*: Charles M. Macal, Decision and Information Sciences Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439; e-mail: macal@anl.gov.

INTRODUCTION

This paper describes investigations into linking computational mathematics systems (CMSs) with the Java-based Repast agent-based modeling toolkit. The goal is to create an interactive and seamless agent simulation environment that can benefit from the strongest points of each system. In general, computational mathematics systems, such as *Mathematica* and MATLAB, are fully integrated development environments and offer linkages to Java in some form. After reviewing the Java linking capabilities of MATLAB and *Mathematica*, we conclude that *Mathematica* is more suitable for linking to Repast owing to its more complete capabilities for linking to Java, as provided by the *J/Link* environment. We go on to describe a new configuration for building agent-based models in which Repast simulation classes are extended to call *Mathematica* programs that model agent behaviors and interactions. Repast is used for its discrete-event (time-stepped) simulation features and interactive simulation controls. The interactive graphical user interface (GUI) is composed of components from both the Repast and *J/Link* GUI classes.

This paper is organized as follows. The second section provides background information on agent-based modeling, computational mathematics systems (specifically, MATLAB and *Mathematica*), and agent-based toolkits (specifically, Repast). The third section presents the agent implementation approach through two examples, the Boids model by Reynolds (2005) and a social agent interaction model. The fourth section presents the architecture of the Repast-*Mathematica* system. The final section draws conclusions and identifies promising future directions.

BACKGROUND

Agent-based Modeling and Simulation (ABMS)

ABMS is an active area of rapid growth and development in modeling and simulation. Social agent simulation offers new possibilities for gaining insights into fundamental social processes (Axelrod 1997; Axtell 2000; Bankes 2002; Bonabeau 2001; Cederman 2002; Epstein and Axtell 1996; Gilbert and Troitzsch 1999; Kohler et al. 2005; Sallach and Macal 2001; Schelling 1971). A social agent is a discrete entity with its own goals and behaviors. Agents have autonomy, and, in general, have capabilities to adapt and modify their own behaviors. Agents are diverse and heterogeneous. The key assumptions are as follows:

- The most relevant aspects of behaviors (at least the most relevant to the model application) can be described.

- Mechanisms by which agents interact can be described and represented algorithmically.

- Complex social processes can be modeled by building a system “from the bottom up.”

Examples of agents include people, groups, and organizations; social insects and swarms; robots and systems of autonomous vehicles, and many others.

An agent-based model consists of three elements:

- Set of agents,
- Set of agent relationships, and
- Framework for simulating agent behaviors and agent interactions.

In an abstract view of an agent-based model, an agent includes a representation for the agent and its behaviors; that is, how the agent processes information (updates its state) on the basis of current events and references to past events (memory). An agent relationship exists for a pair of agents and includes a representation of the relationship's characteristics and the mechanisms that act on the relationship to update its state. A framework for simulating agent behaviors and interactions is the mechanism that updates the universe of agents and their relationships in the model; the simulation framework is provided by agent-based modeling toolkits.

A defining principle of ABMS is the assumption that agents have access to only local information. The visibility of agents is constrained to be within an agent's local neighborhood, where the extent of the neighborhood is open to the discretion of the modeler. This characteristic has important implications for the structure of ABMS. A more complete overview of ABMS can be found in Macal and North (2005).

Agent-based Modeling Toolkits: Repast

Substantial public research and development investments have produced many ABMS software environments that are now freely available. These include Repast, Swarm, NetLogo, and MASON, as well as numerous others. Large-scale ABMS toolkits, such as Repast, extend agent modeling beyond simple desktop environments and allow thousands to millions of agents to engage in sophisticated interchanges. Proprietary toolkits are also available. For a recent review and comparison of Java-based agent modeling toolkits, see Tobias and Hoffman (2004). Large-scale agent-based modeling environments generally support features specific to ABMS, including:

- Availability of sophisticated time schedulers,
- Facilities for storing and displaying agent states,
- Agent communications mechanisms, and
- Flexible agent interaction topologies (grids, networks, and geographical information-system [GIS]-based structures)

Repast (REcursive Porous Agent Simulation Toolkit; see Repast 2005, http://repast.sourceforge.net/) is the leading free and open-source, large-scale ABMS library (Collier and Sallach 2001; Collier et al. 2003). Repast seeks to support the development of extremely flexible models of agents with an emphasis on social interactions. It has been used extensively in social simulation applications. North and Macal (2005) has an overview of Repast applications in the social sciences.

Repast is a pure Java modeling environment that supports the development of large-scale agent models and linkages to other software systems. It includes a variety of features, such as a fully concurrent discrete event scheduler, a model visualization environment, integration with GISs for modeling spatially situated agents on real maps, and adaptive behavioral tools (e.g., neural networks, genetic algorithms). Repast includes a full range of time-stepped simulation run controls that are useful in driving simulation models written in other languages, such as *Mathematica*, which is the focus of the remainder of this paper. Users build simulations by incorporating Repast library components into their own programs or by using the visual Repast for Python Scripting environment (Collier et al. 2003). More information on Repast, as well as downloads, can be found at the Repast home page (Repast 2005). Repast is maintained by the Repast Organization for Architecture and Design (ROAD).

Computational Mathematics Systems: *Mathematica* and MATLAB

Mathematica (see Wolfram Research, Inc. 2005, www.wolfram.com) and MATLAB (see MathWorks 2005, www.mathworks.com) are examples of CMSs, which allow users to apply powerful mathematical algorithms to solve problems through a convenient and interactive user interface. CMSs supply a wide range of built-in functions and algorithms. Their origins go back to the late 1980s.

CMSs are structured in two main parts: (1) a user interface that allows dynamic user interaction and (2) an underlying computational engine, or kernel, that performs the computations according to the user's instructions. Unlike conventional programming languages, CMSs are interpreted rather than compiled, so there is immediate feedback to the user, but some performance penalty is paid. The underlying computational engine is written in the C programming language for these systems, but the user does not see the C coding. The most recent releases of CMSs are fully integrated systems that combine capabilities for data input and export, graphical display, and the capability to link to external programs written in conventional languages such as C or Java by using interprocess communication protocols. The powerful features of CMSs, their convenience of use, the need for the user to learn only a limited number of instructions, and the immediate feedback provided to users make CMSs good candidates for developing agent-based social simulations. However, it should be noted that unlike dedicated agent-based toolkits like Repast, CMSs have not, to date, provided specific capabilities or resources for modeling agents.

A further distinction can be made among CMSs. A subset of CMSs — called *computational algebra systems* (CASs) — are interactive programs that, in contrast to numerical processing systems, allow mathematical computations with symbolic expressions. Computations are carried out exactly, according to the rules of algebra, instead of numerically with approximate floating point arithmetic. In contrast, a numeric processing language requires that every variable have a value assigned before it is used. Typical uses of CASs are equation solving, symbolic integration and differentiation, exact calculations in linear algebra, simplification of mathematical expressions, and variable precision arithmetic. Computational mathematics systems consist of numeric processing systems or symbolic processing systems, or possibly a combination of both. Especially when numeric and algebraic capabilities are combined into a multi-paradigm programming environment, new modeling possibilities for developing sophisticated agent-based social simulations with minimal coding open up. The core MATLAB system is strictly a numeric CMS, with limited symbolic processing capabilities provided by an add-on module; *Mathematica*, however, combines numeric and extensive symbolic processing capabilities. The symbolic processing capabilities in *Mathematica* greatly extend its ability to represent abstract data types and generalized data structures. The flexibility of data types plays an important role in developing large-scale, extensible models for agent-based social simulation. Macal (2004) has a more extensive review of CMS applications to agent simulation.

In-depth reviews of the system documentation and testing concluded that MATLAB's current implementation allows Java GUIs and graphics to be driven by MATLAB programs, but it does not allow the two-way linkage (i.e., only the MATLAB-to-Java linkage is implemented). Furthermore, *Mathematica's* underlying architecture and *J/Link* class library permit it to be extended to allow the two-way linkage requirements to be met (Gayley 2004a,b). Therefore, the remainder of this paper describes the Repast-*Mathematica* system that was successfully implemented.

Mathematica

Mathematica's symbolic processing capabilities allow programming in multiple paradigms, either as alternatives or in combination (Wolfram 2003). Programming paradigms include functional programming, logic programming, procedural programming, rule-based programming, and object-oriented programming. *Mathematica* is an interpreted language with the C-based kernel of *Mathematica* running underneath the notebook interface. In terms of data types, everything in *Mathematica* is an "expression." An expression is a data type with a head and a list of arguments in which even the head of the expression is part of the expression's arguments.

The *Mathematica* user interface consists of a notebook. A notebook is a fully integratable development environment plus a complete publication environment. The *Mathematica* application programming interface (API) allows programs written in C, Fortran, or Java to interact with the kernel. The API has facilities for dynamically calling routines from *Mathematica* as well as for calling *Mathematica* as a computational engine.

Capabilities of Mathematica for Linking to Java

Computational social scientists have identified the need for a rapid social science discovery process in which computational experimentation and electronic laboratories would be key instruments facilitating rapid progress (Sallach 2003). A linked Repast-CMS system would provide a unique set of benefits for rapid prototype development of social agent models, including these:

- Access to Repast's scheduler and simulation controls (SetUp, Initialize, Step, Run, Pause, Stop, Exit), which allow users to interactively control a time-stepped simulation;

- *Mathematica's* mathematical modeling libraries, visualization capabilities, and statistical analysis routines;

- Database capabilities available to the real-time agent-based modeling environment and for post-simulation analysis;
- Spatial and geodesic data and analysis capabilities such as geographical information systems (GIS) provided through Repast and *Mathematica*;
- Access to the full range of Repast-provided Java classes, as well as the *Mathematica* Java classes provided through *J/Link*;
- Interpreted environment, requiring no compilation or linking steps needed by general programming languages and agent-based modeling toolkits alone; and
- A real-time, interactive animation system for agent-based simulations.

The linkage is created by developing a software component — a Java class — that wraps communication protocols between Java and *Mathematica*.

AGENT IMPLEMENTATION

Boids Model Example

We next implement and describe the "Boids" model, developed by Craig Reynolds (Reynolds 2005). Boids is an interesting example of a very simple agent simulation that demonstrates emergent group behavior based on simple rules for agent behaviors. (Note that the rules do not include learning or adaptation on the part of the agents.) There are three main rules (Rules 1–3) and one rule that has been added for the demonstration example (Rule 4). The rules, which treat the agents as a "flock," are:

- *Rule 1 (cohesion).* An agent steers to move toward the average position of local flockmates.
- *Rule 2 (separation).* An agent steers to avoid crowding local flockmates.
- *Rule 3 (alignment).* An agent steers toward the average heading of local flock-mates.
- *Rule 4 (containment).* An agent heads toward in-bounds if it strays out of bounds.

The agent behaviors implied by the four rules are illustrated in Figure 1.

Agents move in continuous space and discrete time. This allows them to move at different speeds according to their individual characteristics, as opposed to a grid-based model, which would enforce movement over discrete space and constant speeds.

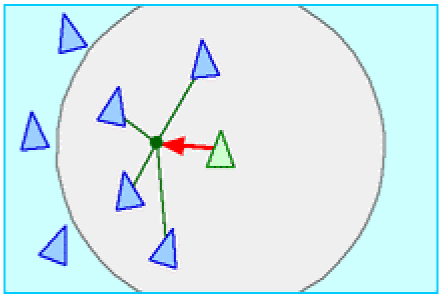

Cohesion: Steer to move toward the average position of local flockmates

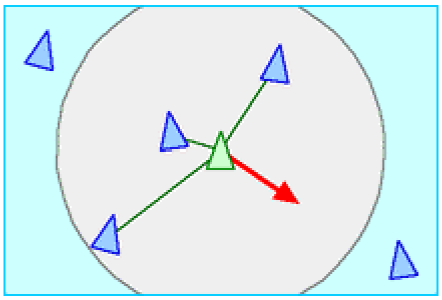

Separation: Steer to avoid crowding local flockmates

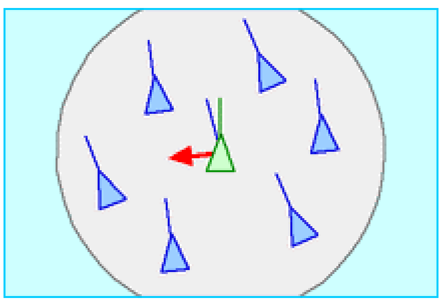

Alignment: Steer toward the average heading of local flock-mates

FIGURE 1 Agent rules in the Boids model

We next demonstrate the Boids model by using the Repast-*Mathematica* system. We may ask some questions in advance of the demonstration, such as:

- How do we expect the agents to behave as they operate according to the simple rules just described?
- Will the agents randomly mill around, stand still, or exhibit other behavior?
- How can we statistically characterize the behavior of the agents that we observe during the simulation?

A snapshot of the Boids simulation is shown in Figure 2.

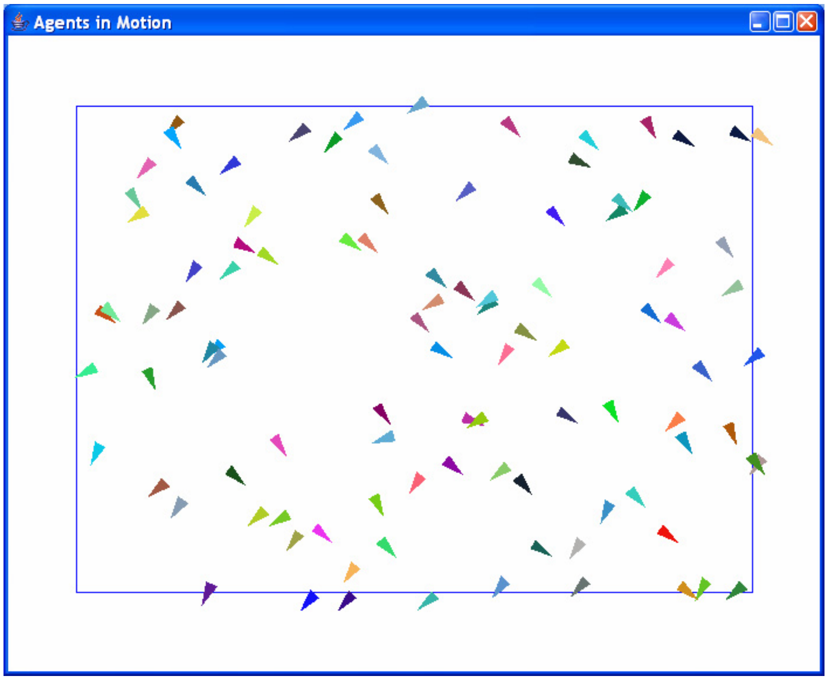

(a) Initial random configuration

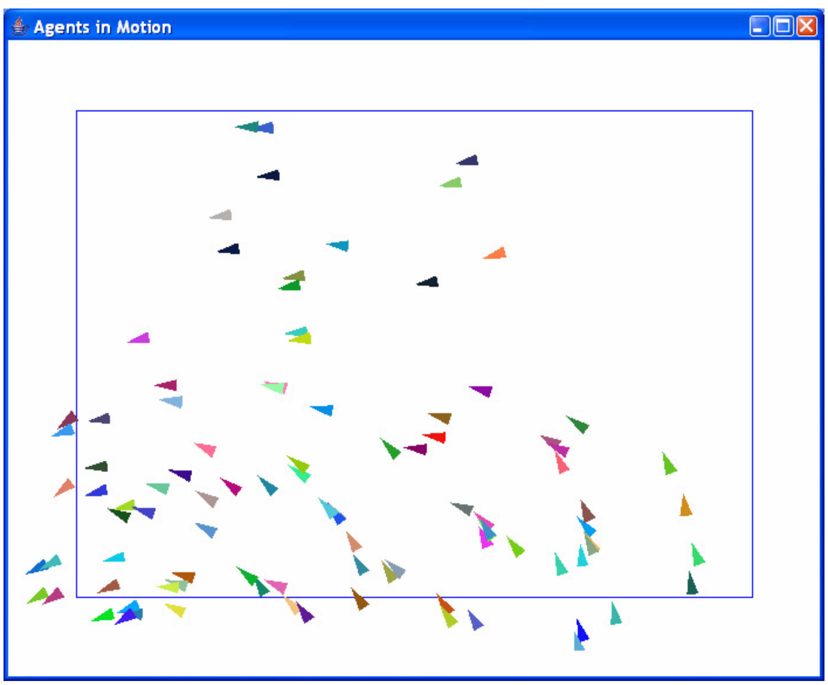

(b) After 650 updates

FIGURE 2 Boids model simulation

Figure 3 shows an analysis of the simulation results. The statistics recorded during the simulation are available for post-simulation analysis in the *Mathematica* notebook from which we launched the simulation. The center of mass (CM) for the agents is plotted over the simulation time along with the distance of the CM from the original CM. Further time-series analysis done in the *Mathematica* notebook (not shown here) on the data series logged during the simulation reveals that the time series is of a chaotic nature.

Agent Representation

Here we briefly describe the agent implementation in *Mathematica* that underlies the Boids model. The representation is based on defining abstract data types and associated attributes and functions, similar to an object-oriented (O-O) implementation. Maeder (2000) describes this approach for implementing models in *Mathematica*. This approach most closely aligns the *Mathematica* implementation of the agent model with the representation that would be most natural to implement directly in Repast or any other O-O agent-based modeling toolkit. This is desirable from the standpoint of having a smooth transition path for the Repast-CMS system to a full Repast implementation, should this transition become advantageous at some point in the development cycle. Furthermore, the *Mathematica* agent model can be easily described via a unified modeling language (UML) representation (Booch et al. 1998), which separates the model specification from its implementation, whether the implementation is in *Mathematica*, Repast, or another O-O programming language.

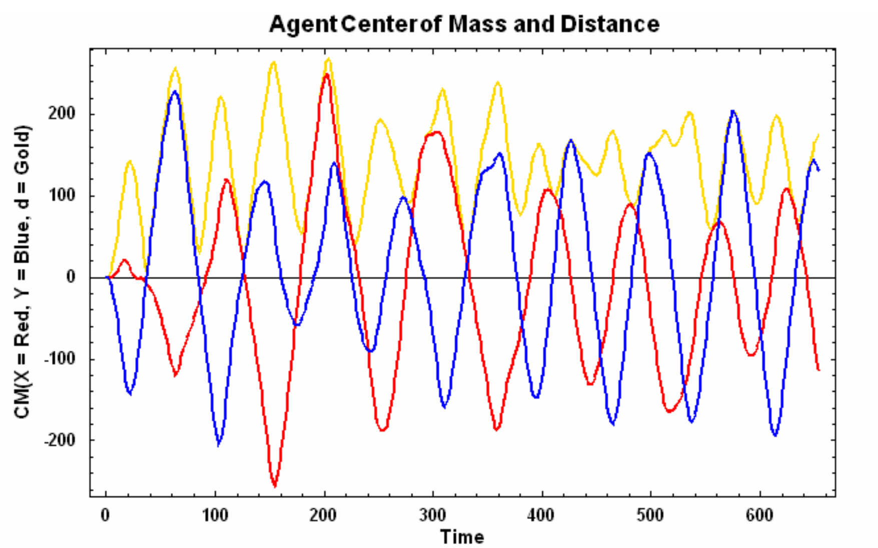

FIGURE 3 Post-simulation analysis of Boids simulation time-series data

An agent type or pattern is represented by an abstract data type, `agent`. This is a *Mathematica* expression with a head `agent` and a sequence of patterns that corresponds to agent attributes. The agent data type is a template for the specific instance of agents in the model and corresponds directly to a class in Java or Repast. For example, in the Boids model, the agent data type is defined as:

`agent`[name (string), location: x coordinate, location: y coordinate,
velocity: x component, velocity: y component, other attributes (e.g., color)].

Agent Operations

Agent Creation

All of the agents in the model are represented by a list called `agents`. Agents are created and updated in user-written *Mathematica* programs. An example of the `agents` list is:

```
agents ={
      agent["#26", 26.0, -50.9, -19.8, 38.7, «JavaObject[java.awt.Color]», 42,
                  False, 16.0, False, 0, 1, TargetNone, 2.6, {}],
      … ,
      agent[…]
      };
```

Agent Updating

Agent states are updated at every time step in the simulation. Updated agent attributes include location and status. *Mathematica* allows for a functional programming style and has the capacity for mapping a function onto a list such that the function is applied to every element in the list in one statement. For example, the function `agentRulesBoid` is a user-written *Mathematica* program that updates the status of an agent. This is accomplished via the following statement, which is applied to `agents` each time period:

```
agents = agentRulesBoid /@ agents;
```

where /@ is shorthand for the mapping function applied to the list `agents`. Note that the order in which the agents are updated is the order in which the agents appear in the list `agents`. This order is arbitrary. Since the order of agent interactions in the Boids model has no effect on the final states of the updated agents, this arbitrary ordering is adequate. If, on the other hand, the ordering of agent updating was a significant factor in determining model results, other agent ordering schemes could be invoked before the `agentRulesBoid` program was applied. For example, the `agents` list could be randomized each time before applying `agentRulesBoid`, as in:

```
agents = agentRulesBoid /@ randomizeList[agents] ;
```

where `randomizeList` is a program that simply randomizes an arbitrary list.

The `agentRulesBoid` is a program structured as follows in *Mathematica* notation:

```
agentRulesBoid[a:agent[id_, x_, y_, vx_, vy_, col_, sz_, fil_, v0_, sel_, nColl_, 0,
                    target_, orient_, leat_]] := Module[...

  ... (*CALCULATE prey exclusive center of mass*

  cmXPrey = If[neighborsXPreyLocs === {}, {0,0}, Mean[neighborsXPreyLocs]
                - {x,y}] (*prey XCM*);

  ... (*RULE 1: COHESION Move rule1Param % way toward XCM.*)

  v1 = rule1Param * cmXPrey;
  xx =…; yy =…; wx =…; wy =…;

  ... (*RETURN updated agent with new position, velocity and orientation*)

  agent [id, xx, yy, wx, wy, col, sz, fil, v0, sel, nColl, 0, target, newOrient, leat]

   ];
```

The `agentRulesBoid` program is structured to take an individual agent as input and to return that agent, with its attributes updated, to reflect the results of agent decision making and interaction. The `agentRulesBoid` program is a modular representation of agent behavior. All the rules of agent behavior and interaction are included in this single program. A similar program can be written to represent the behaviors of each of the other types of agents in the model. For example, a predator agent was developed and instantiated with behaviors to chase and prey upon the other agents in the Boids model.

Social Agent Interaction Model

We next describe a social agent interaction model (SAIM) and implement it in the Repast-*Mathematica* system. Agents in this example have two types of social behaviors: movement and influence. The movement rules are motivated by the "mobile heterogeneous agent model" of Gaylord and D'Andria (1998). The influence rules are based on opinion change as applied to social influence (Friedkin and Johnsen 1990). The SAIM combines both of these behaviors into a single model at the agent level. This model turns out to be an interesting example that demonstrates the highly nonlinear behavior of a complex system composed of agents with very simple rules for agent behaviors. The social agent mobility update rules are as follows:

- *Rule 1 (sequencing).* Agents move (or not) in random order at each simulation time step. (This assumption eliminates contention for moving into unoccupied spaces.)

- *Rule 2 (stationarity).* If the nearest neighbor site that the agent is facing is occupied by another individual, the individual remains in place and chooses a new random direction to face.

- *Rule 3 (mobility).* If the nearest neighbor site that the agent is facing is unoccupied, the individual moves into that space and chooses a new random direction to face.

The social agent influence update rules are as follows:

- *Rule 1.* If an agent is surrounded by more red neighbors than blue neighbors, the agent changes to red

- *Rule 2.* If an agent is surrounded by more blue neighbors than red neighbors, the agent changes to blue.

- *Rule 3.* Ties are resolved randomly.

SAIM is a grid-based (lattice) model in which agents move in discrete space and discrete time. In each time period, agents either remain stationary or move to an adjacent grid cell.

Next we demonstrate the SAIM by using the Repast-*Mathematica* system. We set up an experimental simulation in which there are two types of agents: blue and red. Initially, there are twice as many blue agents as red agents. Each side tries to convince the other side of its position. An agent surrounded by more agents of one color than another adopts the position of the dominant agent. There is one stipulation: Each red agent is *twice* as convincing as each blue agent. An agent surrounded by only half as many red agents as blue agents is just as likely to adopt red's position. Thus, the total convincing power of the red and blue agents is equal. In the event that an agent is surrounded by an equal number of red and blue agents in terms of their relative convincing power, the agent randomly chooses a color.

We begin an experimental simulation by randomly placing the red and blue agents across the grid. Simulation experiments suggest some interesting questions: Will one type of agent have its opinion dominate, or will there be a stable mix of positions in the long run? What agent and system behaviors will we observe during the simulation? A snapshot of the agent simulation is displayed in Figure 4.

Figure 5 shows the results of a single simulation run, which are available for post-simulation analysis in the *Mathematica* notebook from which we launched the simulation. The numbers of red and blue agents are plotted over the simulation. The plot for this case indicates an extensive initial period of give-and-take between the red and blue agents, which ultimately gives way to what appears to be some kind of phase transition that leads directly to the complete dominance of the red agent population. It should be noted that other random initial placements of red and blue agents may result in different outcomes, such as blue agents winning.

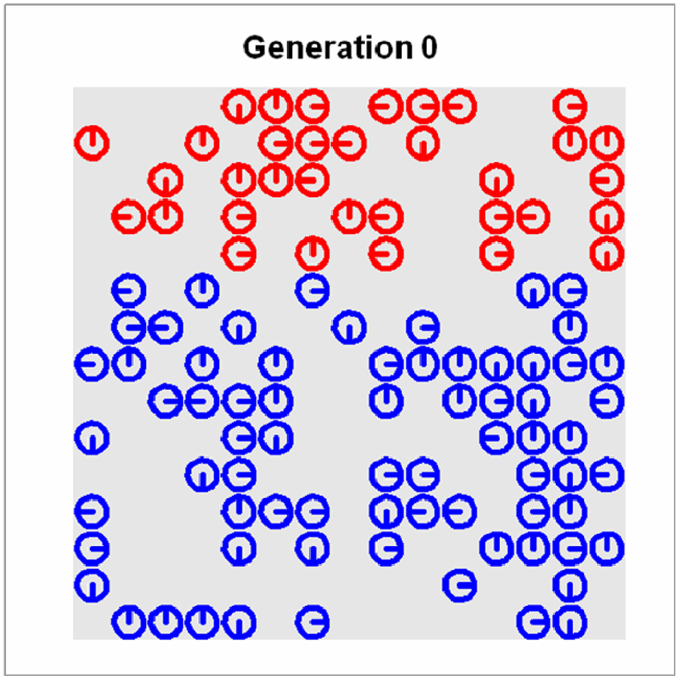

(a) Initial agent distribution

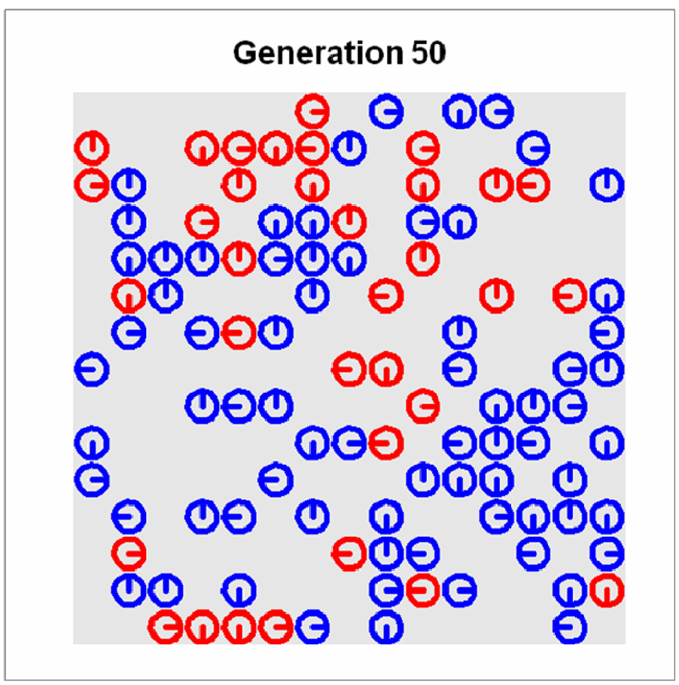

(b) After 50 generations

FIGURE 4 SAIM

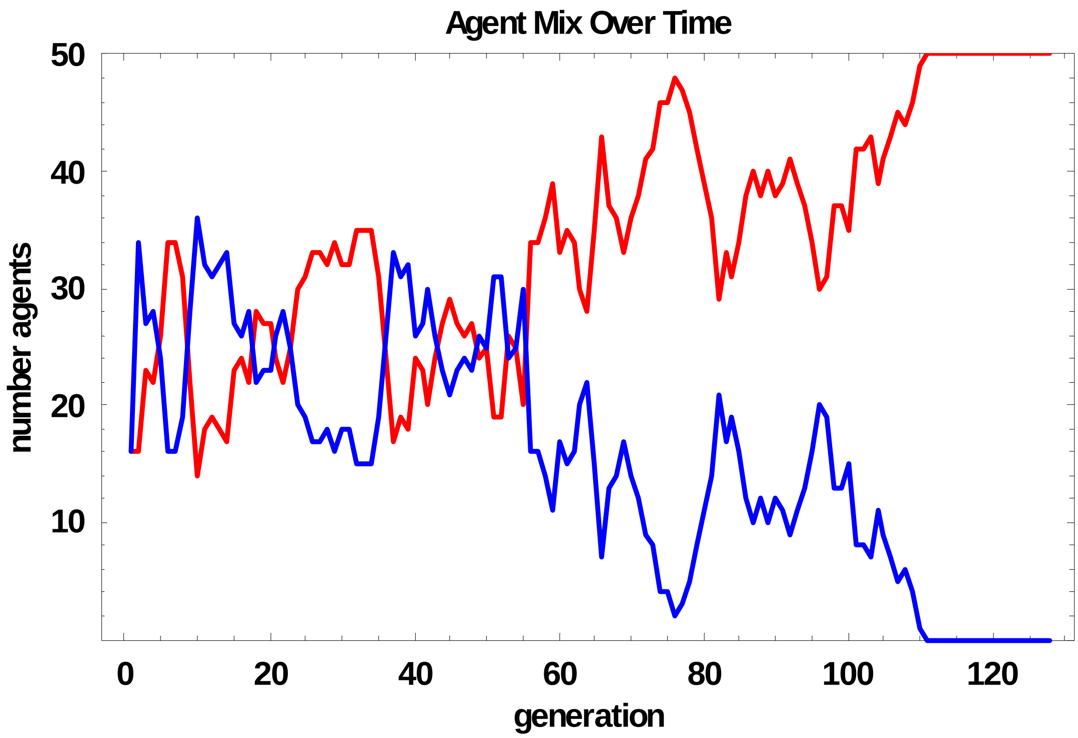

FIGURE 5 Post-simulation analysis of SAIM time-series data

ARCHITECTURE OF REPAST-*MATHEMATICA*

The interactive Repast simulation controls (buttons of the user interface) for controlling the course of a time-stepped simulation include `Setup`, `Initialize`, `Step`, `Run`, `Pause`, `Stop`, and `Exit`. *Mathematica's J/Link* classes provide functionality to the user for interacting with and controlling the simulation as it progresses via the Repast controller through the Repast GUI. A single Mathematica notebook corresponds to a single agent simulation model and includes these programs. The simulation user interface for the Repast-*Mathematica* system is shown in Figure 6. The details of the Repast-*Mathematica* architecture for a typical agent simulation model are shown in Figure 7.

User-written *Mathematica* Simulation Control Programs

The following user-written *Mathematica* programs are the core of an agent simulation application: `Simulate`, `init`, `step`, `stop`, `exit`, and `plotCMX`**.**

- `Simulate` is a *Mathematica* program that creates an instance of the Repast class controller. `Simulate` loads any needed Java classes, creates the GUI objects and event handlers, and assigns *Mathematica* simulation parameters applicable to all simulation runs. Once called, `Simulate` hands over program control to the Repast controller for the duration of the simulation. The controller calls various *Mathematica* programs in response to user actions on the Repast control panel.

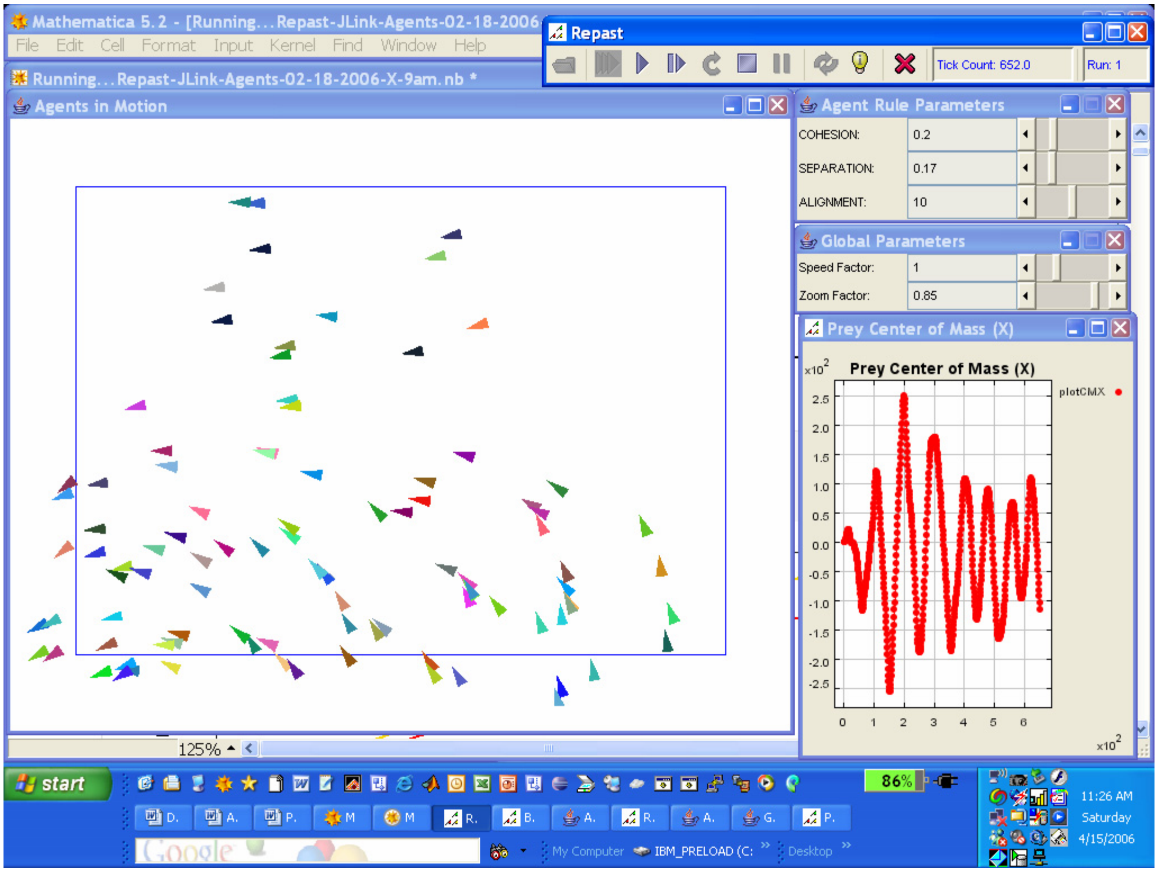

FIGURE 6 User interface for the Repast-*Mathematica* simulation system

- `init` initializes parameters for the *Mathematica* simulation, creates the agents, and initializes the GUI. The `init` program is called when the user presses the `Initialize` button.

- `step` is called by the Repast controller at each time-step. `step` calls *Mathematica* programs that are the heart of the agent simulation model. These programs update agent locations and status. The `step` program is called once when the user presses the `step` button or called repeatedly when the user presses the `Run` button. The calls to `step` are interrupted when the user presses the `Pause`, `Stop`, or `Exit` buttons.

- `stop` resets the simulation and GUI object parameters. `stop` is called when the user presses the `Stop` button.

- `exit` disposes of user-created GUI objects (in *Mathematica* and Java) and returns control of the program to *Mathematica*. `exit` is called when the user presses the `Exit` button.

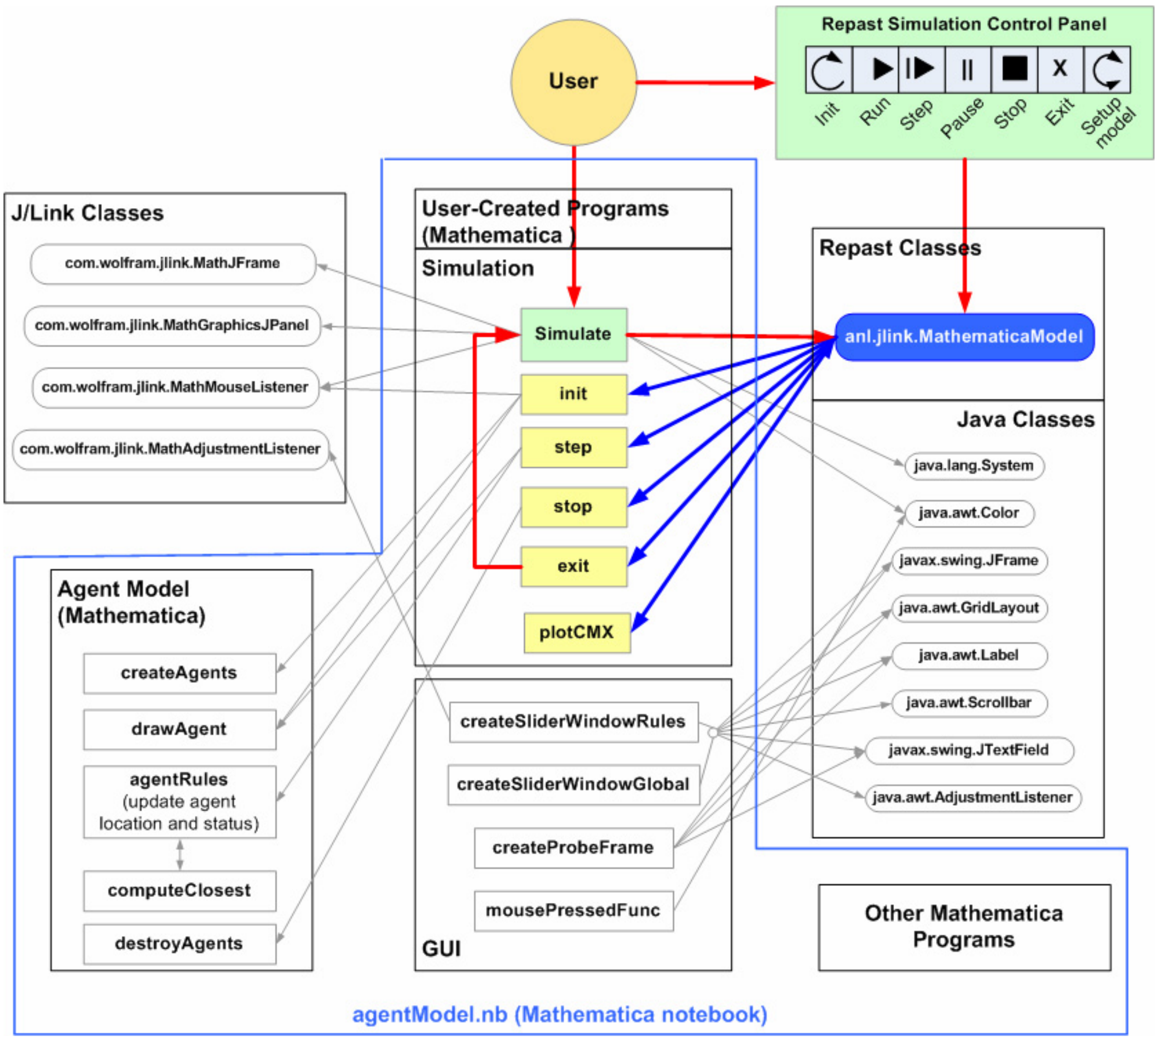

FIGURE 7 Repast-*Mathematica* agent model architecture

- `plotCMX` is called at the end of each time-step to return the data to the controller for plotting on the Repast/Java side. For the Boids example, `plotCMX` returns the x and y coordinates of the center of mass of the flock.

User-written *Mathematica* Interactive GUI Programs

In addition to the user-written *Mathematica* programs involved in the simulation, there are also user-written *Mathematica* GUI programs. This part of the *Mathematica* GUI allows the user to interact directly with the agent model. For the Boids Model, there are four main GUI components that are written in *Mathematica* using *J/Link*:

- Agent Action frame displays the agents and their animation.

- Rule Parameter window allows the user to dynamically vary the agent rule parameters during the simulation.

- Global Parameter window allows the user to dynamically vary global parameters, such as zooming.

- Probe window allows the user to view the dynamically changing attributes of a selected agent as the simulation progresses.

An example GUI is shown in Figure 6. These GUI components are implemented via Java and the *J/Link* GUI classes, which include versions of slider panels and mouse event handlers. It should be noted that a great deal of flexibility exists in terms of which GUI components are implemented as Repast/Java classes versus *J/Link-Mathematica* classes.

Java Classes for Linking Repast and *Mathematica*

- `JLink-Repast.jar` is the main entry point for using Repast from *Mathematica*. It provides a simple class called `MathematicaModel` along with some helper classes that allow the user to expose his *Mathematica* functions to the Repast toolkit. It also allows the user to exploit the simulation scheduling facilities, the user interface facilities, and the charting and data collection facilities present in the Repast toolkit. It uses the *J/Link* library to allow the user to expose his functions and data to Repast.

- `Repast.jar` includes the classes that make up the full Repast library. It stores the `Scheduler`, which is the heart of the Repast library. The scheduler manages the execution of actions that drive the simulation. It also provides tools for data collection and analysis, reusable agent components, and the user interface. These wrapper classes put all of these pieces into one simple “facade” class that can be easily used by *Mathematica*. The wrapper exposes many of the useful tools from the `Repast.jar` without requiring the user to understand how they all relate to one another.

- `Plot.jar` is a set of two-dimensional signal plotter components written in Java to provide real-time animated plots, as used in the Boids model simulation.

- `Trove.jar` is a library that generates Java byte-code that allows Repast to create new classes (create and alter compiled Java code) at runtime on the basis of simple specifications provided by the user.

- `Colt.jar` consists of support libraries for Repast numerical computations and for simulating concurrency.

To complete the linkage of Repast and *Mathematica*, the five jar files described above are placed in a directory, whose location is added to the `JavaClassPath`, which makes them accessible to *Mathematica*. The *J/Link* libraries are automatically accessible, since they are a standard part of the *Mathematica* distribution.

SUMMARY AND CONCLUSIONS

We have demonstrated that the linked Repast-*Mathematica* system is a viable candidate for a rapid prototyping agent modeling environment. This environment is flexible enough to model different agent interaction topologies (grids, networks, free space, etc.). The system allows the user to observe agent behaviors in real time for simulations over extended time horizons, and it allows real-time user interaction to explore agent behaviors and interactions. Results from the simulation are readily available for post-simulation analysis in *Mathematica*. Future activities include exploring scale-up issues associated with the simulation and display of very large numbers of agents and implementing the Repast-*Mathematica* agent modeling system on cluster or grid computing platforms.

ACKNOWLEDGMENT

This work is sponsored by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Axelrod, R., 1997, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton University Press, Princeton, NJ.

Axtell, R., 2000, "Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences," Working Paper 17, Center on Social and Economic Dynamics, Brookings Institution, Washington, DC.

Bankes, S.C., 2002, "Agent-based Modeling: A Revolution?," *Proceedings of the National Academy of Sciences* 99(Suppl. 3):7199–7200.

Bonabeau, E., 2001, "Agent-based Modeling: Methods and Techniques for Simulating Human Systems," *Proceedings of the National Academy of Sciences* 99(Suppl. 3):7280–7287.

Booch, G., J. Rumbaugh, and I. Jacobson, 1998, *The Unified Modeling Language User Guide*, Addison-Wesley, New York, NY.

Cederman, L.-E., 2002, "Endogenizing Geopolitical Boundaries with Agent-based Modeling," *Proceedings of the National Academy of Sciences* 99(Suppl. 3):7796–7303.

Collier, N., et al., 2003, "Onward and Upward: The Transition to Repast 2.0," in *Proceedings of the First Annual North American Association for Computational Social and Organizational Science Conference*, Pittsburgh, PA.

Collier, N., and D.L. Sallach, 2001, *Repast, REcursive Porous Agent Simulation Toolkit;* available at http://repast.sourceforge.net/.

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom Up*, MIT Press, Cambridge, MA.

Friedkin, N.E., and E.C. Johnsen, 1990, "Social Influence and Opinions," *Journal of Mathematical Sociology* 15:193–206.

Gayley, T., 2004a, *J/Link User Guide, Version 3.0;* available from Wolfram Research, Inc., Champaign, IL, Sept.

Gayley, T., 2004b, "Building User Interfaces Using J/Link," *The Mathematica Journal* 9(1).

Gaylord, R.J., and L. D'Andria, 1998, *Simulating Society — A Mathematica Toolkit for Modeling Socioeconomic Behavior*, TELOS/Springer Verlag, New York, NY.

Gilbert, N., and K.G. Troitzsch, 1999, *Simulation for the Social Scientist*, Open University Press, Buckingham, UK.

Kohler, T.A., G.J. Gumerman, and R.G. Reynolds, 2005, "Simulating Ancient Societies," *Scientific American* 293:77–84.

Macal, C., 2004, "Agent-based Modeling and Social Simulation with *Mathematica* and MATLAB," in C.M. Macal, D. Sallach, and M.J. North, eds., *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Macal, C.M., and M.J. North, 2005, "Tutorial on Agent-based Modeling and Simulation," in M.E. Kuhl, N.M. Steiger, F.B. Armstrong, and J.A. Joines, eds., *Proceedings of the 2005 Winter Simulation Conference*, Orlando, FL, Dec.

Maeder, R.E., 2000, *Computer Science with Mathematica*, Cambridge University Press, Cambridge, UK.

MathWorks, 2005, MATLAB home page; available at www.mathworks.com.

North, M.J., and C.M. Macal, 2005, "Escaping the Accidents of History: An Overview of Artificial Life Modeling with Repast," in A. Adamatzky and M. Komosinski, eds., *Artificial Life Models in Software*, Springer-Verlag: Dordrecht, Netherlands.

Repast, 2005, Repast home page; available at http://repast.sourceforge.net/.

Reynolds, C., 2005, *Boids*; available at www.red3d.com/cwr/boids/.

Sallach, D., 2003, "Social Theory and Agent Architectures: Prospective Issues in Rapid-Discovery Social Science," *Social Science Computer Review* 21:179–195.

Sallach, D., and C. Macal, 2001, "The Simulation of Social Agents: An Introduction," *Special Issue of Social Science Computer Review* 19(3):245–248.

Schelling, T.C., 1971, "Dynamic Models of Segregation," *Journal of Mathematical Sociology* 1:143–186.

Tobias, R., and C. Hofmann, 2004, "Evaluation of Free Java-Libraries for Social-Scientific Agent Based Simulation," *Journal of Artificial Societies and Social Simulation* 7(1), Jan. 31.

Wolfram, S., 2003, *The Mathematica Book: 5th Edition*, Wolfram Media, Champaign, IL.

Wolfram Research, Inc., 2005, *Mathematica* home page; available at www.wolfram.com.

AGENTS IN SPACE: BUILDING UPON THE GEOGRAPHIC INFORMATION SCIENCE INITIATIVE

B. PIJANOWSKI and K. ALEXANDRIDIS,*
Purdue University, West Lafayette, IN

ABSTRACT

Simulating how humans behave within a spatial context is challenging. Integrating the spatial context into agent-based modeling approaches adds yet another important dimension to social simulation that allows researchers to provide a more geographic and perhaps a more realistic context to their simulations. However, most social scientists are not trained in the tools, methods, and theory that now form a huge area of research called geographic information science. The purpose of our presentation is to summarize the major spatial relationships that might be considered in social simulations; review some of the available tools (e.g., Repast, MASON, ArcGIS, FRAGSTATS) and spatial data sources that could be used; and highlight important considerations to modeling agents in space. We conclude with a quick summary of how we have used spatial analysis tools to model agents in our spatially explicit Multi Agent-based Behavioral Economic Landscape (MABEL) model.

Editors' Note: The full paper was not received in time for publication. The abstract is included to provide a frame of reference for the discussion that follows this session.

* *Corresponding author address*: Konstantinos Alexandridis, Human-Environment Modeling and Analysis Laboratory (HEMA), Department of Forestry and Natural Resources, Purdue University, 204 Forestry Building, 195 Marsteller Street, West Lafayette, IN 47907; email: ktalexan@purdue.edu.

TOWARD A GRAPHICAL ABM TOOLKIT WITH GIS INTEGRATION

W. RAND,* D. BROWN, R. RIOLO, and D. ROBINSON,
University of Michigan, Ann Arbor, MI

ABSTRACT

Agent-based modeling (ABM) has proved useful in a number of fields. Many of the early successes of ABM were due to its ability to represent the processes of a phenomenon. However, there is less emphasis within ABM on developing its ability to replicate spatial patterns of phenomena. In order to accomplish that, more powerful spatial modeling techniques, like those within geographical information systems (GISs), are necessary. The integration of ABM and a GIS into a cohesive package would allow for elegant modeling of both process and pattern. One problem with an integrated toolkit is that most GIS users are not programmers. However, most GIS users are familiar with the use of detailed graphical user interfaces (GUIs) in order to create complex visualizations of data. Thus, providing a GUI to access an integrated ABM-GIS toolkit would vastly expand the number of users for such a toolkit. This paper is a first step toward describing such a toolkit. It first outlines several design principles for an ABM-GIS toolkit and then describes a survey of extant toolkits (Repast Py, NetLogo, and MobiDyc) that were selected on the basis of design principles. The toolkits were surveyed to see how well they fulfill some of the design principles. This survey was not meant to be a comparative review of these toolkits; rather, it was conducted to determine what useful design principles could be gathered from them that might inform a new "ideal" ABM-GIS toolkit. Finally, the paper concludes with some design recommendations for such a toolkit.

Keywords: Agent-based modeling, toolkit, GUI, GIS, design

INTRODUCTION

Agent-based modeling (ABM) has proved useful in a number of fields, from population biology, ecology, and epidemiology to international relations, economics, and urban planning (Epstein and Axtell 1996; Axelrod 1997). However, as this modeling technique continues to mature, it will often be useful to integrate it with more powerful data-handling methods like a geographical information system (GIS) (Gimblett 2002; Parker 2005). To date, typical agent-based models of spatially embedded systems use very simplistic representations of space, spatial patterns, and spatial processes. Where ABM has excelled is in its ability to represent the processes underlying a particular phenomenon, but it does not have a rich representation of the patterns of phenomena. On the other hand, while GISs are regularly used to build complex and interesting spatial models that clearly represent the patterns of a phenomenon, these models tend either to be static models of pattern or to be statistical (e.g., Markovian) models of process and therefore do not contain algorithmic processes to generate the phenomenon. Thus, easy access to ABM techniques would enhance the range of models that GIS users could employ, by making it

* *Corresponding author address (current affiliation)*: William Rand, Northwestern Institute on Complex Systems, Northwestern University, 600 Foster Street, Evanston, IL 60208-4057; e-mail: wrand@northwestern.edu.

possible to combine individual (bottom-up) models of processes with sophisticated spatial models of pattern.

However, to make it possible to define arbitrarily complex agent behaviors, general-purpose agent-based modeling packages rely, more or less, on universal computer programming languages like Java, NetLogo, Python, Objective-C, and so on. But most GIS users are not programmers by training; instead, they have learned to use the powerful graphical user interfaces (GUIs) now available on most GISs. Thus the motivation for our project is to explore how to make it easier for GIS users to employ ABM techniques in combination with standard GIS tools by using standard GUI interfaces and frameworks. We believe one way to move toward that goal is to design a conceptual architecture for ABM toolkits that specifically facilitates the definition of combined spatial and agent-based process models within a GUI framework. We feel that by doing this, we can greatly expand the range of ABM applications and bring this technology to a new group of users.

Since a number of existing systems have already been designed to make it easier for nonprogrammers to create agent-based models, we began by reviewing these systems and their intended scope. In this paper, we examine three ABM GUI toolkits and evaluate their capabilities on several dimensions related to their functionality, interface, and primary intended audience. We chose these systems on the basis of their explicit use of a GUI, their capability to support spatially explicit ABMs, and their ability to minimize programming requirements. The first is NetLogo (Wilensky 1999), which has an easy-to-use GUI for developing the interface of the ABM. The second is Repast Py (Collier and North 2004), which has a GUI for model development as well as strong GIS integration. The final toolkit we examined was MobiDyc (Ginot et al. 2002), which has one of the most comprehensive GUIs for model development and also has an ecological focus that aligns well with the interests of many GIS users.

We carried out a systematic characterization of the functionality of all three platforms. In this paper, we answer a list of questions devised to categorize and describe the capabilities of each platform. After describing the results of our review of these systems, we discuss what was learned about each toolkit’s contribution to the development of ABM architecture design, and then we distill these lessons into a list of desired capabilities for a GUI-based ABM-GIS toolkit. In short, we found that each toolkit had its own strengths and weaknesses, and we summarize these in order to create a picture of a more ideal toolkit. We conclude this paper with a presentation on the desired capabilities of our “ideal” toolkit and on general lessons gleaned from experience with existing systems.

DESIGN PRINCIPLES FOR AN ABM-GIS-GUI TOOLKIT

Having established that there are reasons why a combined ABM-GIS toolkit with a GUI would be useful (Brown et al. 2005), we felt it would be useful to systematically characterize what we would want in such a toolkit. By creating a list of desired capabilities, we can start to understand how such a toolkit could be put together. Since there are three main elements to the combined toolkit (ABM, GIS, and GUI), we created separate lists of desired characteristics for each of those areas; these characteristics are summarized in Table 1.

TABLE 1 Desirable characteristics

ABM	Event scheduling, heterogeneous entities, environmental processes
GIS	Multiple layers of spatial data, rapid spatial queries, GIS to ABM correspondence
GUI	Simple to use, drag and drop interaction, ability to create complex queries

To begin with, from the world of ABM, we want the full power of scheduling and heterogeneous entities that are normally available in ABM. Thus, we want the ability to schedule an event at any time in the future, and when it occurs, to allow it to trigger other events. This aspect of ABM comes out of work in discrete event simulation (DES) (Cassandras and Lafortune 1999). It is one of the main components in the creation of rich models of process and events. One of the hallmarks of ABM is the ability to create large numbers of heterogeneous agents and combine those agents into arbitrary groups. This gives the modeler the ability to describe agent heterogeneity as (1) variation in properties and methods and (2) the categorization of different agent types (Brown and Robinson in review). Moreover. the modeler can combine these agents and ask all of them to carry out an action simultaneously. Another feature of ABM that has proven useful is the ability to use multiple representations of the environment within the same model. In the case of this toolkit, different GIS maps could be utilized to situate the model in heterogeneous locations and assess model performance under different environmental conditions or landscape patterns. Similarly, ABM also has a powerful representation of environment processes. The environment has the ability to carry out its own processes and interact with the agents in autonomous ways. For instance, the ecological processes of photosynthesis, nutrient cycling, water use efficiency, and succession may interact locally and autonomously with climate conditions within a predator-prey or grazing model (Wilensky and Reisman 2005).

Of course there are also some capabilities from GIS that would be desirable. First of all, the ability to store multiple layers of data in one collection, which is tied together by the physical location of those layers in the world, is a powerful data model tool that would be useful within an ABM. For instance, residents moving around in a residential location model should be able to access information — such as the amount of open space, distance to a central business district, and proximity to schools — for one location in an easy and effective manner. Moreover, the ability to do rapid spatial queries would be useful. For instance, in the residential location model, developers should be able to quickly determine which lots are available within a hundred meters of a main arterial road. Another desired capability would be the transformation of GIS objects into ABM agents. For instance, a store in a GIS database could be reified as an ABM agent that buys and sells products with its neighbors. Of course the ability to export the GIS data about the environment to the ABM is also very important.

Finally, a GUI for constructing integrated ABM and GIS models, such as the GUI that is available within ArcGIS for building spatial data models (i.e., model builder), would be very useful. Model builders should be able to create agents, processes, and data reporters by doing nothing but pointing, clicking, and typing a few names. However, just because the GUI would be simple does not mean that it would necessarily only involve the creation of simple methods. Traditional GIS systems (like ArcGIS) use drop-down menus to construct detailed and rich structured query language (SQL) queries into the GIS database. These query systems are easy to use in part because they are graphical and in part because they require little (if any) formal knowledge of programming.

SURVEY

On the basis of these design guidelines (see Table 1), we undertook a qualitative survey of toolkits that have been built with one or more of these guidelines in mind. Our overall goal was to understand better whether or not extant toolkits had already integrated the aspects of a toolkit that we desired, and, if so, how they accomplished this integration. The specific objectives of this survey were therefore twofold: (1) evaluate the toolkit in terms of how well it accomplished the task we had set before us and (2) examine the basic ideas of the toolkit and see if there was anything useful we could incorporate into our design of an ideal toolkit. To accomplish this task, we created a list of questions about the capabilities of each toolkit and sought to answer those questions by examining the toolkits. However, in order to carry out this survey, we first had to determine which toolkits to examine and then determine what questions we would answer about each of the toolkits. Finally, we had to actually answer the questions and summarize the results.

Selection of Toolkits

There exists a myriad of ABM toolkits: Repast, Swarm, MAML, Ascape, AnyLogic, MASON, CORMAS, NetLogo, and MobiDyc, among others. As a result, narrowing down the toolkits to a reasonable number that we could survey was daunting. However, since a number of existing systems have already been designed to make it easier for nonprogrammers to create agent-based models, we began by reviewing these systems and their intended scope. We developed a list of criteria for determining which toolkits we would examine. The toolkit had to have a strong GUI, powerful ABM tools, strong support for the toolkit, and be provided for free. In addition, we thought it would be very useful if the toolkit already had some GIS integration and ability to model ecological systems (since that is one of the major uses of GIS data).

Given the above criteria, we selected three ABM GUI toolkits and evaluated their capabilities on several dimensions related to their functionality, interface, and primary intended audience. The first toolkit we examined was NetLogo, which was developed by Wilensky as a pedagogical and research tool (Wilensky 1999). NetLogo has an easy-to-use GUI for developing the interface of the ABM. It also has an interesting programming paradigm (everything happens in parallel) and was built with a "low-threshold, high-ceiling" language paradigm (Tisue and Wilensky 2004). The second was Repast Py (Collier and North 2004), which was developed at Argonne National Laboratory in order to make Repast easier to use (Collier et al. 2003). Repast Py has a GUI for model development that utilizes a drag-and-drop interface, and Repast Py also has strong GIS integration. The final toolkit we examined was MobiDyc (Ginot et al. 2002), which was developed at the National Research Center in Avignon, France, and was primarily built for ecological modeling. The basic concept of MobiDyc is that everything is an agent, including tasks and the environment. MobiDyc has one of the most comprehensive GUIs for model development and requires the use of only drop down menus to build a model. It also has an ecological focus that aligns well with the interests of many GIS users.

Design of the Survey

We carried out a systematic characterization of the functionality of all three platforms. In this paper, we answer a list of questions we devised to categorize and describe the capabilities of each platform. These questions are of the form “Can the system...?,” referring to specific capabilities. Besides providing detailed responses to these questions (Appendix 1), we also graded the ability of each system to carry out the particular function by using a simplified scale (Appendix 1 and Table 2). The system receives a ‘G’ (color-coded as green) if it was possible to carry out the entire task using the (G)raphical interface, a ‘P’ (color-coded as yellow) if there were specific (P)rimitives in the toolkit for carrying out the task, a ‘C’ (color-coded as red) if (C)oding was required to carry out the task, and an ‘N’ (color-coded as black) if it was (N)ot possible (without extreme measures) to carry out the task.

In order to clarify our thinking about capabilities that we desired in the integrated toolkit, we distinguished the following six modeling topics crucial to any ABM development used for rigorous scientific purposes and publication: (1) agents, (2) agent groups, (3) environment, (4) experiments, (5) reports, and (6) interoperability. We determined these topics were relevant on the basis of our experience with building and utilizing ABMs in the past. Some of these topics are not specific to use for GIS users. However, we believe that the design of experiments, software interoperability, and model output through reports and graphs are important topics relevant to the design, use, and interpretation of any ABM, and we therefore included them in the overall survey.

Once we had the six major groups established, we developed a list of questions within each group that detailed the functionality we desired in any integrated toolkit. Within each group of questions, we also found it useful to create subcategories that helped to classify the question. Finally within each of these subcategories, we listed the questions in approximate order of difficulty, moving from the least difficult to the most difficult goals to accomplish.

One word of warning: Many of these questions were very difficult to answer in any objective sense. However, we did attempt to create standards within the grading so that even if the answers are not absolute grades in any sense, they are at least a decent relative comparison of the three toolkits. In the end, because of the subjective nature of these results, they may not be as applicable for one particular project as they are for another one.

It is also important to remember that surveys like this one only make sense within the context of the questions being asked. Our questions and answers were designed specifically to inquire about the construction of an integrated ABM-GIS toolkit with a strong GUI. There are many criteria that we could have utilized that we did not. For instance, we did not ask “Are the primitives easy to use? Is the architecture of the toolkit intuitive? Is there a wide base of support for the toolkit?” It may very well be impossible to carry out a truly comprehensive survey of toolkits that would be appropriate for all users; hence, all such surveys are going to be subjective and thus at least partially controversial. There have been several other surveys of toolkits that have had other goals; some of these are more general surveys (Gilbert and Bankes 2002; Tobias and Hoffman 2004; Fedrizzi 2005; Wiedmann and Girardin 2005; Railsback et al. in review).

RESULTS

We present the results of our survey in two different formats. In the more extensive format (Appendix 1), we present all of the questions and the exact answers that we gave to those questions. The answers are presented both in terms of a quick description of an answer and in terms of the grading system described above. In addition, for quicker reference and to provide a higher-level summary of our results, Table 2 presents the letter grades that we gave to each toolkit for each answer (G, P, C, N) and is color-coded to reflect these grades (green, yellow, red, black). In addition, the questions are not listed in full in Table 2, but the categories, subcategories, and keywords referencing the question are listed.

DISCUSSION

The results of our survey were mixed. It seems obvious that none of these packages measure up to our ideal toolkit in terms of ABM-GIS integration with a strong GUI. However, we were able to update our design principles by looking over these results.

For instance, NetLogo has a programming paradigm (enforced parallelism) that causes the programmer to write code for the model in a specific way. MobiDyc also makes use of a particular paradigm (everything is an agent). As we went through the questions in the survey, we realized that this paradigm had a dramatic effect on the answers to some of the questions for these toolkits, but it did not necessarily have a negative effect. In some cases, it probably had a positive effect. In the end, it was clear that the programming paradigm utilized by a toolkit will force trade-offs to be made in the toolkit; thus, choosing a paradigm requires careful thought before designing a new toolkit.

NetLogo probably has one of the best GUIs for designing the look of the ABM, but it has little to no GUI for actually creating the model. This was an interesting result, and we realized that being able to design the look of the ABM enhances the model development experience for novices. Having to specify screen coordinates and sizes within code is very daunting; being able to drag and drop graphs and sliders around the world is much more natural.

Instead of providing the ability to design many (if any) of the model components graphically, NetLogo relies on a long list of primitives that can be used to carry out most of the basic operations that an ABM developer would desire. This emphasis on primitives, as opposed to visual programming, may not specifically address the goals we had in this survey, but it does seem to aid novice programmers in learning how to program. In fact, NetLogo and MobiDyc together (for opposite reasons) caused us to reassess our desire for a strictly graphically based language. It may, in fact, be easier to use a large graphical component with some simple coding than to design a fully functional GUI-only system.

NetLogo has also made recent strides in being able to run experiments from the GUI (i.e., BehaviorSpace) without ever having to control the model from the command line. This is a feature that will likely be appreciated by novice model users who simply want to see what the effect of a particular range of values is on the overall model performance. Part of the “ease of use” of NetLogo is a result of the fact that it has good support and the development team has included new features requested by users on a regular basis. Though support was not an explicit part of our survey, it does have a positive impact on many of the questions that we asked in our survey.

Repast Py has, by far, the best GIS integration of any of the toolkits we examined. It allows the model developer to read GIS data within the drag-and-drop environment and the click of a button. In addition, since it works with both OpenMap and ESRI products, it is usable by a

TABLE 2 Summarized and color-coded results

Agents			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
Creation			
basic	P	G	G
types	P	G	G
Properties			
basic	P	G	G
values	P	P	G
type-based	P	G	G
Methods			
basic	P	C	G
Initialization			
external	G	C	G
GIS	C	G	N
Scheduling			
parallel	P	N	G
agents	N	N	G
schedule	P	G	G
properties	C	C	G
Sensors			
other agents	P	C	G
environment	P	C	G
Effectors			
other agents	P	C	G
environment	P	C	G
Termination			
die	P	N	G
kill	C	N	G

Groups			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
Creation			
groups	P	P	G
het. groups	C	C	G
Scheduling			
schedule	P	C	G

Reports			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
world display	G	G	G
agent stats	G	C	G
envt. Stats	G	C	G
Graphs	G	G	G
output files	G	G	G
GIS	N	C	N

Environment			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
Initialization			
Values	P	C	G
External	P	C	G
GIS	C	G	N
statistical	C	C	G
non-Euclidean	C	G	N
Properties			
Global	P	G	G
Raster	G	G	G
Vector	C	G	N
GIS methods	N	C	N
layers	N	C	N
Methods			
basic	P	G	G
independent	P	G	G
topology	N	N	N
Scheduling			
schedule	P	G	G
independent	P	G	N

Experiments			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
batch	G	C	G
monte carlo	G	C	G
sweep par.	G	G	G

Interoperability			
Question?	**NetLogo**	**RepastPy**	**MobiDyc**
called from	C	C	N
calls to	C	C	N
analysis	G	G	G
experimental	C	C	N

Legend	
N = No	N
C = Code	C
P = Primitive	P
G = Graphical	G

wide variety of GIS practitioners. There is still work that needs to be done in terms of incorporating topological vector data and multiple layers and being able to easily carry out spatial queries, but, in general, Repast Py is a good first step toward GIS integration into an ABM toolkit.

Repast Py also used different GUIs for different types of models. For instance when work was being done with a vector-based model, a different GUI was required from the one used when work was being done with a raster-based model. In fact, these two worlds are so different that it may be impossible to reconcile them within one GUI.

MobiDyc seemed to be the toolkit closest toward achieving our goal of having a truly GUI-driven ABM toolkit. It had selectable menus for everything. However, the interface seemed a little confusing at times, and sometimes it was inefficient to select three or four menu items just to write a simple equation like "z = x + y." In addition, MobiDyc lacks GIS integration and, because it is written in SmallTalk, is not easily extensible.

However, in MobiDyc, it is possible to write very complicated expressions with just a few primitives. The entire MobiDyc "language" can fit on one sheet of paper with brief descriptions and yet has been used to build some fairly complicated and complex ecological models. Therefore, it seems clear that designing a good system of primitives is critical to the development of a good toolkit.

CONCLUSION

In the future, we hope to make use of this survey to design a toolkit that would meet the goal of integrating ABM and GIS while still being usable by a novice model builder. A large component of this design will involve the identification and description of the primitives of the language. A "primitive" is a basic command that is easily identifiable and can be used by a model builder without an explicit knowledge of the internal implementation of that primitive. In particular, one group of primitives that would be useful for us would be those related to the modeling of land use dynamics (i.e., land-use modeling primitives [LUMPs]), which would be tailored to allow GIS users who are interested in land-use and land-cover change to build models of real systems.

The design of primitives is very important to the eventual realization of such a toolkit. If the primitives of the toolkit are chosen carefully, then it is possible for novice users to build complicated models. NetLogo provides a clear example of that, having been used, for example, by elementary school students to build models of traffic simulation. However, the primitives also dictate what is hard and what is easy in a given language. For instance, because of the particular parallel paradigm chosen in NetLogo, it can be difficult to build a true discrete-event simulator within that toolkit.

In order to move forward toward the design of such an integrated toolkit, we plan to refine and reconsider our goals. As mentioned above, maybe it is not necessary to have every aspect of the toolkit be built around visual programming aspects. Of course, one of the major components of this design process will be the development of a set of ideal LUMPs. Making a concise but effective list of primitives will facilitate the development of a prototype of an ideal ABM-GIS.

Ultimately, there does appear to be a trade-off between the ease of use and power of the modeling environment, but based on our analysis of these three toolkits, we believe that we have not yet hit the pareto-optimal front of that trade-off and that it is possible to continue to make improvements in both areas.

REFERENCES

Axelrod, R. 1997. "Advancing the Art of Simulation in the Social Sciences," pp. 21–40 in R. Conte, R. Hegelsmann, and P. Terna, *Simulating Social Phenomena*, Berlin, Germany: Springer-Verlag.

Brown, D.G., et al. 2005. "Spatial Process and Data Models: Toward Integration of Agent-based Models and GIS," *Journal of Geographic Systems, Special Issue on Space-Time Information Systems* **7**(1): 25–47.

Brown, D.G., and D.T. Robinson. In review. "Effects of Heterogeneity in Residential Preferences on an Agent-based Model of Urban Sprawl," *Ecology and Society*.

Cassandras, C.G., and S. Lafortune. 1999. *Introduction to Discrete Event Systems*, Springer.

Collier, N., et al. 2003. "Onward and Upward: The Transition to Repast 2.0," presented at First Annual North American Association for Computational Social and Organizational Science Conference, Pittsburgh, PA.

Collier, N., and M. North. 2004. "Repast for Python Scripting," in C.M. Macal, D. Sallach, and M.J. North (editors), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Epstein, J., and R. Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*, Cambridge, MA: MIT Press.

Fedrizzi, A. 2005. *Evaluation of Tools for Agent-based Simulation*, Technische Universität München, p. 218.

Gilbert, N., and S. Bankes. 2002. "Platforms and Methods for Agent-based Modeling," in *Proceedings of the National Academy of Sciences* **99**(7), Supplement 3, pp. 197–198.

Gimblett, R.H. 2002. *Integrating Geographic Information Systems and Agent-based Modeling Techniques for Simulating Social and Ecological Processes*, New York, NY: Oxford University Press.

Ginot, V., et al. 2002. "A Multi-agents Architecture to Enhance End-user Individual-based Modelling," *Ecological Modelling* (157):23–41.

Parker, D.C. 2005. "Challenges and Prospects for Integration of Geographic Information Systems and Agent-based Models of Land Use," in D.J. Maguire, M.F. Goodchild, and M. Batty, *GIS, Spatial Analysis and Modeling*, ESRI Press.

Railsback, S.F., et al. In review. "Agent-based Simulation Platforms: Review and Development Recommendations," *Simulation;* available at http://www.humboldt.edu/~ecomodel/documents/ABMPlatformReview.pdf.

Tisue, S., and U. Wilensky. 2004. "NetLogo: Design and Implementation of a Multi-agent Modeling Environment," in C.M. Macal, D. Sallach, and M.J. North (editors), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Tobias, R., and C. Hoffman. 2004. "Evaluation of Free Java-libraries for Social-scientific Agent-based Simulation," *Journal of Artificial Societies and Social Simulation* **7**(1).

Wiedmann, N.B., and L. Girardin. 2005. "Technical Note: Evaluating Java Development Kits for Agent-based Modeling," *Journal of Artificial Societies and Social Simulation* **8**(2).

Wilensky, U. 1999. *NetLogo*, Center for Connected Learning and Computer-based Modeling, Northwestern University, Evanston, IL.

Wilensky, U., and K. Reisman. 2005. "Thinking Like a Wolf, a Sheep, or a Firefly: Learning Biology through Constructing and Testing Computational Theories — An Embodied Modeling Approach," *Cognition & Instruction*.

APPENDIX 1: FULL SURVEY RESULTS

Agents			
Question? Can it...	**NetLogo**	**Repast Py**	**MobiDyc**
Creation			
create agents?	yes P	yes G	yes G
create different types of agents?	yes, by creating different breeds P	yes, you create different agent classes graphically and then instantiate them G	yes, called entities, and there can be different “stages” within entities G
Properties			
create agent properties?	yes, using -own predicate P	yes, agents have fields G	yes, agents have attributes G
set agent properties to heterogeneous values?	yes, you can specify all agents properties uniquely P	yes, you can specify all agents properties uniquely P	yes, attributes can be initialized via a “series” G
create different properties for each agent type?	yes, different breeds own different properties P	yes, all agents must have their fields specified G	yes, each entity has different properties, though some are automatic (“age,” “location”) G
Methods			
create agent methods?	yes, but all methods are available to all entities P	yes, agents have individual methods C	yes, agents have tasks; some are built in, and others can be created G
Initialization			
initialize agents from external sources?	yes, using “import-world” and from text files G	yes, you can read in data using standard file I/O C	yes, there is a standard initialization file format and you can write Smalltalk I/O code G
initialize agents from GIS data?	yes, using standard GridASCII and file I/O C	yes, agents can be read directly from shapefiles via the GUI G	no, though you could convert GIS data into the proper MobiDyc format N
Scheduling			
have agents take actions in a distributed, parallel fashion?	yes, this is the standard method of executing actions P	no, agents take actions in asynchronous fashion N	yes, you can switch the scheduler between synchronous and sequential modes G
create events as agents?	no, events are methods that are requested of agents N	no, methods are something that agents and the environment have and are not agents themselves N	yes, all tasks and events are actually considered agents and thus treated in the same way as other agents G
schedule agents to take actions?	yes, though there is no discrete event system, you can ask a turtle to do anything at any time P	yes, there is a full dynamic event system scheduler G	yes, though there is no schedule, you can ask an agent to perform a task conditionally G
schedule agents to take actions on the basis of their properties?	yes, though there is no schedule, you can ask them to take actions on the basis of properties C	yes, you can determine in the code if the agent should actually take the action C	yes, though there is no schedule, you can ask them to take actions on the basis of properties G

Agents			
Question? Can it...	**NetLogo**	**Repast Py**	**MobiDyc**
Sensors			
have agents learn about other agents?	yes, all agents can access anything about all other agents, though it can be difficult to single out agents that do not have particular properties or spatial nearness P	yes, agents can access information about other agents C	yes, all agents have access to all other agents G
have agents learn about their environment?	yes, agents can ask questions of the patches P	yes, agents can access information about the environment C	yes, all agents of any entity type can access all other agents, and since MobiDyc uses an agent to represent the environment, that includes the environment G
Effectors			
have agents which affect other agents?	yes, any agent can ask any other agent to set a particular value P	yes, agents can force other agents to change fields or execute methods given the right permissions C	yes, “modify an attribute” is one of the most common tasks G
have agents which affect the environment?	yes, agents can set attributes of patches and can “stamp” their environment P	yes, agents can change environmental values C	yes, agents can modify attributes of the environment G
Termination			
destroy agents?	yes, “die” is a primitive P	no, Python’s “del” is not supported, though you can create workarounds that “imitate death” usually N	yes, “die” is a built-in task G
have agents destroy each other?	yes, agents can be asked by other agents to die C	no, see above N	yes, “kill” is a built-in primitive G
Creation			
create groups of agents?	yes, breeds are a great way to do this P	yes, there are even primitives to get many basic groups like neighbors P	yes, you have entities, stages, and some basic queries like neighbors G
create groups made of heterogeneous agent types?	yes, but very constrained; you can create a property shared by two different groups and then create an agent class using a filter based on that property C	yes, you can create lists of different types of agents fairly easily C	yes, but these groups are calculated each time you perform a task and are not preserved over time G
Scheduling			
schedule agents to take actions on the basis of a group that is independent of the type?	yes, though there is no general scheduling mechanism, you can ask different breeds to take different actions P	yes, agents in a group can be asked to perform an action, but it cannot be explicitly scheduled C	yes, though there is no general scheduling mechanism, you can ask different groups to take different actions G

Agents			
Question? Can it...	**NetLogo**	**Repast Py**	**MobiDyc**
Initialization			
create environmental values?	yes, patches have an "-own" predicate P	yes, you can create underlying grids like in regular Repast, but there is no way to just set up a grid with values from the GUI C	yes, cells have attributes since they are also agents G
initialize the environment from external sources?	yes, you can read in values from files using standard I/O and then set patch values based on that; "import-pcolor" lets you do this from the menu G	yes, you can read in values from files using standard I/O C	yes, there is a standard initialization file format G
initialize the environment from GIS data?	yes, but there are no specific GIS I/O primitives C	yes, a GIS environment is a specific type that allows you to read in shapefiles to define the environment G	no, there is no standard way to read in GIS data, though you could write a script to turn GIS data into the MobiDyc file format C
initialize the environment from statistical distributions?	yes, most standard distributions can be generated C	yes, most standard distributions can be generated C	yes, but the initialization only happens once and thus is always the same every time you start the model G
create non-Euclidean environments?	yes, standard techniques exist to create network-based topologies C	yes, networked environments are another built-in environment type G	no, it is all grid based N
Properties			
create global properties for the entire model?	yes, the "globals" command defines properties for the whole world P	yes, the environment has fields that can be set for the whole world G	yes, these are considered nonlocated agents G
create properties in the environment on a raster basis?	yes, rasters/grids are the basic environment G	yes, the normal grid data can be a raster, but there is no way to import GIS raster data G	yes, rasters/grids are the basic environment G
create properties in the environment on a vector basis?	yes, by using the network-like agents to demarcate areas; moreover, you can describe a raster with a vector to a very fine level C	yes, either through the use of a network or through GIS data G	no, though agents are smaller than the grid N
create properties based on GIS methods (i.e., buffering, intersections)?	no, though there are some things like neighbors and the like that could be used to generate similar results N	yes, you can use either OpenMap or ArcObjects to manipulate GIS objects C	no, though there are some things like neighbors and the like that could be used to generate similar results N

Agents			
Question? Can it...	**NetLogo**	**Repast Py**	**MobiDyc**
create properties of the environment in multiple layers?	no, but patches can have multiple properties that might be equivalent to multiple layers N	yes, you can create multiple layers in a grid model, but not in a network model or GIS model C	no, but cells can have multiple properties that might be equivalent to multiple layers N
Methods			
have the environment take action?	yes, patches can determine that certain things should be done, and diffuse is a basic command P	yes, the environment has its own actions; in GIS models, the environment is even identified with agents G	yes, the cells can perform tasks just like any other agent G
have the environment act independently of the agents?	yes, "diffuse" is one clear example of this; more important, patches are agents independent of the turtles P	yes, though sometimes the environment is an agent as described above G	yes, and you can even modify whether the grid or the agents act first G
have the environment enforce topological rules?	no, agents are responsible for checking that they are not violating any topological rules N	no, agents are responsible for checking that they are not violating any topological rules N	no, agents are responsible for checking that they are not violating any topological rules N
Scheduling			
schedule the environment to take actions?	yes, though there is no schedule, any patch can be asked to do anything at any time P	yes, the environment can use the same scheduler as the agents G	yes, though there is no schedule, any cell can be asked to do anything at any time G
schedule the environment to take action independently of the agents?	yes, patches/the environment are independent agents P	yes, the environment can use the same scheduler as the agents G	no, though the whole world can be asked to perform actions at a different time than the agents G
generate graphical output of the world?	yes, this is all manipulated from the interface G	yes, you drag and drop a viewer into the model G	yes, you can define your own visualization options G
calculate statistics about agents?	yes, BehaviorSpace makes this easy, G	yes, you can write code to calculate just about any statistic C	yes, you can perform many standard statistical calculations G
calculate statistics about the environment?	yes, BehaviorSpace makes this easy G	yes, you can write code to calculate just about any statistic C	yes, cells are just like agents in this environment G
output statistics to graphical displays?	yes, the graphs themselves are designed graphically and are linked to report values in the code G	yes, you can select from a drop-down menu what variables you want to graph G	yes, they have line graphs and histograms; unfortunately, these are not real time; they can be examined only after the experiment G
output statistics to an output file?	yes, BehaviorSpace makes this easy G	yes, this is part of each graph you create and is specified in the GUI G	yes, you can save the text used to generate any display G

Agents			
Question? Can it...	**NetLogo**	**Repast Py**	**MobiDyc**
output data to a GIS server?	no, but the data can be written to a text file and then imported N	yes, you can write back to ShapeFiles C	no, there is no way to write to a GIS file, though you could use the output text file as a GIS input N
run the model in batch mode (i.e. run the model one or more times without the GUI)?	yes, by turning off the update display and using BehaviorSpace, or it can be called from another Java program, or it can be run from the command line by using the headless version G	yes, you can turn off the GUI output and just have the controller come up C	yes, you define it through the GUI and run it from there, but you can turn off visualization G
run the model automatically with different random number seeds in batch mode?	yes, by using BehaviorSpace G	yes, though you have to create a random number seed input that varies as one of the parameters in multi-run C	yes, part of the standard batch mode is to select the number of times to replicate the experiment G
sweep parameters while running the model multiple times?	yes, by using BehaviorSpace G	yes, by using multi-run, though it has never worked in our installation G	yes, and there are even multiple ways that MobiDyc will sweep the parameters for you G
be called from Java or C?	yes, there is a Java API that allows you to call a NetLogo model C	yes, you can export to Java and then compile it in any way you want C	no, since the code is in SmallTalk, it would be hard to access from anything but SmallTalk N
call Java or C standard programming libraries?	yes, you can use the extension API to call out to other Java libraries and even create new primitives in the language C	yes, it supports all Python and Java objects C	no, it could read other SmallTalk libraries, but that is all N
generate data for use with other analysis tools?	yes, you can output data to text files and then analyze them, or you can use the CSV files generated by BehaviorSpace G	yes, you can output data to text files, and supposedly multi-run will output data to XML files C/G	yes, in fact, they are working on an interface with R G
be run using third-party experimental tools (e.g., execute a model via a shell process from a third-party software platform)?	yes, you can run NetLogo with the command line and pass in an arbitrary parameter list via the BehaviorSpace files C	yes, since you can create a standard Repast model, but this takes work C	no, since there is no way to run it from the command line N

CLUSTERED COMPUTING WITH NETLOGO AND REPAST J: BEYOND CHEWING GUM AND DUCT TAPE

M.T.K. KOEHLER* and B.F. TIVNAN, The MITRE Corporation, McLean, VA
S. UPTON, Referentia Systems, Inc., Tampa, FL

ABSTRACT

We have developed a methodology to run NetLogo in an automated fashion in a cluster computing environment or on a single machine, varying both parameter values and/or random seeds. We utilize the Java application program interface (API) of NetLogo, as well as those of Condor, and three custom software programs: a NetLogo XML parser, XStudy, and OldMcData (OMD). This paper describes how to set up a NetLogo program to run in this environment, use the XML parser and XStudy to create the other necessary files, and then use OMD to perform the runs and collect the output data files. The software currently supports automation of statistical experimental designs (such as a nearly orthogonal Latin hypercube), evolutionary algorithms (utilizing a user-defined, potentially very complex fitness function), and full factorial parameter sweeps. All of the software discussed in the paper will be made available to the research community. The current system represents a fully functional prototype, although additional development work is ongoing to improve the robustness and accessibility of the system. The system described herein will also work with Repast J. However, the system requires writing some additional Java code that is specific to the Repast J model. It is our hope in the future to make the current system work with Repast J more seamlessly. We hope the release of the software and methodology will be seen as an invitation to collaborate to improve the system as a whole and enhance its utility to the vibrant modeling communities of NetLogo and Repast J.

Keywords: NetLogo, Repast, cluster computing

INTRODUCTION

Many agent-based models have some inherent randomness associated with them. This means there will likely be a distribution associated with output from the models. One must be cognizant of where on the distribution of possible outcomes a particular model run may lie. This is difficult to know a priori, however. Therefore, it can be particularly useful to run agent-based models many times, varying not only model parameter values but also random seeds. Done manually, however, this task is tedious in the extreme. This paper will discuss a methodology to run NetLogo many times in an automated fashion in a cluster computing environment that is more flexible than, although not as user friendly as, NetLogo's internal BehaviorSpaces.

The work described in this paper is the outgrowth of work started as a part of Project Albert known as data farming and operational synthesis; see Brandstein (1998) and Horne (2001) for a more complete discussion. One of the many foci of Project Albert was the creation of a data

* *Corresponding author address:* Matthew Koehler, The MITRE Corporation, 7515 Colshire Dr., Mail Stop H305, McLean, VA 22102; e-mail: mkoehler@mitre.org.

farming environment (DFE); this is a cluster computing environment designed to run simple, stochastic models many times (tens of thousands to millions) as a way to explore the dynamics of the model and understand the possible outcomes, especially potential outlier events; see Barry (2004a,b) for a discussion of the importance and utility of understanding outlier events. This was done with an eye toward decision support, so the system was made as easy to use as possible in the hope that it would be utilized by decision makers and subject matter experts. The models that are currently part of the DFE include ISAAC, MANA, Pythagoras, Socrates, PAX, and NetLogo. ISAAC, MANA, Pythagoras, and Socrates are all agent-based combat models. PAX is an agent-based peace-support, peace-keeping model. NetLogo is a general-purpose agent-based modeling environment (Wilensky 1999). The incorporation of NetLogo represents a major step forward for the generalization of the DFE, making it useful not only to military analysts but also to the greater academic and analytic communities.

The DFE is based on XML input and study definition files. This presented some challenges with respect to NetLogo, since it does not use XML for its input files. We overcame this by creating a few standards for the NetLogo program and a parser that breaks the NetLogo program into a series of XML blocks on the basis of each part's functionality. This paper briefly describes the major steps necessary to "data farm" a NetLogo model.

GENERAL FLOW OF THE DFE WITH NETLOGO

The general flow of the system is as follows: (1) create a NetLogo model following a set of conventions; (2) parse the NetLogo file into an input XML file; (3) by using the XStudy tool, pick the sliders, choosers, or switches that will be varied during the runs; (4) use OMD and Condor to kick off the runs and collect the data (this can de done on a single machine or multiple machines). All of the software is written in Java and should work on any machine with a Java Virtual Machine. This system, although it is not perfected, is robust enough to handle the pressure of workshop demands, including thousands of runs done remotely on clusters in different countries. We have successfully run Netlogo in two different cluster computing environments: the Maui High Performance Computing Center and the Singapore Defense Science Organization. The system is capable of handling any sort of experimental design, from full factorial to nearly orthogonal Latin hypercube (NOLH). Furthermore, OMD has post-processing capabilities that can be used with evolutionary programming algorithms and other types of user-defined algorithms to create a more dynamic study.

In the following discussion, we examine the conventions needed to put together a NetLogo program, and we provide general instructions for using the other software used for the multiple runs; however, we assume the reader is familiar with Condor. The software discussed in this paper is, or soon will be, available on SourceForge. Alternatively, the software is available from the authors. Condor is available from its developers at http://www.cs.wisc.edu/condor/. NetLogo is available from its developers at http://ccl.northwestern.edu/netlogo/.

SETTING UP THE NETLOGO MODEL

The current system requires certain features within the NetLogo model. These requirements, which are discussed below, have minimal impact on the structure of the program or the speed of execution and are designed to allow an external Java program to start the model, set parameter values (sliders, choosers, and switches), start and end a run, and collect output data (both end-of-run and time-series data). In general, the wrapper starts NetLogo and loads the model, and then it tells NetLogo to iterate a certain number of times. At the end of the requisite number of iterations, output data are collected, and the NetLogo run is terminated.

Global Variables

The model needs three global variables: `stopped`, `filename`, and `clock`. These are used by the external program to run NetLogo, keep track of output data, and allow the modelers to control the behavior of their model separately from the Java wrapper.

Setup

First, the NetLogo model must have a procedure called `setup` to instantiate the model and to prepare the output files. At a minimum, it needs the following lines of code:

```
to setup
set clock 0
set stopped false
setup-file
end
```

Every time the model is run, it will be in a newly started instantiation of NetLogo; therefore, you are not required to set variables (unless they need to be something other than zero). However, you may want to clear values and set others so that you will know exactly how the model is starting up. If you do clear values, DO NOT use the command `clear-all` or `ca`. If you want to clear values, use commands such as `clear-turtles`, `clear-patches`, `clear-all-plots`, or `clear-output`; then manually set the variables. If you use `clear-all`, you will set the variable `filename` to 0. This will cause problems later on, when the output from all the runs is collected, because all the files will have the same name. The batch version of NetLogo is run by a Java program that will set certain parameters; among them is `filename`. Once NetLogo is started, the Java program will call the `setup` procedure. If `setup` then resets the value of `filename`, Condor and OMD will have trouble keeping track of the output files because all the files will have the same name. A more comprehensive version of the setup procedure that includes resetting of values is shown below:

```
to setup
ct
cp
clear-output
clear-all-plots
;; manually set all variables
```

```
set clock 0
set stopped false
setup-file
end
```

The setup-file procedure is very short and could be called from within the setup procedure. It is recommended to keep them separate for clarity. A sample of this procedure is shown below:

```
to setup-file
 ifelse filename = 0
 [file-open "Your_BackUp_Name_Here.csv"]
 [file-open filename]
end
```

This procedure allows you to run the NetLogo program inside the cluster computing environment or in the standard NetLogo program for testing purposes. This works because it checks to see if the variable `filename` has been set by the Java wrapper program. If it has not been set by the Java wrapper, it will open a default file of your choosing.

Go

All models must also have a `go` procedure. The `go` procedure is a little different than the usual NetLogo program. First of all, the procedure must be called "`go`." Second, the wrapper runs the NetLogo program by asking it to step a certain number of times. Because of this structure, it is important to "protect" your runtime code by nesting it inside an `if` statement that returns true if stopped is false. Sample code for the `go` procedure is given below:

```
to go
set clock clock + 1
if not stopped
  [
  ;;runtime code goes in here

  if 'stop condition is true'
[do-file-print close-files set stopped true]
  ]
end
```

By nesting the runtime code inside the `if` statement, the wrapper can run the model any number of times without any potential damage to the output after the stop condition is met. For example, if you have set up the wrapper to run your model 6,000 times, but you have a stop condition that is triggered at time-step 3,500, the wrapper will continue to tell your model to step another 2,500 times. If you generate output at every time step and do not protect it, then you will end up with another 2,500 lines of output. Since your stop condition could be triggered at different times, it could be very difficult to fix your data after the run. It is also important to segregate any end-of-run printing procedures from the file close procedure. Once the wrapper is done stepping the NetLogo program, it will tell the program to `close-files`. Therefore, you must have a procedure in your program that is called `close-files`. If this procedure includes anything other than file closing code, it may cause a problem, since it will be run any time files are closed. If you close files at any time that the stop condition for your model is true, then any other code will be run every time the wrapper steps your program once the stop condition is met. (This is not an issue if you protect your runtime procedures in the aforementioned `if` statement and make the `close-files` procedure exclusively devoted to closing files.) However, this does require that your model have a stop condition that will be triggered at least one time-step before the wrapper ends the run, because the wrapper will simply stop telling the program to step and then call the `close-files` procedure. Sample code for the `do-file-print` and `close-files` procedures are provided below:

```
to do-file-print
file-print "output goes here"
end

to close-files
file-close-all
end
```

Also, there is no post-processing currently associated with NetLogo runs, so if you want something in the output file, such as input parameters, you must write it there in the program (in something like the `do-file-print` procedure). This file will be a single line if you are only collecting end-of-run data. If, however, you are collecting time-series data, this file may be very large.

The above represents all the requisite code for a NetLogo program to set it up for cluster computing. Now, part of the utility of cluster computing is being able to run a model many times with different parameter values. The system we have developed can run NetLogo programs many times and change parameter values. However, the parameters that will change need to comport with a set of standards. First, they must be sliders, choosers, switches, etc., so they must therefore appear in the “Interface” tab of the NetLogo environment. Second, the parameters must not contain any special characters, like ?, %, $, or *. Third, they may be set to numeric values only — no strings. For example, a chooser with the values high, medium, and low would not be acceptable. The chooser should have values such as 1, 2, and 3, which could then be mapped to high, medium, and low in the procedural part of the NetLogo model. This does not preclude other parameters from taking on any values you wish or from having special characters in their name; these standards apply only to parameter values you wish to change in an automated fashion.

NETLOGO TO XML PARSER

Once the NetLogo model is completed, it is time to prepare the system for multiple runs. To do this, first create a new folder to use for file preparation and to collect the output created by the multiple runs. Place the NetLogo model in this folder, and make sure that it writes output to this folder. (NetLogo defaults to the user folder, which can cause problems later on, when OMD tries to collect output that is not in the right folder.) Next use NetLogo2XML to parse the NetLogo file into an XML format usable by XStudy. The process basically pulls the variables that are declared as sliders, choosers, and switches and places them within XML tags that

XStudy will recognize as variables that can be manipulated from one run to the next. To accomplish this conversion, place the NetLogo2XML.jar and the dom4j-1.5.jar into your working folder, then launch the parser from a command prompt within your working directory with the following commands:

```
java -cp NetLogo2XML.jar;dom4j-1.5.jar

albert.datafarm.netlogo.GenerateXMLScenario

Your-NetLogo-Model-Name.nlogo
```

with spaces in between the separate lines above. The parser then outputs an appropriate XML file based on your model name (e.g., Your-NetLogo-Model-Name.xml) into the same working directory. This file can then be used by XStudy to define your experimental design.

THE XSTUDY TOOL

Once the NetLogo XML file has been created, start XStudy by double clicking the `xpath.bat` file. Once XStudy is running, open the NetLogo XML file. Per Figure 1, in the box on the left, you will see all of the variables associated with the sliders, choosers, and switches of your model. Next, create a parameter group by clicking the “Add Param” button and typing in a name. You must have at least one group, but you do not need more than one group. To add a variable to the group, click the variable in the box on the left, then click the “Add Current Selection” button. Do this for all variables you wish to change during the run. Every variable that will take on different values during the runs will need to be in a different group. For example, if you have sliders associated with sight for three types of agent and you wish to vary them in the same way across all three groups, then place them in a single group — perhaps called vision. If there is another slider in your model for the number of obstacles within the environment and you wish to change it in a way that is dissimilar to the values for agent sight, then it will need to go in a different group — perhaps one called obstacles.

Per Figure 2, once you have created all of the necessary parameter groups, the Gridded tab can be used to create a simple, full factorial or gridded experimental design. In this tab, you will see a list of all the parameter groups you created. Simply enter in a minimum, maximum, and a delta for all the groups to create your study. This is very similar to the BehaviorSpace experimental setup found within NetLogo and the structure of the multi-keyword value definitions in Repast parameter files.

If you click the DOE radio button at the bottom of the Select Parameters tab (see Figure 1), the CSV_DOE tab becomes active. Figure 3 is a screen shot of that tab. This option allows you to create any study design you wish. It simply requires a CSV file of the values you wish to run for each parameter. The structure is simply this: columns are parameter groups created in the Select Parameters tab (order here is important, the left to right order of the columns in the CSV must be the same as the top to bottom order of the parameter groups), and each row is a run. To load the CSV file you created, browse to its location, tell XStudy how many lines to skip (if you have any header information in the file, tell XStudy which line has the labels in it), and then click Parse. Once that has been accomplished, click the cell in the Column

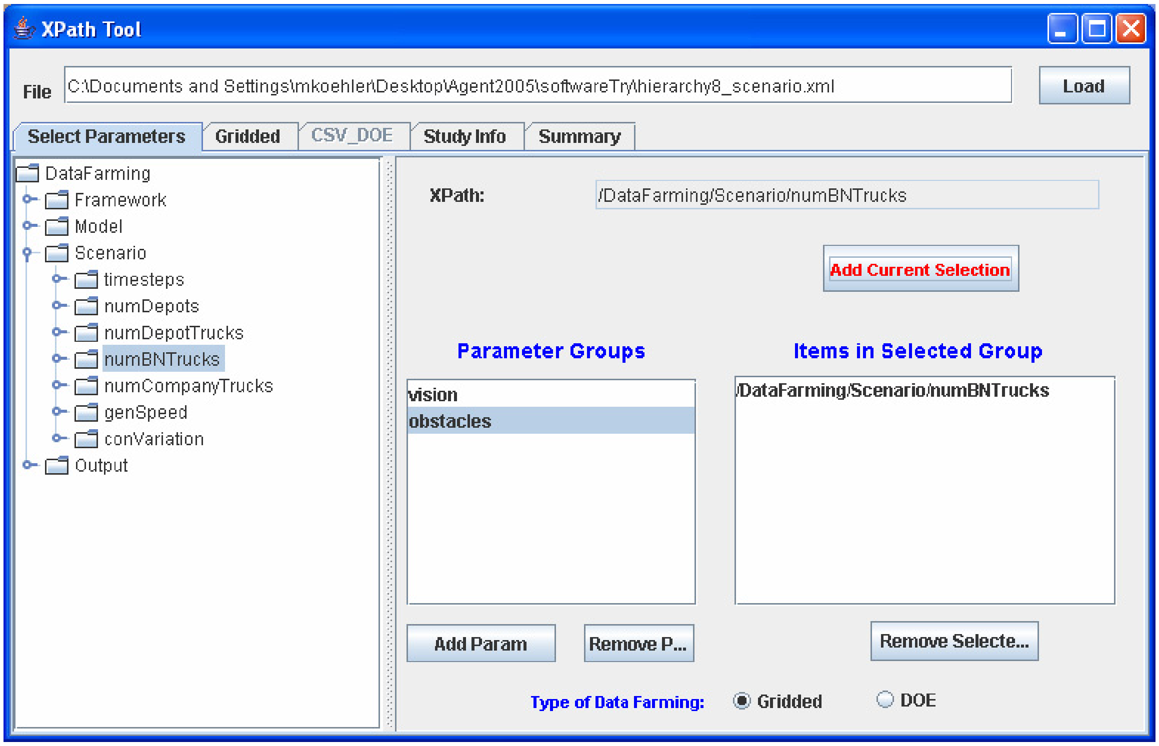

FIGURE 1 Select Parameters tab

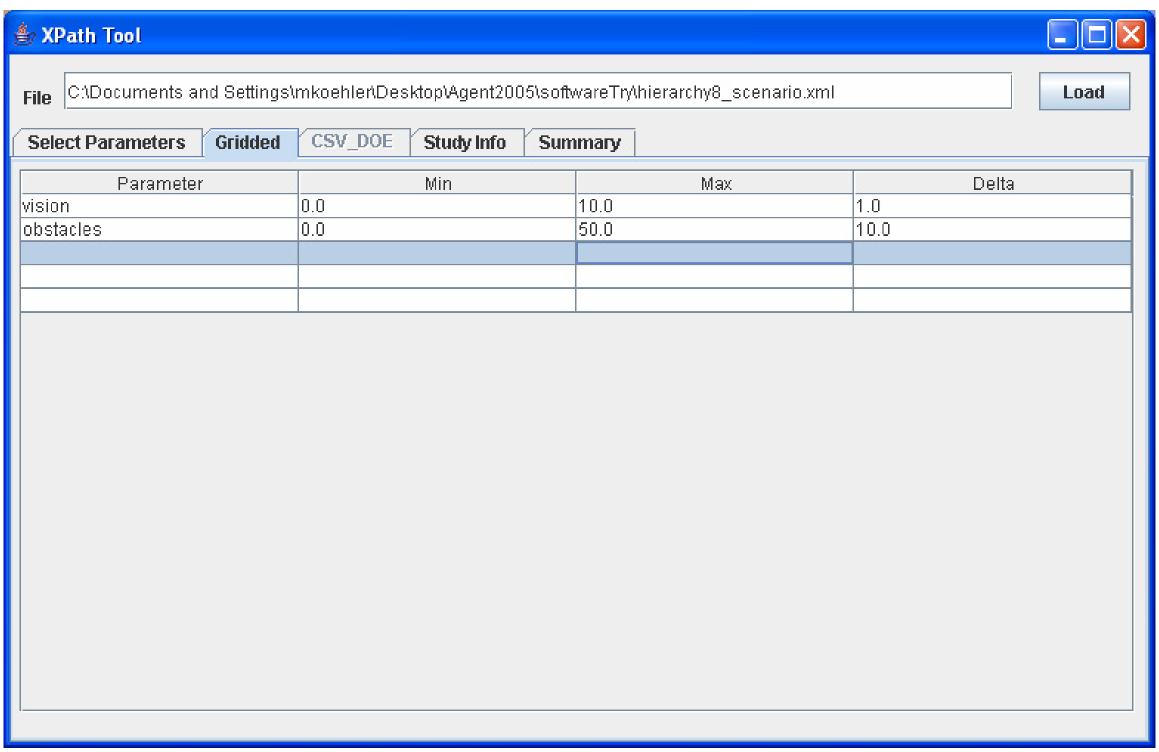

FIGURE 2 Gridded tab

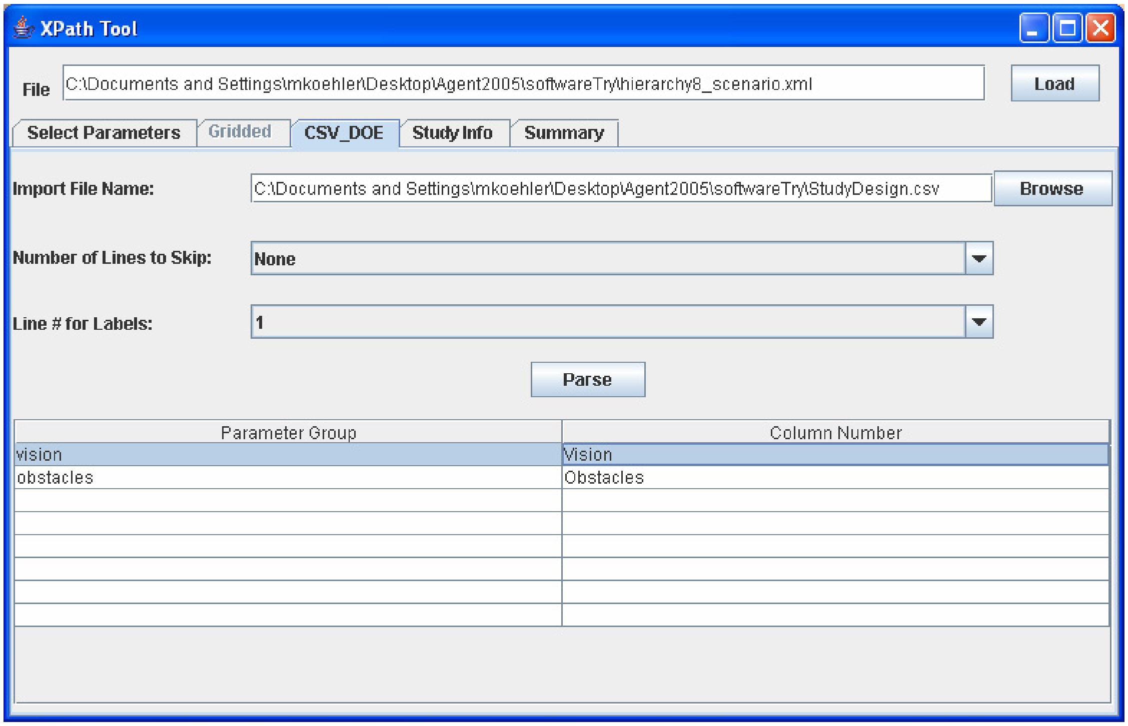

FIGURE 3 CSV_DOE tab

Number column adjacent to a parameter group and choose the appropriate column for the values from the drop-down menu.

It is in the CSV_DOE tab that you would create a more sophisticated study design, such as NOLH. Software to generate a NOLH study design in an Excel spreadsheet, as well as a wealth of other information, is available from the Naval Postgraduate School's Simulations of Experiments and Efficient Designs (SEED) Lab at http://diana.or.nps.navy.mil/~susan/SeedLab/index.html, under "Software Downloads." As the parameter space you wish to search gets larger, these study designs become more and more important. A simple, gridded design will quickly outpace available computing resources, even for fairly small numbers of parameters. NOLH represents an alternative to gridded designs that still affords a statistically valid sample of the parameter space. More information about NOLH designs to explore high-dimensional simulations can be found at Lucas (2002) and at the SEED Lab Web site, under "Papers."

Once the study has been defined in either the Gridded tab or CSV_DOE tab, there are only a few last items to specify before one can begin the runs. These last items are specified in the Study Info tab. An example of this tab can be found in Figure 4. Here you may enter information about who is performing the study, as well as a narrative about the study. All of this information is optional. What are important on this tab are Model Info and Model Run Info. In these areas, you specify the model that OMD should use to run the program and the number of replicates to run for each parameter combination. The number of replicates to run is usually driven by how many processors you have on which to run the program, how long it takes to run your model, and what level of statistical significance you wish to achieve. In the past, we have used anything from 10 replicates for very cursory, fast analyses to 25,000 replicates when we are particularly concerned with very low base-rate phenomena.

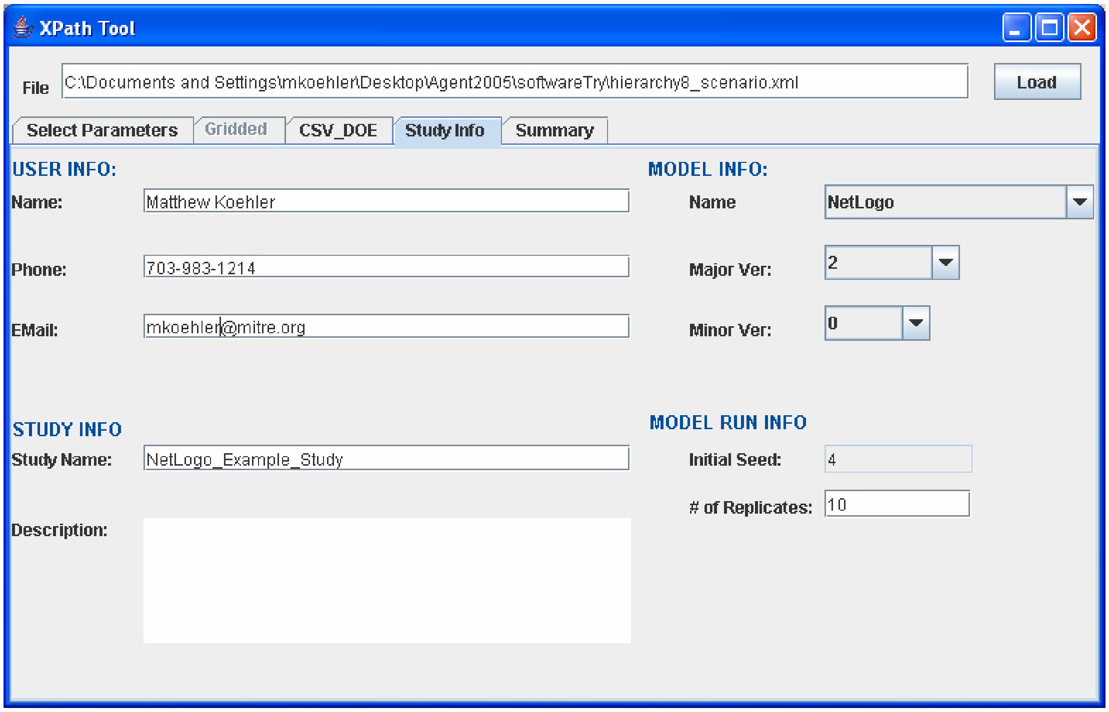

FIGURE 4 Study Info tab

The last step is to move to the Summary tab (Figure 5). Here you are presented with a summary of the information that you have entered thus far. At the bottom of the tab are two buttons. One button (Make Maui Study) is used to create a study.XML file for the Maui High Performance Computing Center; the other (Make OldMcData Study) is used to create a study.XML file for OMD. The only button that is of use in this case is the Make OldMcData Study button. This button will produce the necessary OMD XML file to run the experiment and place it in the same directory as the NetLogo XML file.

USING OLD MCDATA

Once you have created the OMD study.XML file, make sure everything is in the same directory as your NetLogo program. This includes, at a minimum, the .nlogo file and the study.XML file. If you imported a CSV file for the study design, that file must also be in the directory with the .nlogo and study.XML files.

Although this paper does not discuss the installation of OMD, the applications needed for OMD to function properly can be found in Table 1. Installation documentation for OMD can be found with the software, which is available from the authors and should soon be available on SourceForge.

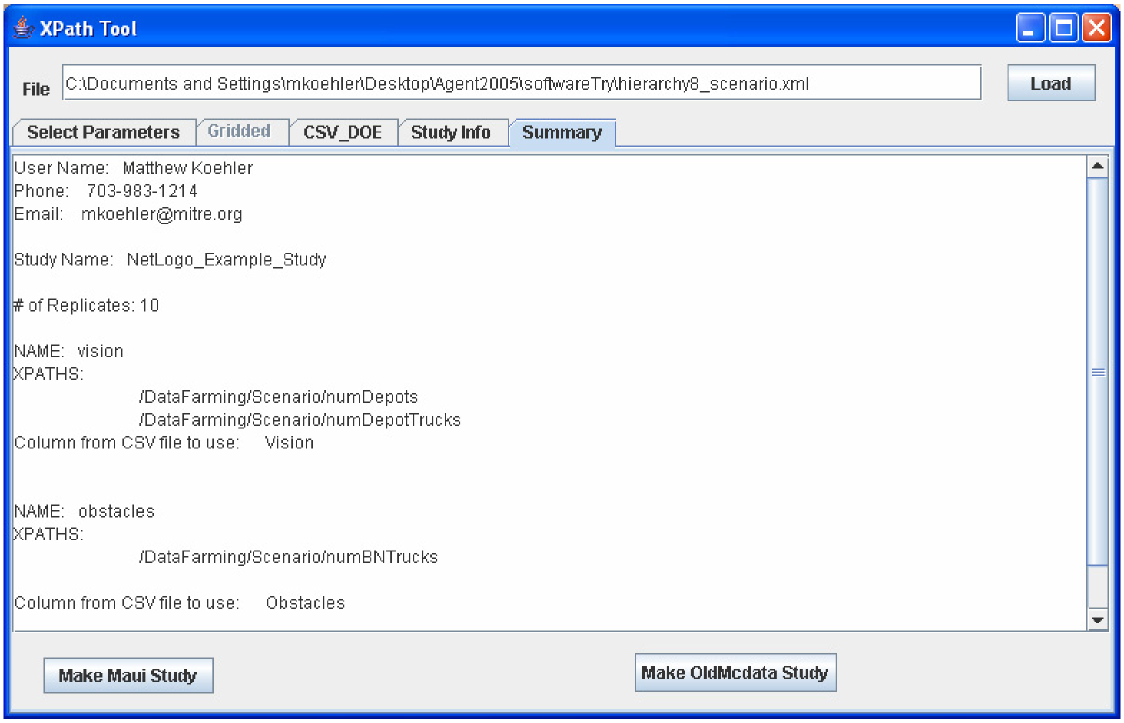

FIGURE 5 Summary tab

TABLE 1 Java applications required to run OldMcData

Application	Purpose	Jar files
Jade 2.5	Agent-based development environment	Base64.jar, iiop.jar, jadeTools.jar, jade.jar
Colt	Scientific code – use random number generators	colt1.0.2.jar
PES	Redirecting output	PES.jar
Xalan	XSLT processor for transforming XML – used by dom4j	xalan.jar
Xerces	XML parsing	Xerces.jar
Jakarta-oro-2.0.7	Regular expression pattern matching and processing	jakarta-oro-2.0.7.jar
NALEX	Natural algorithms, such as simulated annealing, genetic algorithms, and evolutionary programming	nalex1.0-20031119b306.jar
dom4j	XML parsing	dom4j-full.jar

Once OMD is installed on your computer, you will have the directory `oldmcdata` in your root directory. Place the NetLogo folder you created previously in the `test` directory found within the `oldmcdata` directory. In the following example, the NetLogo directory is called `NetLogoTest`. Once you have accomplished this, you should have the following directory tree: `c:/oldmcdata/test/NetLogoTest`. To run the experiment, open a command prompt window, navigate to the `oldmcdata` directory, and type in the following line:

```
oldmcdata.start c:/oldmcdata/test/NetLogoTest study.xml
```

The above is mostly for convenience and bookkeeping. You may put your NetLogo study directory anywhere you wish, thus making the command line statement:

```
oldmcdata.start <path to your study directory> <your study filename>
```

Unless you have manually changed the name of your study file that XStudy created, that last argument in the statement will remain `study.xml`.

Upon completion of the run, there will be three new folders in the `NetLogoTest` folder: `Excursions`, `Output`, and `Playback`. The `Excursions` folder contains the XML files used as input for each NetLogo run. The `Output` folder contains the output data that NetLogo created. The `Playback` folder should be empty and is created to hold output produced by some of the Project Albert models. At this point, you can use whatever tool you choose for analyzing the output data.

CONCLUSION

Project Albert has created a number of tools for running simple models for a large numbers of times in a cluster computing environment. Initially, these tools were specifically for Project Albert agent-based combat models. Recently, however, we have spent time trying to generalize this capability for use by the greater analytic and academic community. This is still very much a work in progress, and we would welcome collaboration with others as we continue to develop this capability and make it more accessible. This paper necessarily glossed over some details with regard to the process described; we encourage interested readers to contact the authors for more information and the latest versions of the software discussed herein. By November 15, 2005, all of the software discussed in this paper should be available on SourceForge. The Web site is http://sourceforge.net/projects/datafarm. The software is currently in an alpha form but will continue to be updated on the SourceForge site.

Although space limitations do not permit its detailed description, the reader should note that a similar XML framework has been developed for Repast J utilizing a custom controller that passes in a random seed and an XML input file. However, this system currently requires more customized code than the NetLogo version. As with the NetLogo case, the authors also wish to share this framework with interested members of the Repast community and will gladly share software and lessons learned with any interested party.

REFERENCES

Barry, P., and M. Koehler, 2004a, “Simulation in Context: Using Data Farming for Decision Support,” in *Proceedings of the 2004 Winter Simulation Conference*, eds. R. Ingalls, M. Rossetti, J. Smith, and B. Peters, Institute of Electrical and Electronics Engineers, Piscataway, NJ; available at http://www.informs-sim.org/wsc04papers/101.pdf.

Barry, P., M. Koehler, and A. Forsyth, 2004b, “The Bigger Hammer Approach: Using Massively Parallel Computation to Address Low Base Rate Problems,” in *Proceedings of the 2004 North American Computational Social and Organizational Science Conference;* available at http://www.casos.cs.cmu.edu/events/conferences/2004/2004_proceedings/Barry_Philip.pdf.

Brandstein, A., and G. Horne, 1998, “Data Farming: A Meta-Technique for Research in the 21st Century,” in *Maneuver Warfare Science 1998,* eds. F. Hoffman and G. Horne, USMC Combat Development Command, Quantico, VA.

Horne, G., 2001, “Beyond Point Estimates: Operational Synthesis and Data Farming,” in *Maneuver Warfare Science 2001,* eds. G. Horne and M. Leonardi, USMC Combat Development Command, Quantico, VA.

Lucas, T., S. Sanchez, Maj. L. Brown, and Maj. W. Vinyard, 2002, “Better Designs for High Dimensional Explorations of Distillations,” in *Maneuver Warfare Science 2002,* eds. G. Horne and S. Johnson, USMC Project Albert, Quantico, VA.

Wilensky, U., 1999, *NetLogo. Center for Connected Learning and Computer-Based Modeling*, Northwestern University, Evanston, IL; available at: http://ccl.northwestern.edu/netlogo/.

DISCUSSION

Methods, Toolkits, and Techniques

(Integrating Agent Modeling Toolkits and Critical Applications, Thursday, October 13, 2005, 8:00–10:00 a.m.)

Chair and Discussant: *Nick Collier, Argonne National Laboratory and PantaRei Corp.*

Charles Macal: Welcome to Agent 2005. My name is Charles Macal. You should all have a program book that lists the sessions along with the times. Today, we have an invited distinguished speaker, Steve Bankes, who will be speaking after lunch. Tomorrow, we have Joshua Epstein, who many of you in the agent-based modeling realm already know. He will speak at 8:45 tomorrow, so please try to be here by 8:30. On Saturday, we have another invited speaker, Lars-Erik Cederman, who will talk between 8:30 and 9:30. Those are the high points.

Today is devoted to modeling methods and toolkits — a methodological day that includes simulation and agent modeling. Validation is an important component of that. Tomorrow will be devoted to agent-based applications, so the theme will be oriented toward applications — people actually building models and, at least conceptually, coming up with results from models. Saturday will be devoted to computational social theory. David Sallach will run the computational social theory day. Mike North is actually in charge today, and I'm in charge of the application program tomorrow. So that's the framework.

To the best of my knowledge, all the papers and all the sessions are as they're printed in the program book. No one has dropped out or been added, and we have not rearranged the schedule. With that in mind, I'll turn things over to Nick Collier for the first session.

Nick Collier: I want to welcome everybody to our first session. Each person will have 25 minutes to talk about each paper and 5 minutes for questions. Please remember to speak into the microphone for questions, so we will have a record for the transcripts.

Linking Repast to Computational Mathematics Systems: *Mathematica* and MATLAB

Collier: Our first speaker is Charles Macal, who will present "Linking Repast to Computational Mathematics Systems: *Mathematica* and MATLAB."

Macal: This talk is about linking Repast to computational mathematics systems, such as *Mathematica* and MATLAB. My co-author on this paper, in terms of the work involved, is Tom Howe. He's one of the Repast developers. The goal is to link Repast and *Mathematica* in a seamless interactive agent simulation environment. In effect, I will describe a new configuration for building agent-based simulation models in which Repast simulation classes are extended to *Mathematica*.

[Presentation]

Collier: Thank you, Chick, for that interesting talk. I have a few questions. First, can you say a few words about this approach for doing simulations in *Mathematica* versus the approach taken by Richard Gaylord in his book? Also, if you could say something about Wolfram's approach, all this CA stuff, in his book as well.

Second, what makes this kind of interaction difficult? Are there any drawbacks? What could be done to improve getting Repast or any kind of Java thing to integrate better with *Mathematica*?

Macal: As for Gaylord's approach compared with this approach, let me say that Richard Gaylord is a professor at the University of Illinois who did a lot of work in social simulation about six or seven years ago. He's still working on things, and he's written books, particularly on *Mathematica* and social simulation in *Mathematica*. I learned how to do social simulation in *Mathematica* from his books and from reading Richard's work as well.

This implementation is a lot different. I would call it beyond, certainly much different than Gaylord's implementation. It is much more akin to object-oriented modeling with respect to modeling abstract data types as agents and with the notion that one could easily create a UML diagram from the way that I've modeled this in *Mathematica*. This type of diagram could be handed off to somebody to implement just as easily in Repast or any object-oriented programming language.

The advantage of my approach, without going into the details of how it differs from Gaylord's, is that there are not embedded assumptions in it. The approach I use, in my opinion, is much more scalable than Gaylord's approach. You can have a million agents, as opposed to 30 like I did. Of course, you would pay a performance penalty.

As for the second question, about Wolfram and CA …. *Mathematica per se* is a general-purpose programming language and includes a programming and publishing environment. It does everything in an interpretive nature, so there is no constraint relative to modeling things as CAs *per se*. This is totally independent of what Wolfram has done using *Mathematica*, so I would speculate that Wolfram would say, "Oh, you could just build all these agent models as CAs if you interpret them correctly." I'd rather not put that much work into the interpretation step.

The final question was about making it easier. The real bottleneck in terms of how this system operates is the link to the processes, the Java processes. Essentially what happens is that JLink calls are made during the *Mathematica* simulation. So each time I reference an object over on the Java side, that's another call, and you pay a penalty for that. So there are some aspects where that can be potentially improved that are very promising relative to the graphics objects, for example. Rather than sending over each agent individually to draw on the graphics panel, one could simply send over the list of all agents at each time step and parse those out on the Java side. Now, that would require either knowing Java or developing perhaps classes to make that a little easier on the Java side. So basically that's the answer.

Collier: I don't know if there was a question from Joanna Bryson.

Joanna Bryson: Lisp is my favorite language, so I really love the way you're remapping your agents. But I build more cognitive agents and I'm worried. Isn't it going to be a hit if you start having more than trivial amounts of data associated with each agent? Do you have to do that kind of remapping or can you keep the objects you already have once you start building up a bigger data structure?

Macal: I guess the answer to that question has to do with the fact that it's the subject of experimentation and speculation. At this stage, obviously being able to observe this system in real time takes a lot of overhead, and we couldn't do the same thing with a 1,000 agents in real time that I was able to do with 30.

I would probably say that for larger-scale simulations for which you'd want to spend the extra effort, that it would be natural to put that over perhaps the whole thing on the Repast side. So there's some transition point where I think that would make sense.

Zhian Li: Zhian Li from Argonne National Lab. Chick, I have a brief comment. I think this is a very innovative way of linking the common tools to Repast. This is very interesting. Would your simulation still be in an interpretive environment once you hand off the simulation to the Repast environment?

Macal: Well, no. Just like Repast, you'd have to compile things and so forth. If you're creating new Java code or new classes, they would all have to be compiled at that point. The issue really is about how far you can go scalewise, or even if it makes sense technically to go that far. Now, if you don't do the graphics part of the real-time animation, then of course that overhead is greatly reduced, although it's all the interactivity and the graphics and things from the Repast side that are really why this is useful to begin with. If you don't do a real-time animation observation of the system, it's probably of less value. But for a lot of the agent models that I've been looking at, and certainly for instructional purposes and helping to observe experimentation, I personally found this to be a useful thing to do. I'm not saying that everyone would find it to be as useful, but I think it's worth adding to the mix of tools and methods and approaches that are out there.

Li: Yes, but I think the way you extend this would be useful to a lot of people who really do not have a great Repast or Java background.

Macal: Yes, in that case, you have to learn *Mathematica*.

Li: A lot of people already know *Mathematica*.

Macal: Well, right. If you don't know *Mathematica* and you don't know Repast, I would recommend learning Repast if you want to do agent modeling, unless you have a particular reason to learn *Mathematica*.

Thomas Howe: The one little piece that I would add to that is to ask yourself if it's useful to you if you don't use the graphics capabilities. I would say that there's actually still a huge benefit in that we exposed the full capabilities of the Repast scheduler. So if you want to do more complex scheduling where you have triggers that are being dynamically flipped so that actions will happen down the road, and they happen in proper order instead of having to develop that whole framework in *Mathematica*, you can just use the stuff that already comes with Repast. So even if you're not using the visualization of the data collection stuff that comes with Repast, I think that having access to the Repast scheduler saves a lot of time in terms of coding in the *Mathematica* side.

Macal: I have a final comment that goes along with what Tom said. If you don't do graphics, you still may want to have a real-time system that is somehow taking in real-time sensor data or information from the outside world. You have a simulation model as well that is planning ahead for some time in the future. Both of those things are going on at the same time. Some of those triggers and things that Tom mentioned could be very useful to access in an integrated planning real-time system for whatever application that might suggest.

Collier: Thank you very much, Chick.

Agents in Space: Building upon the Geographic Information Science Initiative

Collier: Next we have Kostas Alexandridis with "Agents in Space: Building upon the Geographic Information Science Initiative."

[Presentation]

Kostas Alexandridis: Now I am glad to answer any questions.

Collier: Thank you, Kostas. I have a comment and then a question related to that comment. The comment, I believe, somewhat anticipates Bill Rand's paper, which is coming up.

First, the comment is somewhat of an endorsement about the importance of spatial stuff, and you can see that everyone seems to be pushing in this direction. You see this in NetLogo and Repast. We're all trying to push GIS stuff into our toolkits to get them spatially enabled in the right way. At the ESRI user's conference this year, I noticed that they also are moving in that direction. They're getting into more time series, and they can do much more dynamic things. In their modeling language, they're adding loops and random numbers to try and build in some dynamics over space, some process in the pattern you might say. So yes, this is important.

With that in mind, I would like to ask you to talk about what you'd like to see in a tool that would do some of this. These are interesting, important requirements or outlines as to why we want to move in this direction, but in terms of a tool, what do you see? Is it like a GIS-based tool … with agents moving on GIS, or is it something different than that?

Alexandridis: We did a lot of work with MABLE with that, and we started discovering that there are boundaries to what you can do in a toolkit. Lately, and I think my talk tomorrow will speak to some efforts to integrate tools developed to couple different …. I think what we need to incorporate in the models is the functionality of GIS software. Right now, ESRI has a number of DLLs, and it's 20 years of research and development unless we spend the first 10 years developing our toolkits, and so far as I know, toolkits do not talk to each other. So somehow we can take that functionality and bring it into models.

Another direction that I personally find very useful is getting out of grid-based, cell-based modeling and starting to look at a vector or continuous data computations and pattern recognition in space and things like that.

The third area is, and I think this is my bias, but I'm always looking for real-world rich applications. I can understand the functionality of the experimentation process, but unless we are able to move into the real-world processes, we cannot convince people outside the discipline or planners or decision makers that we can improve their decisions and decision making.

Collier: Are there any other questions?

László Gulyás: I would like to ask whether you are aware of the "OBUL" system. It was developed by Itzhak Benenson and his group Tel Aviv University. The name stands for, well, it's objective-based something for urban simulation, and it's built on top of a GIS and the database. It's especially very powerful in finding your neighbors in the geographic system. It's also a mix of a GIS and an agent-based simulation system. It's fairly new. I'm not even sure how much it's advertised or published, but the first version that I looked at was released some time this summer. I would suggest looking at it because it's pretty much in line with what you are talking about. [Editors' note: Gulyás is referring to the Object-based Environment for Urban Simulations (OBEUS) software developed at the Environment Simulation Laboratory of the University of Tel Aviv; see *Geographic Automata Systems: A New Paradigm for Integrating GIS and Geographic Simulation*, by Itzhak Benenson and Paul M. Torrens.]

Alexandridis: Well, we did, and there's a forthcoming paper we have from MABLE that looks at a sample of about 6,000 simulations of two counties in Michigan with Monte Carlo techniques. We start looking at different pattern recognitions, splitting rules for agents, and start revealing a huge level of complexity. I would like to talk to you later and learn more about this.

Collier: Any more questions? All right then. Thank you, Kostas.

Toward a Graphical ABM Toolkit with GIS Integration

Collier: Next we have William Rand with "Toward a Graphical ABM Toolkit with GIS Integration."

William Rand: That last talk is a great lead-in to my paper, so I'm going to skip the slides and go straight to the end. I talk a lot about the desire and the need for spatial integration. Actually, though, I'm going to talk primarily about how to proceed to the point of doing the integration.

First, I'd like to give some background on this project. This work was performed while I was at the University of Michigan. Since that time, I've moved to Northwestern University. I also presented a workshop on NetLogo, which is what I'm now doing. You can still reach that group [at Michigan] if you have any questions about this paper.

I'd like to thank my collaborators: Dan Brown, who's in the School of Natural Resources and the Environment in Michigan; Rick Riolo, who is in the Center for the Study of Complex Systems; and Derek Robinson, who is a GIS graduate student there.

[Presentation]

Rand: Do you have any questions? I'll answer them now, but there's also a website for all the work for this research project.

Venkatesh Mysore: My question relates to something you said at the beginning of your talk, or at least, my question *begins* with that. You had mentioned agent analyst and that was something we had developed based on Repast Py. The point is that it integrates right into ArcMap. So you have ArcMap up and you click on something, and you see something that looks like Repast Py and it uses the ArcMap — visualization application — to show what's going on. The power, not of the agent analyst, but what we're trying to leverage there, as Kostas [Alexandridis] said, is 20 years of DLL [dynamic link library] development — or I guess it started on Unix, particularly for spatial analysts. If you've managed to sort out your licenses and forked over your zillions of dollars, you can use all that kind of stuff. The question then is, "How do you see approaching things from that direction, trying to leverage everything to spatial analysts, for example, the modeling stuff? That's GIS for you right there. It doesn't have to be ArcMap; it could be GRASS [GRAphics Symbiosis System], too."

Rand: Let me clarify the question. Are you asking about building models within the GIS side or something else?

Mysore: Well, no. How would you make an ideal ABM GIS? It seems that much of the work is trying to get the real GIS stuff, which means things like spatial analysts.

Rand: You're completely correct. I guess that wasn't as highlighted here as it should have been, but that is exactly what we're trying to do. We want to be able to, for instance, create a road agent in a model and learn how many other agents are on 50-meter sides of each side of this road agent. That's one of the powerful things. You can do a quick filtering — all these spatial transactions — use spatial analysts.

That, in essence, is what we would like to do. I would guess that we have much more work that we talk about, and we have another paper on this subject that was published in a GIS journal about where that line should be. Where do you draw the line between the ABM toolkit and the GIS toolkit in the end? That's a difficult question, and this is all premised on the idea that you have this agent-based modeling toolkit, and you have some sort of middleware connection that exists between the ABM setup and the GIS setup to allow them to communicate efficiently. You could, obviously, do it the other way. You could move one inside the other, or you could move the other around.

What we would like and what we really want is the ability to have a control bar — one of these standard agent-based modeling interfaces — where you have your GIS all open at the same time, which agent analysts take a first step toward doing, and be able to run your model and see it happening on the GIS, but also collect all the data, do all the statistics, do all the cluster analysts, do all that stuff within the GIS. Right now, half the time when we do the Swiss project, we wind up exporting the data into a great ASCII format and then analyzing it using GIS tools. It would be nice to do that automatically, so that is what we're trying to get at.

Luis Fernandez: Luis Fernandez with the EPA, and in the interest of full disclosure, I worked with Bill a few years back, though I wasn't part of the group that actually developed this tool. My question is more conceptual. How do you see this "ideal" toolkit fit into the new paradigm that Kostas outlined in the previous presentation, and how can that toolkit expand or adapt that new paradigm?

Rand: I think in many ways that this toolkit embraces that paradigm, to take the idea of spatial knowledge and special information and bring it directly into the agent-based modeling paradigm. You have these toolkits that have evolved simultaneously over time. It would be nice to merge them to enfold that whole paradigm together.

I didn't spend a lot of time talking about why this would be a good thing, partially because Kostas' talk did such a good job of explaining that. I think that agent-based modeling can gain from having accurate representations within geographic information systems, within spatial worlds, but I think that spatial worlds do a great job, or GIS tools do a great job of modeling pattern, but they don't necessarily do such a good of job modeling the process. That's something that agent-based modeling toolkits are very good at doing. So you have pattern on the GIS side, you have process on the ABM side, and together I think they make a nice match that could really be quite powerful.

Gulyás: I would like you to comment as to why you excluded Swarm from your survey and briefly give your feeling on how it would perform for this kind of comparison.

Rand: Swarm, I think, is quite powerful. I originally started coding in Swarm long ago. The thing about Swarm — all three of these toolkits at least contain some aspect of graphical model development — is that it is primarily driven through code model development, and so we excluded it because it has no easy graphical model development aspect. We wanted to have something in there along those lines.

Some great work, I think by Paul Box, is being done integrating Swarm with GRASS. So there is some good work on doing a tight integration. But without the graphical element, we want this ideal toolkit to be aimed toward GIS users, and so we think it would take a lot more work to take Swarm out of that. Again, we're also not necessarily trying to evaluate all the toolkits. We're just trying to learn lessons to help develop this ideal toolkit. And so Swarm might very well do quite well in this. I haven't thought through it too much, except for the comments I just made, but our goal wasn't necessarily to encompass all those, but rather to take a set of them, to have employed a lot of these attributes and see what we can learn from them.

Collier: If there are any other questions, please talk to Bill offline.

Clustered Computing with NetLogo and Repast J: Beyond Chewing Gum and Duct Tape

Collier: Next we have Matthew Koehler with "Clustered Computing with NetLogo and Repast J: Beyond Chewing Gum and Duct Tape."

Matthew Koehler: You'll notice there's a slight change in the title. I leave it to the listener to figure out if it is in fact better than getting poked in the eye. As you'll see, that was primarily based on coming up with a nice, easy-to-use graphical user interface. We took a slightly different tack with this one and came up with as many different command lines as we possibly could and a dizzying variety of software.

[Presentation]

Koehler: I'm happy to take any questions.

Steven Bankes: Could you talk a little about OldMcData? I know you said something, but if you could expand a little, it would be helpful. I've heard of it, but could you touch on the basics of it?

Koehler: Okay. Basically, OldMcData takes the study XML file, which spells out the way you've designed your experiment, and looks and says, "What model do I need?" It then makes sure it has the model. It asks you for the input file. It grabs the input file and asks for the parameter values for that run. Then it turns it over to, in the case of NetLogo, a Java wrapper. It sets the parameters in the NetLogo model and runs the NetLogo model. After a certain number of time steps, it kills the NetLogo model, grabs all the output files and sticks it into a folder, kills off NetLogo, grabs the next study XML file, figures out what the next one will be, and runs it again. Is that what you're looking for?

Bankes: Yes.

Koehler: Outstanding.

Collier: Is there another question?

Bankes: This is a leading question, so I don't know if there's time or not. I have to say that I've been very puzzled by the Project Albert stuff every time I heard it briefed.

Koehler: Understandably.

Bankes: I'm tempted to jump to the conclusion that a lot of the arcane stuff we see, which always seems somewhat impenetrable and backward-looking, is a consequence of Project Albert being a visionary effort that started quite a few years ago. You have constraints in the environment that you have to work around. Not to dwell on that, but to look forward, many of us have been inspired to run lots of iterates across lots of machines. In your opinion, to what extent is some of this infrastructure mineable? If I don't want to run cases in Maui, but I want to run cases on a bunch of machines that I acquired, and I'm looking to acquire software that is going to stand the test of time rather than redevelop it, would you care to opine about which of the things that have been developed under Project Albert should be imported into a new environment?

Koehler: I will ask Steve to comment on that when I'm done because he's been more intimately involved with the actual software creation. In my opinion, OldMcData is really good because it's very modular, very general purpose. I wouldn't necessarily re-create the Maui infrastructure for this, and in fact they're working on a new, more generalized system that is nearly complete. It will be a 'multi' and will be an operating system with a neutral job distribution and collection system.

The system I just discussed actually won't work at Maui, so the Maui side of the house is a nonissue with what's discussed in our paper. Definitely, some of the archaic stuff that was in here is an outgrowth of Maui, since we started there, and all the systems we've been trying to create need to continue working there. As the military has become more interested in this type of thing, they've also become more interested in less specific questions that are not geared perfectly toward the classic combat modeling. That's required NetLogo and Repast to be usable in these systems. It forces us to make them much more general, and as they become more general, they're more useful to more people. Now we're actually trying to make it nonspecific to Maui, and so, yes, there are certainly some historical accidents still in there. Steve, do you agree?

Steven Upton: Yes, I'd just say, that some things are still being worked on. This year people are trying to devise a set of standards that will specify terms. For example, study XML file is just an XML file that has some structure that says, "What kind of algorithms do you want to use to generate your experiment?"

Koehler: In some sense, as I said, the study XML file is probably something that could be mineable, and we've got copies of that. You can look at that, and there's no schema to it, but there's probably an implied schema.

OldMcData is a set of classes in Java code. Basically, it takes what's in that schema and generates the excursion from your model. That's the other thing we were trying to do — put in an XML front for the parameters. For all your parameters, your agent specifications are in that XML file, and all you're doing is sampling over that and making changes to that XML file. That's some of the vestigial parameters in NetLogo — making an XML file for that and then having the wrapper to call NetLogo with the XML file and using the workspace provided by the NetLogo guys.

Let's see, what else? I was trying to make OldMcData so you could actually have different plug-ins for distributed mechanisms, so you could use Sun's grid or whatever. At this time, Condor is the only one we've got a couple wrappers on that actually generates OldMcData from your parameter specification; generates the submission data file, which is what Condor requires for a job; and has all those parameters out there to send the stuff off to the code. This means you can actually change another evaluation mechanism to say that you want to run a local or Condor. Of course, there are some issues, but basically, you could pass that off to the distributing mechanism. You could make it generic enough so it really doesn't matter. That was the intent. To sum up, there are mineable things in terms of, say, the standards, if that answers your question.

Bankes: We'll take this offline. I have a lot more I'd like to ask about.

Koehler: Okay.

Collier: Are there any more questions?

Seth Tisue: Seth Tisue, Northwestern University from the NetLogo team. I think software for setting up runs on clusters is something that's in its infancy, and so there really aren't a lot of easy-to-use tools out there for doing that, particularly when you're trying to develop a tool that's generic across different toolkits and different cluster environments. That's actually a really ambitious project, so you seem uncomfortable with the level of ease of use that you've achieved so far. I just wanted to say that there really isn't a lot of competition out there, so the fact that you're doing this at all is a step that needs to be taken. You can worry about making it easier later.

Koehler: Well, that's clearly been the path we've chosen.

Tisue: Yes, and I have a brief question. You said that this is more flexible than the behavior space tool that's built into NetLogo for specifying your experiments. What's more flexible about it?

Koehler: Well, the main thing is just that it can interface with Condor and run on a cluster. I'm not an overly sophisticated user of the behavior space, but one thing that I believe you can't do on the behavior space is, in essence, import a CSD file that specifies all the parameters you want to use over the course of the entire experiment so that you can use some other piece of software to create your design or experiment. That might be troublesome to input by hand.

Tisue: I see. Well, maybe we can work together on making behavior space work with the system so that if that level of flexibility isn't needed, people could specify their experiments in our software and then use yours to do their runs.

Koehler: Absolutely. Yes, that'd be great.

Collier: Do we have any more questions? No. Thank you, Matthew.

Macal: I'd like to thank all the speakers in the first session. I'd especially like to thank Nick Collier for being the chair and the discussant.

Behavioral Modeling Methods

INTEGRATING LIFE-LIKE ACTION SELECTION INTO CYCLE-BASED AGENT SIMULATION ENVIRONMENTS

J.J. BRYSON,* T.J. CAULFIELD, and J. DRUGOWITSCH,
Department of Computer Science, University of Bath, Bath, UK

ABSTRACT

Standardized simulation platforms such as Repast, Swarm, MASON, and NetLogo, are making agent-based modeling (ABM) accessible to ever-widening audiences. Some proportion of these modelers have good reason to want their agents to express relatively complex behavior, or they may wish to describe their agents' actions in terms of real time. Agents of increasing complexity may often be better (more simply) described by using hierarchical constructs that express the priorities and goals of their actions and the contexts in which sets of actions may be applicable. Describing an agent's behavior clearly and succinctly in this way might seem at odds with the iterative, cycle-based nature of most simulation platforms. Because each agent is known to act in lock-step synchrony with the others, describing the individual's behavior in terms of fluid, coherent long-term plans may seem difficult. This paper describes how an action-selection system designed for more conventionally humanoid artificial intelligence, such as robotics and virtual reality, can be incorporated into a cycle-based ABM simulation platform. We integrate a Python-language version of the action selection for Bryson's Behavior-oriented design (BOD) into a fairly standard cycle-based simulation platform, MASON. The resulting system is currently being used as a research platform in our group, and has been used for laboratories in the European Agent Systems Summer School.

Keywords: Behavior-oriented design, agent-based modeling, POSH, action selection system, Python, social behavior

INTRODUCTION

Standardized simulation platforms, such as Repast, Swarm, MASON, and NetLogo, are making agent-based modeling (AMB) accessible to ever widening audiences. Some proportion of these modelers have good reason to want to express the actions of their agents in terms of real time. For example, they may want to describe continuous, durative actions, such as walking or fighting wars. Such actions can either have prespecified durations, or they can continue occurring until some event consummates or interrupts them. Modelers may also need to increase the complexity of their agents to include a relatively complex variable state, such as memories of past interactions with other agents or conflicting theories of relationships between third parties.

Bryson (2003a) describes how agents of increasing complexity may often be better (that is, more simply) described by using hierarchical constructions that express the priorities and

* *Corresponding author address*: Joanne J. Bryson, Department of Computer Science, Artificial Models of Natural Intelligence (AmonI), University of Bath, Bath, BA2 7AY, United Kingdom; email: jjb@cs.bath.ac.uk; website: http://www.bath.ac.uk/comp-sci/ai/AmonI.html.

goals of their actions and the contexts in which sets of actions may be applicable. Describing an agent's behavior clearly and succinctly in this way can seem at odds though with the iterative, cycle-based nature of most simulation platforms, which seem to specify that each agent proposes its own action in lock-step synchrony.

This paper describes a Python-based version of the action selection for Bryson's Behavior-oriented design (BOD). BOD splits an agent's control into two sorts of entities: modular *behavior libraries*, which are written in conventional object-oriented languages such as Python, and parallel-rooted, ordered, slip-stack hierarchical (POSH) *action selection scripts*. Behavior modules enable action by:

1. Containing the code describing *how* an action is performed,

2. Providing a place to store any state/memory needed to conduct those actions, and

3. Containing code describing any sensing that must be conducted to acquire that state/knowledge.

Action selection scripts (also called *plans*) then order these actions and also the sensing events necessary to support decision making in a way that reflects the individual agent's priorities. In other words, action selection determines *when* an action should be expressed.

The Python version of the core action-selection system, jyPOSH, is currently being used for two different research projects. One is as a real-time system controlling a virtual-reality game agent in Unreal Tournament (Partington and Bryson 2005; Bída et al. 2006), and the other is for agent-based social simulations. We are currently using MASON (Luke et al. 2003) for the social simulations, but our solution should be easily extendible to other platforms. We have used an early demo of the BOD/MASON system involving dogs herding sheep to teach BOD at the European Agent Systems Summer School in July 2005, where students were encouraged to extend the behavior of the dogs and sheep, including turning the dogs into wolves that ate the sheep. We are now also building research projects on primate social behavior on this BOD/MASON platform, although earlier versions of this work were run in NetLogo and SmallTalk (Bryson et al. 2006).

In this paper, we begin by describing the MASON platform and why we are using it. We then give a brief overview of BOD and its POSH action-selection framework. The main purpose of this paper is to outline the relatively simple steps necessary in the abstract for integrating POSH action selection into a cycle-based simulation tool, and the relatively more complicated process of actually integrating a Python planning system and behavior modules into MASON.

BACKGROUND: COMPONENT TECHNOLOGIES

MASON and Other Platforms

MASON (Luke et al. 2003) is multi-agent simulation environment written in Java. It is designed to provide a generic platform with core functionality upon which many different types of simulation can be built. The philosophy behind MASON is to provide a fast, portable, and non-domain-specific framework that contains features commonly required by different types of multi-agent models. Other simulation environments, such as NetLogo and Repast, are more domain-specific but may also provide a great deal more support for writing simulations in their respective domains. NetLogo, for instance, is designed for simulating social interactions, and this is supported by the existence of fairly high-level, competent agents (called turtles) natively within NetLogo and a large number of functions (moving, turning, finding other turtles, etc.) to manipulate these agents. NetLogo provides its own high-level language that features the turtle's actions and ways to manipulate the turtles as special primitives. It is a relatively simple and constrained language intended to help inexperienced programmers build simulations quickly.

In contrast, MASON does not restrict itself to a particular type of agent but provides the flexibility to completely customize its agents. As a result, the primitives provided by MASON are on a lower level. If developers want to manipulate the agent's behavior at a higher level of abstraction, they are required to program custom high-level primitives themselves. The group that developed MASON has a particular interest in evolutionary simulations and, as such, were willing to sacrifice greater development time in exchange for rapid execution time. All entities must be implemented in MASON's native Java language.

Over the last four years, we have had approximately 12 student dissertations (undergraduate and one-year taught masters) on social simulation conducted in our group. Each of these projects has contained a phase where students evaluate a number of platforms and then choose the one on which they build their project. By far the most popular platform has been NetLogo (Wilensky 1999). This is primarily because of the ease of developing both agents and UIs for assisting in running experiments (and, more recently, behavior-space parameter sweeps.) MASON is the second most popular tool. When it wins over NetLogo, it does so because it runs faster and is more programmable (e.g., it allows linking to Java libraries).[1] Two students selected Repast and one student selected SeSAm.[2]

MASON is divided into separate layers (see Figure 1). The inner layers — *the core* — can function without the outer layers, which provides visualizations and domain-specific functions. At the core are two essential layers — the model and utilities layers. The model layer contains classes that represent two-and three-dimensional fields in both bounded and toroidal forms, both discrete and continuous, along with methods to place, locate, and determine the distance between objects in the field. A scheduler is also supplied, allowing events to be run at various frequencies of iterations and even in a particular order during one iteration. The other core layer implements classes that are useful when writing simulations, such as optimized

1 Although NetLogo now also has a Java API, we have not yet tried to exploit this.

2 The SeSAm student did not complete her dissertation, but this was probably not a consequence of the platform.

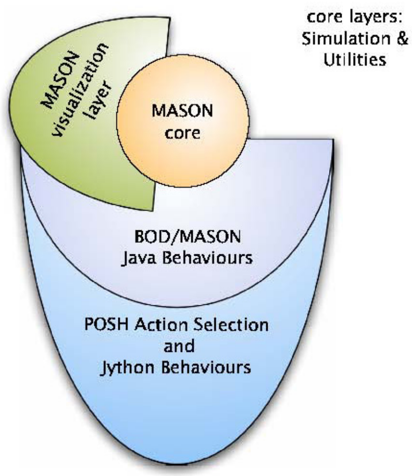

FIGURE 1 Basic BOD/MASON architecture

collections and a random number generator. We have added extra layers of functionality on top of MASON's default layers (in keeping with MASON's design principles) that provide the user a simpler programming interface, including access to specialist action-selection code.

Besides the core layers, MASON includes a visualization layer. This layer is designed in the basic MASON architecture as optional — simulations can be run with or without it, and performance increases when it is turned off. The layer allows for 2D and 3D visualizations by using a sophisticated class model, making the display of model data easy to customize. This method also isolates the simulation in a separate thread from the GUI and displays, increasing performance, and enabling the visualisations to be turned on or off as the simulation runs. Information on the state of objects in the model can also be observed and altered at runtime through the use of the inspector classes in this layer.

Beyond these layers, MASON does not provide anything. It is intended as a generic simulation environment and leaves the implementation of domain-specific layers and of the actual models up to others. Our contribution is that we have extended the MASON environment by integrating it with a POSH engine in a compatible language. This allows modelers to use the BOD development methodology in MASON and, as a by-product, supports their describing their agents in Python rather than Java. This is an advantage because Python is a far quicker and easier language to code in.

BOD and POSH Action Selection

Behavior-oriented design is a methodology for constructing relatively complex, animal-like agents (Bryson 2001, 2003b). The purpose of the methodology is to help ordinary programmers produce agents that generate behavior to meet multiple, possibly conflicting, goals that operate continuously and in parallel. Most programmers — and indeed, most ordinary people who simply make plans — are used to thinking in terms of sequences of actions. The challenge of making an autonomous agent is that many goals are held at the same time and must be met concurrently (e.g., one agent can simultaneously entertain the desire to get a promotion, have dinner, and wear a clean shirt in the morning). Further, these goals must be achieved in the face of an unpredictable, dynamic environment, which may provide both unexpected opportunities and challenges. Steps that might ordinarily be considered necessary for achieving a goal may not be required in some contexts but may have to be repeated or even abandoned for other strategies in others.

BOD splits the problem of agent intelligence into two parts. One part is the construction of a *behavior library*. The behavior library consists of relatively ordinary code for describing individual actions, including active, goal-directed sensing and evaluation of the environment. The behavior library is typically modular, based on the established development methodologies of object-oriented design (Parnas et al. 1985; Coad et al. 1997) and behavior-based AI (Brooks 1991; Matarić 1997). The other part of a BOD agent is a hierarchical structure known as a *POSH plan*, which organizes primitive actions, including sensing.[3] The POSH plan determines an individual agent's capabilities and priorities. Many very different agents can share the same library of behavior modules, provided each has its own POSH action selection. A single behavior library is typically developed for any one research platform, but multiple types of agents can be designed that reuse aspects of this code yet have different priorities or "personalities" specified in their POSH plans.

The primitives of POSH plans are *acts* and *senses*. These are the interface to the behavior library: each primitive is typically a method call to one of the objects used to represent a behavior module. Senses report on conditions in the simulation to inform decision points in the plan, while acts actually change some aspect of the simulation (possibly just the agent).

Picking appropriate primitives is a challenge. Deciding the correct granularity at which action selection should occur is half of the problem of action selection for AI. BOD does this through a set of heuristics that are run iteratively over the development period. Essentially, if a POSH plan is becoming too complicated, current granularity may be too small and new, more abstract primitives should be built. On the other hand, if there is redundancy in the behavior library code, a primitive probably needs to be decomposed into smaller elements to facilitate reuse. If there is redundancy in a plan, additional memory should be introduced into the agent in the form of state in a behavior module, which the plans can the refer to, making them more general. Details of the BOD methodology can be found elsewhere (Bryson 2001, 2003b).

Besides the plan primitives, POSH also provides three aggregate types that provide the plans' order and hierarchy. These are *action patterns*, which are simple sequences; *competences*, which are prioritized sets of productions (pairs of sensory preconditions and their action consequences); and a *drive collection*, which is a special competence that serves as the root of the action-selection hierarchy and specifies the main drives or motivations for the agent. Details of POSH can also be found elsewhere (Bryson and Stein 2001; Bryson 2001, 2003a).

Action selection is where the programmer or planner's established ability to sequence actions and prioritize goals can be expressed. The underlying sequential structure can be encoded in the plan hierarchy, then the action selection mechanisms can manipulate the actual order of action expression in response to motivational and environmental context. These acts and context-checking senses are then the plan primitives, which must be supported by the behavior library.

---

3 POSH is an adjective that stands for "Parallel-rooted, Ordered, Slip-stack, Hierarchical."

The version of POSH we use with MASON is a derivative of pyPOSH. Originally implemented by Kwong (2003), pyPOSH is a Python version derived from the (Bryson 2001) lisp version of POSH action selection. As documented by Kwong (2003), we chose to implement a version of POSH in Python because Python is (like lisp) a high-level loosely typed language that allows rapid code development. However, Python is also a scripting language with relatively familiar syntax and structure, making it more accessible for programmers familiar with languages such as Java, C, or perl. Python is also strongly object-oriented, which makes it amenable to both contemporary software engineering in general and to BOD in particular. Previous to this project, the main pyPOSH behavior library was for Unreal Tournament (Kwong, 2003; Partington and Bryson 2005). Figure 2 shows a BOD/MASON simulation. The MASON simulation window shows the agents, while the MASON control panel (upper right) shows the console tab; the panel controls the simulation. The smaller windows are inspectors. The simulation can be started from the command line, or from the jyPOSH GUI (overlapping the simulation window on the lower left.)

EXAMPLE: SHEEP AND DOGS

Before detailing the technical issues involved in bringing POSH to MASON, we start by describing a complete BOD/MASON simulation. Let us consider a simple example: we want to simulate a bunch of sheep that are herded by one or more dogs. This gives us two types of agents to model, the sheep and the dog, both situated in the same environment.

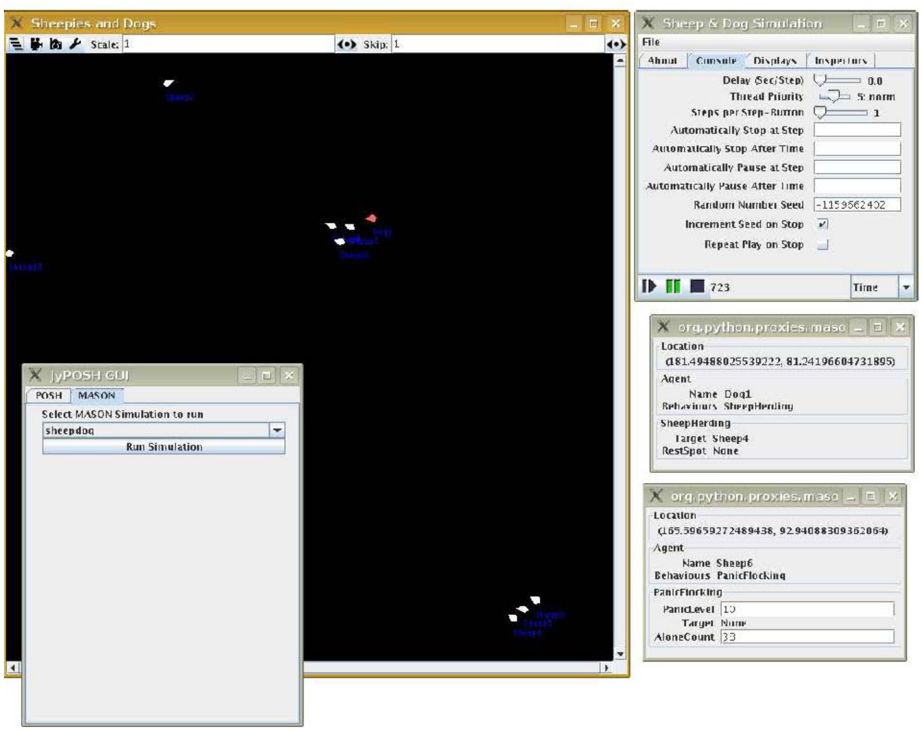

FIGURE 2 Screenshot of BOD/MASON running the sheep/dog demo

Even though we think of the behavior of sheep and dogs as very different, when we consider the problem of describing intelligence from first principles (e.g., physics), we realize that all agents in a particular environment share a large amount of basic behavior. For example, consider navigation. Currently, we restrict ourselves to using the continuous toroidal environment MASON provides since this simplifies navigation. MASON is itself in Java and does not provide many primitives, and so we have implemented some Java primitives on top of MASON that simplify the control of agents in these environments (e.g., the actions `move` and `setDirection` and the senses `location, orientation`, and `distanceTo`). These primitives are implemented in Java to provide sufficient speed in the simulation. An additional layer in Python links the MASON agents to POSH action selection and provides all the primitives implemented in the Java layer, like, for example, information about the agent's location. Once implemented, all these primitives can be used and reused by all agents that reside in that environment.

The sheep's behavior is determined by the combination of its POSH plan and its behavior library. The plan tells them to graze (i.e., do nothing) as long as they are calm and close enough to fellow sheep. If a sheep is not close enough to other sheep (where *enough* is determined in a behavior and depends on just how calm they are feeling), it will move toward the others. When moving to other sheep agents, their motion is determined by simple flocking behavior, balancing cohesion and aversion between sheep, and some momentum and random motion (Reynolds 1987). If the sheep sense a nearby dog, their plan says that this is a higher priority concern than grazing. When a dog approaches, the sheep's panic level rises and influences the weight of the different motion components, causing the sheep to increase their speed and flock more densely. Grazing, moving toward other sheep, and flocking motion are implemented as a POSH behavior library, written in Python. The act/sense primitives provide a thin layer on top of the MASON primitives. The different behaviors are then integrated by the POSH plan.

The dogs' behavior is designed in the same manner, by linking a set of POSH primitives provided by the dog behavior library by a POSH plan. The plan tells a dog to rest as long as the herd width does not exceed a certain threshold, and otherwise to approach the closest sheep to make them move closer together. The dogs can share the same resting behavior as the sheep, but their approaching behavior is different, since they pick just one sheep to hassle first. The agents' plans determine which behavior from the library is expressed by those particular agents.

Having implemented the sheep and dog behavior, all that is left for the user to do is to specify the environment and the location of sheep and dogs. A simulation master file, also written in Python, takes this role by defining a set of agent classes by their POSH plan, behavior libraries, and their appearance in the GUI. It also specifies the size of the environment and how many sheep and dogs are initially randomly located in the environment.

The sheep/dog demo has been implemented for teaching purposes to show the power of the BOD approach. Students are encouraged to think about how to change the behavior of the agents. For example, how would you change the sheep into deer, which are smart enough to scatter if a predator attacks rather than to clump? How would you change a dog into a wolf, which catches sheep rather than just frightening them? Could you make a smarter dog that herds the sheep more quickly? Can you use a team of dogs?

To summarize, the only parts that have to be written by the users are:

- *POSH behavior libraries* containing a set of actions and senses that are usable by POSH plans. These actions and senses provide agent-specific higher-level functionality than what is provided by the simulation environment. The behavior library for sheep, for example, provides among others the actions `set_panic_to_max`, `move_to_target`, `flock_move`, and senses `alone`, `predator_close`. The behavior libraries are written in Python.

- *POSH plans* linking the actions and senses of the behavior libraries to form the behavior of the agent. The syntax of these plans is POSH-specific; the complete grammar can be found on the POSH web page or in the pyPOSH documentation. There is also an integrated development environment (IDE) for developing these plans, called the Advanced Behavior Oriented Design Environment (ABODE). Figure 3 shows a screenshot of ABODE editing a sheep's POSH plan.

- *Main simulation script* to set up the simulation, specify the environment size, and initialize the agents of the environment. In our sheep/dog example, this script defines the sheep and dog agents by linking their behaviors with their plans, specifies their appearance in the simulation, and sets the number of sheep and dogs in the simulation, and their initial location. The simulation script is written in Python.

The simulation is displayed and controlled by MASON, which can in turn be started from the main pyPOSH GUI. MASON uses pyPOSH and the user-provided behaviors and plans to let the sheep flock and the dog hunt them. The details of how exactly POSH is controlled from within a simulation environment like MASON — all hidden away from the user — are given in the next section.

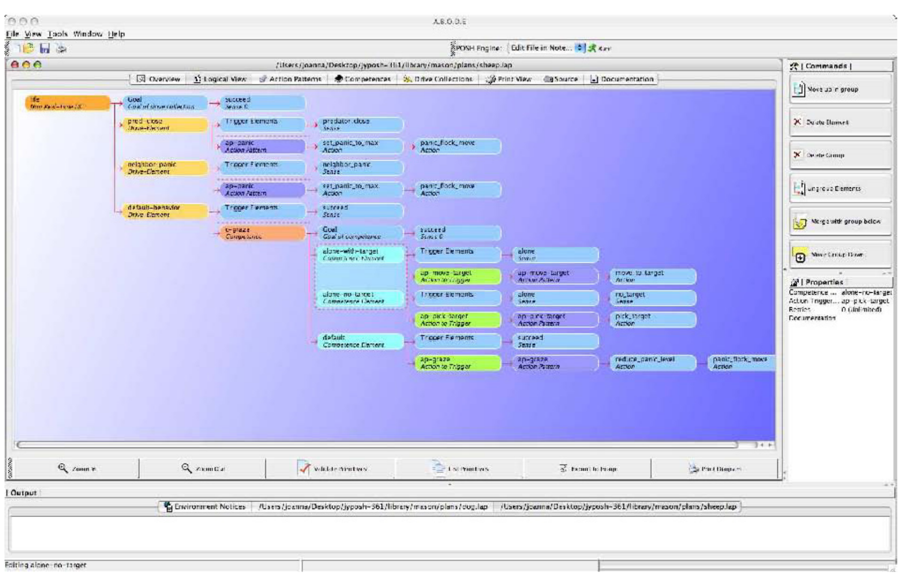

FIGURE 3 Screenshot of ABODE editing a sheep's POSH plan

UNDER THE HOOD

Our aim was to keep the tasks of the user developers simple while still maintaining the flexibility of POSH action selection and the power of the MASON simulation toolkit. All the user needs to do is to write or extend Python behavior modules, POSH scripts, and a master simulation file. How we have provided such simplicity, and how MASON actually interfaces with pyPOSH, is described below.

Language Issues

The POSH system was already written in the Python language. Rather than making it compatible with MASON by rewriting it in Java, we decided to use Jython to integrate the two. Jython is an implementation of Python written in Java, which compiles the Python code to Java byte code, which is then run on the Java virtual machine. Jython has the extra advantage over ordinary Python that Java classes can be used from within Python code, and vice-versa (with a few caveats). Jython also supports multiple-inheritance, and more important for us, multiple-inheritance from Python and Java classes simultaneously.

Jython implements an earlier version of the Python language than the one that POSH was written in, meaning that the code initially would not run under Jython. To fix this, it was necessary to go through the POSH code and change any parts that used features only found in the newer version of the language. For instance, in more recent versions of Python, there are "true" and "false" keywords in the language, whereas in earlier versions these were represented by "1" and "0" respectively. Another problem that was at first more difficult to solve was the fact that the way Python did variable scoping changed between versions. This presented a problem, particularly with nested functions, where inner functions could not access the members of the enclosing class. Luckily, the new scoping was accessible in the version of Python that Jython implements, so instead of rewriting all of the existing POSH code, we only needed to add an import statement (from_future_import_nested_scopes) at the top of each Python file to enable the new scoping.

Layer on Top of MASON

Jython allows the direct use of Java classes from Python scripts. Hence, MASON can be directly used from Python scripts, but then the user still has to cope with the complexity of the MASON framework. To avoid that, we have written a thin layer around MASON that removes some of its complexity and tailors the API specifically for a certain sets of environments to be used in combination with POSH actions selection.[4]

For now we have committed ourselves to continuous toroidal environments of arbitrary size. Consequently, our simulation class only provides such environments and a simulation GUI

4 Early users of BOD/MASON complained a great deal that they needed to look in both Java and Python behaviors. That was never really our intention, but more of a teething problem. In the current version of BOD/MASON, we seem to have been successful at completely encapsulating the Java away from the user.

class that visualizes them. In the master simulation scripts, the user is only required to perform the following only tasks:

- Override the simulation class to specify the size of the environment and initialize the agents at the start of the simulation. That is, the user has to specify the agent type, its initial state, and where it is located in the environment.

- Override the simulation GUI class to set the name of the simulation and the color of the environment.

All the rest of the set up for the MASON simulation, and GUI is performed in the thin layer between the user and MASON.

Entities and Agents

Because MASON provides only very limited agent functionality, we have constructed an additional layer for entities and agents that is similar to that which can be found in NetLogo. MASON provides the concept of *fields* of the environment that can be seen as a set of different layers, each of which covers the entire environment. Objects in the environment are located in one or several of those fields, and all the functions we have provided for these objects can be restricted to certain fields. Going back to the sheep/dog example, all the sheep are located on the sheep field, and all dogs are on the dog field. If a sheep wants to find the closest other sheep, it just needs to look for the closest object in the sheep field. Hence, using different fields allows grouping of agents and entities. This speeds up certain functions, particularly those exponential on the number of agents, such as comparisons. However, it adds a level of complexity for the MASON developer, since the fields need to kept track of.

The simplest objects we provide are *entities*, which have a location and a certain appearance in the environment, but no orientation. They can represent any stationary object in the simulation, for example, a tree or pond. In addition to its properties, we provide a set of sense primitives for these entities, like its distance to other objects, or the approximate "center" of all objects in a certain field. These senses apply either to all fields, or just to a given subset.

Another class *oriented entity* inherits the basic entities and provides an additional "orientation" property and some methods that allow it to move in the environment. Even though it does not yet provide autonomous control, this class provides all that is required from the MASON side to interface it with POSH action selection, as described in the next section.

Synchronous Asynchrony

BOD agents are asynchronous by design. New actions are performed in response to events in the environment or changes of internal motivation of the agent, not at regularly scheduled intervals. On the other hand, most simulation environments are stepped, and all agents have to choose their actions at each step. How can we combine those two modes of operation?

In fact, this problem is simpler to solve than it appears. Although BOD agents may have multiple threads or even devices supporting the behavior generated by the behavior modules, the actual POSH action selection is sequential and cycle-based. The responsiveness is a consequence of the rapid cycle rate of POSH, which typically can operate at hundreds of cycles per second, providing that no primitive action it calls has blocked it for any length of time. In fact, to sustain the illusion of asynchrony, one of the requisite properties for BOD behavior libraries is that the method calls to behaviors should not block while they wait for a protracted action to happen. Where protracted actions occur (e.g., motion in a robot) methods should only serve to initialize or reparameterize an action, which is then sustained by the behavior module itself.

When a BOD agent appears to be engaged in a sustained activity, what actually happens with respect to action selection is the following. At each cycle, POSH checks if the conjunction of its sensed internal or external states is unchanged, which causes it to perform the same act as in the last cycle. The expressed behavior is continuous because the POSH cycles are performed at a high frequency, and the action primitives are carefully designed so they show no disruption in expressions of behavior when they recur.

This architecture was originally designed to allow an agent with hierarchical action selection to still be a fully responsive and reactive real-time autonomous agent; in fact, POSH was originally designed for and implemented on autonomous robots. Fortunately, this structure also allows POSH to be easily integrated into stepped simulation environments. Rather than having the POSH agent continuously call the internal cycle, the control is now given to the simulation environment. Its role is to signal the POSH agent at each simulation step to perform one internal cycle, which causes the agent to perform actions according to its POSH plan. The agent's environment becomes the simulation environment, and its behavior libraries use the simulation environment primitives to sense and act within this environment. Because some cycles in POSH are dedicated to details of the decision making, a new expressed action may not be chosen on every cycle of the simulator. But this is not really a problem if we are trying to simulate realistic real-time agents. In fact, faking asynchrony in cyclic environments has been shown to be important in simulating animal-like behavior (Hemelrijk 2000).

What does that mean for the specific case of using MASON as a simulation environment? First, we need to modify the MASON agent class we have described in the last section to use POSH to perform its action selection. A new class, *AgentBase*, which inherits from both the POSH agent and our MASON agent, is the foundation for all agents that perform POSH action selection in a MASON simulation (e.g., the sheep and the dog in our example). This class overrides the standard POSH cycle control and instantiates a separate MASON control class. At each step of the simulation, a method of this control class is called by the MASON-stepped scheduler, and that method subsequently calls the POSH core cycle. Hence, each step of the simulation causes one POSH cycle to be performed. Having several agents in the same simulation, the MASON scheduler calls the POSH cycle of each agent one after the other.

The final question is what implementation of POSH the agent will use. There are currently two versions of POSH action selection encoded in jyPOSH — a scheduled version, *ScheduledAgent*, which maintains a little more decision state and therefore is more cognitively plausible, and a strict slip-stack, *StrictAgent*, which runs faster. Bryson (2001, Section 4.6) details the differences between these POSH implementations.

The only thing that is left to the user is to specify for each agent class the set of behavior libraries to use, and which plan to run. The agent class has to inherit out combine POSH/MASON agent class, and the behavior/plan information is given by overriding its constructor.

Intra-Module and Intra-Agent Communication

The behavior library for BOD agents is usually modular. Modular decomposition driven by the sorts of memory an agent needs, just like in standard object-oriented design (Coad et al. 1997). BOD, however, does not require the modules to be fully encapsulated. Just as for the action selection, one behavior module may poll another for information the second module is a specialist in. Similarly, sometimes agents will need to communicate with each other in the simulation. Even if in the real world that communication would have been tactile or visual (for example, if one agent hugs another or glares at another), in a simulation, this information has to be transmitted between agents. Since in a BOD system all sensing is done in behavior modules, in a BOD agent such information needs to be transmitted between two agents' behavior modules.

BOD/MASON framework provides a unified solution for both of these problems. All `Behaviour` instances for a single agent are stored in a behavior *dictionary* (the Python version of a hash table), which is maintained by the POSH `Agent` object. The `Agent` object also provides the method `getBehaviour` which allows any process having access to the agent object to get the state of any of the agent's behavior module objects. Additionally, each behavior object keeps a reference to the agent object that owns that behavior. Thus, different behaviors can exchange state information by calling `self.agent.getBehaviour (behaviour_name)`, and can get information about the internal states of other agents (like the closest other object on the agents field) by `self.agent.closest (agents).getBehaviour (behaviour_name).`

Agent Visualization

The idea of separating the visualization from the simulation core in MASON is also applied on the visualization of agents. While the agent only exists in the simulation, it is linked to a portrayal object that is responsible for its visual appearance in the MASON GUI. Directly translating that to POSH/MASON, the user would have to perform several additional (nontrivial) steps to define the appearance of an agent.

We have simplified this procedure by removing the separation between the agent object and its visualization. The user now has to call a class method in the overridden constructor of the POSH/MASON agent to set the appearance of the agent. This method creates the agent's portrayal object and links it to the simulation visualization as soon as it is created. This is another measure for the user to simplify the use of MASON.

Inspecting Agent States

In complex simulations with many agents, it is hard to keep the overview over the variables that specify the memory and identity of the agents. Fortunately, MASON provides some tools that allow the display of changes in the agent's internal states while the simulation is running. Its implementation is based on using Java *reflection* on the methods of the agent to identify accessor and mutator methods and display them in the GUI's inspector window.

In BOD, the state of an agent is determined by the state of the behavior modules it uses from the behavior library. These behaviors are implemented in Python. Unfortunately, Jython creates Java proxy objects for Python objects, so its methods are not accessible through Java reflection. Hence, the standard inspectors do not work for POSH/MASON agents.

To regain access to the MASON inspectors for BOD/MASON agents, we have added an additional method to the `Behaviour` class (the base class for all behaviors) that lets the user register Python accessor and mutator methods to be used by MASON inspectors. These methods are collected in the agent's inspection objects. Whenever the user requests inspection of a certain agent, some Python code in the BOD/MASON agent's base class adds the inspection GUI components and queries the state accessor methods to display the current state of the agent's behaviors. The user can then modify that state through text fields in the GUI, and these modifications are communicated to the agent's behaviors via the provided mutator methods.

Again, all the implementation detail has been hidden from the users. The only thing they have to do is to provide and register accessor/mutator methods for all the behavior she defines.

SUMMARY

This paper has presented the BOD/MASON agent-based modeling tool suite. Both are described in general in terms of incorporating complex agent action selection into a stepped ABM simulator, and in particular the problems we encountered and the solutions we have implemented in building POSH capabilities into MASON. We have also briefly outlined the BOD methodology and provided an overview of MASON, including a brief comparison between it and other simulation platforms.

Since BOD/MASON is one of the development platforms for our research into the evolution of social behavior, we expect it will continue to grow and improve. We have now created a mailing list and a bug-tracking system as well as having alpha released an IDE for POSH plans. We encourage our colleagues to consider using this system.

ACKNOWLEDGMENTS

BOD/MASON was originally built by Tristan Caulfield and has been under continuing development and refinement by Jan Drugowitsch. Development time for both researchers was funded from a grant from The UK Engineering and Physical Sciences Research Council (EPSRC), Grant GR/S79299/01 (AIBACS). ABODE is being funded as contract work by an anonymous industrial benefactor. Thanks also to Ivana Čače and the students of the 2005 European Agent-based Systems Summer School (EASSS) who provided useful feedback. The greatest thanks goes to our primary user, Hagen Lehmann, for suffering through all stages of BOD/MASON development — an interesting way to learn to program for the first time.

REFERENCES

Bída, M., Burket, O., Brom, C., and Gemrot, J. (2006). Pogamut — platform for prototyping bots in unreal tournament (*pogamut — platforma pro prototypování botøu v unreal tournamentu*). In J. Kelemen, and V. Kvasnicka (editors), *Proceedings of Sixth Czech-Slovak Workshop on Cognition and Artificial Life*, pages 67–74. Slezská Univerzita v Opave (in Czech).

Brooks, R.A. (1991). Intelligence without representation. *Artificial Intelligence*, 47:139–159.

Bryson, J.J. (2001). *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. PhD thesis, MIT, Department of EECS, Cambridge, MA. AI Technical Report 2001–003.

Bryson, J.J. (2003a). Action selection and individuation in agent based modeling In D.L. Sallach, and C. Macal (editors), *Proceedings of the Agent 2003 Conference on Challenges in Social Simulation*, pages 317–330, Argonne, IL. Argonne National Laboratory.

Bryson, J.J. (2003b). The behavior-oriented design of modular agent intelligence. In R. Kowalszyk, J.P. Müller, H. Tianfield, and R. Unland (editors), *Agent Technologies, Infrastructures, Tools, and Applications for e-Services*, pages 61–76. Springer.

Bryson, J.J., and Stein, L.A. (2001). Architectures and idioms: Making progress in agent design. In C. Castelfranchi, and Y. Lespérance (editors), *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*. Springer.

Bryson, J.J., A. Yasushi, and H. Lehmann (2006). Agent-based models as scientific methodology: A case study analysing primate social behavior. *Philosophical Transactions of the Royal Society, B*. (in press).

Coad, P., D. North, and M. Mayfield (1997). *Object Models: Strategies, Patterns and Applications*. Prentice Hall, 2nd edition.

Hemelrijk, C.K. (2000). Towards the integration of social dominance and spatial structure. *Animal Behaviour*, 59(5):1035–1048.

Kwong, A. (2003). A framework for reactive intelligence through agile component-based behaviors. Master's thesis, University of Bath. Department of Computer Science.

Luke, S., Balan, G.C., Panait, L., Cioffi-Revilla, C., and Paus, S. (2003). MASON: A Java multi-agent simulation library. In D.L. Sallach, and C. Macal (editors), *Proceedings of the Agent 2003 Conference on Challenges in Social Simulation*, pages 49–64, Argonne, IL. Argonne National Laboratory.

Matarić, M.J. (1997). Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2/3):323–336.

Parnas, D.L., P.C. Clements, and D.M. Weiss (1985). The modular structure of complex systems. *IEEE Transactions on Software Engineering*, SE-11(3):259–266.

Partington, S.J., and J.J. Bryson (2005). The behavior oriented design of an unreal tournament character. In T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist (editors), *The Fifth International Working Conference on Intelligent Virtual Agents*, pages 466–477, Kos, Greece. Springer.

Reynolds, C.W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34.

Wilensky, U. (1999). *NetLogo*. Evanston, IL, USA.

AGENT-BASED CONTROL FOR DYNAMIC CONFIGURATION OF SPATIALLY DISTRIBUTED NETWORKS

E. TATARA, A. ARTEL, F. TEYMOUR, C. HOOD, and A. CINAR,*
Illinois Institute of Technology, Chicago, IL
M. NORTH, Argonne National Laboratory, Argonne, IL

ABSTRACT

Large-scale spatially distributed systems provide control challenges because of their nonlinearity, spatial distribution, and generally high order. The control structure for these systems tends to be both discrete and distributed. A layered control structure interfaced with complex arrays of sensors and actuators provides a flexible supervision and control system that can deal with local and global challenges. An adaptive agent-based control structure is presented whereby local control objectives may be changed in order to achieve the global control objective. Information is shared through a global knowledge environment that promotes the distribution of ideas through reinforcement. The performance of the agent-based control approach is illustrated in a case study where the interaction front between two competing autocatalytic species is moved from one spatial configuration to another. The multi-agent control system is able to effectively explore the parameter space of the network and intelligently manipulate the network flow rates such that the desired spatial distribution of species is achieved.

Keywords: Agent-based control, networks, generic algorithms

INTRODUCTION

Large-scale spatially distributed systems provide a difficult control challenge because of their nonlinearity, spatial distribution, and generally high order. The control structure for these systems tend to be distributed and contain discrete and continuous elements. Hybrid control systems that combine process dynamics and discrete control elements and include multiple models for different operating points are one way to develop control systems for spatially distributed systems (Christofides and El-Farra 2005). An alternative approach is based on a hierarchical agent-based system with local and global control structures (Tatara et al. 2005) that has been demonstrated on a network of interconnected continuous stirred tank reactors (CSTRs). Reactor networks exhibit highly complex behavior with multiple steady-state operating regimes and have a large pool of candidates for manipulated variables (Tatara et al. 2004).

The operation of highly nonlinear systems like autocatalytic replicator networks may benefit from evolutionary self-organizing control because the optimal operating regime and the required control strategies may not be known a priori. Agent-based control systems provide the capability for localized and global control strategies that are both reactive in controlling disturbances and proactive in searching for better operational solutions (Jennings and Bussmann 2003). An adaptive agent-based control system for a CSTR network is proposed. The

* *Corresponding author address*: Ali Cinar, Department of Chemical and Environmental Engineering, Illinois Institute of Technology, 10 West 33rd Street, Chicago, IL 60616; email: cinar@iit.edu.

performance of the agent-based control approach is illustrated in a case study where the interaction front between competing autocatalytic species is moved from one spatial configuration to another.

AGENT-BASED CONTROL FRAMEWORK

Multi-agent control system architectures have several properties that make them particularly attractive for use in supervising large, complex systems. The first property, which is usually the most important in critical systems, is a high level of reliability. Modularity, scalability, and adaptability are also attractive features of multi-agent systems. The adaptive and self-regulatory nature of agent systems has only recently been investigated for solving control problems that are normally solved with traditional methods.

Design Process

The design procedure used is a derivative of recent agent design methodologies based on the concept of the agent-services-acquaintance model (Wooldridge et al. 1999) and the application to manufacturing control (Brueckner et al. 1998). The goal of the design process is to develop an agent-based control system for physically distributed industrial processes. Certain parts of the agent-based control system are generic because they are based on general concepts of industrial control systems and the operation of distributed processes.

Comprehensive studies of the physical process domain provide information regarding the expected normal operating conditions of the processes, types of faults and disturbances that may occur, and control strategies. In addition, the desired process operation and/or optimal conditions are expected to be known by the designers. Required agent types and roles are identified on the basis of the requirements for controlling the physical system. The details of the hierarchical agent-based architecture (Tatara et al. 2005) are not repeated in detail here. The focus is rather on the specific agent synthesis and instantiation for the presented examples.

Agent Synthesis

There is nearly a one-to-one mapping of roles to agent types. The number of control agents is variable, depending on the number of reactors in the network as well as the complexity of the control actions being performed. Generally, a single control agent can be used for controlling each reactor. While multiple control agents can be applied to each reactor, from a software design point of view, it makes more sense to encapsulate the functionality of several control concepts (temperature, level, etc.) into a single software agent, as long as the control algorithms are not extraordinarily complex.

The number of arbitration agents is probably the most flexible variable in the agent model. As with the control agents, the run time environment will set the number of arbitrators required. Larger networks with more control agents will subsequently require more arbitrators to handle the setpoint change requests coming from the controllers. Ultimately, the number of realized arbiters will be determined by the supervisory level agents.

The simplest implementation will have at least one supervisory-level agent to coordinate efforts between the control and arbitration agents. While local interactions between agents are intended to serve as the primary driving force in the control system, the supervisor needs to maintain an overview at all times. More complex control schemes can include multiple supervisors for each high-level function, such as help in mediating disputes, setting spatial concentration patterns, and supervising process recovery from disturbances that are too complex for the local controllers.

Finally, there will be exactly one instance each of the data collection and data acquisition agents. The data acquisition essentially functions as a bridge between the agent system and the physical domain. The acquisition agent will read values from either a simulator or a hardware data acquisition system and write the numeric values to objects that can be read and manipulated by the control and arbitration agents. The data collection agent encapsulates the roles for both data collection and file input/output (I/O) since these roles share very similar tasks. Any agent in the system will likely have a small memory space for storing local information relative to its specific tasks. The data collection agent will, however, be responsible for cataloging relevant data for the entire network, such as average concentrations, or possibly even the concentration histories in each reactor. These data will be written to a file stream in chunks at some variable rate.

Global Knowledge Environment

Considering the nonlinearity of reactor networks, it is difficult to predict how the behavior of the system changes when the system parameters are manipulated. Consequently, one cannot easily predict how to change the operating conditions of the network by manipulating the flow rates, or what the localized operating conditions should be, in order to satisfy a global objective. Several methods can be used to guide the decision agents in planning their control strategies, including dynamic exploration of the parameter space, rule-based heuristic models, and first-principles-based models.

Although information is exchanged between agents via arbitrators, these interactions are local and limit the amount and quality that can propagate through the system. The global knowledge representation (Figure 1) serves as an environment for indirect communication between agents. This concept builds upon the hierarchical structuring of the control system by adding a mechanism for communication and reinforcement of ideas. The information in the knowledge space is divided into categories, including local control objectives, control heuristics, and data-based models.

Information exchange occurs indirectly between agents because agents asynchronously read/write information from/to the knowledge space. For example, a particular agent may discover a local control strategy that works particularly well in meeting an objective set by a supervisor. This strategy is cataloged in the knowledge space by the originating agent. Other agents may read this strategy from the knowledge space and implement it to satisfy their particular control objective. The value of the strategy is then rated by the agents that adopt this new strategy such that its value relative to others is promoted. Similarly, outdated information in the knowledge space continuously decreases in value and eventually may be deleted from the knowledge space.

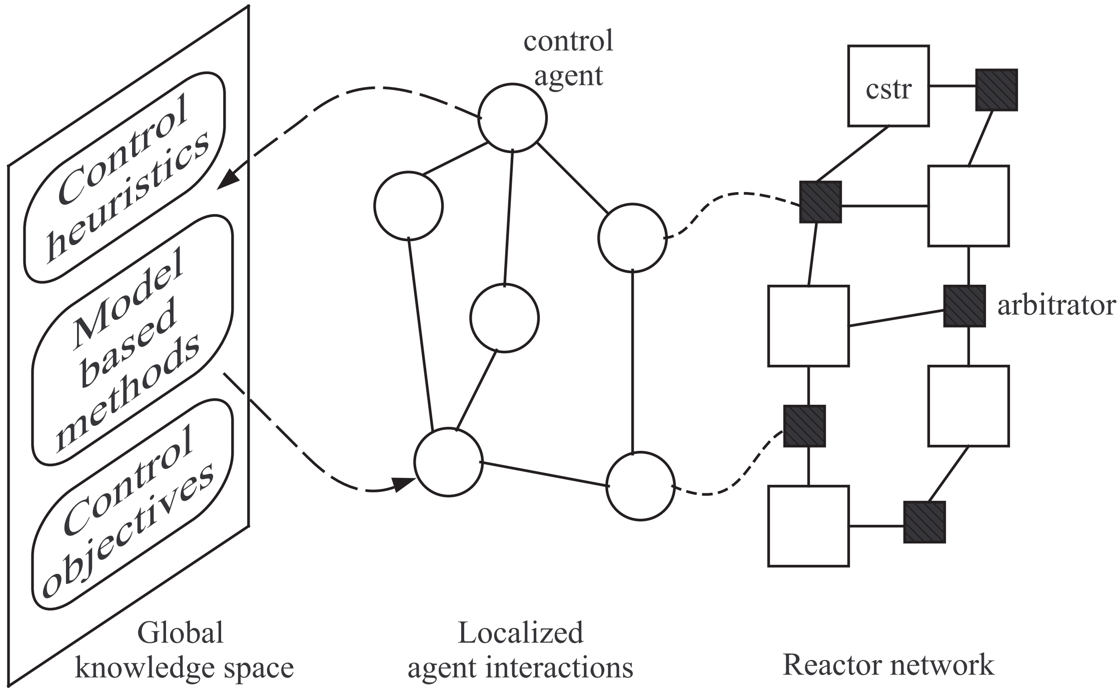

FIGURE 1 Coupling of control agents with global knowledge space

Although the stability of the agent dynamics cannot be guaranteed for every scenario, this methodology helps to reduce or prevent the emergence of dysfunctional agent dynamics by reinforcing "good" agent behavior while punishing undesirable agent behavior. Furthermore, the agent system has been designed under the assumption that the agents' decision delays are small compared to the time scale of the physical process. The importance of this assumption becomes apparent when the consequences of its impact on process performance are examined. If the agents' computing time is very long with respect to the process time scale, control of a continuous process becomes difficult because of the reduced data acquisition, control action computation, and implementation rates. This assumption generally holds for chemical processes in which operating changes are introduced infrequently and process dynamics are represented with time scales of tens of minutes or even hours. Traditional controllers are normally used in the event of very rapid, localized dynamics, and, although the agents may modify the setpoints of such controllers, the time-critical (first-response) control actions are strictly outside the domain of the higher-level decision-making agents.

SOFTWARE IMPLEMENTATION

The agent-based control system is implemented by using the open-source toolkit, Repast (REcursive Porous Agent Simulation Toolkit) (Collier et al. 2003). Repast is a Java-based framework for agent simulation and provides features such as an event scheduler and visualization tools. Software events are executed when autonomous agents register them with the scheduler.

Agents in Repast are implemented as Java classes, and the overall system implementation is the result of combining these various agent classes according to the system design criteria. The agent classes are organized in Java packages according to their functional relationships, including the core system, user interface, network objects, data acquisition, and simulation packages.

The Java agents created with Repast interact with virtual representations of the physical reactor network, as shown in Figure 2. The virtual network objects map the states of the physical system to objects that can be manipulated by the control agents. The interface between the physical network and the agent environment can take the form of a data acquisition system in the case of a real-world control application, or, in this case, a simulation of a chemical reactor network. The ordinary differential equations that describe the reactions in each CSTR are solved numerically by using the CVODE solver (Cohen and Hindmarsh 1994). The solver code is written in C and linked with Repast via the Java native interface (JNI).

PRODUCT GRADE TRANSITION IN A CHEMICAL REACTOR NETWORK

Product grade transitions may be used to schedule the production of various composite compounds at different points in time. The schedule can vary on the basis of a planned transition between product grades or on the basis of product demand. The overall product quality is determined by selecting one or more exit streams from the network and mixing them in the desired proportions. Therefore, if one would like to produce a grade consisting of the majority of one chemical species and only a small amount of another, it would make sense to mirror this grade composition in the network. The undesirable alternative would be to produce equal amounts of all products in the network and then combine only the fractions necessary to achieve the grade composition. However, this approach would be wasteful with respect to the less-used species as well as limit the production of the species that represent the majority of the grade makeup.

Supervisory-level agents set the desired network product grade, although the details of the control strategies are left to the local control agents, again as a result of the fact that network scale-up would reduce the performance of a more centralized control scheme. Each control agent

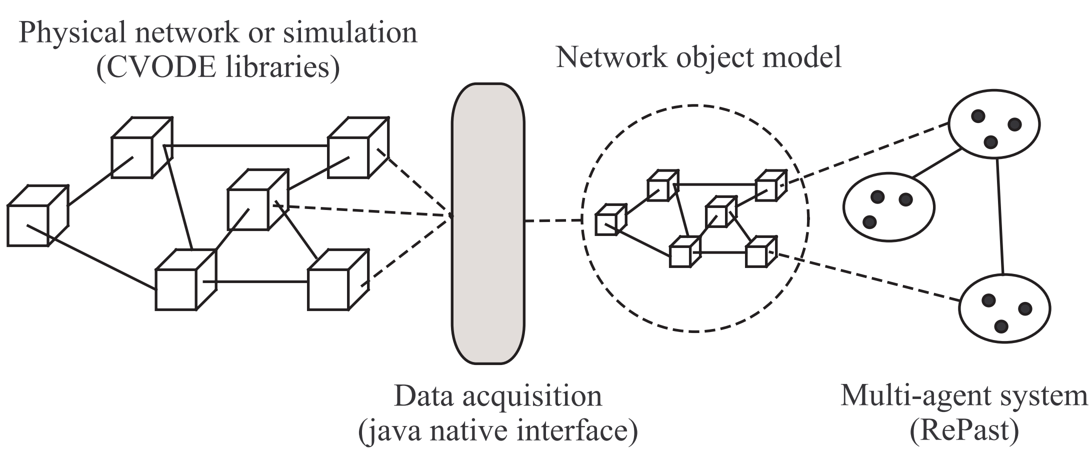

FIGURE 2 Overview of software components

evaluates a utility function that determines how the agent behaves relative to its neighbors. Specifically, the control agent's behavior refers to its desire to change from one dominant species state to another. This concept is adopted from the classic Schelling model of socioeconomic behavior, in which agents segregate in a discrete network on the basis of the makeup of the agents' local neighborhood (Pancs and Vriend 2003).

The supervisory-level agent(s) set the grade composition by specifying the fraction of each species desired in the network. The composition is then transformed into a set of behaviors via the utility function that determines each control agent's goal. The fraction of desired species determines the agent's willingness to change or remain the same. If the agent's average composition value is set very high and if the agent is surrounded by competitor agents, then its desire to remain unchanged is high. Conversely, if the agent's composition value is set very low and the agent is surrounded by cooperator agents, its desire to remain unchanged is low. This behavioral programming results in the network self-organizing to meet the global composition goal. A few isolated agents controlling reactors with a trace compound will contribute a lesser amount of their species to the product grade, while many closely clustered groups of the majority will contribute to the primary product in the desired grade. This self-organizing behavior arises from some aspects of the rules governing the local interactions as well as the open-loop behavior of the network, since local clusters of one autocatalytic species will be more stable than single, isolated reactors.

The performance of the agent-based control architecture is demonstrated in a case study to control the distribution of autocatalytic species in a network of 49 (7 × 7 grid) reactors hosting three autocatalytic species by using the interaction flow rates as the manipulated variables. The species that populate the reactor network are characterized by identical growth and death rates such that one species does not have an unfair advantage over the others.

Figure 3 shows the resulting changes in the species concentration profiles when the agent-based control system is tasked with creating different product grades starting from a random initial condition. After the desired product grade for species (1, 2, 3) = (0.25, 0.02, 0.73) is entered, the system goes through a series of transition states [Figure 3(a)] before successfully settling on the grade setpoint [Figure 3(b)]. A second grade transition (0.3, 0, 0.7) demonstration is also successfully executed by the control system [Figure 3(c)].

ADAPTIVE SPATIAL RECONFIGURATION WITH GENETIC ALGORITHMS

If a control agent desires to change the dominant species within the reactor under its control, it must transport some amount of the desired species from a reactor to which it is connected. If none of the immediate neighboring reactors contain the desired species, then the agent must negotiate with several other agents in the network in order to move the species from a more distant reactor. To minimize the disturbance to the operation of the network as a whole, the control agent attempts to find the shortest path between itself and a reactor containing the species that it needs in order to change the dominant species of its reactor. A genetic-algorithm-based (GA-based) (Mitchell 1998) spatial reconfiguration technique has been implemented and embedded into control agents in order to provide the ability to transport the desired species across the network to the target reactor.

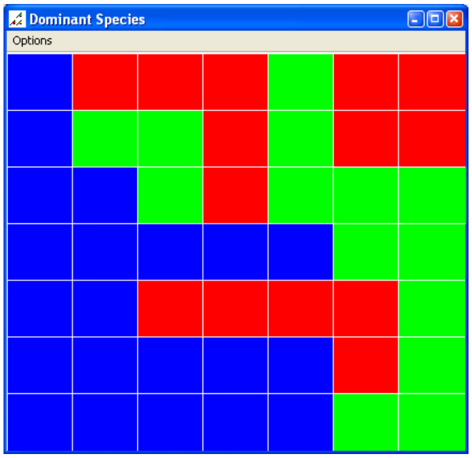

(a)

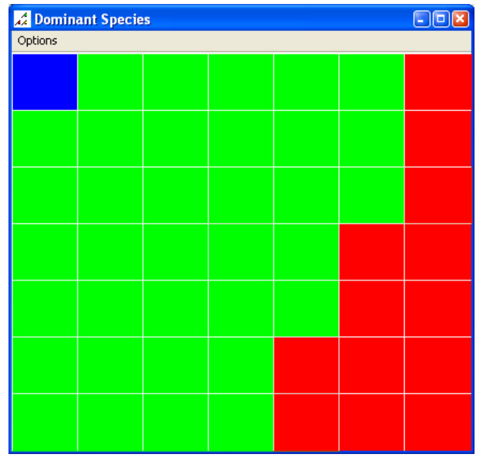

(b)

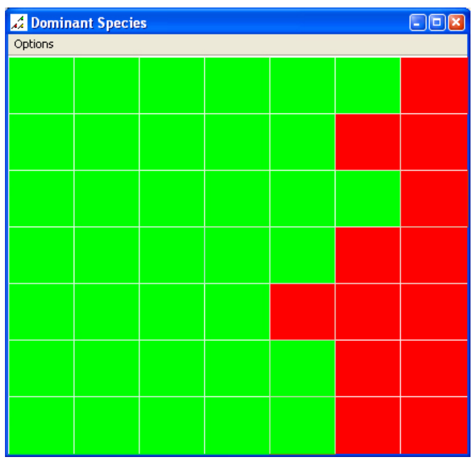

(c)

FIGURE 3 Evolution of two-dimensional spatial concentration profile of the dominant species in each reactor (Red [dark gray] represents species 1, blue [black] represents species 2, and green [light gray] represents species 3.)

The GA-based agent attempts to identify the most favorable source reactor, both in terms of the availability of the desired species and the proximity to the target reactor. The fitness value is calculated on the basis of the distance that the species has to travel in order to reach the target reactor, the magnitude of required increase in the feed flow rate to the source reactor, the rate of increase in interaction between each of the reactors, and the difference between paths from the source reactor to the target. The source reactor and the path to the target reactor are chosen according to the best fitness value calculated to meet the objective of changing the dominant species in the target reactor.

Figure 4(a) shows the initial conditions in a 3 × 3 network of reactors hosting three autocatalytic species. The objective set for the GA to achieve is to change the dominant species

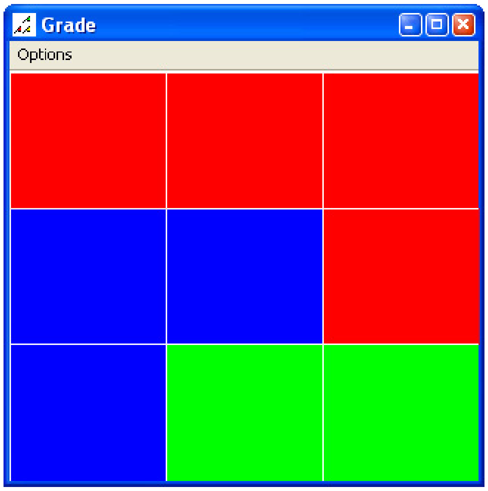

(a)

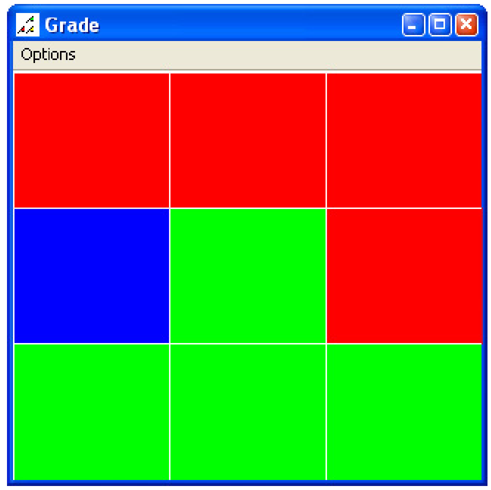

(b)

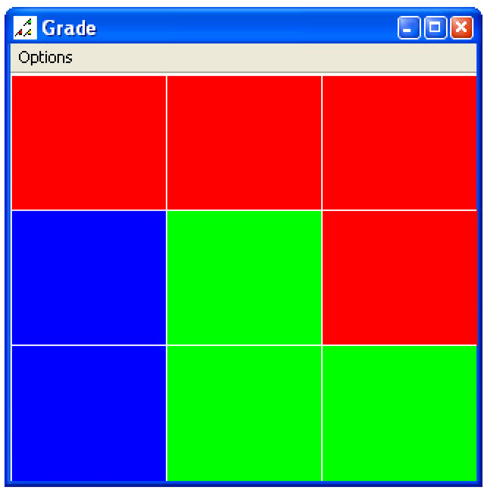

(c)

FIGURE 4 GA-based evolution of two-dimensional spatial concentration profile of the dominant species in each reactor (Red [dark gray] represents species 1, blue [black] represents species 2, and green [light gray] represents species 3.)

of the center reactor from species 2 to species 3. The time series plots for the concentrations are shown in Figure 5. After the network settles down from the initial startup at time 0, the GA agents are activated at tick 1,000. The GA identifies a source reactor as the reactor just below the center reactor. Figure 4(b) shows the system after the parameters found by GA are applied to the system to change the dominant species of the center reactor. After the objective has been changed, the network parameters are reset to the initial conditions, resulting in a new steady-state concentration profile in the network (Figure 4(c)).

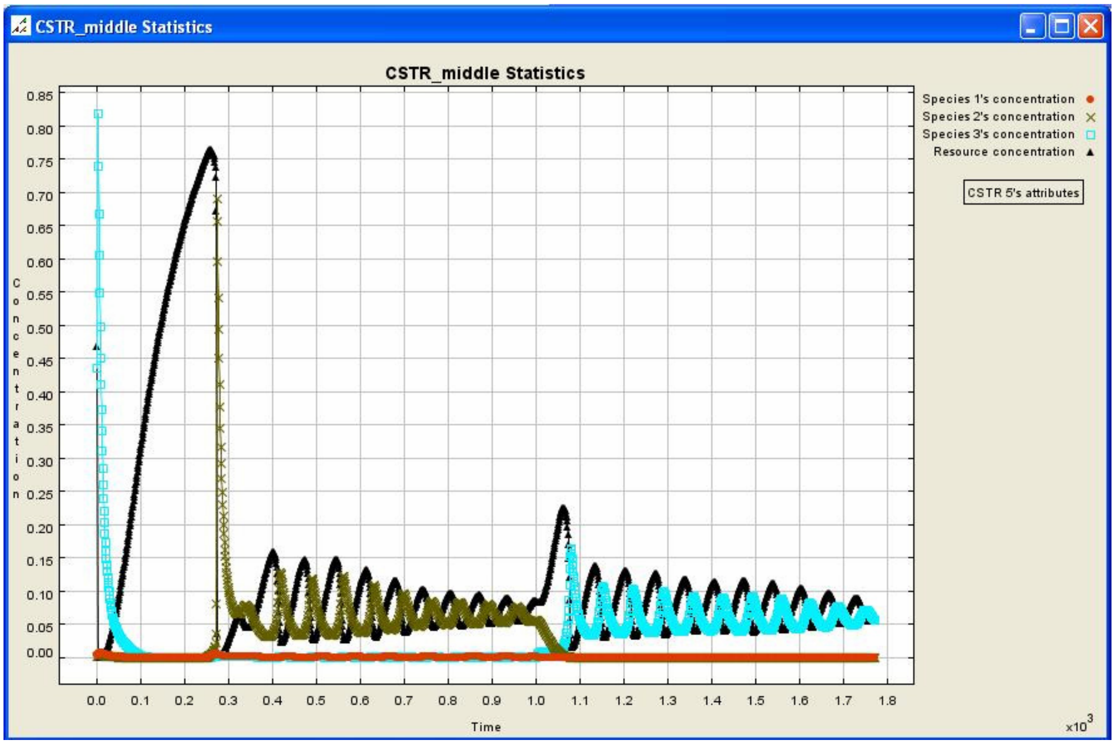

FIGURE 5 Resource and species concentrations in the center (target) reactor

CONCLUSIONS

An adaptive, intelligent agent-based control system has been implemented to control the spatial distribution of autocatalytic species in a reactor network by manipulating the interconnection flow rates. This methodology has been proposed as a real-time alternative to traditional nonlinear control schemes involving predetermined controller configurations or computationally expensive optimization techniques. Controlling the spatial distribution of autocatalytic species in a network of reactors requires the simultaneous manipulation of interconnection flow rates within the system. The multi-agent control system is able to effectively use local rules of interaction combined with an adaptive GA-based approach to intelligently manipulate the network flow rates such that the specified goal is achieved. Product grade transitions can be achieved by setting the overall desired network product qualities, while the local control agents regulate the reactors in a semi-autonomous fashion.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CTS-0325378 of the Information Technology Research (ITR) program. This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Brueckner, S., J. Wyns, P. Peeters, and M. Kollingbaum, 1998, "Designing agents for manufacturing control," in *Proceedings of the 2nd AI & Manufacturing Research Planning Workshop*, Albuquerque, NM.

Christofides, P., and N. El-Farra, 2005, *Control of Nonlinear and Hybrid Process Systems*, Berlin, Germany: Springer-Verlag.

Cohen, S., and A. Hindmarsh, 1994, *CVODE User Guide*, Lawrence Livermore National Laboratory (LLNL) Report UCRL-MA-118618.

Collier, N., T. Howe, and M. North, 2003, "Onward and upward: The transition to Repast 2.0," presented at First Annual North American Association for Computational Social and Organizational Science Conference, Pittsburgh, PA.

Jennings, N., and S. Bussmann, 2003, "Agent-based control systems," *IEEE Control Systems* **23**:61–73.

Mitchell, M., 1998, *An Introduction to Genetic Algorithms*, Cambridge, MA: MIT Press.

Pancs, R., and N. Vriend, 2003, "Schelling's spatial proximity model of segregation revisited," Working Paper No. 487, Department of Economics at Queen Mary, University of London.

Parunak, V., and S. Clark, 1997, "Toward the specification and design of industrial synthetic Ecosystems," in *Proceedings of the Fourth International Workshop on Agent Theories, Architectures, and Languages (ATAL '97).*

Tatara, E., I. Birol, F. Teymour, and A. Cinar, 2004, "Static and dynamic behavior of autocatalytic replicators in reactor networks," *Industrial and Engineering Chemistry Research* **43**:3972–3993.

Tatara, E., I. Birol, F. Teymour, and A. Cinar, 2005, "Agent-based control of autocatalytic replicators in networks of reactors," *Computers & Chemical Engineering* **29**:807–815.

Wooldridge, M., N. Jennings, and D. Kinny, 1999, "A methodology for agent-oriented analysis and design," in *Proceedings of the Third International Conference on Autonomous Agents (Agents '99)*, Seattle, WA.

DYNAMIC SOLIDARITY AMONG AGENTS: AN ARRAY LANGUAGE IMPLEMENTATION

V.S. MELLARKOD, Texas Tech University, Lubbock, TX, and
Argonne National Laboratory, Argonne, IL
C.M. MACAL,* Argonne National Laboratory, Argonne, IL

ABSTRACT

Several tools are currently available for building complex agent-based models (Repast, Swarm, NetLogo, *Mathematica*, etc.). In this paper, we introduce the array programming language J (Thomson 2001; Peele 2005) as one of the languages for building complex systems models. The mathematical approach of the language allows fast computations and efficient programming. J is a scripting language that uses very few lines of code to write complex programs. The idea is to use J as a rapid experimentation language to test and code specifics of a model during the design phase. We implemented an agent-based model of "dynamics of solidarity among occupying authorities and counter authorities" in J. Here we explain the specifics of implementing the model and focus on how array programming makes it simple to code. We implemented a preliminary model in another mathematical language, *Mathematica*, before using J. In this paper, the model is explained briefly, and the implementation details are presented by considering modeling specifics from both languages (*Mathematica* and J).

Keywords: J, computational mathematics systems, agent-based model, array programming

INTRODUCTION

We implemented an agent-based model of "dynamics of solidarity among occupying authorities and counter authorities" in J. This paper explains the specifics of implementing the model and focuses on how array programming makes it simple to code. An early version of the solidarity dynamics model was implemented in *Mathematica* language. This version was used as a reference model for building the model in J. Several mechanisms from the model were implemented and used during the design phase to discuss probable strong points and drawbacks of the mechanisms. This paper gives a brief introduction to the solidarity dynamics model, then introduces *Mathematica* and J, two mathematical array programming languages, to build agent-based models, and discusses J as a promising language for building early prototypes of models.

SOLIDARITY DYNAMICS MODEL

The baseline solidarity dynamics model is a simple model that attempts to capture various dynamics involved when competing authorities try to win the support of an occupied

---

* *Corresponding author address*: Charles M. Macal, Decision and Information Sciences Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439; email: macal@anl.gov.

public. This section gives a brief description of the model; for a detailed description, see Ruby et al. (2005a,b). The model defines two types of actors: authorities and individuals (also referred to as agents). Authorities apply sanctions — in the form of rewards and punishments — to individuals in an attempt to shape public support. Depending on the sanctions, the individuals are affected by economic, social, emotional, and psychological factors, which, in turn, shape their support toward the authorities.

Authorities

The baseline model assumes the presence of two authorities in the occupied land: an occupation authority (OA), which represents the occupier, and a counter authority (CA), which represents a resistance group. Although, in principle, there may be many CAs, the baseline model makes a simplifying assumption and posits the presence of only one CA. Both the OA and CA engage in a mix of coercive and noncoercive strategies (punishments and rewards) to shape support within the occupied public. Rewards may be thought of as various material benefits, such as contracts, jobs, and schools. Punishments may be thought of as detention, injury, or the destruction of resources.

Individuals

The individuals, or agents, of the model are the occupied public. These individuals have two attributes associated with them: valence and alignment. Valence represents an individual's emotional and/or inner feelings toward the authorities. Alignment indicates an individual's outward show of support to the authority. Both the alignment and valence range between values of –1 and 1. A value of 1 represents an extreme liking of the particular authority and –1 denotes extreme dislike. For simplicity, individuals' have only one alignment and valence value. This value represents the intended preference of the agent toward the occupation authority. The negation of this value is understood as their preference toward the counter-authority. The agents also have a static social network of neighbors, representing family/kinship and they have an emotional internal valence toward each of them also ranging between –1 and 1. If an agent's valence toward a neighbor is greater than zero, then either the agent likes the neighbor or dislikes the neighbor.

Executing Sanctions

Selection of a sanction policy by an authority requires a consideration of two elements. The authorities first decide which method they want to use to reward the agents. They have three options to choose from: friends, foes, and egalitarian methods. The friends method focuses on applying higher rewards (punishments) to individuals with higher alignment toward the authority while punishing those with lower alignments. In the foes method, which is the opposite of the friends method, the individuals who show less public support are the focus, and they are rewarded (punished) to attain their support, while friends are taken for granted and rewarded less. In the egalitarian method, all agents are given equal attention, and everyone is rewarded (punished) equally. After selecting the method of sanction, the authorities need to select a point (called Y-spline) on an alignment continuum that divides the reward and punish zones. Though Y-spline can be anywhere on the continuous scale of [–1, 1], the baseline model accepts only three strategic locations. The Y-spline can be among one of –0.5, 0, or 0.5. After the two elements have been selected by the authority, it rewards and punishes the agents by using the each agent's public opinion expressed by its alignments. For instance, if OA selects the friends method along with a 0.5 Y-spline, the agents with alignments greater than 0.5 are rewarded in proportion to their alignment. The higher the alignment, the larger the reward. The agents with an alignment of less than 0.5 are punished in proportion to their alignment.

Agents' Response Mechanisms

The agents respond to the sanctions of authorities via four mechanisms: (1) emotional response, (2) rational choice or expected future mechanism, (3) social conformity, and (4) dissonance constraint. With regard to the first mechanism, the sanctions of an authority on an agent and its neighbors affect the internal emotion (valence) of the agent toward the authority. This change in valence can be positive or negative toward the authority. An agent's valence toward an authority is changed positively when the agent is rewarded or when its friends (neighbors with positive valence or "liked neighbors") are rewarded or, alternatively, when its enemies (neighbors with negative valence or "disliked neighbors") are punished. Similarly, the agent's valence toward the authority changes negatively if it or its liked neighbors are punished or its disliked neighbors are rewarded. The change is proportional to how much the neighbors are liked or disliked. The equations governing the net change of valence toward the authorities use a gap mechanism that restricts the new valence to the range [–1, 1].

The agent's second mechanism governs the change in public opinion of the agent toward the authority. The agent infers the net benefits and costs of the agents that are similar in alignment, then rationally calculates the best possible public opinion that would give higher rewards in the future, and change its alignment accordingly. The similar agents are identified by using an alignment band percentage (currently 5%), which denotes that the agents whose alignments are above or below 5% of a given agent are similar. Researchers can change the alignment band percentage initially.

The third mechanism captures the agent's social conformity issues. The agent adjusts its alignment according to the alignments of its neighbors. The conformity mechanism generates a process by which the alignment choices are imitated, propagated, and/or disseminated throughout the network.

The fourth mechanism is a preference falsification mechanism and constrains the extent to which an agent tolerates the dissonance resulting from varied alignment and valence toward an authority. Dissonance refers to the psychological stress experienced by an agent whose public alignment diverges from its private emotional valence. The dissonance constraint mechanism adapts alignment to valence, or vice versa, depending on the actor's location on the alignment spectrum. When valence is adapted to alignment, it is interpreted as an appreciation of benefits and/or a conformity effect. When alignment is adapted to valence, it is in circumstances where valence is regarded as more authentic and thus less susceptible to adjustment. If the discrepancy is slight, the adjustment is minimal, but the greater the dissonance, the stronger the correction.

MATHEMATICA AS A LANGUAGE FOR ABMS

Mathematica (Wolfram 2003; Wolfram Research, Inc. 2005) is an example of a computational mathematics system (CMS), which allows a user to apply powerful *Mathematica*l algorithms to solve mathematics problems through a convenient and interactive user interface. MATLAB (MathWorks 2005) is another example of a popular commercial CMS. Each CMS has its own limited-syntax programming language, supplies a wide range of built-in *Mathematica*l algorithms and libraries of programs, and is structured in two main parts: (1) an interactive user interface that allows users to submit programs for execution and (2) the underlying computational engine, or kernel, usually written in C, that performs the computations. The CMS programs are translated into C for execution in a separate step that is unseen by the user. Unlike conventional programming languages, CMSs are interpreted rather than compiled, so there is immediate feedback to the user, but some performance penalty is paid. The powerful features of CMSs, their convenience of use, their need for the user to learn only a limited number of instructions, and the immediate feedback provided to users make CMSs good candidates for developing agent-based social simulations. CMSs tend to be array processing systems; that is, data are stored internally in contiguous arrays, and operations are done over large segments of the data with single statements. A full discussion of the application of computational mathematics systems to social simulation is given in Macal (2004).

Mathematica is also a fully functional programming language with symbolic processing in addition to numeric processing capabilities. In contrast, strictly numeric processing languages (such as MATLAB or J) require that every variable have a value assigned to it before it can be used in the program. *Mathematica*'s symbolic processing capabilities allow programming in multiple paradigms, either as alternatives or in combination. Programming paradigms include functional programming, logic programming, procedural programming, rule-based programming, and even object-oriented programming. *Mathematica*'s main data type is that of an expression — a generalization of an array. An expression is a data type with a head and a list of arguments in which even the head of the expression is part of the expression's arguments.

Mathematica has been used extensively for modeling social systems and social agents. See Gaylord and D'Andria (1998) to learn about an application of *Mathematica* to model a number of classic social simulations by using a cellular automata approach. Canty (2003) uses *Mathematica* for a number of classic game theoretic models of strategic interaction. The style of agent-based social simulation used here is much different than what has previously been done when *Mathematica* was used for social agent simulation. It consists of a two-step process: (1) defining agents as abstract data types, which are independent of the particular implementation language, and (2) defining functions and methods that operate on the agents and data types defined in the model. This modeling approach is similar to that taken in object-oriented programming. Maeder (2000) provides an extensive discussion of abstract data types and object-oriented programming in *Mathematica*. This approach separates the modeling (i.e., the model design and specification, Step 1) from the programming implementation (Step 2). For example, the design from Step 1 could easily be captured and communicated to others via a unified modeling language (UML) representation (Booch et al. 1998). This could, in turn, be the basis for programming the model in any object-oriented language, such as Java, C++, or Python, or in a variety of object-oriented, large-scale agent-based modeling toolkits, such as Repast (Collier and Sallach 2001).

In the general modeling approach based on abstract data types, an agent is represented explicitly as an expression that includes a head named agent, a sequence of agent attributes, and a list of the agent's socially connected neighbors:

```
agent[sequence of agent attributes,
 {neighbor 1,…, neighbor i,…, neighbor n}
];
```

The list of neighbor references in the agent expression consists of pointers to the expressions for the agent's neighbors. Agent attributes and pointers could be numeric, strings, or symbols. Social mechanisms between agents, including interactions mediated by the social network, are defined in functions that operate on the agent expression. Access to an agent's neighbors and attributes (including the neighbors of the neighbors) is provided by the list of pointers to the agent's neighbors. Dynamic social networks, which are networks that are formed and change during the simulation, are implemented by manipulating the list of neighbors during the simulation on the basis of the current agent states and the environment.

For the Occupational Dynamics Model, the population of individuals in the model is represented by a list (a list is denoted by {...}) of individuals. Each individual consists of a pattern expression, that is, an abstract data type with head individual and named patterns (indicated by the underscore "_") denoting the individual's attributes:

```
population = {
individual[
        individualIndex_Integer,
        individualID_String,
        coord[loc:{x, y, z}],
        resources_,
        culturalAttL:{
                {ethnicity, ethnicitySalience},
                {religion, religionSalience}
                },
        networkL:{
                networkKinL:{…, {individualID, valence},…},
                networkFriendsL{…, {individualID, valence},…},
                networkSpatialL{…, {individualID, valence},…},
                },
        authorityL:{…, {authorityID, valence, alignment},…}
                }
        ],…,
individual[…]
};
```

where

individualIndex_Integer = the individual identifier number, an integer;

individualID_String = the individual identifier name, a string;

coord[loc:{x_, y_, z_}] = the individual's location in three-dimensional (3D) space;

resources_ = the individual's resource level (assumed to be between 0 and 100, where 0 indicates death and 100 indicates maximum wealth or health);

ethnicity_ = the individual's ethnicity (a symbol);

ethnicitySalience_ = the individual's relative strength of ethnic identification (assumed to be between 0 and 1);

religion_ = the individual's religion (a symbol);

religionSalience = the individual's relative strength of religious identification (assumed to be between 0 and 1); and

networkL = a list of the social networks in which the individual is a member (e.g., networkKinL is the individual's list of kin and the valences that it holds toward them, networkFriendsL is the individual's list of friends and the valences that it has toward them, networkSpatialL is the individual's list of spatial neighbors and the valences it has toward them, and authorityAlignL is the individual's list of authorities that affect it and the individual's valences and alignments toward each).

The authorities in the model are represented by a list of authority expressions, that is, an abstract data type with head authority. Each authority consists of a pattern expression, that is, an abstract data type with head authority and named patterns (indicated by the underscore "_") denoting the authority's attributes:

```
authorities = {
authority[
        authorityID_String,
        coord[loc:{x_,y_,z_}],
        actions_,
        alignAttractor_,
        coerceIndex_,
        coerceNum_
        ],…,
authority[…]
};
```

where

authorityID_String = the authority identifier; a string;

coord[loc:{x_,y_,z_}] = the authority's location in 3D space;

actions_ = a list of the possible actions by the authority and the associated colors that individuals who are the subject of that action (actions consist of the symbols coerce, convince, or deter) take on when displayed;

alignAttractor = the position in alignment space of the authority;

coerceIndex and coerceNum = coercison parameters (the authority coerces up to coerceNum the number of individuals in the population each period that have an alignment index below the coerceIndex threshold value).

The world is then created by joining the list of the population and the list of authorities into a single expression with head society by the statement:

```
world = society@@Join[authorities, population]
```

The world is the abstract data type society with attributes that consist of the two authorities — Authority and InsurgentAuthor — and the full population of individuals. For example:

```
society[
    authority[Authority, coord[{1, 2, 20}], {{coerce, GrayLevel[0]}, {convince,
        RGBColor[1,0,0]}, {deter, RGBColor[1.,0.84,0.]}}, –1, 0, 20],
    authority[InsurgentAuthor, coord[{26., 26., 20}], {{coerce, GrayLevel[0]},
        {convince, RGBColor[1,0,0]}, {deter, RGBColor[1., 0.84, 0.]}}, 0, 0, 10],
    individual[1, AB, coord[{1, 2, 13.2}], 43.8, {{1, 0.51}, {1, 0.82}}, {{{DG,
        0.981197}, {KK, 0.95}, {TN, 0.57}, {YP, -.71}}, {}, {}}, {{Authority,
        0.53, –1}, {InsurgentAuthor, -0.36, 0}}],
    …,
    individual[4, AY, coord[{1, 25, 11.0}], 37.4, {{0, 0.69}, {–1, 0.16}}, {{{OM, –
        0.18}, {YN, 0.70}, {ZN, –0.88}}, {}, {}}, {{Authority, 0.62, –1},
        {InsurgentAuthor, 0.15, 0}}]
    ];
```

The simulation model updates the state of the world at each time-step. Various functions (not shown here) are associated with individual and authority expressions, similar to get and set methods in object-oriented programming. For example, the getResources[i_individual] function takes an individual as an argument and returns the resource level of the individual. Each time period, the state of the society is updated by updating the states of all the individuals and authorities in the society. The updateWorld[s_society] function takes the world as input and applies a chain of functions (social mechanism) to the individuals and authorities in the world, returning an updated world:

world = updateWorld[world]

This single *Mathematica* statement updates the entire society by using the user-defined functional operator updateWorld. To do this, updateWorld calls the function actionPublic[a_authority, world_society, action_String], which takes an authority (i.e., the world) and an action as input and applies the action to the population, updating the states of the affected individuals in the population in the process. An entire simulation can be run by using the functional programming construct Nest:

worldList = NestList[updateWorld, world0, numGen];

where world0 is the starting world to which updateWorld is recursively applied for a total of numGen times (the simulation length) to return worldList (i.e., a list of all of the updated worlds computed during the simulation).

J AS A LANGUAGE FOR ABMS

A baseline solidarity dynamics model is implemented in an array programming language called J (Thomson 2001, Alvord and Thomson 2002, Rich 2004, Hui and Iverson 2005). The *Mathematica* implementation was used as a reference from which the design for the model was drawn. While one of our goals was to implement the model, another was to check whether the J language is suitable for building agent-based models. The main advantage of using this language is that it has a rich set of primitives (functions) that can be used to write complex programs in fewer lines of code. The idea, therefore, is to use this language for building simple and small prototypes of a model while the mechanisms of the model are actively under development. Such speedy implementation would help in eliciting an intuitive understanding of the mechanisms and thereby a recognition of the various dimensions and aspects of particular mechanisms. Unlike the impact that results after waiting for several days (or weeks) to implement a mechanism, if the modelers implement the mechanisms while actively discussing them or soon thereafter, the impact of the implementation, although somewhat primitive, might help advance the model in the right directions.

From our experience in implementing this model, we found that J can be successfully used to build ABMS applications and would naturally fit in as a candidate for building prototypes of the model during the early stages of its development. This section deals with the introduction of the language and its strong points with regard to building agent-based models. The next section deals with the implementation details of the solidarity dynamics model in J.

J is a mathematical language with its roots from APL. J is a scripting language and could even be cryptic for a new learner. Although the learning curve for the language is very steep, the essence of the language is simple and consistent. Most users learn to use parts of the language and can manage quite well with that knowledge. One of the authors learned the language and was able to implement the classic heat bugs agent simulation model within a week. This shows how easy it is to pick up at least parts of the language and still manage to build real and interesting applications. Like any language, there are advantages and disadvantages to using J. It is unlike any procedural language, although it has several constructs that come from procedural programming languages. For instance, loops like *if, for, while,* and *case* are all supported.

As mentioned earlier, there are several built-in functions called *primitives* that can be used to write complex code in few lines. For instance, to calculate the mean of an array of numbers, we can use the primitives +/, and %, where +/ is defined as adding all numbers from the input, # is a function that returns the number of elements in the input, and % is the division operator. Therefore, *(+/ % #) 2 3 6 7 10* calculates the mean of numbers 2, 3, 6, 7, and 10. Finding all input numbers whose value is less than five can be written as *(] #~ <&5) 9 8 3 2 1* returns 3 2 1. J works with arrays efficiently; for instance, if the input s for finding the average of numbers is a two-dimensional array, then *(+/ % #) s* gives the row means, *(+/ % #)”1 s* gives the column means, and *(+/ % #) ,s* gives the average of elements in the array s. (The comma operator before s “unravels” the array and makes it one dimensional). Here is another example for stressing the terse code feature: The heat bugs model contained one file with 70 lines of code, apart from variables defined and comments. The main mechanisms coded were not more than 10 to 15 lines altogether. The rest was written toward building graphical tools for the model.

Though learning J at the beginning is difficult, writing functions would come naturally with the definitions of the primitives. One main advantage of using these primitives over writing self loops and code is that the primitives have been implemented very carefully, with space and time efficiencies kept in mind. For instance, writing a sort function to sort a sequence of numbers y is as simple as writing */:~ y* and results in sorting the numbers y. The input y can be a number, a set of numbers, an n-dimensional array, or a sparse array. The nature of the input is considered to select the algorithm used in the sort function. The algorithm almost always returns in linear time. This feature lets the user to think of the problem in detail, without worrying about computational-efficiency-related problems in implementation.

Designing a model to implement in J lets the user think about the problem in different ways from the standard methods used in other procedural languages. To implement a mechanism, it is important to think of the mechanism as a whole and not in parts. The data structures are different from the standard object-oriented sense. All data expressed are in arrays and arrays of boxes. A box can be seen as a way of putting multiple data types together, like structures in C. Mostly, an array would be composed of one or more attributes of agents/objects put together, which is unlike an object-oriented paradigm, in which each object contains its attributes, and an array of such objects is built. J has classes and objects to support object-oriented programming, but its efficiency lies is using/working with a bulk of attributes/arrays at the same time.

J allows easy connection/interaction with other programming languages, allows easy handling of files, has a project builder that can be used to build huge applications by using an interface, and has a form builder that can be used to build graphical user interfaces (GUIs; see www.jsoftware.com). While building a model is one part of developing a complete system, the other part is to understand the mechanisms involved and experience the model visually; agent-based models depend heavily on visualization techniques. J’s form builder could be effortlessly used to design user interfaces and other visualization tools. All the GUIs for the solidarity dynamics model were built by using the form builder. Working with charts and plots, recording data, and calculating statistical measures are straightforward in J.

As noted earlier, J works with huge multi-dimensional arrays and sparse arrays. It is quite efficient at working with them as a whole instead of working with each of the elements in the arrays separately. For this reason, any model that requires huge calculations in bulk can be easily handled in J. Models that work with minute details and whose mechanisms cannot be expressed clearly without various “if” loops are less efficient in J. For such models, J can be used to build early prototypes, which are easy to build but whose final implementation would be efficient (in time) by using a different language.

The answer to the question “Can J be used as an effective language for agent-based modeling?” is “Yes.” For models in which agents or groups of agents are given common tasks, goals, etc., J can be used and will work efficiently. During the design phase, when there are multiple mechanisms in question, J can be used to build early prototypes in much less time, to see the behaviors of the mechanisms or to see their combinations in general. The ‘sys_time’ function in J can be logically used to implement the step function in the simulation. The visualization tools (charts, plots, GUIs, etc.) can be built easily, which is a necessity in simulation models. Data can be stored to files, and even statistical data can be calculated by using the agent data and stored to be directly used in an analysis. Input data for the simulation can be directly imported by using Excel or other data formats. The next section talks about implementation of the solidarity dynamics in particular.

IMPLEMENTATION IN J

The baseline solidarity dynamics application was implemented in J (Ruby et al. 2005). This section discusses the details of implementing the model. The initialization phase consists of extracting user inputs from the Initialize GUI and setting up the model. The initialization interface is shown in Figure 1. Several parameters can be changed from the default values shown on the interface. The parameters and their default values and descriptions are given in Tables 1 and 2. The main data structures involved with the individuals are shown in Table 3.

The value INum is the number of individuals in the model. The network neighbor relation is symmetric (i.e., if agent *x* is a neighbor of *y*, then *y* is a neighbor of *x*). The other data structures relate to rewards and punishments received by individuals from both the authorities. All the data structures are arrays. J can work with arrays very efficiently. Since an array always

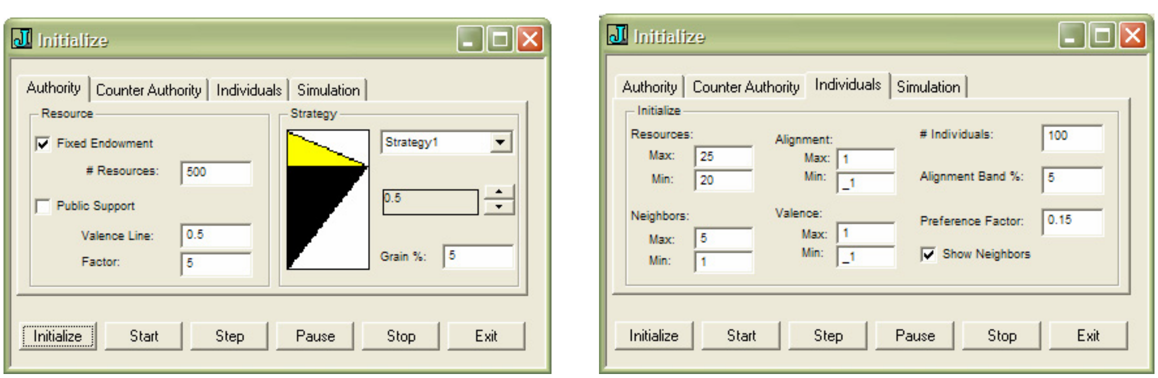

FIGURE 1 Authority and individual initialization interfaces

TABLE 1 Initialization defaults for authorities

NB. Initialization defaults
NB. Parameters that can be changed on the authorities interface
Fixed resources = 500 NB. Authority’s fixed endowment – number of resources (every tick)
Fixed endowment = 1 NB. Default is receiving fixed endowment every tick
Public support = 0 NB. Authority public support; default is NOT receiving public support
Valence line = 0.5 NB. If public support is used, default is valence line = 0.5
Valence factor = 5 NB. If public support is used, authority receives resources = valence factor * # individuals above the valence line
Strategy number = 0 NB. Main strategy selection (foe, friend, egalitarian)
Y-spline = 0 NB. Spline at which strategy splits between reward and punishment (0.5, 0, –0.5)
Grain = 5% NB. Percent grain size used when executing strategy

TABLE 2 Initialization defaults for individuals

NB. Parameters that can be changed on the individual interface
individuals = 100 NB. Number of Individuals
Max resources = 25 NB. Maximum initial resources given to an individual
Min resources = 20 NB. Minimum initial resources given to an individual
Max alignment = 1.0 NB. Initial maximum alignment toward authority
Min alignment =–1.0 NB. Initial minimum alignment toward authority
Max valence = 1.0 NB. Initial maximum valence toward authority
Min valence = –1.0 NB. Initial minimum valence toward authority
Max # neighbors = 5 NB. Maximum number of neighbors
Min # neighbors = 1 NB. Minimum number of neighbors
Alignment band % = 5 NB. % grain for alignment bands; used during rational choice mechanism
Show neighbors = 1 NB. Show the neighbor network during model execution (GUI)
Preference factor = 0.15 NB. Alpha factor used during preference falsification (constraint factor)

TABLE 3 Data structures for individuals

AlignmentIdx	Alignment index is an INum x 1 array of individual alignments
ValenceIdx	Valence index is an INum x 1 array of individual valences
Resources	Resources is INum x 1 array of individual resources
NeighborsL	Neighbors list is 1 x INum boxed array of neighbors; each box in the array consists of an agent’s neighbors list
NeighborVal	Neighbor valence is 1 x INum boxed array of valences toward neighbors

contains the same number and type of elements in each row, sometimes it is not very efficient to represent structures like network neighbors. The simplest form of representing the network neighbors by using an array is an adjacency list, but it would require a lot of memory. Boxed arrays in J are used to represent dissimilar numbers of elements and even elements of different types. Each box is a unit and is a container. The boxed array is a list of containers. The containers can contain any type of element. For instance, a string, a number, or a class can be in different containers in the same boxed array. This is analogous to structures in C programming language. In the neighbors list boxed array, each box *b* contains a set of network neighbors of agent *b*.

After initialization, the step function is executed at every tick. (See flowchart in Figure 2.) Depending on the initialization parameters, the OA and CA get resources at every tick. The total resource for an authority is calculated as the sum of fixed endowment and the resources from public support:

Total resources = fixed endowment + public support.

The fixed endowment is set at initialization. The default is 500 for the OA and 150 for the CA. The public support is calculated as follows:

Public support (OA) = valence factor * number of Individuals with valence above valence line.

Public support (CA) = valence factor * number of Individuals with valence below valence line.

The valence factor and valence line are set at initialization. To find the number of individuals with valence above a particular valence *v* in J, we use (+/ *v* < valence) and multiply by the valence factor to give the public support for the authority. This code shows how simple and concise J code is; it removes the need for loops in most cases. The OA and CA strategies are selected by the user at initialization time. The user can also change its strategy by using the interface at any time during the simulation. The execution of OA and CA strategies occurs in a random order (see Figure 2).

For the execution of the strategy, the continuum space of alignment [–1, 1] is divided into zones by using alignment grain % (which is set at initialization). The area under each zone is calculated, and the resources of the authority are divided among these zones with a ratio proportional to their respective areas. For each reward zone, *Rz* (reward allocated to the zone) is distributed to individuals in that zone randomly in chunks of *rUnits*.

Total rewards *R* = total resources * reward area/total area

Reward for Zone *Rz* = *R* * zone area/reward area

Multiple units *Mu* = floor (reward units/#individuals) where reward units = *Rz/rUnit*

Single units *Su* = floor (reward units – multiple units * # individuals)

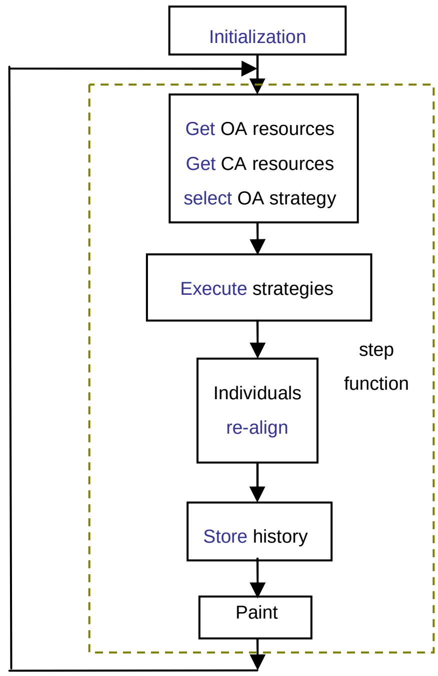

FIGURE 2 Solidarity dynamics application flow chart

Multiple units are distributed among all the individuals (i.e., each individual is given *Mu* resources), and single units are given to random individuals (i.e., each individual is given one extra resource). This is like distributing reward units/packets to individuals one at a time until the packets are consumed. When an individual gets some reward units, the individual's resources increase by this number of units. Punishments follow the same formula, except the resources of the agent are decreased by the calculated amount. Both authorities follow the same rules in applying sanctions. The strategies, the alignment grain percent, and the area of reward and punishment zones (can) differ between the authorities. The *rUnits* represent the minimum number of reward units that an individual can get (and consider as reward). In J, calculations of rewards and punishments are implemented in a band-wise fashion (i.e., the set of individuals of a particular band are found, and all are given the reward or punishment at the same time). "At the same time" does not mean that the code runs parallel but rather that code is written for a band without the use of loops.

Individuals realign their orientations by using the four mechanisms defined earlier. As specified before, the valence orientation of an agent depends on the rewards and punishments received by the agent and its network neighbors. For instance, if an agent i likes a neighbor j, and j gets rewarded by authority a, then the equation used to calculate the new valence of i toward authority a is $nvia = via + vij * (1 – via)$, where vij is the valence of i toward j. Similarly, if neighbor j is punished, then the new valence is calculated by using $nvia = via – vij * (1 + via)$. The other equations follow the same principles. Since the number of neighbors, the valences toward them, and their rewards and punishments are different for each individual, this part of the code in J looks like a procedural programming code as written in many languages.

In the rational choice mechanism, an agent takes other agents' rewards/punishments into account, calculates the future costs/benefits of its public position, and moves rationally toward the maximum benefit zone. In the baseline model, the direction of movement is chosen, and the movement in that direction is random within a limit. The direction is chosen as follows:

$\Delta\uparrow$ = average net change of resources for the band above the current alignment of self,

$\Delta\downarrow$ = average net change of resources for the band below the current alignment of self, and

$\Delta\equiv$ = average net change of resources for individuals having the same current alignment as self.

Alignment band % is used to determine the width of the bands considered and can be changed at the start of simulation.

$\Delta\uparrow=$ = total change of resources for individuals in the band above "divided by number of individuals" in the band.

The other alignment bands, are calculated similarly. The alignment band selected is chosen by max ($\Delta\uparrow$, $\Delta\downarrow$, $\Delta\equiv$), and the directions are $\uparrow$, $\downarrow$, or $\equiv$ (move up, down, or same in alignment), respectively. After an agent chooses a direction to move, the alignment band % is again used as the limit to which the agent can move randomly in that direction. In the J implementation, the net change of resources for all individuals is calculated once, and it is used for all the individual calculations. Since the agent's similar bands above and below depend on its alignment, this last calculation works on every agent separately.

Social conformity effects of an agent are calculated by using two components: the valence toward neighbors (denoted by v_{ij}), and the distance between self and neighbors (denoted by daij), where i is an individual and j is a neighbor. The effects follow the principle that an individual imitates a friendly neighbor and tries to become unlike a neighbor it dislikes. The strength of an imitation effect is directly proportional to v_{ij} and d_{ij}. The more an agent likes a neighbor, the more strongly it wants to imitate the neighbor; the more dissimilar their public opinions (alignment) toward an authority are, the stronger is the agent's need to decrease this dissimilarity. Similarly, the more an agent dislikes a neighbor, the more strongly it wants to be unlike the neighbor; and the more similar their opinions toward an authority are, the stronger is the agent's need to become unlike the neighbor. All these effects are captured in the following formula:

$$\Delta ia = (0.5/\#\text{ neighbors}) * \{\sum \alpha * signd * vin * [(daij * ind) + (2 - daij) * (1 - ind)]\},$$

where

Δia = change in alignment of individual i toward authority a;

vin = valence of self toward neighbor;

α = the imitation factor set to 0.5 in the model;

dain = the absolute difference in alignments of self and neighbor, calculated as daij = |Align of Neigh – Align of self |;

signd = the sign (positive or negative) of the difference in alignments, where signd = 1 if (Align of Neigh – Align of Self) is ≥0 and signd = –1 if (Align of Neigh – Align of self) is < 0; and

ind = an indicator function that is calculated as ind = 1 if valence toward neighbor vij is ≥0 and ind = 0 if otherwise.

The implementation works for every individual, since each has a separate set of friends and enemies. For a single agent, finding the friends from its neighbor list n with valences nval, is (0 < n {nval). To find daij, we can use (a -~ n{Alignment), which gives the difference in alignments between neighbors and self (a). Similarly, the other equations can be stated in simple terms in J. The code looks cryptic to a newcomer to J, but a user can easily read and understand it.

The baseline uses a quadratic model for the preference falsification mechanism. When there is a large dissonance between valence and alignment of an individual, then the alignment or valence is pulled (changed) to decrease the dissonance. The diagonal valence = alignment line is used as the ideal state preferred by an agent. If an agent does not meet this preference, then either alignment or valence is changed. The change is proportional to the square of the perpendicular distance between the agent's current location in alignment-valence coordinates and to the diagonal alignment = valence. Valence is changed when an individual's alignment is >0, and alignment is changed otherwise. The formula for calculating the new valence or alignment is given as follows:

via= = via + sign * I *(α_p/2) * (vminusa)^2 if alignment is >0 and

aia = aia + sign * I *(α_p/2) * (vminusa)^2 if alignment is <0,

where

vminusa = via – aia, which is the perpendicular distance;

α_p = preference factor, which can be set at initialization and denotes the amount of preference falsification effects seen in the individual; and

I = indicator function equal to 1 for alignment >0 and –1 for alignment <0.

The value sign * I gives the direction of change of valence/alignment and is always toward the diagonal valence = alignment,

where

sign = 1 if vminusa is ≤0 and sign = –1 if vminusa is >0.

The calculations for this mechanism are straightforward in J, as the above equations stated. The values for vminusa for all agents can be calculated with (Valence – Alignment). The sign for all agents can be calculated as (1&+[+[) – 0 < vminusa. The final equation to calculate the new valence for individuals with an alignment of >0 is written as:

via= = via + sign * (α_p/2) * vminusa * vminusa * (0 is <Alignment).

For all individuals whose alignment is >0, the new valence is calculated, and the others have the same valence as before. The calculations for alignment change are similar. The preference falsification mechanism has undergone several modifications. During the design phase, each modification on the mechanism was implemented during the discussions.

The store history function stores information regarding rewards/punishments issued by the authorities to the individuals. This is used for saving input and simulation data to files for further analysis and also for plotting a selected agent's change in alignment, valence, and resources with respect to time. The paint function paints the GUI. The main GUI (Figure 3) is a distribution of individuals on a coordinate plane. The horizontal axis is the valence of the individual toward authority, and the vertical axis is the alignment. The authority considered is the occupation authority. The green box in the figure denotes the occupational authority OA, while the red box denotes the counterauthority CA. The strategies applied by both authorities are shown on their respective sides. The lines connect the network neighbors. The color of the agent denotes the level of resources that the agent possesses.

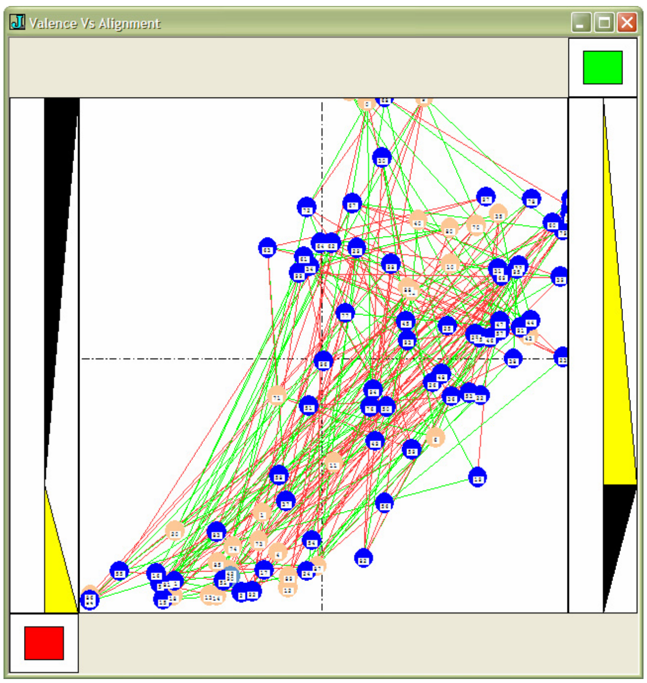

FIGURE 3 Valence versus alignment of individuals

CONCLUSIONS

A new language J was experimented with to assess whether it may provide a particularly appropriate language for building faster prototypes for agent-based models. The language has proven to be successful in building faster and fairly complex implementations and can be used during the design phase, when the modelers can implement mechanisms in their model to get a deeper understanding that ultimately leads them to move in productive directions.

ACKNOWLEDGMENT

This work is sponsored by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Alvord, L., and N. Thomson, 2002, *Easy J;* available for download at www.jsoftware.com (Oct. 2005).

Booch, G., J. Rumbaugh, and I. Jacobson, 1998, *The Unified Modeling Language User Guide*, Addison-Wesley, New York, NY.

Canty, M.J., 2003, *Resolving Conflicts with Mathematica: Algorithms for Two-person Games*, Academic Press.

Collier, N., and D.L. Sallach, 2001, *Repast, REcursive Porous Agent Simulation Toolkit;* available at http://repast.sourceforge.net/.

Gaylord, R.J., and L. D'Andria, 1998, *Simulating Society — A Mathematica Toolkit for Modeling Socioeconomic Behavior*, TELOS/Springer Verlag, New York, NY.

Hui, R., and E.K. Iverson, 2005, *A Note on Programming Style*; available at www.unb.ca/web/transpo/mynet/jpgmstyle.htm (Oct.).

Macal, C., 2004, "Agent-based modeling and social simulation with *Mathematica* and MATLAB," in C.M. Macal, D. Sallach, and M.J. North, eds., *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Maeder, R.E., 2000, *Computer Science with Mathematica*, Cambridge University Press, Cambridge, UK.

MathWorks, 2005, MATLAB home page; available at www.mathworks.com.

Peele, H., 2005, *Mathematical Computing in J: Introduction,* Research Studies Press, Hertfordshire, UK.

Rich, H., 2004, *J for C Programmers*; available at www.jsoftware.com/jwiki/Books.

Ruby, K.G., D.L. Sallach, C.M. Macal, V.S. Mellarkod, and A. Wendt, 2005a, "Occupation Dynamics: Contesting Hearts and Minds," *Proceedings of the North American Computational Social and Organization Science (NAACSOS) 2005 Conference*, South Bend, IN, USA, June 26–28, 2005; available on CD.

Ruby, K., et al., 2005b, "Contesting Hearts and Minds: A Baseline Model of Occupation Dynamics in Military Occupations," in C.M. Macal, M.J. North, and D. Sallach, eds., *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms*, ANL/DIS-06-1, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 13–15.

Thomson, N., 2001. *J: The Natural Language for Analytic Computing*, Research Studies Press, Baldock, UK.

Wolfram, S., 2003, *The Mathematica Book: 5th Edition*, Wolfram Media, Champaign, IL.

Wolfram Research, Inc., 2005, *Mathematica* home page; available at www.wolfram.com.

DISCUSSION

Methods, Toolkits, and Techniques

(Behavioral Modeling Methods, Thursday, October 13, 2005, 10:15–11:45 a.m.)

Chair and Discussant: *Thomas Howe, Argonne National Laboratory*

Integrating Life-like Action Selection into Cycle-based Agent Simulation Environments

Michael North: I'd like to welcome everyone back to the conference. Now we will hear about behavioral modeling methods. The chair and discussant will be Tom Howe, so I'd like to turn things over to Tom.

Thomas Howe: Okay. As Mike said, this section is on behavioral modeling methods, and our first paper is called "Integrating Lifelike Action Selection into Cycle-based Agent Simulation Environments" by Joanna Bryson, Tristan Caulfield, and Jan Drugowitsch from University of Bath.

[Presentation]

Howe: Thanks, Joanna. I have a couple questions, and then I'll open it up. One question is rather trivial. Is the editor for the BOD stuff available or is it proprietary?

Joanna Bryson: Believe it or not, it is part of a contract that will be open-source like everything else. It was just delivered a week ago, and this is my first time playing with it. Basically, it's an alpha mode. If you ask me for it now, I have it on my stick. I'll give it to you now. It should be on the web page soon so that the Boeing people can start playing with it.

Howe: Yes, I think your idea of trying to make complex agents and thicker agents easier is obviously something that most people in this room are probably interested in. Anybody who's tried to do an agent-based simulation is stuck in the world of heatbugs, because they just don't have the programming skills to go beyond that, would probably attest to that. I think that's the motivation behind things like the Mathematica work, Repast Py work, and NetLogo work obviously.

How tightly coupled is the work that you've done here with MASON? I'm sure that there are other toolkits that would be interested in working with it, but it seems that if it's just Jython, it wouldn't be that hard to extract, maybe to a level where it could be used in multiple toolkits.

Bryson: I would honestly say that it's not tightly coupled enough. I wish it were smoother. But no, the whole intention is to find other groups to work with, and the Jython stuff is totally stand-alone. It used to be Python because some roboticists asked for the .NET language basically, and for the Unreal tournament thing. So like I said, the stuff that's running on Unreal Tournament, which is probably our most complicated agent so far, is completely stand-alone. It's not MASON or anything, and this was just our first foray. I am open to talking. Part of the reason

I'm here is that I was hoping there'd be a bunch of MASON developers to talk about the problems we are having, but also anyone else who wants to talk, that's fine.

I'd like to say one more thing though about the complexities. I ran a workshop this summer on modeling natural action selection, and Mick Laber was absolutely brilliant. I recommend him for next year if you want to get a good invited speaker. He just started doing agent-based modeling —taught himself — and programming. He's leading the *Political Science Journal*, and he's the first person to get in with an agent-based model that not only shows you the end points, but one of the very few people — I don't know if Lars has gotten into APSR [*American Political Science Review*] yet, he'll tell us — but also, he also shows the track, so over time he's matching the voters' behavior.

He says, though, that he gets constant flack from his colleagues in political science and economics because he has too many variables in his models. When he saw what the rest of us were doing at this conference, he just completely flipped out. He said, "You know, if I had — what my colleagues say to me is, if I had seven free variables, I could model anything." Now, I think all of us who are programmers here know that that's not entirely true. But I do think we have a big methodological issue to go up against, and certainly we've all seen students shoot themselves in the foot. Certainly, when I go through this code, I think they don't know how to simplify; they haven't got the heuristics down, so I think we will have to do a lot of methodological work if we give people more power. With great power comes great responsibility for methodology.

Howe: Are there any other questions? Thank you very much, Joanna.

Agent-based Control for Dynamic Configuration of Spatially Distributed Networks

Howe: The next speaker is Arsun Artel who will be presenting "Agent-based Control for Dynamic Configuration of Spatially Distributed Networks," which is actually by a whole bunch of people.

[Presentation]

Howe: I have one question at this point. Do you see these techniques as having a wider reach? In what ways do you think that they could be used for control of agent behavior and agent distribution for more general-purpose simulation models?

Arsun Artel: I'm not sure that I get your question, but our design goal was to map those things to real systems. Eventually, we will try to do those, but currently the bottom layer is only a simulator from which we are getting the data. I'm not sure whether I answered your question.

Howe: I guess the types of simulations that you've done had to do with very autocatalytic kind or species distribution.

Artel: Yes.

Howe: Have you thought about other kinds of data and other kinds of behaviors that this would be capable of supporting?

Artel: Yes, but some of the methods are based on heuristics, which are also based on the autocatalytic system. If we change system, if we change the simulator under that agent-based system, we have to also modify those heuristics. Basically, though, the methods will work fine.

Eric Tatara: I'll just add to that. Basically, the goal was to simulate the people who are operating industrial facilities. Essentially, in all industrial facilities, you have human operators. We tried to mimic their behavior and automate it. We basically set out to improve the overall operating efficiency.

You could apply that to anything where human intervention toward another type of system takes place. For example, say you have natural gas pipelines, electrical grids, or ecosystem management. You could even stretch it and say prison populations, where you have a distributed control in terms of the police and things like that. There are a lot of possible application areas.

Howe: Are there any other questions? Okay, thank you very much.

Dynamic Solidarity among Agents: An Array Language Implementation

Howe: The last paper in this section, and the last paper before lunch, is called "Dynamic Solidarity among Agents: An Array Language Implementation" by Veena Mellarkod and Charles Macal.

Veena Mellarkod: I'm Veena Mellarkod from Texas Tech University, Computer Science Department. I've been working with Charles Macal, David Sallach, and Keven Ruby for quite some time now toward building an agent-based model for capturing dynamics and solidarity among agents. Today, I'm going to give some highlights on the implementation. This implementation was done by using an array language called J.

[Presentation]

Howe: Thank you. I have one question, partially because of the time. There has been a lot of discussion today about what it takes to build a rapid prototyping agent-based framework, something that you can actually use to build models very quickly. It looks like you can build things in J in a very small amount of code. As you say, though, there are many things that you have to construct on your own. For this to become a truly viable rapid prototyping tool, what needs to be done by the agent-based community to enhance J in order …

[Editors' note: Tape malfunctioned.]

Invited Speaker

Steven Bankes

SUPPORTING THE MODELING LIFE CYCLE

S.C. BANKES,* Evolving Logic, Los Angeles, CA

ABSTRACT

Computational science has for the most part been focused on creating appropriate representations of problem domains of interest and reporting the results of modeling experiments based on those representations. This focus has caused other equally important issues to be neglected. In order for computational science to achieve its potential, methods and technologies are needed to support the entire modeling life cycle, including model creation, modification, and exploitation. In this paper, I discuss these challenges and provide examples where initial steps have been taken to meet them.

Keywords: Computational science, model exploitation, exploratory modeling, computer-assisted reasoning, robust inference

INTRODUCTION

The idea of support throughout a life cycle has a substantial history. Our current understanding has grown out of a natural but naïve concept of life cycle as a linear series of stages. For nearly any artifact, it seems natural to begin with the concept for the product and proceed to analysis, design, and implementation, arriving finally at a functional system. That system will get used and maintained and finally be retired.

Such a linear process, while intuitively comfortable, suppresses the iterative nature of the design and construction process. And as an ideal, it can lead to many disasters, as it implies a process that contacts the world only at beginning and end. Often, this can result in a product that is rejected by the user when it is fielded, even though users were fully consulted during the specification and design phases.

A notorious version of the linear approach is the waterfall development methodology, which has been used for software engineering and has been blamed for many software engineering failures. Although it is still in use here and there, the waterfall methodology is identified in many texts on software engineering as a discredited approach. Certainly, it cannot be recommended for programming in general, as experience has shown it to lead to systems that do not track shifting requirements and that can be difficult to maintain. There are a variety of methods that are preferred, but a prominent example designed to counter the natural tendency towards a waterfall approach was the spiral development method (Boehm 1988 and 2000).

Spiral development encourages us to consider a cyclic process in which specification, design, implementation, and testing occur repeatedly, with new features appearing and design errors being fixed with each cycle. Related ideas such as rapid prototyping also encourage engagement with the user throughout the development process. It is an implication of such an

* *Corresponding author address*: Steven C. Bankes, Evolving Logic, Inc., 2118 Wilshire Blvd., #423, Santa Monica, CA 90403; e-mail: bankes@evolvinglogic.com.

approach that development is never truly completed, and that the distinction between development and maintenance is not hard and fast.

Computational science has its own version of the waterfall process, and just as with the early years of software engineering, it is viewed as so obvious, it is seldom questioned. Roughly, the stages of a linear process of computer modeling are:

1. Gather all relevant data and theory.

2. Design the model.

3. Implement the model, including debugging and verification.

4. Load the inputs corresponding to some question.

5. Run the model once or a series of times in order to answer the question. If it is a policy question, search for the policy that produces the best model outcomes.

The waterfall approach to computational science has led to many problems, including both specious research and a blindness to opportunities for rigorous research that are overlooked in the search for predictive accuracy. This leads one to ask the question, “What is the equivalent to spiral development for computational science?” A different way to express this need is to say that we must move beyond the focus on models, to provide better support for the activities of modeling. Models by themselves provide no value. The value arises in the context of a web of relationships that must be supported and utilized. These include:

- *Beliefs that motivate and explain the structure of the model.* Of special note are assumptions built into the model that must be noted if inference based on the model is to be rigorous.

- *The experts who are the source of these beliefs and assumptions.* They will often be the model’s authors, but could include others. Since these experts may have beliefs and assumptions that they implicitly hold, this meta-information is also important to rigor in model-based science.

- *The data used in model construction.* These data inform the model just as much as beliefs do, so its provenance is should also be captured. Increasingly, model data will be updated regularly, perhaps with real-time feeds. Technology to make this easy, to provide configuration management and to verify the process, is needed as part of the infrastructure for modeling.

- *The relationship of a model to its users.*

- *The relationship of a given model to alternative models as well as to previous versions.* An audit trail or change history is one example of how this relationship can be made explicit and supported. A deeper need is to represent the differences between models in terms of the assumptions they embody, so that joint inference can be supported.

- *The relationship to other, complementary models.* Great benefits could be achieved if models of different phenomenology could be used together to understand the union of the information they embody. Simply gluing the models together may not be the best means for achieving this, however.

- *Data from cases run on the model.*

- *Analytic materials derived from the model.* Examples are statistical analyses or graphical visualizations.

- *Arguments and conclusions based on the model.*

- *Other software systems and applications to which the model may be linked.* Examples are decision support tools or control systems.

SOFTWARE TO SUPPORT COMPUTATIONAL SCIENCE

An example of a system that provides at least some categories of support for "modeling" is the computer assisted reasoning system (CARs™). The original design for CARs was inspired by the concept of exploratory modeling (Bankes 1993). That is to say, the design makes computational experiment the central concept, and provides various facilities for using models to conduct computational experiments, and for using the results of computational experiments to support reasoning. CARs' design emphasizes providing options for interactive use, and for managing large ensembles of models and cases. It has been used as a basis for implementing methods for robust decision analysis, and has been applied to a variety of planning and decision problems. In the discussion below, examples from the design for CARs and its use on various projects will be provided.

CARs is a Java-based system that links to virtually any type of model, treating it as a "scenario generator." In particular, CARs can be harnessed with models implemented in C++ or Java, or any language that supports sockets or Micosoft's COM. CARs provides a variety of services that can be applied to any model once it has been harnessed, including interactive use of search and visualization to create, explore, compare, and understand very large ensembles of scenarios with ease.

In CARs, a scenario generator (aka model) is encapsulated in an object that augments its input/output behavior with a variety of other services, including databases of cases (input/output pairs) and methods that help other parts of the system customize their behavior appropriately. Software clients for these services can interact with objects of this class (known inside of CARs as a "context") transparent to the means that are being used to run cases. Aside from response time, other objects are free to regard contexts as repositories of every possible case that could be run, virtually representing ensembles of what-if questions that can be of infinite cardinality or infinite dimensionality. Keeping a database of previous runs can improve response time, but does not fundamentally alter the logical requirements for making inferences about an infinite set from a finite sample.

On the basis of software architectures such as that of CARs, a variety of avante garde methodologies can be entertained. For some applications, these methodologies may be more appropriate than a realist strategy of model use that has been so successful for engineering models used for computer assisted design. For example, Figure 1 is taken from work in which a coevolutionary strategy has been employed for robust inference. Here virtual ensembles of plausible future challenges are allowed to evolve in parallel to a virtual space of possible coping strategies. Sampling out of each of these sets is guided by the other, seeking assumptions about the future that are maximally stressing for the leading candidate strategies and strategies that perform well across the challenge landscape.

THE CHALLENGE OF MODEL EXPLOITATION

Computational science can be seen as a sequential (and iterative) process of divergent and convergent thinking. Instead of having only a single image of the future in mind (left side of Figure 2), the idea is to confront uncertainty and recognize as much as possible of the highly diverse futures that are in fact plausible. There are two universes of cases in the center of the figure. The larger one is the virtual ensemble of all cases that might possibly be run, while the smaller is the corpus of cases that actually have been run.

Just generating cases is not enough; the next challenge is to make use of the results — to apply more convergent methods for sense making and for the discovery of insights. Various techniques can be used to accomplish this, including interactive visualization, search techniques to run cases that are most likely to be salient to the problem at hand, and the use of data mining algorithms to extract information from the resulting database of cases. (Also interesting is the hybrid set of algorithms called active learning or adaptive sampling that combine search and data mining properties.)

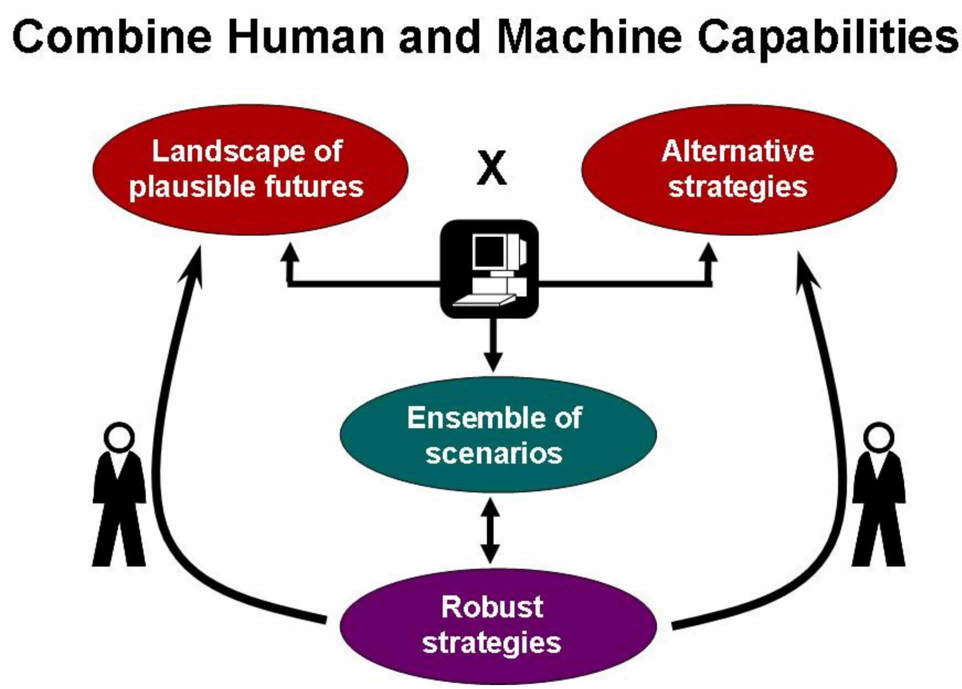

FIGURE 1 Coevolutionary analytic methodology

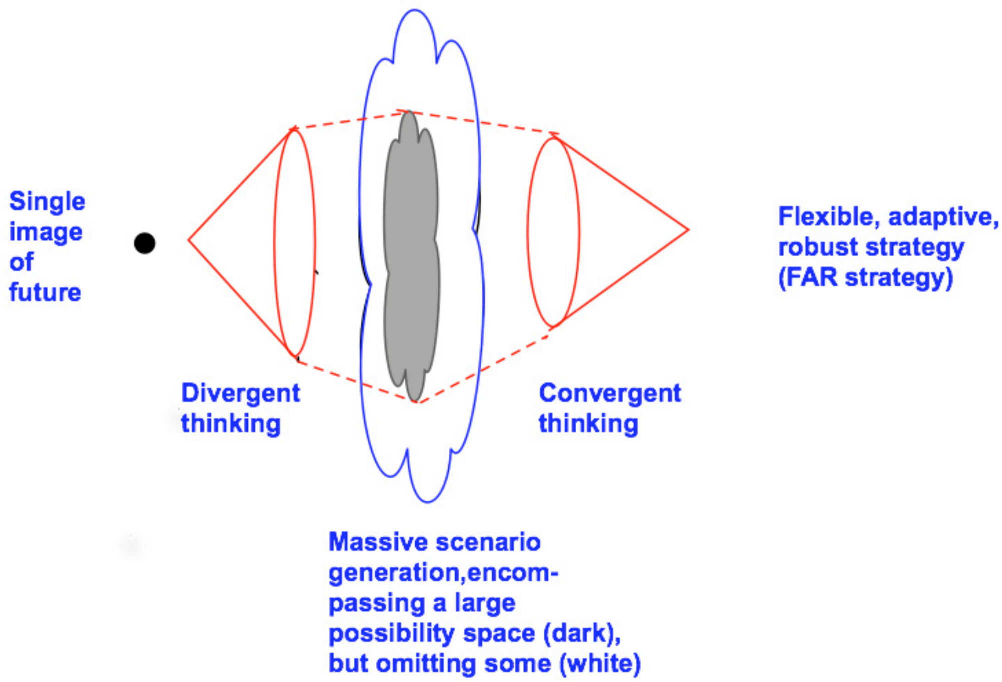

FIGURE 2 Diverging and converging

This approach suggests a series of implications regarding how best to contend with difficulties that have confounded computational science in the past, and opportunities to improve its impact on the future. In the remainder of this paper, I give a few examples that I and colleagues have begun to address.

CONFRONTING THE CURSE OF DIMENSIONALITY

The first of these challenges is the so-called “curse of dimensionality.” In brief, the number of cases required to assess the response of a model rises geometrically with the dimensionality of the input space. In a rectangular region of an N-dimensional input space, assessing the range of response requires 2^N cases if the response is monotonic throughout the region. If the response surface is more rugged, no upper bound on the number of cases exists, and even for the monotonic case, 2^N cases can become impractical for a relatively modest N.

Sensitivity analysis as typically practiced does not attempt 2^N cases, but merely 2^N excursions around a baseline case. This approach assumes a linear model. Its application to nonlinear models amounts to an asymptotic argument of linearization for small enough excursions, meaning that one must assume that the uncertainty range in the inputs is quite small. For nonlinear models with significant uncertainty in the inputs, making a best estimate prediction with sensitivity analysis is not a coherent research strategy. It has been the only strategy available to researchers, so it has been used, with or without apologies. The curse of dimensionality is typically only mentioned in the context of providing an excuse for leaving out the sensitivity analysis part of this method.

Robust inference provides a superior basis for reasoning based on computational experiments on problems with deep uncertainty (Bankes and Lempert 2004). Robust inference is based on identifying a class of experiments and a proposition such that the proposition has been true for all experiments conducted that satisfy the class definitions. The strength of such an inference depends both on how many cases have been examined and how cleverly they were crafted to try to invalidate the proposition. All of the following discussion is based on this idea, and the following provides a few examples of its use.

REASONING FROM MULTIPLE MODELS

Often, there is not a single parameterized model structure that is sufficient to express all of the possibilities consistent with our knowledge. The ability to reason across model structure is useful in a variety of contexts. Figures 3 and 4 exhibit two different Bayes nets, built from expert elicitation, that capture two different hypotheses. In this case, the two hypotheses have to do with the nature of a terrorist threat, and these two diagrams capture for a notional counter-terrorism problem the difference between a threat that is organized hierarchically and one that is more networked.

In order to make robust inferences across this sort of model resident knowledge, it is necessary to find conclusions (about policies in this case) that are robust to uncertainties about both the structure of the model and the parameters that characterize any particular model instance.

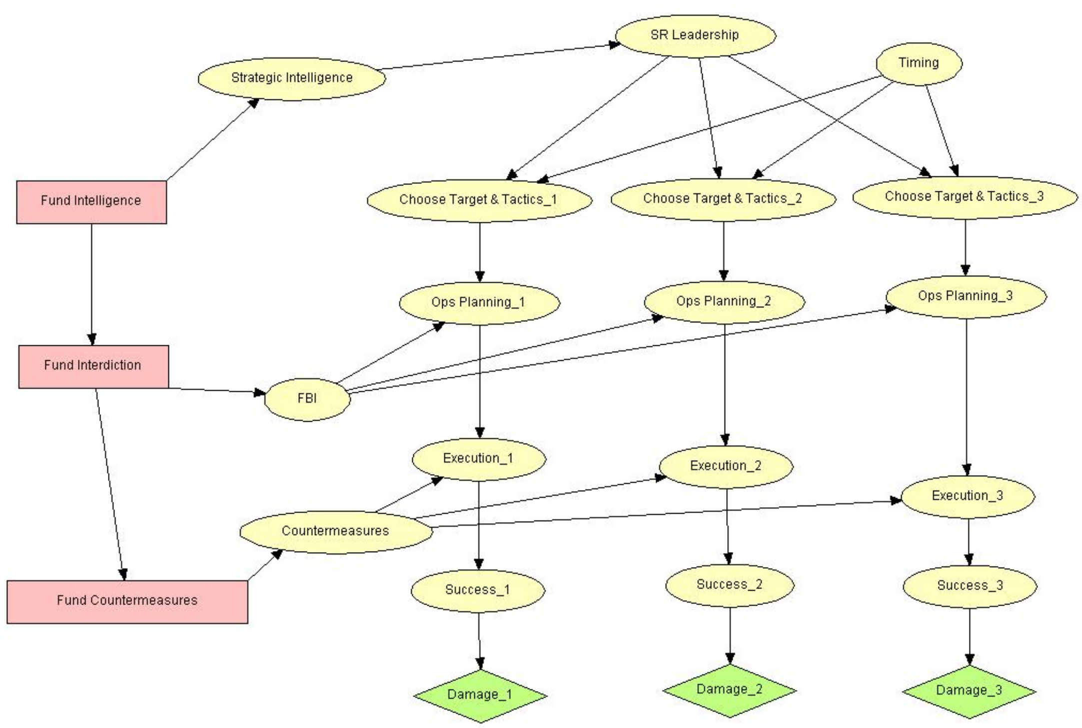

FIGURE 3 Counter-terrorist model for network threat organization

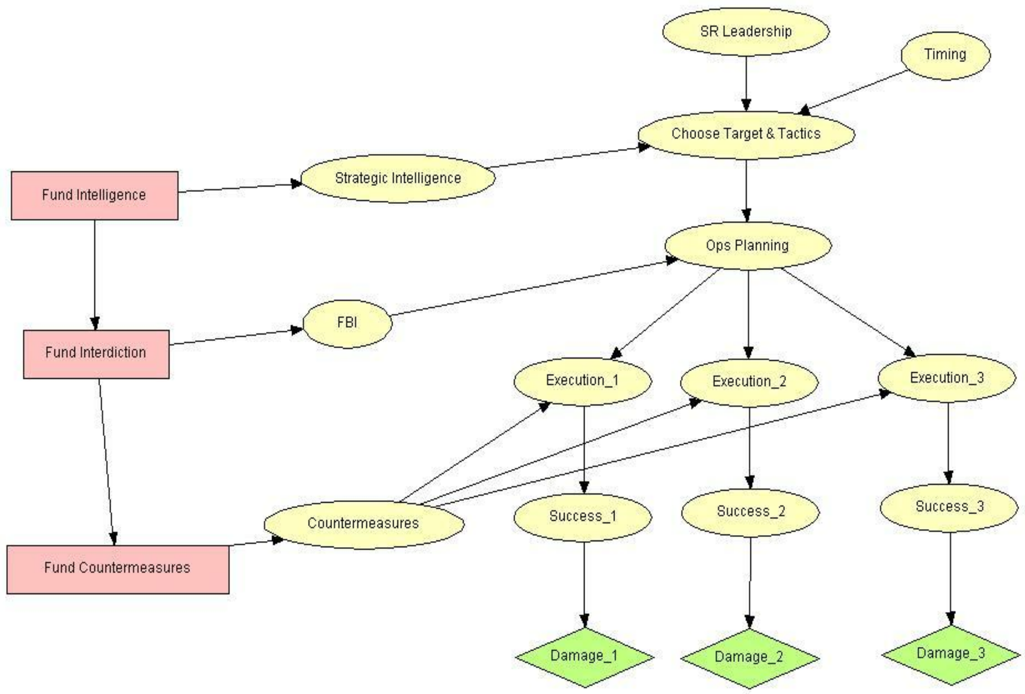

FIGURE 4 Counter-terrorist model for hierarchical threat organization

Multiple Models from Data

In recent years, the machine learning community has been moving to ensemble-based methods. Techniques such as bagging, boosting, and stacking have proven their value across a wide range of applications and constituent model types. From this, it can be inferred that ensembles of models often contain more information than any single model does, and a growing portfolio of means for exploiting ensembles of models are becoming available. I and colleagues have been using this approach to create ensembles of models from data. Instead of using model averaging as in common in machine learning, however, we use the resulting ensemble as a foundation for robust inference, as described above. In this section, I present two examples of this. In the first example, the ensemble consists of models with different structure whose parameters are estimated from the same data set. In the second example, the ensemble is created by bootstrap resampling of the training dataset, as is done in bagging.

The first example uses six regression models created by Paul Collier (Collier and Hoeffler 2000). Collier fit a wide variety of such models to a data set that contained historical data regarding internal conflict across 120 countries over the past forty years. Six of these models had high measures of goodness of fit. Though there is a maximally likely model, differences in likelihood across the six do not provide a statistical basis for rejecting any of the other five. These six models include four models whose predictors are consistent with a theory that conflict results from opportunity for overthrowing the government, aka “greed,” and one model whose predictors are consistent with the view that conflict results from “grievance.” The most likely model used combinations of these predictors.

In this exemplary analysis, these six models are used to forecast the probability of conflict in Pakistan in 2015. These forecasts involve assumptions regarding the continuation of observed trends and possible U.S. policy over the intervening years. By coupling these models with a policy effect model (based on expert elicitation), the impact of alternative U.S. policies on the probably of conflict in Pakistan can be assessed for each model.

This ensemble of models demonstrates that a strategy devised to be robust over all six models has better performance than the "optimal" strategy based upon the "best estimate" model. Figure 5 shows the structure and parameters of these six models.

Figure 6 displays the regret of the candidate strategy (the optimal strategy for the nominal case on the best estimate model) for random cases on all six models. The sampling technique used was a Latin hypercube (space-filling) design. As can be seen on from this graph, the candidate strategy has zero or small regret on all cases sampled from five of the models, and for some cases on the sixth. However, there are also some cases for the grievance model in which the regret is substantial, as much as a 60% increase in the probability of conflict.

Further analysis of the high regret cases can provide a great deal of insight. Figure 7 shows only the high regret cases from the previous figure, displaying for these cases the two most influential variables in producing them (based on an analysis of variance). This view allows us to say that the "optimal" strategy performs well unless it turns out that the grievance model is the right way to think about conflict in Pakistan. If the grievance model holds, there is a significant decline in democracy in the coming years, and the terrorist infrastructure is relatively proficient, then the optimal policy will perform badly compared with alternatives. If this strategy represented our best option, perhaps the decision maker would be willing to wager against this failure mode happening. But, as it turns out, further analysis can provide a better candidate and avoid the necessity of placing that bet.

ALL BETAS	Greed-NoPeaceDur	Greed-PeaceDur	Greed-GDP	Greed-Diasp	Grievance	Greed-Grievance (Base)
constant term	-15.56	-12.16	-3.704	0.7460	-1.688	-13.07
male schooling	-0.0336	-0.02470				-0.03156
ln GDP per capita			-0.8434	-1.237		
GDP growth-3*Pop growth	-0.1110	-0.1146	-0.0955			-0.1152
primary comm exports	19.91	18.93	16.58	17.57		18.94
primary comm exports^2	-31.79	-29.28	-23.42	-28.82		-29.44
ln population	0.8370	0.6775	0.4700	0.2949		0.7677
social fractionalization	-0.0001374	-0.0001590	-0.0002090		0.00005952	-0.0002135
geographic dispersion	-2.238	-2.115	-0.8242		0.2507	-2.487
mountainous terrain	0.01603	0.01326	0.007856		0.01144	
peace duration		-0.003636	-0.003678	-0.002010	-0.004566	-0.003713
diaspora/peace duration				700.9		
ethnic dominance					0.2361	0.6704
democracy					-0.1033	

FIGURE 5 Basis of work of World Bank's Paul Collier using data for several dozen conflicts for 120+ countries over 40 years

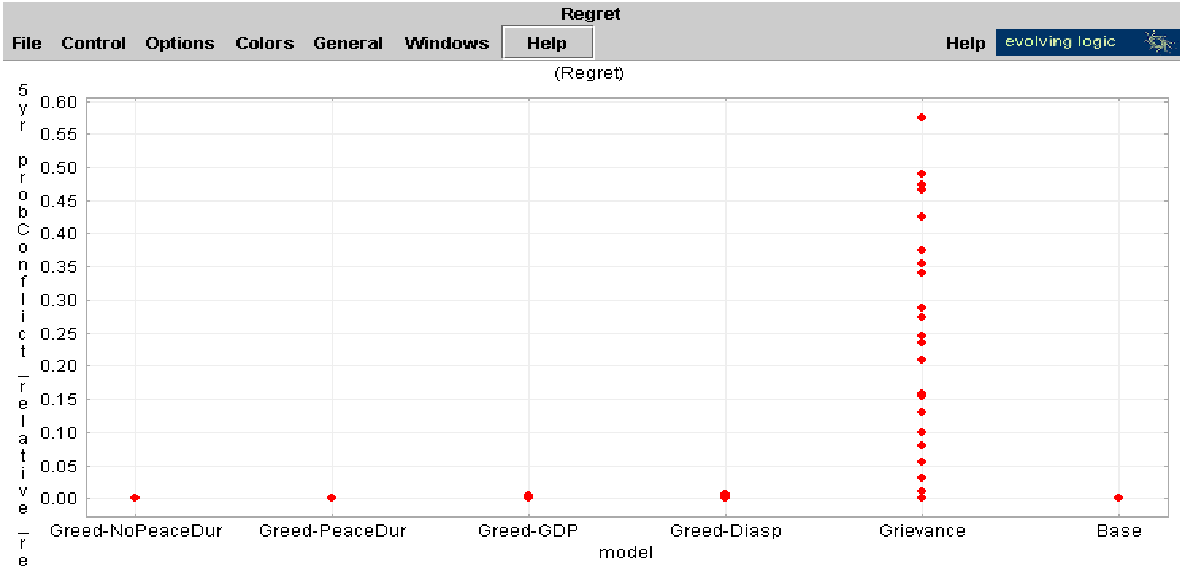

FIGURE 6 Optimal policy for best estimate model, which is vulnerable in the grievance model

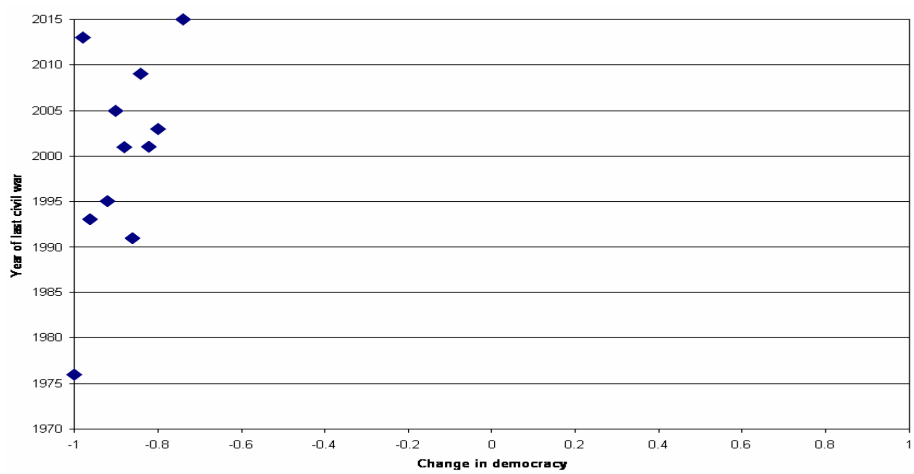

FIGURE 7 Characterizing the cases where optimal policy is vulnerable

A search over strategies reveals one that has the best overall performance on the high regret cases, as shown in Figure 8. Note that this strategy is not optimal for the nominal case for any of the models. Repeating the vulnerability analysis done for the best estimate strategy, the regret of this new strategy is calculated across a series of cases on each of the six models. Those results are displayed in Figure 8. This new strategy has non-zero regret for all of the models. But its maximum regret is two orders of magnitude less than the maximum for the best estimate strategy. (Note that the scale of Figure 8 is expanded compared to that of Figure 6.)

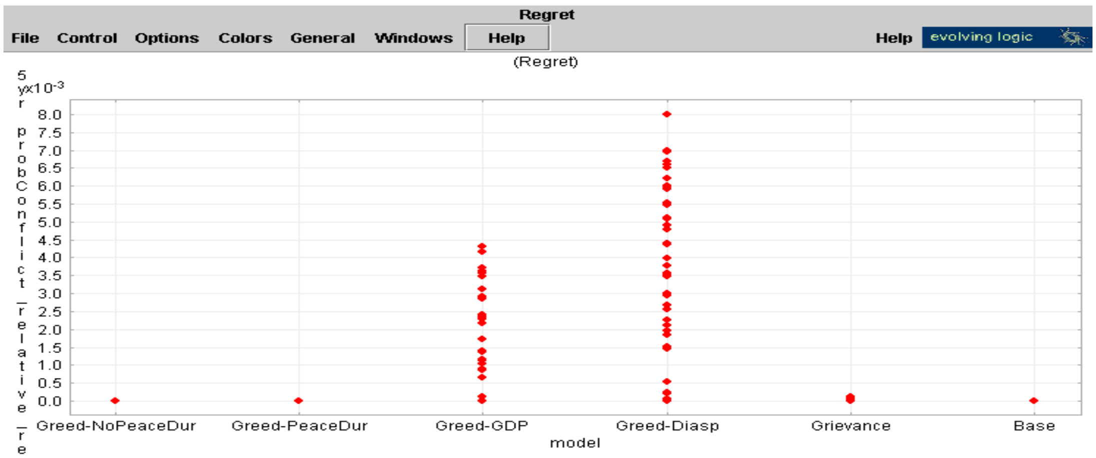

FIGURE 8 Vulnerabilities of candidate robust policy

The worst case for the hedging strategy has a relative regret of less than 1%. While further searching might reveal an even more robust strategy, this one is clearly much more robust than the original candidate, which was optimal for the best estimate model but had failure modes revealed by the ensemble of plausible models.

Had we used only the best model as the basis for our analysis, we would have recommended a strategy that was more vulnerable to surprise than the hedging strategy. The discovery of the hedging strategy was made possible by our use of the suite of alternative models. These different views, all derived from the same data, provide us with more information than any single model would have provided. If we are interested in robust strategies, the use of multiple alternative models provides a clear advantage.

This basic approach can be used with a wide variety of modeling structures. Figure 9, displaying the second example, is taken from a study in which an ensemble of feed-forward neural nets was created by bootstrap resampling among the training exemplars. This figure demonstrates that for new test instances, a variety of possibilities can be observed, from all of the members in the ensemble agreeing on the classification of this input to a broad distribution of classifications being observed across the network. Such an ensemble can be used as a challenge set to help us craft robust strategies.

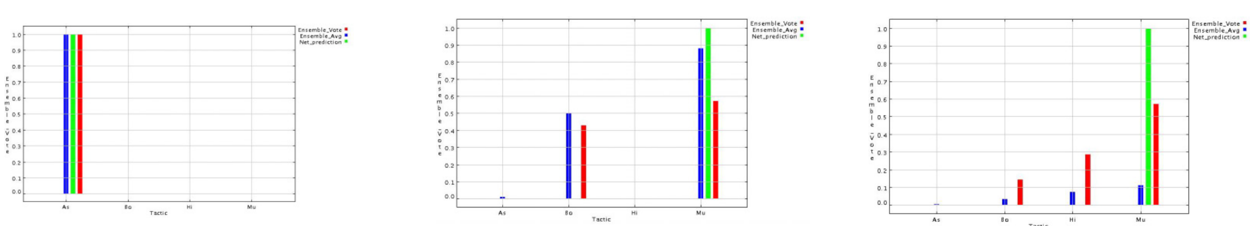

FIGURE 9 Predictions of model ensemble for a selection of novel test cases

Inference with Complementary Models

In addition to reasoning across ensembles of alternative models, computational science could benefit from an established method for joint inference across complementary models. Many efforts have been made historically to combine models by piping the output of one into the input of another, perhaps with some transformation of coordinates, such as miles to kilometers. However, such efforts have often been much less successful than hoped. This is because all models are approximations, and so are based on various assumptions. When the models used have unstated assumptions that contradict each other, essentially any conclusion can potentially result when they are combined. (Recall that from A and not-A, you can deduce anything.) Thus, combining models is in itself an act of modeling, requiring that assumptions of the combining be explicitly addressed and rectified against the pieces.

Treating models as representing ensembles of possible computational experiments can provide a much better basis for approaching the problem of model fusion than is possible through "wiring the models together." An environment for reasoning from modeling experiments (such as CARs) allows the "wiring" between the models to be set at "run time," and hence varied in response to changing hypotheses and goals. This allows strategies for specifying groups of computational experiments to account for uncertainties inherent in the model fusion itself. This approach allows the problem of model fusion to be transformed into a version of the machine learning problem.

To demonstrate this, the six Collier models are combined with two additional models of different type. The first is CAST, which digests daily news reports, and combines them using an expert-designed weighting scheme to produce 12 indicators of potential conflict. The second is a model based on structured interviews with military officers and government officials. This model represented the impacts differing strategies might have on various aspects of the future situation in Pakistan, along with their uncertainty about the size of these effects.

The inputs and outputs of these various models are not unrelated. In fact, the lever inputs to the Collier models and to the Effects model (Expert Elicitation model) are identical by design. The 12 indicators that are outputs of CAST can be related to various inputs to the Collier models. And outputs of the Collier models, particularly the probability of conflict by year, serve as an input to the Effects model.

CARs allows us to avoid expressing this relationships between the models as "hard wiring" created by revising the software. Such hard wiring often results in difficulties when models of aspects of a deeply uncertain problem are combined in this way. The scheme for joining two models in this way is itself an act of modeling, and, as with other models of deeply uncertain phenomena, the assumptions embodied in the wiring itself represent an uncertainty whose implications should be explored. Failure to contend with this difficulty has at time led to nonsensical results in past efforts at model fusion.

Instead of hard wiring, we continue to treat each model as a separate platform for computational experiments, but allow the strategy for case generation for experiments on various platforms to be coupled or correlated in arbitrary ways. This allows us to treat the constellation of models as a single unified model if we wish to, but also allows the confederation of models to be exercised in other ways, to answer other questions than "what would happen if."

REFERENCES

Bankes, S., 1993, “Exploratory Modeling for Policy Analysis,” *Operations Research*, Vol. 41, No. 3, pp. 435–449.

Bankes, S., and R. J. Lempert, 2004, “Robust Reasoning with Agent Based Models,” *Nonlinear Dynamics, Psychology, and Life Sciences*.

Boehm, B., 1988, “A Spiral Model of Software Development and Enhancement,” *IEEE Computer*.

Boehm, B., 2000, *Spiral Development: Experience, Principles, and Refinements*, Special Report CMU/SEI-00-SR-08, ESC-SR-00-08, Software Engineering Institute, Carnegie Mellon University; available at www.sei.cmu.edu/cbs/spiral2000/february2000/BoehmSR.html.

Collier, P., and A. Hoeffler, 2000, “Greed and Grievance in Civil War,” World Bank Policy Research Working Paper No. 2355; available at http://ssrn.com/abstract=630727.

DISCUSSION

Methods, Toolkits, and Techniques

(Invited Speaker, Steven Bankes
Thursday, October 13, 2005, 1:00–2:00 p.m.)

Chair and Discussant: *Charles Macal, Argonne National Laboratory and The University of Chicago*

Supporting the Modeling Life Cycle

Charles Macal: We've been concerned with model validation and associated issues for many years. Many of us have read Steve's thoughts on the subject as we've matured in our use and our development of models — even before agent-based modeling came to fruition as a field. Steve is well-published in the area. He is the coauthor of an influential book called *Shaping the Next 100 Years: New Methods of Quantitative Methods for Policy Analysis*. Steve also has a publication in the proceedings of the National Academy of Sciences that is of interest, called "Agent-based Modeling: A Revolution?" In April (2005), Steve coauthored an article ("Shaping the Future") that appeared in *Scientific American*. With these thoughts in mind, I'd like to turn things over to Steve Bankes.

Steven Bankes: Thank you, Chick. Could I get a show of hands to see how many people have heard me talk before? [Editors' note: The show of hands was more than 50%.] I was afraid of that. Well, I will do my best to say enough about things that are fundamental to the approach — the rather eccentric approach — for those who are having their first experience with this stuff. The main focus of my talk, however, will be directed toward relatively new issues for the people who were here two years ago.

In general, I want to go through a smorgasbord of methodology "quickies" that hopefully will leave you with a good selection of things that will stimulate your future work; however, I want to try and unite it with a little bit of polemic, since this is the tool builders' day and this is the plenary talk. I will give just enough description of the tool that we have been building to provide some basis for understanding how we pull off some of the methodological things I want to talk about. We'll probably run out of time, but that's the way my talks always go.

[Presentation]

Macal: Thank you very much, Steve, for that provocative presentation. Let's begin the question and answer session with questions from Joanna Bryson.

Joanna Bryson: I have two questions. First, you seemed to suggest that basically we should be running, or supporting, lots of different models in some sense.

Bankes: We should be as free to do so as we can.

Bryson: And I totally buy that idea, especially since we were just talking about the extreme programming perspective, that is, that you don't know when you're done until you keep working. On the other hand, the problem that I've seen a lot is being able to document where you've been to keep people from making a big mess. In my own system, you can keep track of previous action selection scripts so that you can always rerun them if you change some of the code and check things. Even then, the consequence is that the code gets messy because it keeps having to support old ways of thinking so that we can perform comparisons. This means that there is a huge trade-off. I wonder to what extent the real ideal is not getting something where you can run a "behavior space" simulation across, so you really want to change parameters rather than supporting lots of models. Now, that won't do the kind of work you're talking about, but I'm worried about the mess.

Bankes: There are several things to say about that, I think, because there are several facets to it. I believe that one facet on that list of relationships is the issue of version control. One of the aspirations I had for my environment was to have an aggressive version control early on because the exploration idea had that as central. In fact, though, we've done a horrible job. You do the things that are going to make you some money and are today's fascination and so forth, and we're not as good at that as we'd like to be. I think the ideal tool ought to be very good at that.

Second, I think that the aspiration I have is that we're very much in model exploitation environments. What frequently happens is that somebody shows up with an awful Excel spreadsheet, and we show that the author implies things he didn't know, which in some cases he never expected. That's kind of what we do for a living. While we do not have the support to aggressively go after the idea, these ideas have a lot of leverage for model construction environments. We've not gotten to do as much of that as possible, but let me say a few things.

I think in knowing you have the ability to exercise the model, and discovering that an input variable doesn't matter, provides the cover for pulling inputs to the faceplate and actually putting them in the model that otherwise people will never do because there's going to be a "giggle" factor. One of the things we find, especially with industrial clients, is that there are things that are crucial to the decision that are left out of the Excel spreadsheet or the agent-based model because nobody has any data; they have no idea about this effect. And so rather than have to assume that the marginal propensity to do *x* given *y* is 5, since nobody knows, they leave it out completely, which means they've actually assumed something, but it's buried. By being able to exercise it, we can get them to say, "Assume it's someplace between 0.1 and 2,000." Okay, right. Now we can find out that in fact it either doesn't matter or that the flip point is 13. If it's less than 13, it's one regime; if it's more than 13, it's another. By doing that, we suddenly can get all the stuff forward that has been hidden as part of the recent models. It's part of leaving your hack in the model to go in and change that assumption that's in the data statement, for example.

I also think that model composability is going to be very important because we end up building models that are too big. If we increased our ability to build something small and concise that really exercises and understands properties, and then combine it with something else and build these things up, we'd be in a different world for better managing this stuff.

What we've accomplished is having a very good mechanism for keeping analysis code and the physics separate. With most people modeling, you'll find there's code for doing Monte Carlo stuff in the same file with stuff that actually has real science in it. To me that's crazy. And then when you want to change the probability distribution, you've got to hack in the code that has the science in it, and you may introduce a bug. There are all kinds of problems that come from not keeping modularity. So I think that one of the answers here is if we have the right scheme, we're only partway there. And somebody else has some other ideas. But the right scheme for keeping the stuff modular enough …, then you change the pieces and you keep good version control. I've got 700 different versions of this model, all off on disk with different version numbers on it, and I've got a notebook, the tree of which one got changed to make which one?

Bryson: Can I ask my second question? It's on a totally different subject — about the API. What I did for my API is to throw all my code online and let people do whatever they want. Do you have a better recommendation? At least that way I'm not charging you more money to get access, but obviously that's not ideal. You say we should have an API, but what do you mean?

Bankes: Well, yes, you're right. I think we all need to make a change in attitude; we're not perfect about this either, but again, it's a different kind of modularity in a way. You have to understand yourself as building an engine that provides some list of services. Plus a GUI that allows the user access to those services gives you a big leg up on providing an API now, which gets the same services by some means. And I actually don't care whether it's XML or remote procedure calls to some Java method, or a little socket guy that accepts some kind of weird coded language.

So it's not a question of the technology, although obviously some technology is better than others. I think it really is a question of concept; it was for us. We didn't start thinking we're building an engine with services. I'm building a system that I get to play with because I've always wanted to be able to do this. And so we've come to the point of realizing I'm providing a thing that may in fact just be middleware in front of a big federation of other people that do better GUIs and people that built the model, and I'm just sitting in the middle exercising. That's been a conclusion we came to after several years of building code. So there is an advantage to always thinking about stuff we build from a service point of view. If we do that, it's not impossible [because] we've got machines all over the country communicating via some kind of XML layer. I don't have to download your code; all I've got to do is follow the rules. And so I think it's possible to get to a world where you don't have to keep downloading somebody's newest version. He can change his website, and it's just seamless.

Things are harder than the facile version of this thing, but I think there is a version that we can get to. We can at least try to stay with the curve — if not ahead of it — in terms of making our stuff ready to work with one another because, frankly, there are a lot of hard problems here, and none of us have enough money. So if you've done a good job, I'd much rather be able to ride on top of your stuff and do something new and different than always reinventing, Everybody's still coding their own statistical routines. You ought to be able to just take SAS and plug it in and work. And so we need to get there.

Lars-Erik Cederman: Steve, it's always a pleasure to see you in action, and there's a lot of wisdom in what you're saying. Very often I realize what I've been doing only after having seen you give a talk. Or sometimes I realize what I should be doing.

I wanted to ask you about the sociology of knowledge and perhaps science because you've been talking about policymakers here, and you have a lot of experience interacting with them. You've also worked with social scientists, and it seems that some of the advice you are giving may be more applicable to the policy environment than to the treacherous social sciences. My experience is that when we're talking about the ensembles of models, you're exactly right. That's the way we should be moving forward, but that's easier said than done from, say, a PR perspective. How do you convince social scientists, or a classical economist who's chasing that holy grail, that you need equilibrium, and nothing short of that is going to convince this guy? How do you convince a person like that that we can be working with this highly dimensional ensemble of models? They think you're from a different planet.

Bankes: I think your point's well taken. It's been one of the disappointments of my career that a paper I wrote in 1993 that I thought would get me some spiff in the next couple of years and I could go back to doing agent-based modeling and other things, has been turned into a telescope-maker for some future Galileo or something. I do think, though, that we're fighting some cultural issues that have been laid in place through a few hundred years of practice, and we may not heat the bathtub with the first couple of matches. So I do take it that it's not automatic to win these things. And it may even be that in some contexts you need to be — duplicitous is too strong a word — clever. So it may in fact be that one does a lot of exploration across an ensemble of models with technologically assisted tools to find the simple picture that you actually show your colleagues.

Certainly we're juggling a lot of models to show that final simple picture to a decision maker. I don't want to show the decision maker 200,000 models; I want to show him this. Okay? But I want to show him this backed up by an environment where decision makers can treat my models like they do their own staff, which is to say, if you watch the general G-2 presents of the thing, or the CEO, "Here's a briefing from the analysis guys." He says, "What happens if interest rates go way high?" If you say that in that case we're hedged, he's happy. But if interest rates go high, it's different, and then he's not so likely to take their advice without more further analysis.

The same thing is true here and that's part of the reason for our tooling. We'd like to provide the ability where people go and be able to get those cases out for them and have them say, "Ah, yes. That confirms my prejudice and I'm beginning to see." And so there needs to be a conversation that happens between the end user and that substrate to get a certain level where the answers are sensible. Then you get to the simple story without understanding how you got there.

In academia (of course, I'm not working in so much these days), it may be the same kind of thing. You have a sales pitch to the guys at the big economic conference. The economists are famous for being doctors. How do you put this in a form that these guys can understand? Well, perhaps these become tools you use in the shop, but what you take out of the shop is much more carefully crafted. The ultimate answer is going to have to be a series of key papers where we demonstrate this stuff is solid. And so then you can cite that key paper, and I'm doing the same thing. The way we get most people to believe our stuff is to show that it's already been proven on this new application area. And we have not gotten those key results out yet. I mean, a few of them are probably starting, but I think in particular the use of data combined with agent-based models to produce something that we understand better now because we combined the data with an agent-based model, it is not true there have been no cases that people have calibrated in agent-based model data, but we don't have a good killer application where after we did this we understood something and everyone went, "Wow!" where I don't know the paper to cite.

And so to be able to say — I mean, partially this little example with the logistic progression models is — I'm reaching out toward kill the statisticians with their own tools, right?

I mean these guys thought up some really great stuff, but to say model averaging and prediction's not the only thing to do; you can do robustness. This is a step that I'm aggressively trying to find the right toy problem that people can quickly understand and believe the story. Some good papers from you would help.

David Sallach: I appreciate many of the elements of it: ensemble strategies, dimensional reduction, value of cluster analysis, and so forth. I wanted to ask you, though, an ontological question, which is: we're at the latter stages of a 50-year period of time when the social sciences have ripped apart all natural entities into variables, so that variables are run against variables without the kind of naturally enfolding structure from which they emerged. And I wonder if that structure, if we had it and were using it, would give us some dimensional reduction in and of itself. I'm wondering whether the ontological experimentation that computational modeling has begun to re-create is not an important part of the process of reassembling some coherence into the domains that we study. Even if it is, it's not sufficient because in sciences that are making progress, it's innovations in the conceptual or ontological view — the innovations like genes quarks, tectonic plates, and so forth — that have allowed a much more coherent view of the processes. That was one element I didn't hear present, and I wonder if you would comment on whether it seems to you that a focus on ontologies and on ontological experimentation advances in the underlying conceptions isn't an important part of the challenge that we face.

Bankes: Well, I suspect it is. This talk is very much shaped by my consciousness that this is the toolkit workshop. And so I tried to give a talk for tool builders. In the end, if it hadn't been tool builders, we probably would have given fewer examples but longer versions of the studies or something.

The slight twist on what you said is that I think that the invention of the computer is one of these huge infrastructure changes that occasionally happens in civilization that takes us a long while to assimilate. I think in the sciences — this is the first one since mathematics that applies across all the sciences, and where the methodological assimilation of computation is, I think, going to take a century. It's a shame for me that I wrote the paper in 1993 and wanted to move on, and I'm just at the beginning of this curve. But I do think that it's going to take redigesting a lot of old truths, re-understanding a lot of old truths.

The thing that the data managers do all the time is this thing we saw with the Collier model, where we start off with 100 possible predictors and we do a lot of specification search and end up picking six because they had the biggest result at the end in terms of whatever score we wanted to use. That's a dimension-collapse strategy. It's not described that way, though. It's usually described in some hyper-realist way, which doesn't hold water. The dimension-collapse strategy then starts to make sense, but if you express it that way, the change in state between paper and pencil when all these tools were invented, when doing one regression was arduous, versus now where regression can be done in milliseconds, depending on the size of the model and data set, and we can do specification. People now have been running GAs on top of specification search. Several years ago the first one happened. There's a lot of this going on. I think our community has not managed to reinterpret truisms proved with the past technology, which had assumptions buried in the way they thought about it that are now no longer true.

And so I think things like this issue of bias versus variance and not overfitting your data, which came up. This is, you know, if you're not aware of this, you make horrible mistakes doing statistics. But I don't believe this is a good argument for excluding variables from simulation models. Until we have the right paper to cite that explains why it's okay to have thousands of variables and then reduce them with these different means, we'll still have economists and statisticians going "Oh! I have a theorem that says that what you just did is wrong." Until we can quickly refute that theorem by showing that one of its assumptions is now violated, we end up not getting to the end of that conversation as a winner. So we have a lot of work to do.

Macal: We have time for one more question.

Pam Sydelko: My question for you relates to the question about APIs and the need for them, which I would argue includes also transparency, because when you talk about ensembles or cross-discipline model ensembles, which is especially interesting if you're looking at wanting to have your agents react to an environment that you may not be the expert in, for instance, if you're trying to have it react to an ecology model. You may not be an ecologist, so you want to use that model, and so APIs need to have that kind of transparency. One of the things that excites me about developments today is this idea of virtual collaborative toolkit environments, where I can actually sit with an expert that may not be at my laboratory, but who is the expert with that model, who can bring that model into my environment along with me, and I can work together at trying to make the best and most clever joining of those two models or three or four models. I just wondered what your thoughts were about that promise.

Bankes: I think it's really important. I think it's a shame we're not smarter, or at least that I'm not smarter. The need to put a transparent or easily-understandable faceplate on all of this stuff lies between me and wealth. If I could take some of the mechanisms — we've got some cool code, but it's "techy," and so there's no market for me to sell this to anybody. We use it for little projects, but we can't get rich doing those types of projects. To wrap this in a thing that I could sell as an Excel add-in would be enormous if I could find the right way to package it.

This leads to a struggle in all parts. I have a lot of room for creativity, but finding the right way to package these things so people go [finger snap]. I think a big part of the reason why simulation didn't become agent-based modeling, but agent-based modeling did, is simply the bitmap graphics and the illusion of a 2D surface with things eating each other. I mean, the visual "Pow! Look, look, look, they're all red!" I mean, it's very [hand claps]. You would think it was trivial, but it's made a big difference. So I think we're looking for easily-approached metaphors of that kind that get us into the meat of some of the stuff, and we need new ones.

One of the things that has emerged from our work that's stupid, but equally good for us at least, is the sliderbar bank thing. It is astonishing to me how content people are to have a graphic up and a bank of sliderbars that correspond with uncertainties. They can pull the sliderbars and see the graph change. If it's responsive enough they don't see the delay, they'll happily sit with a bank of 25 sliderbars and pull them around say it's interesting because one particular thing never makes a difference. I know that a bank of 25 dimensions is too big, but they're happy to pull these sliderbars around. A recent discovery I've made is that I've had all these nice search mechanisms. I can talk about search and searching over ensembles, and it doesn't connect. But by realizing that I can use that same sliderbar bank and have a find button, or I fix some of the sliderbars with a find button and the rest snapped to a certain place that is the best case I could find or the worst case I could find, gives them whole new specs and makes things vary, that rather than, a dialog that comes up saying, "Which variables do you want to bind and which method do you want to use to search?" All of a sudden, I'm finding much more that less technical users find, "Oh, so the computer can pull the sliderbars for me." Right? It's a natural metaphor, but it's taken us five years to stumble across it. The reason we come to these conferences is to show each other stuff. I'm looking for these good ideas. You're absolutely right, it's very important.

Macal: We're going to have to wrap the session up. Thank you very much, Steve, for that very provocative presentation.

Steve referred to his talk that he gave at this agent conference in 2003. I want to point out that the CD on the table outside [this auditorium] consists of the conference proceedings for the Agent 2003 conference. Steve's paper is on PDF page 169. It's in the table of contents as page 155 [Editors' note: page numbers vary due to numbering of front matter.] The title of this talk was "Improving the Validity and the Rigor of Agent-based Modeling through Ensembles of Models."

Toolkits

FABLES: FUNCTIONAL AGENT-BASED LANGUAGE FOR SIMULATIONS

L. GULYÁS* and S. BARTHA,
Loránd Eötvös University, Budapest, Hungary

ABSTRACT

This paper introduces a novel programming language dedicated to the development of agent-based simulations. The Functional Agent-based Language for Simulations (FABLES) attempts to follow the language of agent-based model publications in order to address the needs of modelers who have limited programming knowledge but sufficient skills in algebra and calculus. The first proof-of-concept layout of the language is discussed, together with concrete workable example models written in FABLES. A summary of FABLES's larger context, the Multi-Agent Simulation Suite (MASS), concludes the paper.

Keywords: Agent-based modeling and simulation, computational tools, programming language

INTRODUCTION

Opinions vary about the level of programming skills to be expected from agent-based modelers. Yet there is apparent general agreement that the more skill candidates have, the better. However, most of today's students lack these capabilities, and developing them requires substantial effort from the adventurous entrepreneur. Therefore, lowering the requirements would help agent-based modeling become a more widely accepted methodology (Gulyás 2002; Gulyás and Bartha 2003).

Swarm-like modeling packages, such as Swarm (Swarm undated), MAML (MAML undated), Repast (Repast undated), Ascape (Parker 2001), or MASON (MASON undated), require the use of general-purpose programming languages (Java, Python, Objective-C, etc.). Thus, their users are forced to learn programming at some level. On the other hand, various model-building tools, such as NetLogo (Wilensky 1999) or AgentSheets (AgentSheets undated), demonstrate that by limiting the "space" of possible models, the task of modeling can be efficiently assisted. The real challenge is to bridge the gap between the potential open-ended nature of Swarm-like modeling environments and the ease of use provided by NetLogo-like frameworks. Graphical model building interfaces for Swarm-like agent-based modeling platforms, such as Repast.py (developed for Repast Py, formerly known as SimBuilder) and Visual Swarm Builder (VSB), are attempts at this (Perrone and Tenuti 2002). Still, they impose certain limitations with respect to the models that are possible to build with them and, at the same time, require a certain level of programming knowledge.

* *Corresponding author address:* László Gulyás, Simulation Centre, Cooperative Research and Education Center, Faculty of Informatics, Loránd Eötvös University, H-1117, Budapest, Pázmány P. sétány 1/c, Hungary; also AITIA, Inc., H-1117, Budapest, Infopark sétány 1, Hungary; email: lgulyas@aitia.ai.

FUNCTIONAL AGENT-BASED LANGUAGE FOR SIMULATIONS (FABLES)

The motivation behind FABLES is to improve upon the current situation. One crucial observation is that a large part of the code of a typical agent-based model (ABM) is concerned with observation (collecting statistics, visualization, etc.). On the other hand, the assembly of the observation machinery can be assisted in interactive ways. Both SimBuilder and VSB, or Ascape and NetLogo, provide examples to this effect. As a consequence, in FABLES, tools for input/output (I/O) are kept at the minimum, and observation is delegated to the integrated modeling environment developed for the language (see below).

Another observation is that part of the difficulty in creating and communicating ABMs stems from the fact that the formalism used to describe models in research papers or in oral presentations is far from the language of implementation. Moreover, the model's actual source code is full of what may be called "accidental representation." These are algorithms and data structures that translate the conceptual model's notions into programming constructs. These are "accidental" elements because they are normally developed without much thought, since modelers tend to focus on model details instead of studying computer science textbooks. They are often based on word-of-mouth information (or more specifically, advice given and taken via e-mail distribution lists) rather than on solid software engineering knowledge. These are the parts of the model where programming skills count the most.

Design Goals

FABLES attempts to improve accidental representations by providing a language in which models can be described as close to the conceptual model as possible. Our starting point when designing FABLES was to follow the language that ABM publications use to describe their models. The typical intended user of FABLES is a modeler with limited programming knowledge but with sufficient skills in algebra and calculus to read a research paper. The design goals of FABLES can be summarized as follows.

1. The FABLES source should be easily readable for readers familiar with the basic mathematical formalism.
2. The language should have precise semantics, and the source should be the exact specification of the model.
3. The FABLES source should be as close to a "publishable" model description as possible.
4. FABLES models should be executable.
5. The model description should focus on the nature of the model and leave implementation to the compiler.
6. The language should be general enough to possibly accommodate any ABM but should focus on the common techniques and methods.

Current State

The design goals listed above partially contradict each other. Nonetheless, we attempt to achieve an optimal balance among them. We believe that the ideal model description language is built upon a functional base that is close to the mathematical, algebraic formalism. Such descriptions are typically more concise than those using imperative languages. To get even closer to mathematics, we replace programming constructs like arrays, lists, etc. with respective mathematical concepts like sequences, sets, and relations. Since agent-based systems are close to the object-oriented paradigm, FABLES also uses object-oriented concepts: agents are objects, agent types are object classes. The simulation's behavior in time, however, is hard to describe in functional terms, and the object-oriented framework does not help much either. This component is best described with imperative tools. Separating dynamic behavior from representation also helps achieve a precise semantics. Therefore, the models' event-based dynamics are described by schedules. The schedules contain imperative elements, like those in Swarm, Repast, or MASON. However, these imperative elements (events or actions) are not embedded in objects that often confuse novice users of other packages.

FABLES thus can be separated into three main parts: (1) an object-oriented part describing the general structure of the model (environment, agents, etc.); (2) a functional part (with mathematical equations, functions, sequences, and sets) providing a standard, concise means to summarize the functional relationships among the various components and their behaviors, and (3) an imperative part (assignments and object creation and destruction) with a schedule that specifies the actual dynamics.

This largely corresponds to what is found in published model descriptions. Typically, object-oriented terminology is used to describe the overall structure of the model (the actors and general concepts), and mathematical language is used to picture the components' mutual dependence, while dynamics are given either in functional form (i.e., difference/differential equations) or pseudo-code (often by using the event-based terminology and/or the concept of the scheduler).

Examples

The design and implementation of FABLES are works in progress. Therefore, the syntax of the language is not finalized yet, and the main concepts may also change in the future. The examples below were prepared by using FABLES v0.2.

Random Walk

Example 1 is a simple random walk on a two-dimensional (2D) lattice, performed by 100 agents.

Conway's Game of Life

Example 2 is Conway's famous Game of Life model (Berlenkamp et al. 1982).

```
model randomwalk;

agentNum = 100;

class Agent begin
  pos ; // is Integer x Integer

  schedule step cyclic 1 {
  2 : pos := pos + discreteUniform( [-1..1, -1..1] );
  }
end

schedule init {
0 : seed(0) ;
0 : display:=load("user.Display3",-20,20,-20,20) ;
1 : [ new Agent[ pos:=[0,0] ] : i is [1..agentNum] ] ;
}

/////////////////////// OBSERVER /////////////////////////////
display ;
schedule Observer cyclic 1 {
2 : display([ a.pos : a is Agent ]);
}

end
```

EXAMPLE 1 Random walk

```
model Life;

worldSize=30 ;
world ;

norm(x) = a < 1 => a+worldSize
          otherwise a where ( a= x mod worldSize ) ;

neighbours(x,y) = size ([ 1 : dx is [-1..1],
                              dy is [-1..1]
                              when not (dx==dy==0) and
                              world(norm(x+dx))(norm(y+dy))
                       ]);

step(n,old) = n==3 or (old and n==2) ;

newWorld = [
             [
               step( neighbours(x,y), world(x)(y) ) :
                 y is [1..worldSize]
             ] : x is [1..worldSize]
           ] ;

schedule Init {
0 : seed(0) ;
0 : display:=load( "user.Display2D",
                   "Game of Life",
                   worldSize,
                   worldSize);
```

EXAMPLE 2 Game of Life

```
1 : world := [
                [
                  discreteUniform(true,false,false) :
                    y is [1..worldSize]
                ] : x is [1..worldSize]
              ];
}

schedule Step cyclic 1 {
3 : world := newWorld ;
}

////////////////////////// OBSERVER //////////////////////////
display;
schedule Observer cyclic 1 {
2 : display( [ (x-1,y-1,1) :  x is [1..worldSize], y is [1..worldSize]
                              when world(x)(y) ] );
}

end
```

EXAMPLE 2 (Cont.)

Mousetraps

Example 3 is the Mousetraps model, known from Swarm distributions, which is a cartoon demonstration of chain reactions. Mousetraps are located at fixed positions on a 2D lattice, with each of them having two "ping-pong" balls placed on it. When a mousetrap is hit by a falling ping-pong ball, it is "triggered": it releases its balls into the air. They will eventually come down at random locations on the grid, hitting and triggering other mousetraps. In the model, tossing balls into the air to trigger other mousetraps is accomplished by picking random locations on the lattice and scheduling events in the "near" future that will trigger them.

```
model Mousetrap;

worldSize = 10; // Model parameter

// Shorthand for the space
world = [1..worldSize, 1..worldSize];

mousetrapsFired; // Counter

// The agents
class Mousetrap begin
  pos;      // is Integer x Integer;
  hasFired; // is Boolean;

  activate = (hasFired == false) => [
                mousetrapsFired := mousetrapsFired + 1,
                printLn("Activated"),
                hasFired := true,
                generateTriggers
              ]
              otherwise =>
                printLn("Dummy");
```

EXAMPLE 3 Mousetraps

```
end

// Initialization
schedule init {
  0: graph := load("user.SequenceGraph",
                   "#Mousetraps Fired",
                   0,200,0,100);
  0: seed(1);
  1: [
       new Mousetrap[ pos := [i, j],
                      hasFired := false
                    ] :
       i is [1..worldSize],
       j is [1..worldSize]
     ];
  1: mousetrapsFired := 0;
}

// Main schedule
schedule mainSchedule {
  2: discreteUniform(Mousetrap).activate;
}


generateTriggers = addEvent( mainSchedule, 1, triggers );
triggers = { a.activate :
             a is Mousetrap,
             b is RndPositions when a.pos(1)==b(1) and a.pos(2)==b(2)
           }
  where (
    RndPositions = {discreteUniform(world) : a is [1..2]}
  );


//////////////////////// OBSERVER //////////////////////////

graph; // variable to store the graph object
schedule observer cyclic 1 {
   2: graph([mousetrapsFired]);
}
end
```

EXAMPLE 3 (Cont.)

Schelling's Segregation Model

Example 4 is Thomas Schelling's famous model of residential segregation (Schelling 1978). In this model, two kinds of agents (henceforth called "reds" and "blues") are placed on a 2D thoroidal lattice. The lattice is interpreted as a city, with each square representing a house or a lot. The agents are families of different ethnicities. The neighborhood of an agent occupying any location on the lattice consists of the eight squares adjacent to this location. The agents' happiness depends on the ratio of different color neighbors, with each agent having a specific threshold. Unhappy agents move to the closest empty location that satisfies their expectations.

```
model Schelling;

// Model parameters
worldSize = 10;
agentNum  = 70;
threshold = 0.6;

// Constants
red = 1;
blue = 2;
color = {red, blue};

// Spatial Constructs
world = [1..worldSize, 1..worldSize]; // The world
occupied = {a.pos : a is Resident};   // Occupied positions
empty = setMinus(world, occupied);    // Empty positions

// Helper function to implement a thorus
norm(x) = a<1 => a+worldSize
          otherwise a where (a = x mod worldSize) ;

// Neighborous relation among locations
neighborous(x, y) = { [norm(x+dx), norm(y+dy)] :
                        dx is [-1,0,1],
                        dy is [-1,0,1]
                        when not (dx==0 and dy==0)
                    };

// Manhattan distance function
d(a, b) =  math.abs(a(1)-b(1)) + math.abs(a(2)-b(2)) ;

///////////////////////// INIT /////////////////////////
schedule initModel {
  0 : seed(0) ;
  1 : [ new Resident[ c   := discreteUniform(color),
                      pos := discreteUniform(empty),
                      t   := threshold
                    ] : o is [1 .. agentNum ]
      ] ;
}

// The agents
class Resident begin
  pos ; // is world;
  c;    // is color;
  t;    // is [0.0..1.0];

  neighbors = {a is Resident when
                 a.pos in neighborous(pos(1),pos(2))};
  sameNeighbors = {a is neighbors when a.c == c};
  utility = try( size(sameNeighbors)/size(neighbors), 1.0 );
  closestEmpty = empty(minPlace({d(pos, loc) : loc is empty}));

  schedule Step cyclic 1 {
    2 :  pos := ( utility >= t => pos
                  otherwise => closestEmpty );
  }
end
```

EXAMPLE 4 Schelling's model

```
///////////////////////// OBSERVER //////////////////////////

display; // Variable for the display object
schedule initDisplay {
  0 : display := load("user.Display2D", "Schelling",
                       worldSize, worldSize);
}


schedule cyclic 1 {
  2 : display([(a.pos(1)-1, a.pos(2)-1, a.c) : a is Resident]);
}

end
```

EXAMPLE 4 (Cont.)

MASS: THE CONTEXT OF FABLES

The Multi-agent Simulation Suite (MASS) is a solution candidate for modeling and simulation of complex social systems. It provides the means for rapid development and efficient execution of agent-based computational models. The aim of the MASS project is to create a general, Web-enabled environment for versatile multi-agent-based simulations. The suite consists of reusable core components that can be combined to form the base of both multi-agent and participatory multi-agent simulations. The project also aims at providing a comfortable modeling environment for rapid simulation development. To this end, the suite will offer a high-level programming language dedicated to agent-based simulations and a development environment with a number of interactive functions that help in experimenting with and finalizing the model.

Mass Components

MASS has four major components, as shown in Figure 1 and discussed below.

Multi-agent Core

Multi-agent Core (MAC) is an execution environment for agents. This J2SE-based module provides the basic infrastructure (time and event management, agent-agent and agent-environment interactions, logging and playback functions, etc.) for multi-agent simulations (Figure 1). MAC differs from common packages for agent-based simulation (like Swarm, Repast, Ascape, or MASON) in several ways. Most important, it is a fully Web-enabled environment. (There will be a standalone version, too.) This means not only that prewritten simulations can be run from a browser but also that the assembly and configuration of models (from preexisting components like agent and environment types) are also possible via the Web (Figures 2 and 3). This may be especially useful in educational settings, where novice modelers can experiment with model templates and prewritten components. Moreover, the Web-enabled nature of the MAC platform lends itself naturally as a base for the participatory extension discussed below.

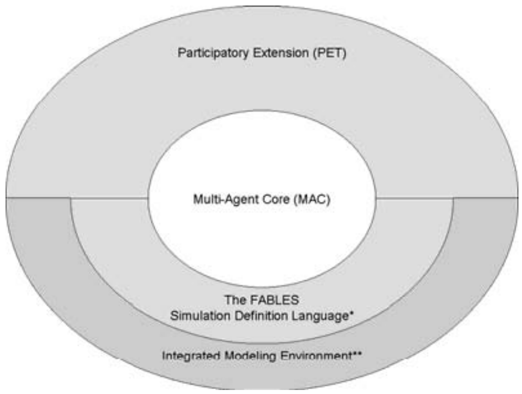

FIGURE 1 Architecture of MASS

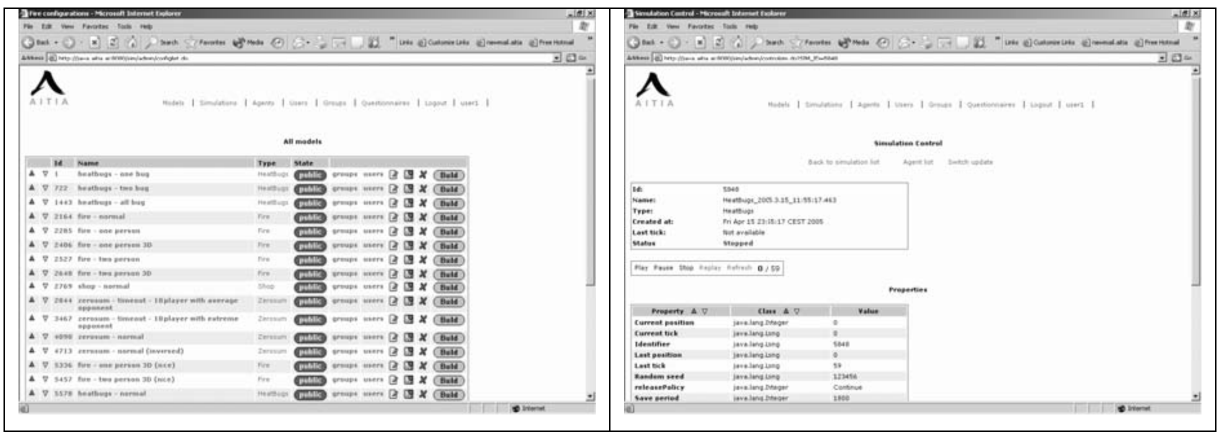

FIGURE 2 Administration interface of MASS

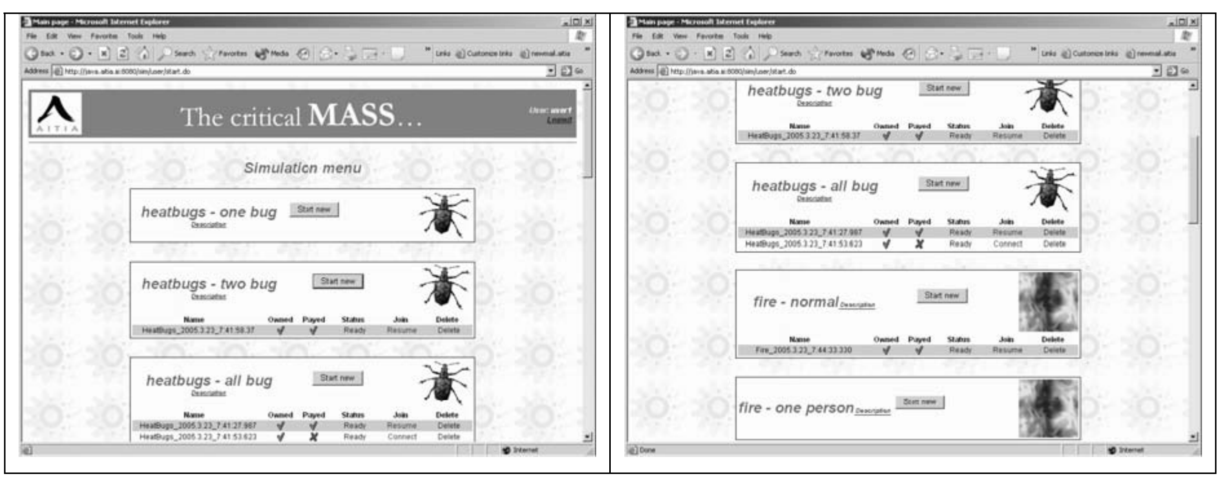

FIGURE 3 Simulation selection interface of MASS

Participatory Extension

Participatory simulation is a methodology that builds on the synergy of human actors and artificial agents and excels in the training and decision-making support domains (Gulyás et al. 2004). In such simulations, some agents are controlled by users, while others are directed by programmed rules. The Participatory Extension (PET) is an add-on to MAC that allows users to take control over agents in the simulation. The J2EE-based extension provides solutions for communications between the client and the main simulation server, including visualization (2D bird's-eye view, 3D world view, etc.) at the client side. The user interface of the simulation client uses standard Web technology, allowing for easily customizable layout and design (Figure 4).

FABLES Simulation Definition Language

As discussed above, FABLES is intended for the concise and efficient definition of agent-based models. FABLES combines the strengths of functional programming with the object-oriented paradigm, providing unique means to implement agent-based simulations.

The FABLES vision is an abstract formalism to describe agent-based models. Models defined in this language could, in principle, be automatically transformed into agent-based simulations in Repast, Swarm, Ascape, etc. Such a description would be ideal for publishing concise definitions of agent-based models. Moreover, with independently developed compilers to different modeling platforms in place, the formalism could also help make the replication and docking of computational models a routine task. In the current, prototype version, FABLES models are interpreted. Our future plans include a compiler that generates (optimized) code for MAC.

Integrated Modeling Environment

The integrated modeling environment makes model development in FABLES more effective by providing a modeler-friendly editor with syntax-highlighting, on-the-fly syntax checking, and a number of exporting options. The environment will also have interactive wizards that help in collecting and charting statistics about the model. This will be completed by wizards that let the modeler interactively set up 2D and 3D displays of the model.

SUMMARY

In this paper, we discuss the FABLES language for simulations and MASS and its development context. In its current prototype version, FABLES is an interpreted language. However, in the long run, models developed in FABLES will be compiled to pure Java code. This way they will be seamlessly integrated in MASS. Obviously, the FABLES compiler may never generate code as efficiently as an experienced programmer. Still, using FABLES may be a viable option for smaller-scale, exploratory models. Also, using FABLES will force making efficiency considerations explicit, especially when they depart from the conceptual model.

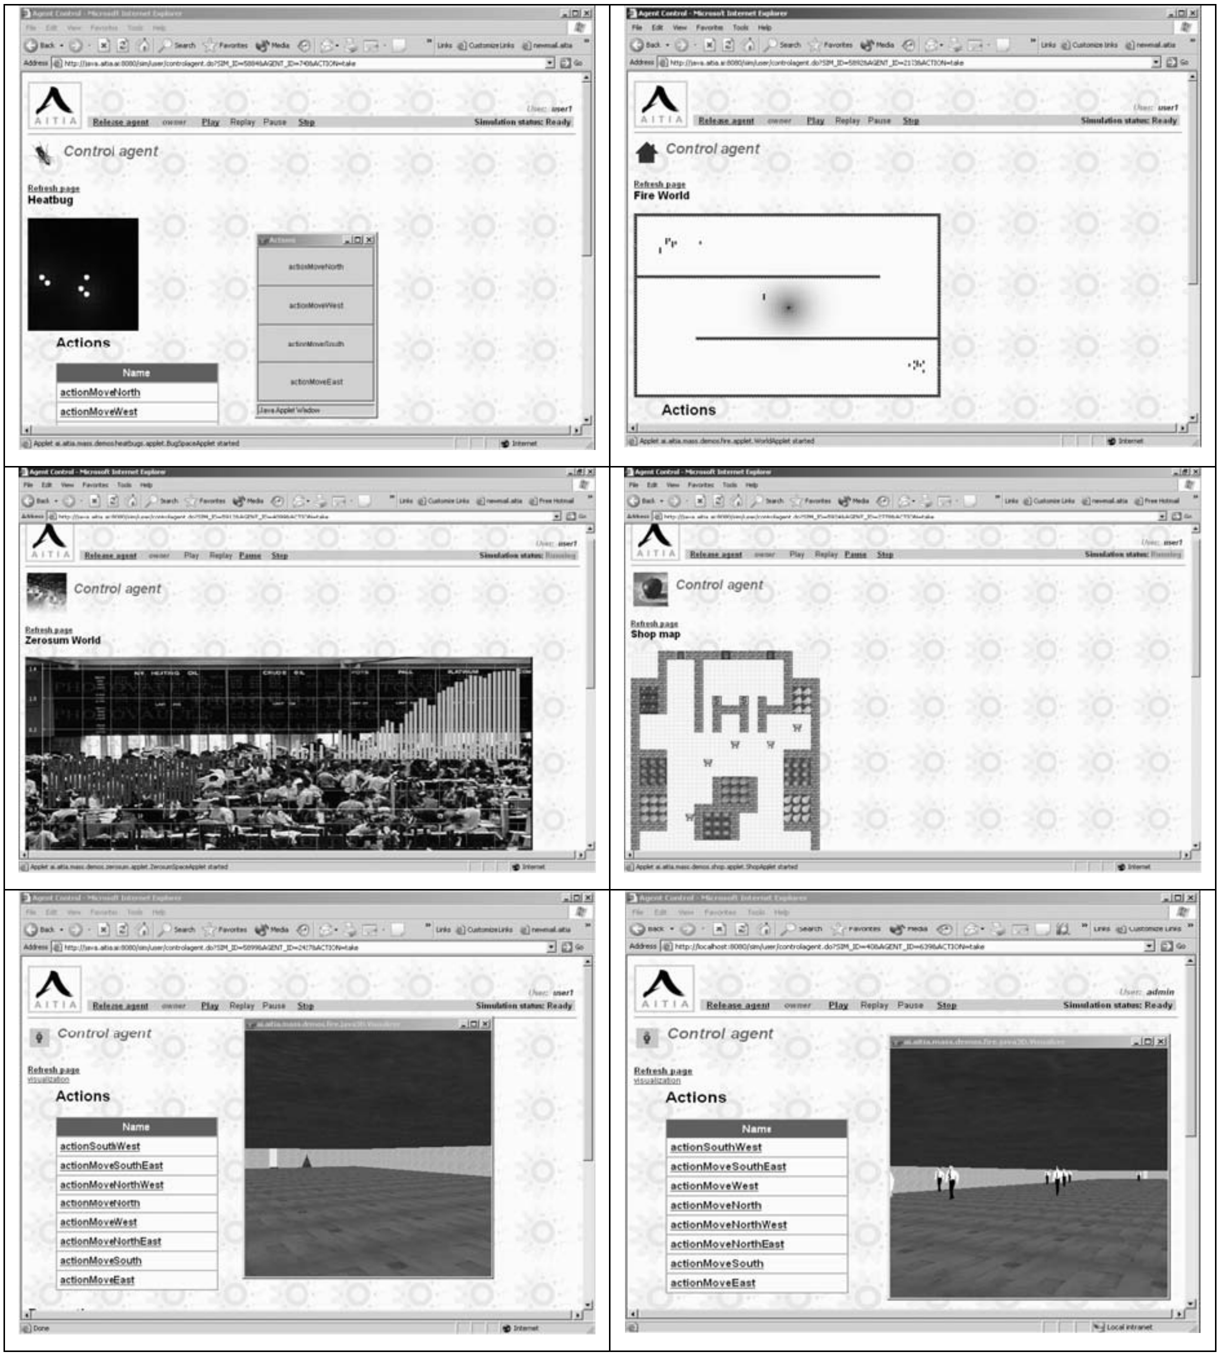

FIGURE 4 User interface of simulations in MASS

REFERENCES

AgentSheets, undated; available at http://www.agentsheets.com/.

Berlenkamp, E.R., J.H. Conway, and R.K. Guy, 1982, *Winning Ways for Your Mathematical Plays*, Academic Press, New York, NY.

Gulyás, L., 2002, "On the Transition to Agent-based Modeling: Implementation Strategies from Variables to Agents," *Social Science Computer Review* 20(4):389–399, winter.

Gulyás, L., and S. Bartha, 2003, "How Much Is Too Much? — What Programming Skills Are Really Needed to Do ABM?" Seventh Annual Swarm Users/Researchers Conference (SwarmFest 2003), Notre Dame, IN, April 13−15.

Gulyás, L., B. Adamcsek, and A. Kiss, 2004, "An Early Agent-based Stock Market: Replication and Participation," *Rendiconti per Gli Studi Economici Quantitativi*, Vol. unico, pp. 47−71.

Gulyás, L., S. Bartha, T. Kozsik, R. Szalai, A. Korompai, and G. Tatai, 2005, "The Multi-agent Simulation Suite (MASS) and the Functional Agent-based Language of Simulation (FABLES)," Ninth Annual Swarm Users/Researchers Conference (SwarmFest 2005), Turin, Italy, June 5−7; available at http://eco83.econ.unito.it/terna/swarmfest2005papers/gulyas_et_al.pdf.

MAML, undated; available at http://www.maml.hu/.

MASON, undated; available at http://cs.gmu.edu/~eclab/projects/mason/.

Parker, M.T., 2001, "What Is Ascape and Why Should You Care?" *Journal of Artificial Societies and Social Simulation* 49(1); available at http://www.soc.surrey.ac.uk/JASSS/4/1/5.html.

Perrone, A., and M. Tenuti, 2002, "VSB — Visual Swarm Builder: A Visual Tool for Swarm Agent-based Environment," in F. Luna and A. Perrone, editors, *Agent-Based Methods in Economics and Finance: Simulations in Swarm,* Kluwer Academic Publishers, Dordrecht, the Netherlands.

Repast, undated; available at http://Repast.sourceforge.net/.

Schelling, T.C., 1978, *Micromotives and Macrobehavior*, W.W. Norton and Co., pp. 147−155.

Swarm, undated; available at http://www.swarm.org/.

Wilensky, U., 1999, *"NetLogo" Software,* Center for Connected Learning and Computer-based Modeling, Northwestern University, Evanston, IL; available at http://ccl.northwestern.edu/netlogo/.

THE REPAST SIMPHONY RUNTIME SYSTEM

M.J. NORTH,* Argonne National Laboratory, Argonne, IL,
and The University of Chicago, Chicago, IL
T.R. HOWE, Argonne National Laboratory, Argonne, IL
N.T. COLLIER, Argonne National Laboratory, Argonne, IL,
and PantaRei Corp., Cambridge, MA
J.R. VOS, Argonne National Laboratory, Argonne, IL,
and the University of Illinois at Urbana-Champaign, Urbana, IL

ABSTRACT

Repast is a widely used free and open-source agent-based modeling and simulation toolkit. Three Repast platforms are currently available, namely, Repast for Java (Repast J), Repast for the Microsoft .NET framework (Repast .NET), and Repast for Python Scripting (Repast Py). Each of these platforms has the same core features. However, each platform provides a different environment for these features. Taken together, the Repast platform portfolio gives modelers a choice of model development and execution environments. Repast Simphony (Repast S) extends the Repast portfolio by offering a new approach to simulation development and execution. The Repast S runtime is designed to include advanced features for agent storage, display, and behavioral activation, as well as new facilities for data analysis and presentation. This paper introduces the architecture and core features of the Repast S runtime system and discusses how Repast S fits within the larger Repast portfolio. A related paper in this Agent 2005 conference proceedings by the same authors, "Repast Simphony Development Environment," discusses the Repast S model authoring system.

Keywords: Agent-based modeling and simulation, Repast, toolkits, model execution, runtime system

INTRODUCTION

Repast is a widely used free and open-source agent-based modeling and simulation toolkit (ROAD 2005; North et al. 2006). Three Repast platforms are currently available, namely Repast for Java (Repast J), Repast for the Microsoft .NET framework (Repast .NET), and Repast for Python Scripting (Repast Py). Each of these platforms has the same core features. However, each platform provides a different implementation environment for these features. Taken together, the Repast platform portfolio gives modelers a choice of model development and execution environments.

Repast Simphony (Repast S) extends the Repast portfolio by offering a new approach to simulation development and execution. Repast S runtime is designed to include advanced features for agent storage, display, and behavioral activation, as well as new facilities for data analysis and presentation. This paper introduces (1) the architecture and core features of the

* *Corresponding author address*: Michael J. North, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439; e-mail: north@anl.gov.

Repast S runtime system, and (2) discusses how Repast S fits within the larger Repast portfolio. A related paper in these conference proceedings (North et al. 2005) discusses the Repast S model authoring system.

It is important to note that Repast S and its related tools are still under development. This paper presents the most current information as of the time of its writing. However, changes may occur before the planned final release.

RELATED WORK

There are a variety of existing agent-based modeling toolkits. Repast J, Repast Py, Repast .NET, NetLogo, and Swarm are just a few examples (ROAD 2005; Wilensky 1999; SDG 2005). The growing agent-based modeling literature suggests that the existing toolkits have been useful for many modelers. However, more is needed:

- There is a need to eliminate the restrictions applied by many existing toolkits such as the need to implement interfaces, extend classes, or manage proxies to access specific toolkit functions.
- There are continuing needs to reduce the distance between modelers and programmers, automate common tasks, and encourage the development of flexible, reusable components.
- There is a need to directly support model and enterprise information system integration.

Repast S is explicitly designed to meet these needs.

DESIGN GOALS

The design goals of Repast S directly address the needs identified in the previous section. The design goals for Repast S are as follows:

- All of the core features and capabilities of Repast J and Repast. NET should be available.
- There should be a *strict separation* between models, data storage, and visualization.
- All user model components should be "*plain old Java objects*" (POJOs) that are accessible to and replaceable with external software (e.g., legacy models and enterprise information systems).
- Common tasks should be *automated* when possible.

- *Imperative* “boilerplate” code[5] should be eliminated or replaced with *declarative* runtime configuration settings when possible.

- *Idiomatic* code expressions[6] should be *simple and direct*.

THE REPAST S MODEL IMPLEMENTATION BUSINESS PROCESS

Based on the design goals in the previous section, the Repast S model implementation business process is as follows:

- The modeler creates model pieces, as needed, in the form of POJOs, often using automated tools.

- The modeler uses declarative configuration settings to pass the model pieces and legacy software connections to the Repast S runtime system.

- The modeler uses the Repast S runtime system to declaratively tell Repast S how to instantiate and connect model components.

- Repast S automatically manages the model pieces based on (1) interactive user input and (2) declarative or imperative requests from the components themselves.

The POJO model components can do anything, but are most commonly used to represent the agents in the model. For example, the model components might represent people, organizations, animals, or markets.

The POJOs can be created using any method. The Repast S development environment plans, in particular, feature a set of Eclipse (Eclipse 2005) plugins that simplify both the creation of model components and the specification of model configurations or settings. This functionality is discussed in these conference proceedings in North et al. (2005).

ATOMS AND ATOMIC BONDS

Components are the “atoms” of Repast S models. The components that make up a Repast S model are managed using settings. Settings play an important role in Repast S, since they link the POJO model components to one another and to the surrounding model. Thus, if components are like atoms, then settings are like atomic “bonds.” These bonds are used to link the atoms or components to form “materials” or models. Just like real atoms, Repast S model components can be combined together in many different ways to form a variety of models. Settings are designed to bond Repast S atoms together by several means, including:

---

5 Much like boilerplate text, boilerplate code is moderate- to large-sized pieces of code that are used again and again in model after model. An example in Repast J is code to setup a master time schedule or configure a network display.

6 Idiomatic code expressions are short pieces of code that act as slang terms or shorthand. An example in Repast J is code to update the weight on a network edge.

- A model and component query capability,
- Network definitions,
- Spatial neighborhood definitions,
- Enterprise data sources, and
- Time scheduling.

Settings are designed to combine with the Repast S remote access and legacy integration tools to allow external models and enterprise information systems to be treated much like any other software component. There are two categories of Repast S settings, both of which are stored in XML files:

- Model descriptor settings are designed to specify the types of components (e.g., Java classes or legacy models) and the kinds of relationships (e.g., networks) that are allowed in a given model, as well as the declarative triggers or "watchers" to be created (e.g., to let human agents know when there is another human agent near them). Model descriptors define what can be present in a model.

- Scenario descriptor settings are designed to specify such things as the data source for the model (e.g., a set of files), the visualizations for the model (e.g., the two- or three-dimensional agent displays), and the logging to be performed (e.g., results files or charts). Scenario descriptors define what is present in a model.

Model descriptor settings are designed to use reactive planning mechanisms to facilitate the declarative specification of models. The reactive planning mechanisms can include rich event triggers or watchers that use queries to define the components to be watched and implicit declarations to specify the component to be notified. Watchers allow components to declaratively monitor changes in other components.

INTERFACE

The Repast S runtime interface is shown in Figure 1. A very basic model called the "Simple Happy Model" is displayed. The lower side panel titled "Scenario Tree" shows the scenario descriptor, while the buttons and menu on the top left show the simulation controls. The model descriptor (not shown) is used to determine what agent types (e.g., simple happy agents), agent relationships (e.g., a network), and watchers (e.g., update our happiness level whenever a linked friend updates its happiness) are present in the model. Finally, a three-dimensional visualization of a small agent friendship network is shown in the lower right.

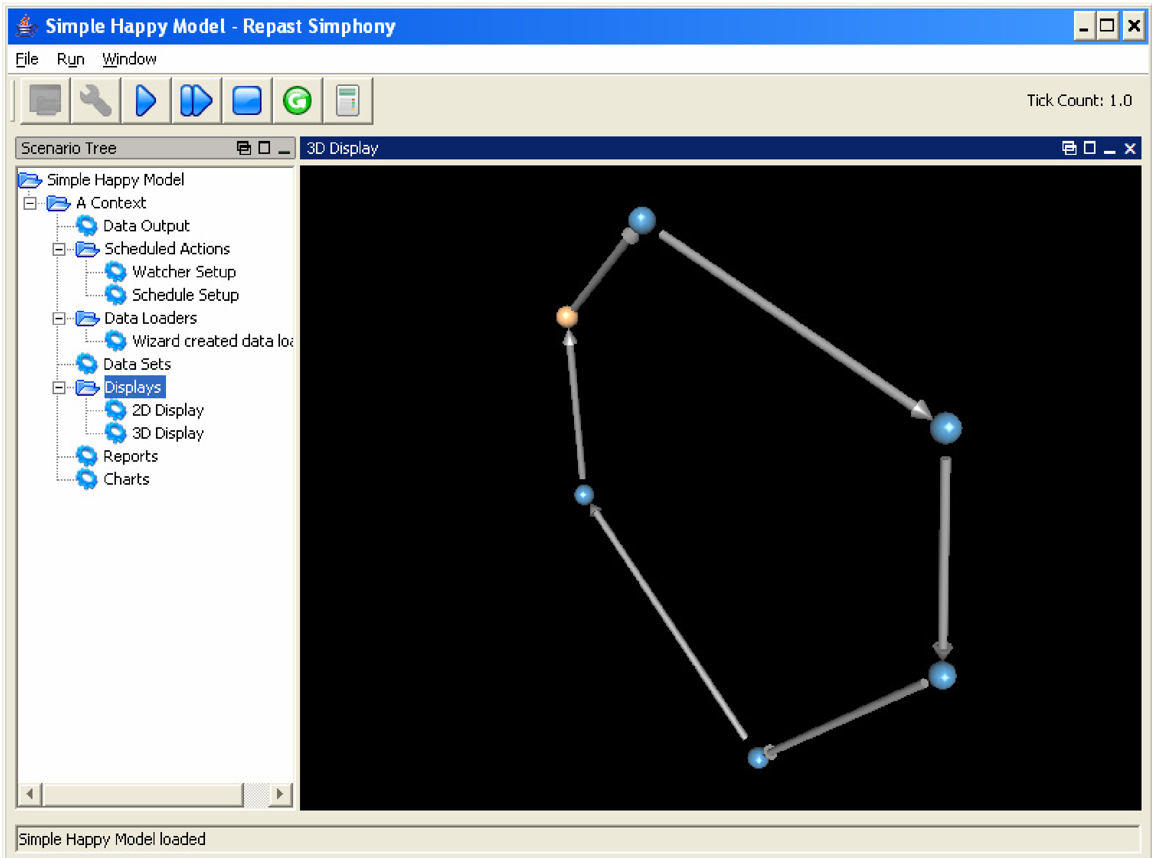

FIGURE 1 The Repast S runtime interface with the Simple Happy Model example

The interface shown in Figure 1 was built using the Simphony Application Framework (SAF). The SAF is a pure Java application development framework intended to simplify the creation of graphical desktop applications. SAF uses Java Foundation Classes, FlexDock (FlexDock 2005), and the Java Plugin Framework (JPF) (JPF 2005) to present and manage application plugins. SAF includes plugin life-cycle management and useful features such as tear-away and dockable windows. Figure 1 shows a set of Repast S windows in the SAF main docking frame. Figure 2 shows a tear-away two-dimensional display. SAF was developed by the Argonne Repast team to support systems such as Repast S. SAF is expected to be released as a separate free and open-source project that supports Repast S.

The Repast S runtime system is designed to include a range of built-in features beyond the previously discussed setting management system. These extended features are organized into Repast S runtime plugins. The planned plugins are expected to work within the SAF-based Repast S runtime interface. The runtime plugins are designed to provide functions such as statistical analysis, interactive charting, and narrative logging. For example, the Repast S plugin

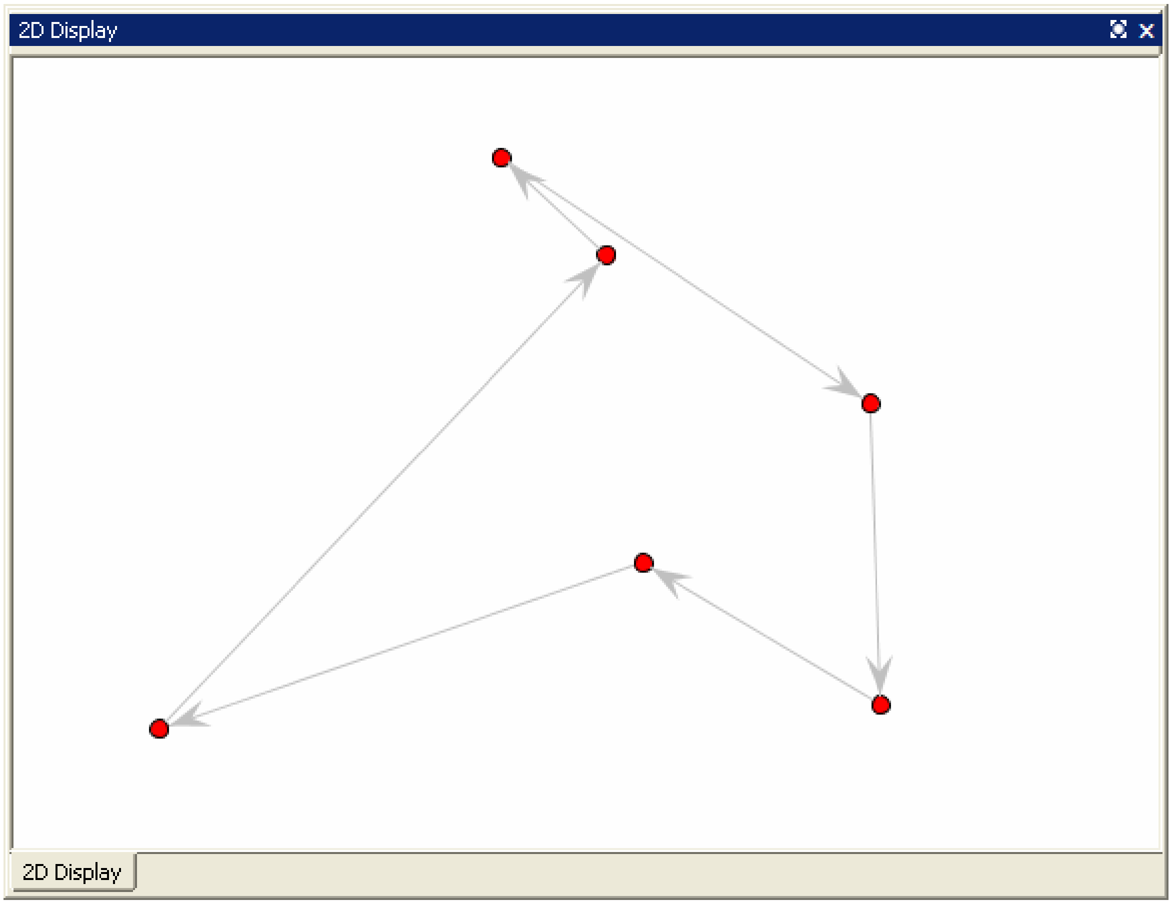

FIGURE 2 A Repast S "tear-away" window with a two-dimensional network visualization

for the R statistics package (R 2005)[7] automatically takes outputs from the Repast S high-performance logging framework and provides them to R Commander[8] (Fox 2005; R Commander 2005) based on interactive user requests. The user provides such requests declaratively in the scenario specification. In this way, Repast S runtime is open to extension, and thus more advanced users are free to add whatever features they may feel necessary.

CONCLUSIONS

The Repast S runtime system is a pure Java extension of the existing Repast portfolio. As discussed above, the Repast S runtime system is designed to provide exciting new features and capabilities to the Repast family. However, Repast S does not replace the existing tools, but

---

7 R and R Commander are both covered by the GNU General Public License (GPL). Software that uses GPL software is itself expected to be covered by the GPL, unless the GPL software is invoked as a clearly identified and separate program. The Repast S plugin architecture can invoke clearly identified and separate programs using a point-and-click interface both for flexibility and to meet licensing requirements such as those for the GPL.

8 R Commander is a point-and-click shell for the R statistical package.

rather is complementary to them. Many users of Repast J, Repast Py, and Agent Analyst may find that porting their models to Repast S is quite straightforward and results in significant simplification and substantial increases in flexibility. Of course, not all existing models should be ported. The users of these models should rest assured that every effort is expected be made to maintain their platforms.

ACKNOWLEDGMENT

The authors wish to thank David L. Sallach for his visionary leadership in founding the Repast project and Charles M. Macal for sustaining involvement in the project. This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Eclipse, 2005, *eclipse*, Eclipse home page, Eclipse Foundation Inc., Ottawa, Ontario, Canada; available at http://www.eclipse.org.

FlexDock, 2005, *flexdock*, FlexDock home page; available at https://flexdock.dev.java.net.

Fox, J., 2005, "The R Commander: A Basic-Statistics Graphical User Interface to R," *Journal of Statistical Software*, Vol. 14, I. 9, Sept.

JPF, 2005, *Java Plug-in Framework (JPF) Project*, JPF home page; available at http://jpf.sourceforge.net.

North, M.J., T.R. Howe, N.T. Collier, and R.J. Vos, 2005, "Repast Simphony Development Environment," in C.M. Macal, M.J. North, and D. Sallach (eds.), *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms*, ANL/DIS-06-1, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 13–15.

North, M.J., N.T. Collier, and J.R. Vos, 2006, "Experiences Creating Three Implementations of the Repast Agent Modeling Toolkit," *ACM Transactions on Modeling and Computer Simulation* **16**(1):1–25, ACM, New York, NY, Jan.

R, 2005, *The R Project for Statistical Computing*, R Project home page, R Foundation for Statistical Computing; available at http://www.r-project.org.

R Commander, 2005, *R Commander Project*, R Commander Project home page; available at http://socserv.socsci.mcmaster.ca/jfox/Misc/Rcmdr.

ROAD, 2005, *Repast*, Repast home page, Repast Organization for Architecture and Design, Chicago, IL; available at http://repast.sourceforge.net.

SDG, 2005, *Welcome to the Swarm Development Group Wiki!*, Swarm home page, Swarm Development Group, Santa Fe, NM; available at http://www.swarm.org/wiki/Main_Page.

Wilensky, U., 1999, *NetLogo*, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

THE REPAST SIMPHONY DEVELOPMENT ENVIRONMENT

M.J. NORTH,* Argonne National Laboratory, Argonne, IL,
and The University of Chicago, Chicago, IL
T.R. HOWE, Argonne National Laboratory, Argonne, IL
N.T. COLLIER, Argonne National Laboratory, Argonne, IL,
and PantaRei Corp., Cambridge, MA
J.R. VOS, Argonne National Laboratory, Argonne, IL,
and the University of Illinois at Urbana-Champaign, Urbana, IL

ABSTRACT

Repast is a widely used free and open-source agent-based modeling and simulation toolkit. Three Repast platforms are currently available, each of which has the same core features but a different environment for these features. Repast Simphony (Repast S) extends the Repast portfolio by offering a new approach to simulation development and execution. The Repast S development environment is expected to include advanced features for agent behavioral specification and dynamic model self-assembly. This paper introduces the architecture and core features of the Repast S development environment. A related paper in the *Agent 2005* conference proceedings by the same authors that is titled "Repast Simphony Runtime System" discusses the Repast S model execution environment.

Keywords: Agent-based modeling and simulation, Repast, toolkits, development environments

INTRODUCTION

Repast is a widely used free and open-source agent-based modeling and simulation toolkit (ROAD 2005; North et al. 2006). Three Repast platforms are currently available, each of which has the same core features but a different environment for these features.

Repast Simphony (Repast S) extends the Repast portfolio by offering a new approach to simulation development and execution. The Repast S development environment is expected to include advanced features for agent behavioral specification and dynamic model self-assembly. This paper introduces the architecture and core features of the Repast S development environment. A related paper in these conference proceedings (North et al. 2005) discusses the Repast S runtime environment.

It is important to note that Repast S and its related tools are still under development. This paper presents the most current information as of the time it was written. However, changes may occur before the planned final release.

* *Corresponding author address*: Michael J. North, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439; e-mail: north@anl.gov.

RELATED WORK

There are a variety of existing agent-based modeling toolkits. Repast J, Repast Py, Repast .NET, NetLogo, and Swarm are just a few examples (ROAD 2005; Wilensky 1999; SDG 2005). The growing agent-based modeling literature suggests that the existing toolkits have been useful for many modelers. However, as outlined in North et al. (2005), more is needed. In particular, with respect to development environments, there are continuing needs to reduce the distance between modelers and programmers, automate common tasks, and directly support model and enterprise information system integration. The Repast S development environment is explicitly intended to meet these needs.

THE REPAST S MODEL IMPLEMENTATION BUSINESS PROCESS

As discussed in North et al. (2005), the Repast S model implementation business process is as follows:

- The modeler creates model pieces, as needed, in the form of plain old Java objects (POJOs), often using automated tools.

- The modeler uses declarative configuration settings to pass the model pieces and legacy software connections to the Repast S runtime system.

- The modeler uses the Repast S runtime system to declaratively tell Repast S how to instantiate and connect model components.

- Repast S automatically manages the model pieces based on (1) interactive user input and (2) declarative or imperative requests from the components themselves.

The POJO model components can represent anything, but are most commonly used to represent the agents in the model. The POJOs can be created using any method. This paper discusses one powerful way to create POJOs for Repast S, namely, the Repast Simphony development environment. However, any method from hand coding, to wrapping binary legacy models, to connecting into enterprise information systems can be used to create the Repast S POJO model components.

Regardless of the source of the POJOs, the Repast S runtime system is used to configure and execute Repast S models. North et al. (2005) details the Repast S runtime system. In summary, the Repast S runtime design includes:

- Point-and-click model configuration and operation;

- Integrated two-dimensional, three-dimensional, geographical information system (GIS), and other model views;

- Automated connections to enterprise data sources; and

- Automated connections to powerful external programs for statistical analysis and visualization of model results.

ANNOTATIONS AND SETTINGS

Repast S uses a new feature in Java 5, namely, annotations, to declaratively mark code for later operations. Annotations are metadata tags that are compiled into binary class files. Like comments, annotations are not directly executed. Unlike comments, annotations are stored in the compiled versions of source code. This storage allows executing Java programs such as the Repast S runtime system to read and act on the encoded metadata. This allows Repast S developers to declaratively mark or annotate code at design time for special processing by the Repast S runtime system. This facility is used for tasks such as declaring “watchers.” The example in Figure 1 shows an agent behavior that is activated any time a connected network neighbor or friend changes its “happiness” attribute. Watchers are considered further in North et al. (2005). Annotations are also used for tasks such as scheduling, as shown in Figure 2.

Repast S is expected to use two major types of settings, namely, model and scenario descriptors, to glue or bond models together. Model descriptors define what *can be* in a model, such as the allowed agent types, permitted agent relationships, and watching information. Scenario descriptors define what *actually is* in a model, such as agent data sources, visualizations, and logging. Model and scenario descriptors are stored in XML files. Descriptors are discussed in North et al. (2005).

```
@Watch(watcheeClassName = “repast.user.models.SimpleHappyAgent”,
watcheeFieldName = “happiness”, query = “linked_from”,
whenToTrigger = WatcherTriggerSchedule.LATER, scheduleTriggerDelta = 1,
scheduleTriggerPriority = 0)
public void friendChanged(SimpleHappyAgent friend) {
     if (Math.random() > .25) {
          this.setHappiness(friend.getHappiness());
     } else {
          this.setHappiness(Random.uniform.nextDouble());
     }
}
```

FIGURE 1 An example “Watcher” annotation for a simple happy agent method

```
@ScheduledMethod(start = 1, pick = 1)
   public void changeHappiness() {
      this.happiness = Random.uniform.nextDoubleFromTo(0, 1);
   }
```

FIGURE 2 An example “Scheduler” annotation for a simple happy agent method (“Start” is the time step to call the method, and “Pick” indicates random selection of one of the available simple happy agents)

Model descriptors are to be created at model development time, while scenario descriptors are expected to be created at runtime. The Repast S development environment is expected to provide both a wizard for creating and a point-and-click editor for modifying model descriptors. The Repast S runtime environment includes a point-and-click panel for creating and maintaining scenario descriptors.

THE WEASELS

The Repast S development environment is expected to use the CodeWeasel Eclipse plugin system. CodeWeasel is a set of Java plugins for the Eclipse development environment (Eclipse 2005). CodeWeasel is being developed by the Argonne Repast team to support systems such as Repast S. CodeWeasel is expected be released as a separate free and open-source project that supports Repast S.

CodeWeasel is expected to work within Eclipse to automate and simplify the creation and maintenance of Java code. Eclipse itself contains powerful tools to create and maintain Java packages and classes. CodeWeasel contains tools that augment and fill in gaps in these existing functions. The guiding design rule for CodeWeasel is that only pure Java files are used. No *separate* non-Java metadata or state files are allowed to store user code. Currently, the main members of the CodeWeasel family are MethodWeasel, FieldWeasel, and LegacyWeasel. Additional tools may also be introduced.

MethodWeasel has a wizard and visual editor for Java methods. The wizard, shown in Figure 3, provides a point-and-click tool for creating new method signatures (e.g., method name, method parameters, and method return type) and specifying method annotations. The wizard currently builds on the free and open-source Eclipse Plug-in Method Wizard (Hawlitzek 2005). The visual editor represents Java code as an editable flowchart as shown in Figure 4. The contents of the flowchart are planned to have a direct correspondence with Java code so the editor can work with almost any standard Java 5 code, regardless of whether or not it was originally created with MethodWeasel. The visual editor also is expected to provide point-and-click tabs for modifying method signatures and changing method annotations.

Much like MethodWeasel, FieldWeasel has a wizard and visual editor. Unlike MethodWeasel, FieldWeasel works with Java fields. The wizard provides a point-and-click interface for creating and initializing new fields, as shown in Figure 5. As before, the wizard currently builds on the free and open-source Eclipse Plug-in Method Wizard (Hawlitzek 2005). The visual editor is expected to allow point-and-click editing of the field signature and annotations, and is expected to represent the field’s initialization code as an editable flowchart.

LegacyWeasel is a model integration tool that is expected to allow developers to use straightforward XML documents to specify the format of legacy files (e.g., existing text and binary files) and the source or destination of the data in Java (e.g., the Java classes that produce or consume the data). These XML files are expected to be creatable using a point-and click editor within Eclipse. Once the appropriate XML files are created, LegacyWeasel is expected to automatically convert the given data sources (e.g., Java objects) into input files; run or activate the legacy model or programs; and then read the resulting output file contents back into the appropriate destinations (e.g., the same or different Java objects). In addition to the goal

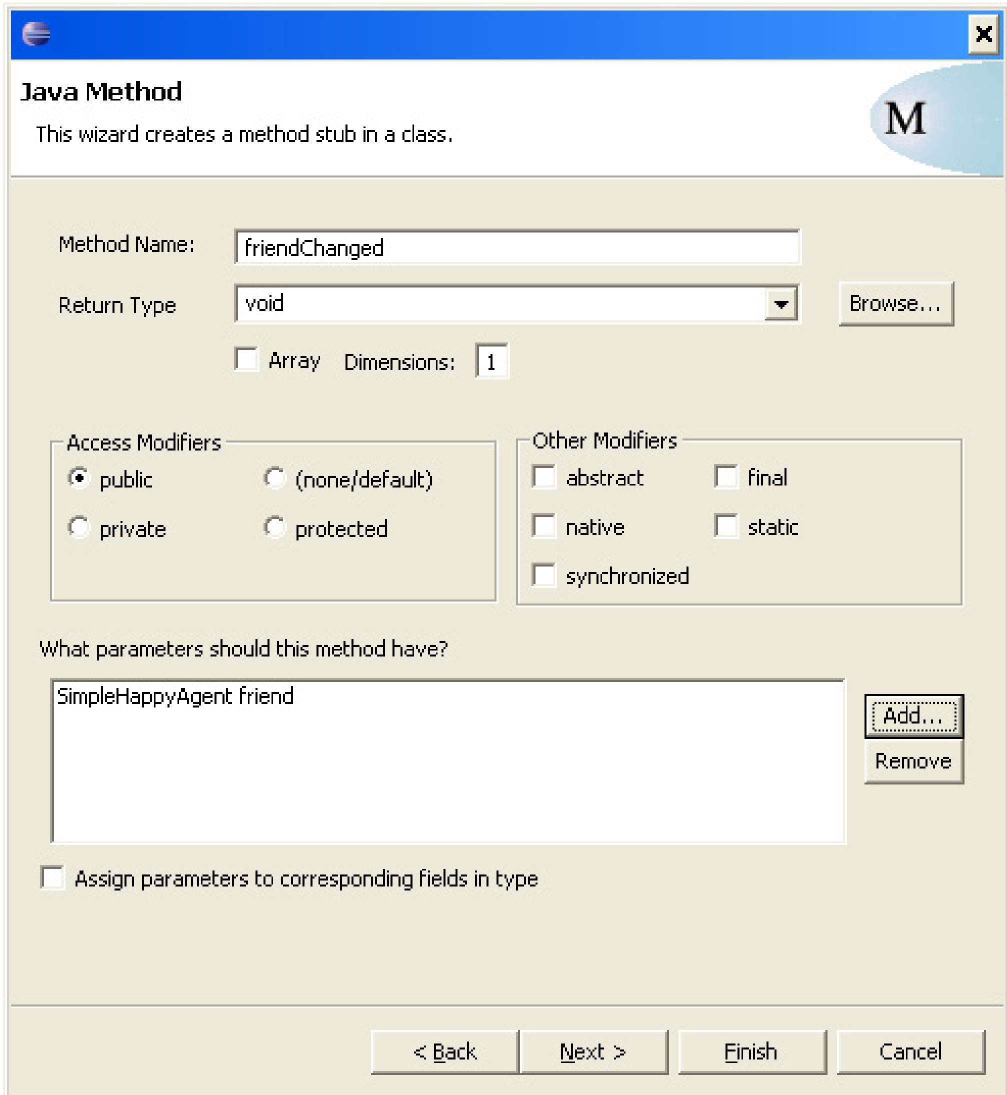

FIGURE 3 One of the pages in MethodWeasel's method wizard

of largely automating and greatly simplifying the often tedious model integration process, LegacyWeasel XML files have the potential to be used as detailed documentation on the format and content of legacy model input and output files.

MODEL TOOLS

As previously stated, CodeWeasel is being developed by the Argonne Repast team to support systems such as Repast S. The various tools within CodeWeasel are expected to provide model developers with a range of useful functions. In addition to these general-purpose tools, the Repast S development environment is expected to include a set of specific support tools. These tools are expected to include a new model wizard, a new model specification file wizard, and a

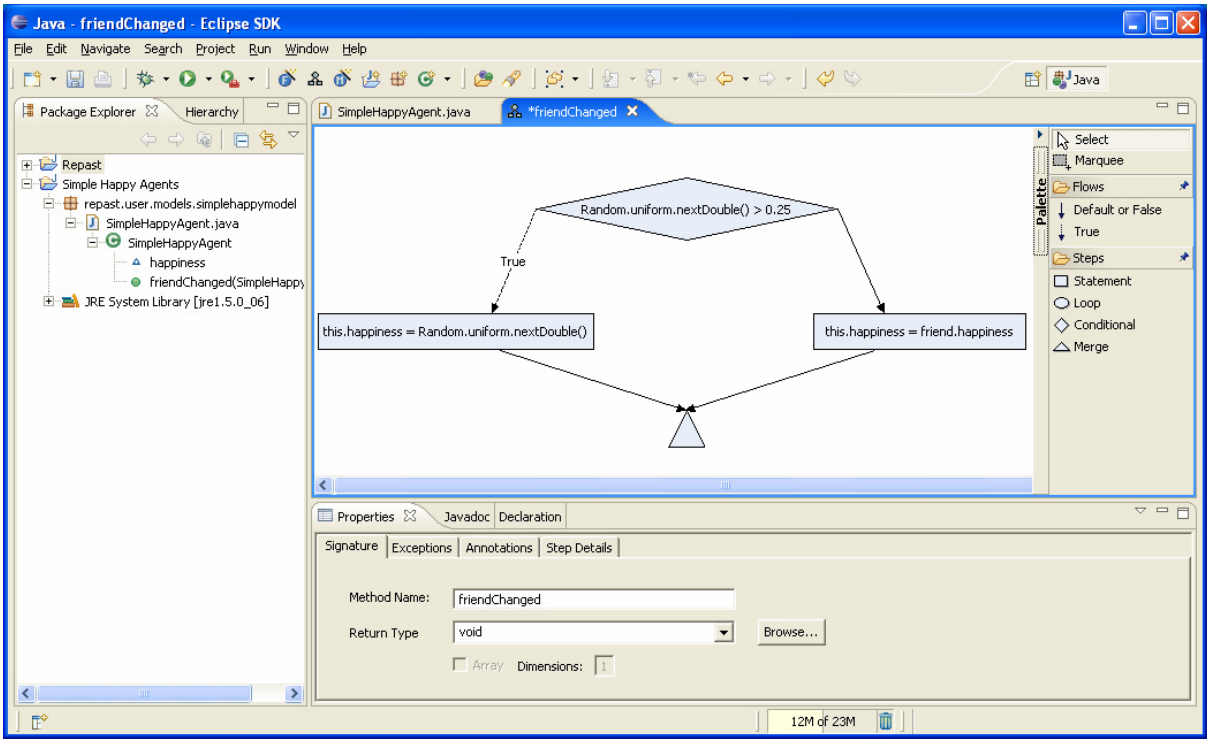

FIGURE 4 The MethodWeasel method editor

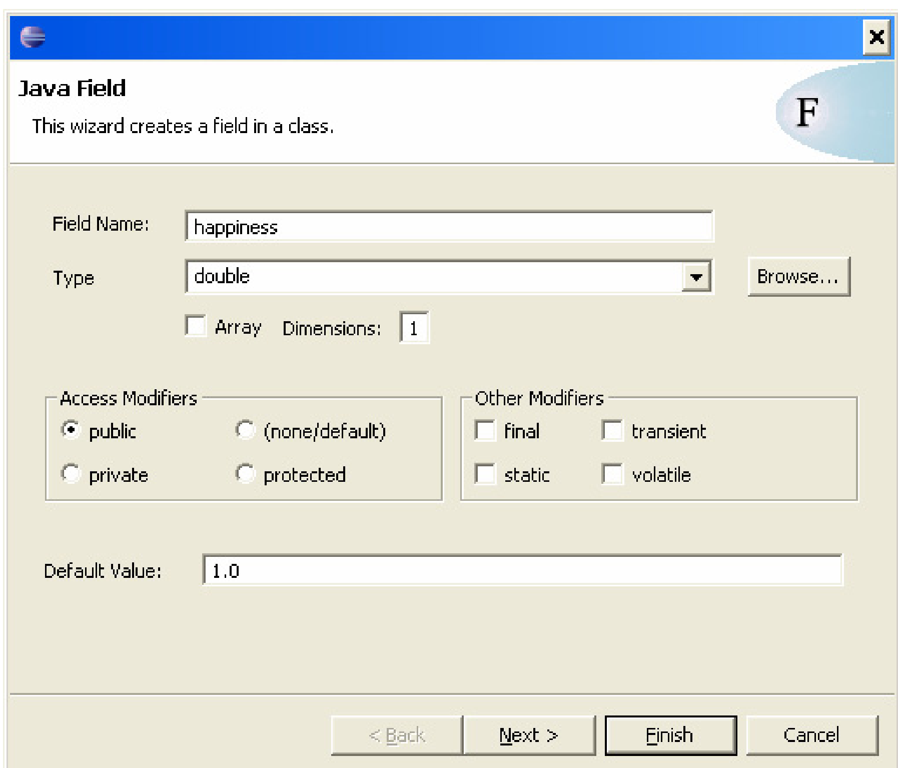

FIGURE 5 One of the pages in FieldWeasel's field wizard

model specification file editor. The new model wizard is expected to be used to create new Repast S models (i.e., Repast S projects) in Eclipse. The new model specification wizard is expected to be used to create model specification XML files, while the model specification editor is expected to allow these files to be updated on a point-and-click basis. Additional tools may also be provided.

CONCLUSIONS

The Repast S runtime is a pure Java extension of the existing Repast portfolio. Repast S extends the Repast portfolio by offering a new approach to simulation development and execution. The Repast S development environment is expected to include advanced features for agent behavioral specification and dynamic model self-assembly. Any POJO can be a Repast S model component. This paper discusses one powerful agent modeling-focused way to create such POJOs, namely, the Repast S development environment. However, any method from hand coding, to wrapping binary legacy models, to connecting into enterprise information systems can be used to create Repast S model components. Once the model components are created, Repast S is expected to provide a set of point-and-click tools for binding the components into working models.

ACKNOWLEDGMENT

The authors wish to thank David L. Sallach for his visionary leadership in founding the Repast project and Charles M. Macal for sustaining involvement in the project. This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Eclipse, 2005, *eclipse*, Eclipse home page, Eclipse Foundation Inc., Ottawa, Ontario, Canada; available at http://www.eclipse.org.

Hawlitzek Consulting, 2005, *Eclipse Plug-in Method Wizard*; available at http://www.hawlitzek-consulting.de/Java_Downloads/Eclipse_Extensions/Eclipse_Extensions_English/eclipse_extensions_english.html.

North, M.J., T.R. Howe, N.T. Collier, and J.R. Vos, 2005, “Repast Simphony Runtime System,” in C.M. Macal, M.J. North, and D. Sallach (eds.), *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms,* ANL/DIS-06-1, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 13–15.

North, M.J., N.T. Collier, and J.R. Vos, 2006, “Experiences Creating Three Implementations of the Repast Agent Modeling Toolkit,” *ACM Transactions on Modeling and Computer Simulation* **16**(1):1–25, ACM, New York, NY, Jan.

ROAD, 2005, *Repast*, Repast home page, Repast Organization for Architecture and Design, Chicago, IL; available at http://repast.sourceforge.net.

SDG, 2005, *Welcome to the Swarm Development Group Wiki!*, Swarm home page, Swarm Development Group, Santa Fe, NM; available at http://www.swarm.org/wiki/Main_Page.

Wilensky, U., 1999, *NetLogo*. Center for Connected Learning and Computer-based Modeling, Northwestern University, Evanston, IL.

DISCUSSION

METHODS, TOOLKITS AND TECHNIQUES

(Toolkits, Thursday, October 13, 2005, 2:15–3:45 p.m.)

Chair and Discussant: *William Rand, University of Michigan and Northwestern University*

FABLES: Functional Agent-based Language for Simulations

William Rand: This session is devoted to toolkits, specifically full toolkits that have been developed by a couple of the participants in the workshop. László Gulyás will do the first session on his FABLES toolkit. Next, Mike North and some of the people from Repast will present their newest Repast version. I'd like to introduce Gulyás, who comes all the way from Budapest to be with us today to give us a talk on his work.

László Gulyás: I'm not in the easiest position to give this talk after Steve's very exciting talk about how we should be doing things. The best I can do is to refer to one of his first slides about the spiral model of development, which required us to get in touch with reality and the real world as often as possible. I will try to get in touch with reality because this is work in progress. We are definitely not following the rotiform model because we have some nice ideas and plans. This is one in a series of talks I'm giving about this topic.

I also want to mention that I have a double appointment, or double affiliation, that might not be clear to everyone. AITIA, Inc., is a small research-oriented IT company in Hungary. I have an appointment with them, and we do joint work with the university on this simulation project. AITIA also has a foothold in the United States, if you want to collaborate with us.

Let me give you the outline of my talk. First, I will say a few words about why somebody would want to develop yet another tool. Then I'll give a short review of the project to show you where we are heading and where we are now. Finally, I will talk about FABLES, the actual language, which is the topic of this presentation. The language, though, is part of a larger project, which is named MASS [Multi-Agent Simulation Suite]. If I have time, I will give you a few screenshots about that, too.

I have one more comment on this functional agent-based language for simulations. Of course, it's a play with words, but functional actually means a paradigm in programming; it's functional programming. It's not really about being very useful for anything, although I do think it's not bad.

[Presentation]

Unidentified Speaker: [Unintelligible]

Gulyás: Yes, this is the syntax of the schedule. I should have mentioned that. It means that at time step 0 you do this, at time step 1 you do this, and so on. If it's cyclic, as it was here, it means that once this agent is created, in the next time step, this is going to be executed. You name the schedule, and you even tell the time step when this should be executed. The line here says that we must add the new evens to the mean schedule, which is going to be the trigger, and this is going to be executed at the actual time step +1.

[Presentation Concludes]

Rand: We can take some questions now. I was going to allude to the very thing that you were talking about — having easier-to-use implementations of agent-based modeling. Of course the problem is that there's a trade-off between ease of use and power, and I alluded to that in my talk earlier. It would be nice to be able to scale up from that easy-to-use version to a more complicated version as you became more familiar with the environment and better at programming or writing computational models. Is there any plan in the future to have FABLES creating the Java code that's the same Java code that your integrated environment is producing on the back end?

Gulyás: Yes, we actually have more plans, but we'll see how far we get. In the long run, we would like not only to comply to Mac or to this core, but also to J or Repast or something. Complying to J and to Repast or MASON should be easy enough, but we'll see.

Michael North: Mike North, Argonne. That was a very interesting presentation. The question I have has to do with the projected license for this. I notice that in the run-time part, you had something about a free demo license. How do you anticipate the licensing to go in the long run?

Gulyás: Yes, that is a good question. I don't know the answer because I'm not the one who is going to give the license. I've had some reassurance that it's going to be a free license, at least for educational use. It could be the same for research use as well, I hope.

Ana Carrie: How is the web interface constructed? Is it written in Java?

Gulyás: It's based on Tomcat, which is basically a Java server-based environment. The code is in Java and the interface is in GSP and servers.

Carrie: Does that mean that it constructs the interface directly from the model?

Gulyás: Do you mean the visualization interface?

Carrie: No, the admin one.

Gulyás: The admin is programmed so it has a database of all the parameters and such, and then it's created on the fly.

Rand: We have time for one more question.

Joanna Bryson: I waited till the end to see if anyone else mentioned this point. This is a little tangential, but at one point you said, and I agree, that the code in a model is much longer than what is in a paper. Obviously, part of the problem may well be something that you're addressing. It might also be more of the duality that got mentioned — that we actually wind up doing quite a lot to make our models work, and then we have to have some simple way to explain it.

North: I think that's a very good point, but I would suggest that maybe it's a combination of things. I definitely think there's a lot of what I call boilerplate code, which is just the same old code again and again — always the same thing like getters and setters if people know what those are. It's always the same; there's nothing new there. There's no information, but I think that it also has to do with the level of detail. There's usually a lot of detail that's missing. If you actually wrote it out, it could turn out that the code is 20 pages, but the real English specification's 10 pages if you get all the details in there. It turns out that there's — and that's the other half — a lot of summarization in the shortest ones. I totally agree with you that there's a lot that we can squeeze out, though, and the code is much more verbose than it needs to be.

Gulyás: I agree. My point, though, is that I think even a model as simple as heatbugs takes in most of the environment in a couple of pages. If you can squeeze that form, I hope it's going to be easier to understand more complex models. It's not going to be as simple as just summarizing it.

Rand: Yes, there's always been some great work on validation. It's nice to see some work being done on verification in the models as well, which is also a very important aspect. With that, we can thank László for his comments.

The Repast Simphony Runtime System

The Repast Simphony Development Environment

Rand: I now have the pleasure of introducing Mike North.

North: How many people here have heard of Repast? Hopefully, everyone has. Today, I'm going to talk to you about the latest work that we've done on Repast. It's not just me. It's also Richie Vos, Nick Collier, and Tom Howe, in no particular order. We've spent the last couple of days tightening up the code and getting it ready to go. We now have a more powerful and easier-to-use version of Repast.

[Presentation]

North: Now I'm going to hand things over to Tom.

Tom Howe: I'm going to talk about the philosophy behind how, why, and what Mike just did. Building on some papers from this conference from a few years ago — one by László Gulyás and one by myself and Mark Diggory — about the necessity of being able to switch the kinds of relationships you're working with. If you write a model that is on a grid, you should also be able to test it in a social network environment. You should be able to test it on a GIS environment without having to record your entire simulation.

[Presentation Continues]

Howe: I want to demonstrate that you can actually initialize your agents from this particular step. I'm going to add a whole bunch of these very quickly. These are instances. Each one of these dots represents an instance of an agent. If I did 100 of them, which I'm not going to do because of the time it would take, I would have 100 agents. I'm actually not instantiating the objects right now. I'm creating a descriptor for my data source so that when I do instantiate the model, it creates all the agents but stores all these data bits.

North: You can use descriptors for things like Monte Carlo simulation where you don't have to re-create the thing from scratch every time. You've got this background description, and you can use that to instantiate as many versions of the model as you want. Normally, we would use human-readable labels. I didn't put the labeling on it. That's why you get the full class name, but we can also add it with the labeling.

Howe: That's all well and good. I have my agents in there, but I also want to set up the relationships between them. There's another little tool that allows you to build the relationships. This works exactly the same way when you're choosing a network, and you can set the strength to whatever you want.

[Presentation Continues]

Howe: Next, Richie Vos is going to do data output, charts, and fun stuff like that. Richie is also a student at the University of Illinois at Urbana-Champaign.

[Presentation Continues]

Richie Vos: To continue, you have file outputters. They write to files on chart outputters; they'll show a chart and basically anything you want. Those are the ones I'm going to show today; we also tested it and have a JDBC [Java Database Connectivity] outputter, so you can write to a database and so forth.

North: To amplify on that, we have what amounts to a data background, so you're using annotations to mark things for collection. Everything collected is sent into the central data stream. On the other end, you can actually divert those data streams to anywhere you want, and you do that through the GUI. On the development end, you can mark the things for inclusion in the data stream. On the run-time end, you can dynamically mark things for logging to various locations.

We're using the underlying log4j system, but it's an enterprise class data-logging framework. It's a free and open source as well, so we have the ability to log the things that Richie mentioned. The log4j also provides it with a huge number of file loggers such as Excel, HTML, comma-separated value. There's a large list of things that come along just in an enterprise system, and many more of them are being developed. Many of the big Java web servers use that, so it's a very high volume system. We basically get everything into that system automatically. You have to mark it as of interest, which means it goes into the stream. Once it's in the stream, it can go anywhere you want visually, as you'll see.

Vos: As Mike said, you just mark something of interest, and then you can automatically specify when you want this information gathered. The data will be automatically pulled out of your agents and out to these outputters.

[Presentation Continues]

Vos: You're going to have some identifiers for the data, so when you generate the data, you generate it in a table because you have to have a way on the outputter end to map if you want data going to a file and data going to a chart. This data goes here and that data goes there, and these are *x* values and these are *y* values and so forth.

North: The basic idea is that these are samples of the stream that is going out. You can imagine a tube, and this is just slicing off a part of that stream and then naming that slice. You can make as many slices as you want, and those named slices are sent to wherever you want them to go.

Vos: So we just add a new data set, and we'll label this one as "Data Set 1." You can see we have listed the classes of agents. This time we have just one because we only had one type of agent. Since we are adding this data set to this root context, it's going to grab every single agent in the model that is of this type. Later, we'll have more advanced filters if you want it. It's going to take every single agent who is a simple, happy agent and gather data off of them.

[Presentation Continues]

Vos: Now Nick will show you some of the display stuff.

North: Between now and the release, we'll be adding drop-down boxes for those things that were already specified, and we'll include other things like the time. The machinery is all there; it's just a matter of populating those boxes.

Collier: Displays refer to the typical heatbugs display. You see the bugs running around on the screen, and the idea behind doing the displays this way is, as Gulyás briefly mentioned, separating the observer from the model, so you, the agent, might set up the watchers and an initial punch to get it going. Then we can add charts and the displays and everything to peek in to see what's going on. I'm going to add the visualization for those right now.

[Presentation Continues]

Collier: We would have a style generator. A simple style generator is easy to make because you're just mapping to some color. If it's greater than 2.5 red and so on, make it yellow, that kind of thing.

North: That's one of the smaller things we'll be adding between now and the release. Plus if you want to make very sophisticated, very complex styles, you could always use the visual environment and draw the style flowchart as well, although normally we'd just use the wizard that we're going to be adding.

Unidentified Speaker: Is that some kind of functional programming style?

North: I'd give you a different answer. That is to say, it's actually a factoring of concerns, and so it's basically separating the model from the appearance of the model because the way the model looks doesn't really matter to the model. The agents don't care that they're blue. It's the happiness that matters, and so it's really a proper factoring of concerns.

Collier: In fact, if you look at most toolkits that have to do with drawing, that use a style-based approach, then its similar because you're separating out, as Mike says, the concerns. If you're going to run a batch model, no one cares if it's blue in a batch model. You never see it anyway.

[Presentation Continues]

Collier: We want to draw the network. Frequency ν refers to how often we update the network. In this case, we're saying whenever any new nodes, any new objects, are added to this context, we'll update the layout. You can change it to ν at update or at some interval, but ν works well because they all get added in the beginning. They get laid out nicely, and you can see what happens.

"Schedule parameters" is the same thing that Richie mentioned with the chart stuff, so I'll just leave it where it is for now. I'm going to make a 2D one. It's the same kind of things, but a different style for 2D. We attempted to unify them, but it didn't go well. We may unify them again in the future.

North: The key thing is that the display is separate from the model. You have one underlying model, and you're just taking these varying views of that model simultaneously. So you can have 2D, 3D, GIS, all, and several of them all simultaneously.

[Presentation Continues]

Unidentified Speaker: I noticed that it automatically created the legend. Based on your little annotations, it looks like it carried over directly.

North: Yes, the system's designed to do all that work. I can't say it removes all the boilerplate, but it takes a lot of it out because it's automated. Most of the functions are obvious in terms of labeling and everything.

Collier: The last thing I want to demonstrate is that we can pull some neat little tricks with the windows. You can pull them out so that you can see them both at the same time. You can probe as well and it shows up there.

North: Richie's going to talk next about the plug-in architecture. The idea is that we wanted to be able to integrate with a variety of existing tools. We're hoping this is a very powerful agent-modeling toolkit, and that's hopefully how you'll see it. Obviously, lots of other tools are available in the world: statistical tools, data analysis, and data mining are just a very short list. We want to be able to naturally integrate with these tools. Since we have these data streams coming out, why not just generate the data streams to sources that can be read by other tools and automatically used by other tools? It's an obvious thing to do once you have the data streams developed.

Vos: Some people want to customize this GUI. One of the key things we were going for on the GUI was to make sure everything is just a JComponent. We passed around a list of stuff, and people can take that out and show it however they want. For example, a chart will be listed with some information about it, and you can put it in your custom map wherever you want it.

North: That's actually a very good point. What that does mean, though, is that if you don't like some of the details of our interface, you can actually take all the parts and embed it into your own Java Swing interface — your own standard Java interface — anyway you want in a straightforward way.

Bryson: This is interesting, but I'm really interested in verification and validation, too.

North: We should be done very quickly. I'm sorry about that.

Vos: Let me show you one last thing. When we created the outputter, it created some files here. People had asked about hooking into more statistical analysis stuff, and we don't want to reproduce that. So through the plug-in architecture, you've connected to *R*. Here it shows the license for *R*. Just because it's GPL, we have to keep everything separate, and this is just the path string.

Now we're just going to pick, and this will have the files generated by the outputters. We click here, pop it open, and up pops *R* and *R* Commander. This will be simplified later. We load up the data set, and you can see the tick, the value for that agent, and the agent's name. We'll click to show that's working, the line graph; the *x* value can be the tick, and the *y* value the happiness. You get a nice little graph there and then just statistics. You can do a linear regression or whatever you want and tick. That's quickly showing the plug-in stuff, and we plan to provide this to simplify life for people after they've run their model and want to analyze things.

Charles Macal: We'll take one or two very quick questions.

Seth Tisue: Seth Tisue, Northwestern University. What does this mean for the future of Repast Py?

North: We'll continue to maintain Repast Py because quite a bit has been done with Repast Py and written in Repast Py. That's going to continue on in maintenance mode. We want to add some features to it; if people need specific features, we can add those as well. For the most part, in the future, it probably would make sense to use this new tool as soon as it's available. Again, we will continue to maintain Repast Py because we have stuff that we need to maintain there for quite a while.

Unidentified Speaker: What does it live in?

North: It also lives in the Agent Analyst tool for the ESRI platform that uses that underlying run-time system. So just as we maintain Repast.Net or Repast J, we'll maintain a variety of different versions of Repast. The trick is to pick the one that matches your specific needs over time, and they all have strengths and weaknesses in terms of what they can and can't do.

Bryson: I liked what you said at the end, but I had a question about what you said at the beginning. It seems that if you know Java, it is clear what to do, all the wizards and things.

North: Yes.

Bryson: In a way, you've constructed a tree where you can only see the leaves, but the branches are still hidden behind it.

North: Yes, sure.

Bryson: Speaking of verification and validation, have you exposed this to nonprogrammers? In my experience with nonprogrammers, NetLogo is a little too much, too little information. They don't get the whole structure of the code.

North: That's an excellent question, and the best way to put it is that we are still working on the development of this system. We've actually talked to nonprogrammers, and they've laid out what they want in a system. They first said that they'd like a native language interface (i.e., talk in English), but we just can't do that now. We're trying to do the best we can with what we have. We're not claiming it's perfect by any means. We expect that it will evolve and improve over time. You have to remember that there's always Java underneath, so in some fundamental sense, it's no worse than Java because it *is* Java. You can always print the Java file. If you don't like the flowchart, look at the Java.

We're showing control flow, and so in that sense, it is exactly the same information in terms of what's available. I think it's really open and depends on what people do. Some people don't like a visual style. You can just use Java, and everything you see here works fine with that. Other people like the visual layout style, being able to put things into flowchart form.

We also think that there's hand-holding involved, with click and select functions, and we may add a little more to that as well. It helps people because one of the biggest problems people have is not with mastering Java; it's mastering this whole mass of API's and all these functions. and digging up all sorts of Java docs, which is overwhelming for people. This provides some hand-holding for that. You don't have to worry about braces and where the braces go to mark blocks. You just say, "On to the left, true or false." Those sorts of things help people a lot. By themselves the flowcharts do not provide verification or validation. It simply provides another way to write your code, and we think that way might be easier for some people.

Verification and Validation

VERIFICATION AND VALIDATION OF SCIENTIFIC AND ECONOMIC MODELS

R.C. KENNEDY,* X. XIANG, G.R. MADEY, and T.F. COSIMANO,
University of Notre Dame, Notre Dame, IN

ABSTRACT

As modeling techniques become increasingly popular and effective means for simulating real-world phenomena, it becomes increasingly important to enhance or verify our confidence in them. Verification and validation techniques are neither as widely used nor as formalized as one would expect when applied to simulation models. In this paper, we present our methods and results through two very different case studies: a scientific model and an economic model. We show that we were able to successively verify the simulations and, in turn, identify general guidelines on the best approach to a new simulation experiment. We also draw conclusions on effective verification and validation techniques and note their applicability.

Keywords: Verification, validation, simulation, natural organic matter, agent-based modeling, Ramsey problems

INTRODUCTION

The use of simulations to model and study scientific and economic phenomena has the potential to be informative; however, the data produced by simulations are most valuable when they can be both verified and validated. In simple terms, this means the data produced are credible and indiscernible from real-world data. This achievement proves to be very difficult, as most real-world systems contain far more constraints and details than computers allow us to reasonably model, and this situation is even more difficult for agent-based simulations and simulations of social and economic phenomena. This leaves most of our simulations as abstractions of real-world phenomena. Their purposes range from helping us to better understand natural phenomena to allowing us to predict the behavior of a system. With the varying purposes of simulations, verification and validation techniques also vary. The problem is that there is no universal verification and validation process that can be applied to all models. The purpose of our work is to explore and apply verification and validation techniques to two very different case studies. The first case study focuses on a scientific problem: the study of natural organic matter (NOM). It has an agent-based backbone and was written first in Pascal then transformed into Java with Repast. The second case study involves an economic problem: solving Ramsey problems in a stochastic monetary economy. It has a more numerical basis and was written first in Matlab and then in C++. We will compare two unrelated simulations that were each written in different programming languages and then compare and verify results. In addition, we will explore some general guidelines to use as an approach to increasing the confidence of a new simulation.

* *Corresponding author address*: Ryan C. Kennedy, 384 Fitzpatrick Hall of Engineering, University of Notre Dame, Notre Dame, IN 46556; email: rkenned1@nd.edu.

The organization of this paper is as follows. The next section outlines what we mean by verification and validation and introduces some general methods. The section that follows describes various aspects of our first case study, including background, validation, and implementation. Then the same is done for our second case study. A conclusion and some general guidelines are provided in the section following that. Finally, a discussion of future work and references conclude the paper.

VERIFICATION AND VALIDATION PROCESS

Simply put, model verification is getting the model right. This means that the code generating the phenomenon being modeled correctly matches the abstract model. Model validation is getting the right model, meaning that the correct abstract model was chosen and accurately represents the real-world phenomenon. It is important to note that verification and validation are key parts of the model development process. Moreover, they must be performed in tandem for the best results. Effective verification and validation of a model will increase the confidence in the model, making it more valuable. An adapted version of Sargent's (1998) and Huang's (2005) verification and validation process diagram is shown in Figure 1. It has been modified for agent-based scientific and economic simulations.

While there have been many verification and validation studies performed for general engineering purposes, verification and validation studies for agent-based and social science simulations are lacking. Some of this can be attributed to agent-based modeling not being as mature as engineering modeling. The point is that we can adapt what has already been done as well as create new tools to fit the needs of agent-based modeling.

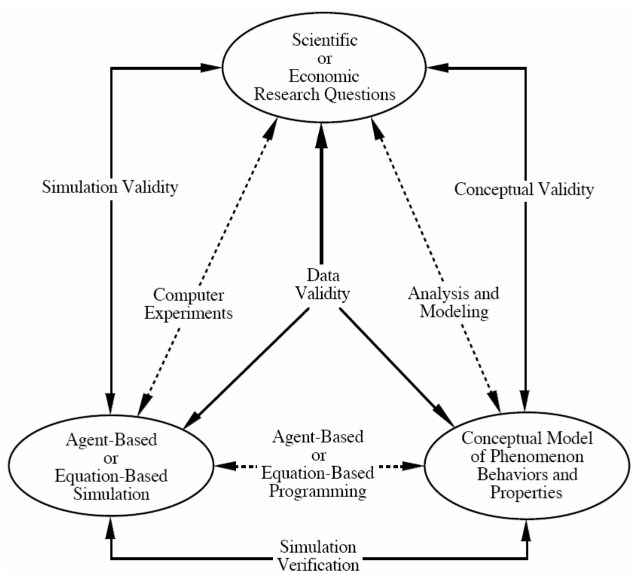

FIGURE 1 A verification and validation process for scientific and economic simulations

Balci (1998) outlined 15 general simulation principles, developed primarily for engineering or management science applications. His principles help engineers and researchers better understand the verification and validation they are performing. This understanding is directly related to model success. A few of his principles that are relevant to scientific and economic modeling are presented next.

1. *The outcome of the simulation model verification, validation, and testing should not be considered as a binary variable where the model outcome is absolutely correct or incorrect.* It is important to realize that models, being abstractions and not absolute representations of phenomena, can never totally and exactly match a system.

2. *Complete simulation model testing is not possible.* As we cannot test all possible inputs and parameters for a system, we must choose the most appropriate ones.

3. *Simulation model verification, validation, and testing must be planned and documented.* Successful planning and documentation are critical and involve the work of many people throughout the lifetime of the system.

4. *Successfully testing each submodel (module) does not imply overall model credibility.* Simply because the modules work well independently does not mean they will work cohesively in a system.

When verification and validation of a model are being performed, it is good to begin by identifying the key principles and techniques to be used for that model. Moreover, planning the verification and validation process, as outlined previously, makes the process more complete and effective. Utilizing Balci's (1998) principles and techniques is a great starting point; from there, model confidence can be improved with further subjective and quantitative methods. We next outline a general verification and validation process that can be adapted to fit many agent-based, social, and economic models. A hierarchy of such methods is shown in Figure 2.

Subjective Methods

Subjective methods largely rely on the judgment of domain experts. They are often used for initial quick-and-dirty validation, but they can also be more formalized. Whatever the purpose, subjective methods typically require less effort than quantitative methods, can detect flaws early in the simulation process, and are often the only applicable verification and validation methods for exploratory simulation studies. We next describe some of the subjective techniques proposed by Balci (1998) that may be applicable to economic and agent-based scientific simulations. His techniques are widely used in validating the models of manufacturing, engineering, and business processes. The following has been adapted from Xiang et al. (2005).

1. *Face validation.* This preliminary approach to validation involves asking domain experts whether the model behaves reasonably and is sufficiently accurate. This is often achieved by evaluating the output or observing a visualization, if applicable.

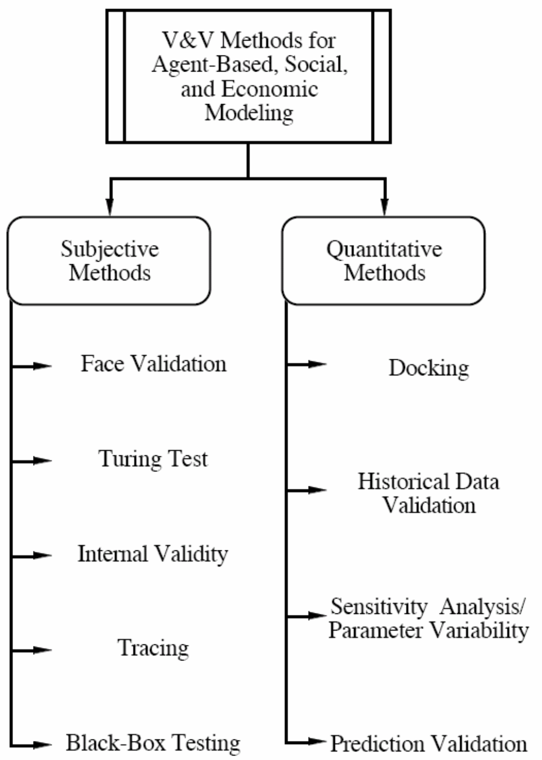

FIGURE 2 Verification and validation methods

2. *Turing test.* This technique is performed by giving domain experts model outputs and real-world outputs and asking them to discriminate them.

3. *Internal validity.* This involves comparing the results of several replications of a simulation, with the only difference being the random seed. Inconsistencies in the results question the validity of some aspect of the model.

4. *Tracing.* Here, the behavior of entities in the model is followed to determine if the logic of the model is correct.

5. *Black-box testing.* This technique involves how accurately the model transforms the input to output in a system.

Quantitative Methods

Incorporating quantitative, or statistical, methods into the validation process can significantly increase the credibility of the model. Model validation is conducted by using statistical techniques to compare the model output data with the corresponding system or with the output data of other models run with the same input data.

The first step to starting quantitative analysis is to determine a set of appropriate output measures that can answer user questions (Xiang et al. 2005). After a set of output measures has been collected, various statistical techniques can be applied to complete the validation process. Time series, means, variances, and aggregations of each output measure can be presented as a set of graphs for model development, face validation, and Turing tests. Confidence intervals and hypothesis tests can be used in the comparison of parameters, distributions, and time series of output data for each set of experimental conditions. These statistical tests can help model developers determine if the model's behavior is acceptably accurate.

The cost of the validation process increases exponentially with the confidence range for a model. There is no single validation approach applicable to all computational models. Choosing the appropriate statistical test techniques and measures of a system is important when conducting a validation process. It is important to note that there is no correct set of statistical tests to use for every simulation; the best results are achieved when tests are carefully chosen according to the model. Some of Balci's (1998) more quantitative techniques that are relevant to our case studies are next described.

1. *Docking*. Docking, or model-to-model comparison or alignment, is used when another model that models the same phenomenon exists or can be created. Docking helps to determine whether two or more models can produce the same results (Axtell et al. 1996). The main idea is that model confidence is significantly improved when two or more models produce the same effective results, particularly if the models were developed independently and with different techniques. In addition, the output from a model can be validated against real-world data.

2. *Historical data validation.* When historical data exist or can be collected, these data can be used to build the model, and the remaining data can then used to determine if the model behaves as the system does.

3. *Sensitivity analysis/parameter variability.* Here, one changes the input values and the internal parameters of a model to determine the effect on the model and its output. Ideally, the relationship in the real-world system should be mimicked in the model. Sensitive parameters that cause significant changes in the model's behavior should be made sufficiently accurate before this model is used.

4. *Predictive validation.* This technique is used to compare the model's prediction with actual system behavior. The system data may come from an operational system or specific experiments, such as from a laboratory or from field experiments.

CASE STUDY I: AN AGENT-BASED SCIENTIFIC MODEL

NOM is a heterogeneous mixture of molecules. NOM plays a crucial role in the evolution of soils, transport of pollutants, and carbon cycle (Cabaniss et al. 2005; Xiang et al. 2005). Its evolution is an important research area in a number of disciplines. NOM is complex; it is made up of molecules with varying molecular weights, reactivity levels, and functional groups. This makes it difficult to model. Performing chemical experiments with NOM is difficult and time-consuming because of its complexity and because of our limited knowledge of its inner workings. The ability to effectively predict NOM behavior as it evolves over space and time would be very valuable to scientists and an accomplishment in the modeling discipline.

Conceptual Model

The NOM conceptual model was based on the work of a chemist working at the University of New Mexico (Cabaniss et al. 2005). He generated his model from extensive observation and experimentation in the laboratory. His basic model outlined the use of the precursor molecules cellulose, lignin, and protein (among others) to be used in a controlled environment where parameters such as light intensity, temperature, and density could be varied. A more detailed description of our model follows and has been adapted from Xiang et al. (2005).

Agents

Our agents are molecules. Each molecule is a representation of its underlying elemental formula, meaning the number of C, H, O, N, S, and P atoms present. This gives rise to a molecular weight for each molecule. Molecules also contain a functional group count, such as the number of alcohol or ester groups present.

Behavior

In our environment, agents can move around a grid, interacting with other molecules and their environment. Molecules undergo chemical reactions on the basis of specific probabilities. Reactions can result in structural changes in the molecule, such as the addition of functional groups. They can also generate new molecules from predecessor molecules, and the predecessor molecules may leave the system. Twelve types of chemical reactions, including first- and second-order chemical reactions, are modeled as described in Table 1. The categories of reactions are as follows:

1. *First-order reactions with a split.* The predecessor molecule A is split into two successor molecules B and C. Molecule B occupies the position of molecule A, while one of the empty cells closest to molecule B is filled with molecule C.

2. *First-order reactions without a split.* The transformation only changes the structure of the predecessor molecule A.

3. *First-order reactions with the disappearance of a molecule.* The predecessor molecule A disappears from the system.

4. *Second-order reactions.* Two molecules A and B are combined to form a new molecule C. Molecule C replaces molecule A, and molecule B is removed from the system.

TABLE 1 Chemical reactions

Name	Type
Ester condensation	Second order
Ester hydrolysis	First order with a split
Amide hydrolysis	First order with a split
Microbial uptake	First order with the disappearance of a molecule
Dehydration	First order with a split
Strong C = C oxidation	First order with a split (50% of the time)
Mild C = C oxidation	First order without a split
Alcohol C-O-H oxidation	First order without a split
Aldehyde C =O oxidation	First order without a split
Decarboxylation	First order without a split
Hydration	First order without a split
Aldol condensation	Second order

Space

In the NOM model, the agents are associated with a location in two-dimensional (2D) geometrical space and can move around their environment. Each cell on the grid can host multiple molecules up to a certain threshold.

Reaction Probabilities

The probability for each reaction type is expressed in terms of intrinsic and extrinsic factors. Intrinsic factors are derived from the molecular structure, including the number of functional groups and many other structural factors. Extrinsic factors arise from the environment and include concentrations of inorganic chemical species, light intensity, availability of surfaces, presence of microorganisms, presence and concentration of extracellular enzymes, and presence and reactivity of other NOM molecules. The intrinsic and extrinsic factors are combined in probabilistic functions.

Molecular Properties

The reactivity of the resulting NOM over time can be predicted on the basis of the distributions of molecular properties, which are calculated from the elemental composition and functional group data. They represent a measurable quantity that can be used as a predictor for environmental function and are useful for the calibration and verification of our conceptual model and simulation.

Simulation Process

The conceptual model is a stochastic synthesis model of NOM evolution, meaning that the state of the system is represented by a set of values with a certain probability distribution, such that the evolution of the system depends on a series of probabilistic discrete events. At each time-step, for each molecule, a uniform random number is generated. This number determines whether a reaction will occur, and if one does occur, its reaction type. After a reaction takes place, the attributes for the current molecule are updated, and the reaction probabilities are recalculated. The molecule structure is changed to alter the outcome of the reaction, and a new probability table entry is added for newly formed molecules, if there are any.

Implementations

The NOM conceptual model was initially implemented in Pascal, resulting in a program for Windows called AlphaStep. Our implementation is coded by using Java (Sun JDK 1.4.2) and the Repast toolkit. Repast is an agent-based simulation toolkit written in Java. It contains a control panel to control and manipulate the model and has rich visualization capabilities. We chose Java for our model because we also incorporate a Web-based front end to the system where users can create and submit simulations, as well as view graphical results.

Validation

We followed the general technique previously outlined when validating the NOM model. We began with subjective analysis and then proceeded with quantitative analysis.

Subjective Analysis

The validation of the NOM model began with face validation of the conceptual model by domain experts. They evaluated the underlying mechanisms and properties and drew their conclusions. After initial face validation was achieved, coding of the agent-based simulation took place. In this step, verification methods, such as code walk-through, trace analysis, input-output testing, pattern testing, boundary testing, code debugging, and calculation verification, were used to verify the correctness of the simulation. Another useful technique used for simulation validation is visualization (Grimm 2002). Visualization is often used in conjunction with face validation. A snapshot of an animated visualization of the flow of molecules through a soil column is shown in Figure 3. A corresponding animated graph shows how the molecular weight distribution shifts with time, initially favoring lower-weight molecules and gradually shifting to larger molecular weights as the simulated time passes. These same behaviors were observed in the laboratory, which increases confidence in the simulation.

A simulation model that uses random number generators must have statistical integrity in that independent simulations with the same input data should have similar results. This is also known as internal validity. If the simulation model produced large variabilities because of the random seeds, there would be a considerable problem with it. To test this, we performed 450 simulations with our NOM simulator, each with a different random seed. We chose the total number of molecules in the system after the simulation had completed as our point for comparison. We found that our simulations produced the expected normal curve upon analysis of the data. Figure 4 shows the histogram for the data. By verifying the independency of the random seeds in the NOM simulator, we were able to conclude that it is statistically robust in terms of repeatability. Further statistical analysis needs to be performed to verify how reliably our simulator conforms to a normal distribution.

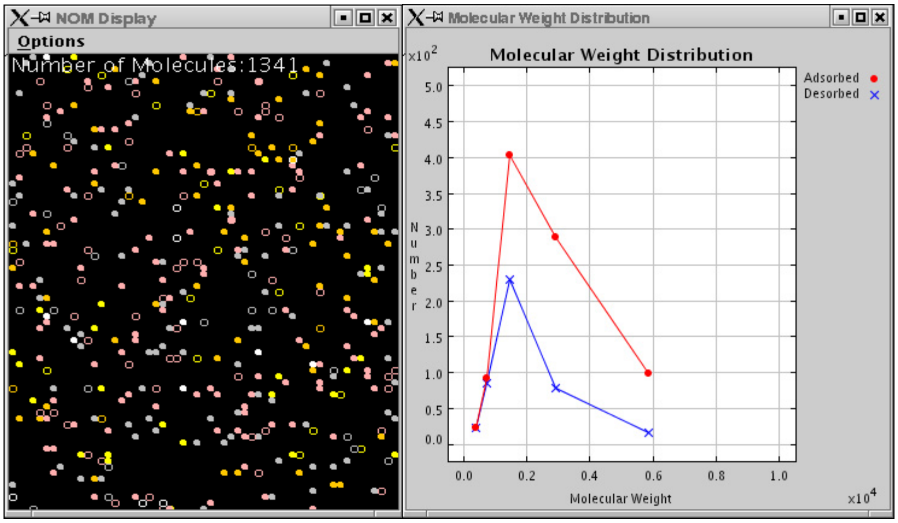

FIGURE 3 NOM visualization

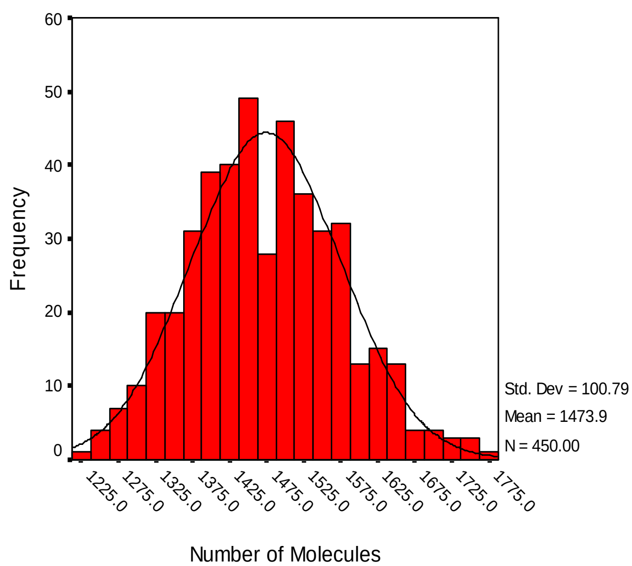

FIGURE 4 Histogram showing distribution after 1,000 simulated hours

Quantitative Analysis

Both our NOM model and the AlphaStep model rely on the same basic conceptual model. However, there are a few inherent differences. First, AlphaStep has no sense of space. Instead, its agents are described as parts of a "well-stirred soup," each equally likely to react with any other molecule. Another key difference is the programming language used in each simulation. We summarize the main differences between implementations in Table 2.

TABLE 2 Implementation differences

Feature	AlphaStep	NOM
Programming language	Pascal	Java
Platform	Windows	Linux
Running mode	Standalone	Web based, standalone
Simulation packages	None	Repast
Initial population	Actual number of molecules	Percentage distribution of molecules
Animation	None	Yes
Spatial representation	None	2D grid
Second-order reaction	Randomly pick one from molecule list	Choose nearest neighbor
First-order reaction with split	Add to molecule list	Find empty cell nearby

To dock these stochastic simulations, we performed 25 replications, each with different random seeds, for both implementations given effectively the same initial conditions. Among the many molecular variables, we chose number of molecules, MWn (number-average molecular weight), MWw (weight-average molecular weight), carbon mass, and carbon percentages as metrics for comparison. We took ensemble averages from 25 replications and compared data points over time. These comparisons are shown in Figure 5. As one can see, visual agreement looks very good; however, statistical testing must be performed to ensure that differences between the models are not significant.

CASE STUDY II: AN EQUATION-BASED ECONOMIC MODEL

Ramsey problems are concerned with setting specific economic variables — money growth and tax rate — to generate the best social welfare for a given economy (Cosimano and Gapen 2005). In our model, nonlinear projection methods are used to solve these problems. The goal is to calculate the real or nominal interest rate for a given economy under the optimal money growth and tax rates. Our model creates a set of residual equations, using bivariate Chebyshev polynomials.

The simulation was initially written for Matlab. It effectively takes advantage of Matlab's built-in functions and capabilities, but execution is slow. The current model works off of four equations and on moderate-sized matrices. The next iteration of the model would include a fifth equation and much larger matrices, making execution in Matlab unpractical. Once the model was verified and validated in Matlab, we converted the code to C++ to make it execute faster.

Conceptual Model and Implementation Differences

Matlab is a matrix laboratory, meaning a rich interactive programming environment that supports many data types best suited for numerical analysis. Matlab, being a high-level language, is very user friendly and has many built-in functions and display options. Matlab is also useful for prototyping. It is, however, inherently slow; it is essentially a software package written in C, and simulations written in Matlab are interpreted, resulting in slow execution. C++, on the other hand, is an object-oriented, lower-level language. Its standard template library incorporates many

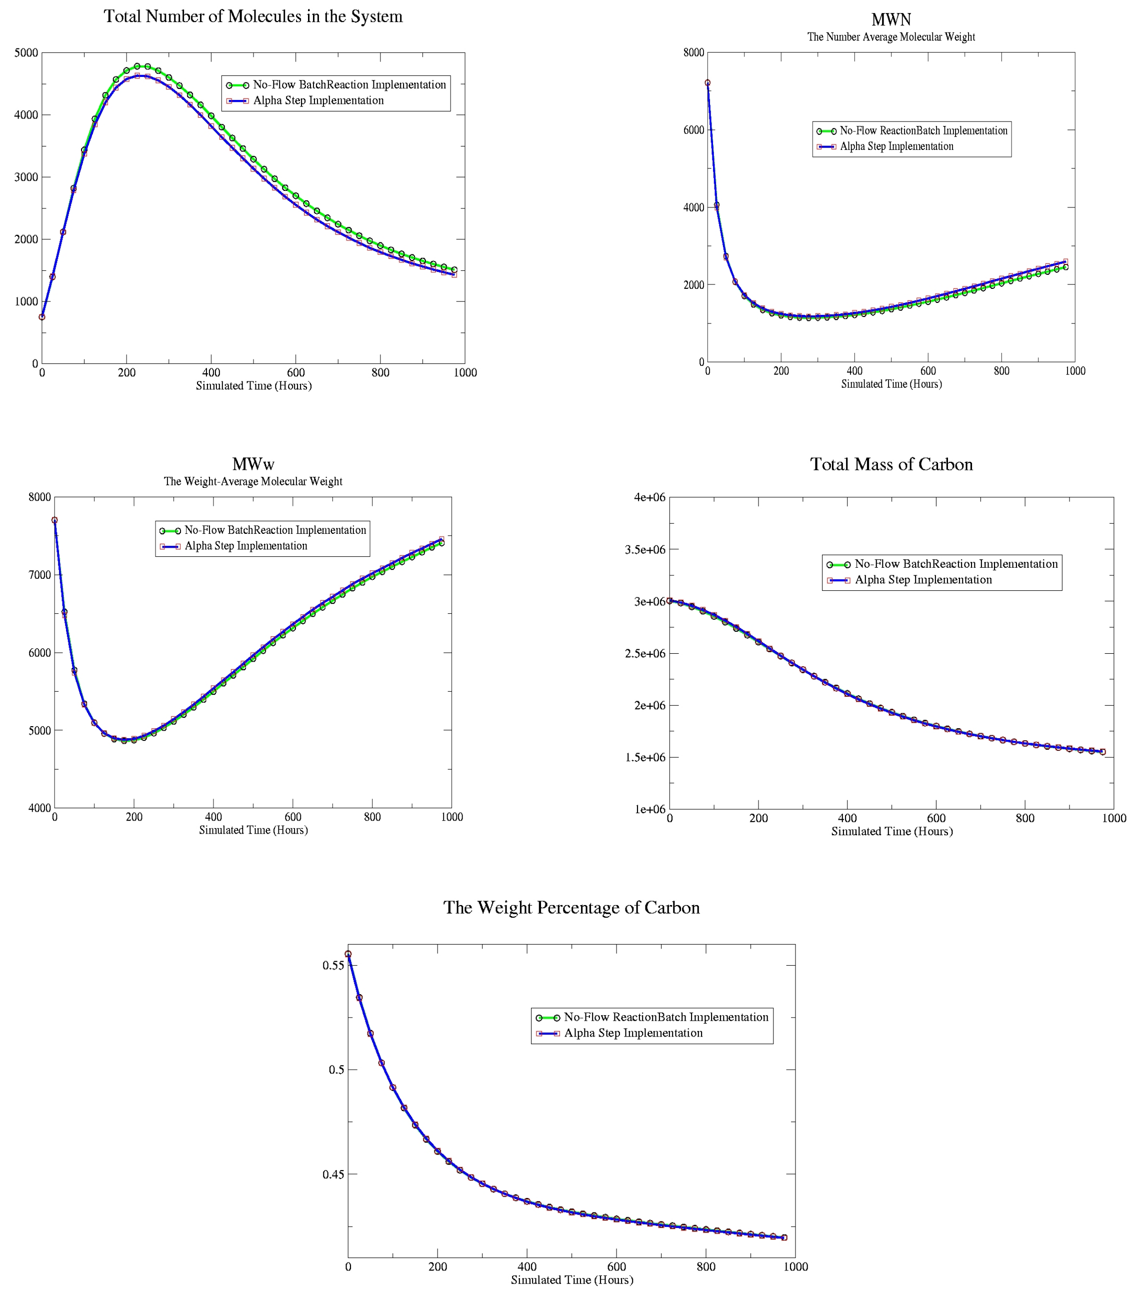

FIGURE 5 Comparisons between AlphaStep and NOM

desirable functions, and it is relatively simple to code. Because C++ has an efficient compiler and is a lower-level language, simulations written directly in C++ run much faster than equivalent simulations written in Matlab.

Converting the simulation from Matlab to C++ was much more difficult than expected. In Matlab, variables are not declared as they are in C++. Instead, most variables are assumed to be arrays. These arrays can contain real numbers, complex numbers, or even other arrays of real numbers, among other things. This creates a big problem in C++, as variables must be declared with a data type. For example, the variable *array1* could represent an array of real numbers at one point in a Matlab program and an array of complex numbers at another point. The idea is that the variable *array1* is, in essence, overloaded to handle many data types. In C++, functions operate and are called differently depending on the type of data passed to them. Overcoming this step was pivotal to porting the code.

Another main difficulty in going from Matlab to C++ was emulating Matlab's built-in functions. The majority of the Matlab functions used in this simulation are part of the LAPACK, or linear algebra package, and include functions such as taking the normal of a matrix or vector, inverting a matrix, conditioning a matrix, etc. Not only are these not inherently included in C++, but they are again overloaded in the sense that you can pass Matlab's max function a vector, a matrix, or a simple set of numbers, and it will give the proper result. While it is possible to call Matlab from within C++ to make use of such functions, the desire of this project was to have everything run in C++ for maximum speed. Determining the inner workings of Matlab's many functions and implementing approximations of them in C++ proved difficult and time-consuming. In the end, the core of the simulations did the same thing, but with an inherently different implementation, requiring a rigorous verification and validation effort.

Performance

Running time for runs of 5, 50, and 500 iterations of the simulation can be found in Table 3 and Figure 6. As evidenced, the speedup is significant, or approximately 30 times faster in C++.

Validation

Validating the economic model was a little different than validating the scientific model described in the first case study. The validation process helped us identify some problems with the C++ implementation, so we were limited in the amount of quantitative analysis that we could perform. In this case, validation served the purpose of identifying what was wrong with our implementation.

TABLE 3 Running time for Matlab and C++ implementations

	5 Iterations	50 Iterations	500 Iterations
Matlab	58 seconds	568 seconds	8,872 seconds
C++	2 seconds	17 seconds	176 seconds

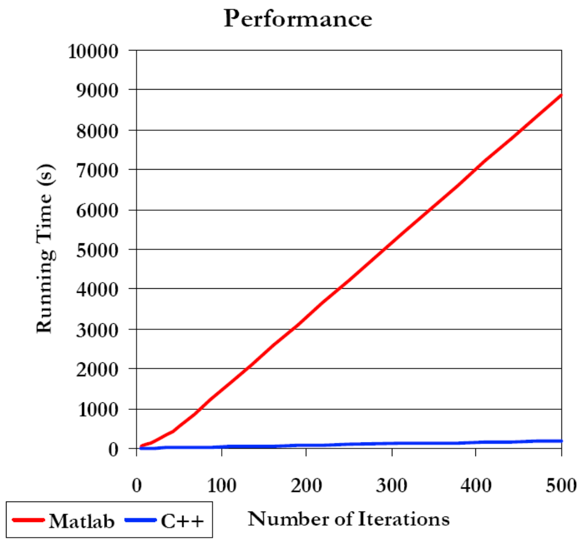

FIGURE 6 Performance comparison of Matlab and C++ implementations

Subjective Analysis

Validation was performed throughout the code porting process. It was important to verify that the control flow was similar in both versions and that the data were handled in the same manner. To accomplish this, a sort of tracing was performed in which the behavior of certain entities in the models, such as the Lagrange multiplier, was followed in both versions of the model. This helped validate the C++ model against the Matlab model and also helped us identify where our code was going wrong. Table 4 shows some sample results from our code. The steady-state data are taken from Cosimano and Gapen (2004). Our results show that the programs produced similar results for the more important variables. However, the variable representing the real interest rate shows a significant disparity. Tracing helped us discover the likely cause: that the matrix inversion function in the C++ code is not as robust as it is in the Matlab version. When presented to a domain expert, face validity was achieved for most of the values shown in Table 4. Because the core of the simulation is equation based, output values should be consistent through both of our versions. This correlates to us performing more of a "face verification" technique in judging the correctness of our results.

Quantitative Analysis

Simple checking, such as outputting key variables as the programs were running, helped validate that the calculations were being done correctly. In essence, the C++ version of the code was validated against the Matlab code in the docking process. In addition, the labor and Lagrange multiplier values were docked against the steady-state data, further increasing confidence. The face validity checking helped identify some errors in our code, while docking helped us isolate and verify the problem. More verification and validation techniques need to be performed for this case study.

TABLE 4 Face verification

Model	Lagrange Multiplier	Labor	Money Growth	Tax Rate	Cash Good	Credit Good	Real Interest Rate
Matlab	0.138	0.309	–0.009	0.188	0.486	0.621	0.009
C++	0.123	0.309	–0.009	0.188	0.486	0.621	–0.659
Steady state	0.138	0.309	–0.009	0.188	0.485	0.620	0.009

CONCLUSION

In this paper, we have shown how we performed verification and validation through docking on two very different models. We have shown that similar techniques can be applied to the models, regardless of the underlying model structure.

General Guidelines

When designing a simulation, it is important to have a concise abstract representation of the model in mind. This abstract representation will help lead to effective programming and implementation choices. On the basis of this, one can choose the correct environment or language for the model. As we have shown here, different environments and languages have their own distinct advantages and disadvantages. It is upon these that our programming decisions must be based. It is important to note that the entire lifetime of the model must be considered when making these decisions. The choices must also be made with consideration given to the verification and validation techniques that will be applied to the model. These verification and validation techniques must also be thought out in advance. We have listed some general ratings for the techniques used in this paper in Table 5. Possible ratings are fair, good, very good, and excellent. It is important to note that the ratings listed are specific to our case studies.

TABLE 5 General ratings for our case studies

	Agent-based	Equation-based
Face validation/verification	Very good	Very good
Turing test	Very good	Good
Internal validity	Very good	N/A
Tracing	Fair	Excellent
Black-box testing	Good	Good
Model-to-model comparison	Very good	Very good
Historical data verification	Very good	Very good
Sensitivity analysis	Good	Good
Prediction validation	Good	Fair

FUTURE WORK

In future studies, it is important that we use and develop more stringent and formalized verification and validation testing methods. Doing this on top of a strong statistical foundation would further increase confidence in the models. Gathering empirical data and generating statistical data would also serve as a better point of comparison when judging our models against real-world systems. In addition, performing some goodness-of-fit tests, such as the chi-square test and the Kolmogorov-Smirnoff test, as well as performing the ANOVA test would help us determine the validity of the models. "Invalidating" our models, meaning performing tests specifically designed to invalidate them, also has the potential to eliminate some of our "validation bias" (Macal and North 2005). Finally, improving the helper functions in the C++ version of the second case study would both speed up and strengthen the results for that model. It would also allow us to do a more in-depth verification and validation study.

REFERENCES

AlphaStep; available at http://www.nd.edu/~nom/Software/software.html.

Axtell, R., R. Axelrod, J.M. Epstein, and M.D. Cohen, 1996, "Aligning Simulation Models: A Case Study and Results," *Computational and Mathematical Organization Theory*.

Balci, O., 1998, "Book Chapter: Verification, Validation, and Testing," *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, John Wiley & Sons, New York.

Banks, J., J.S. Carson III, B.L. Nelson, and D.M. Nicol, 2005, "Discrete-Event System Simulation," Prentice Hall International Series in Industrial and Systems Engineering, Fourth Edition.

Cabaniss, S.E., G. Madey, L. Leff, P. Maurice, and R. Wetzel, 2005, "A Stochastic Model for the Synthesis and Degradation of Natural Organic Matter. Part I. Data Structures and Reaction Kinetics," *Biochemistry*, Vol. 76, No. 2, pp. 319–347, Springer, Nov.

Cosimano, T.F., and M.T. Gapen, 2004, "Optimal Fiscal and Monetary Policy with Nominal and Indexed Debt," working paper; available at http://www.nd.edu/%7Etcosiman/Optimapolicy.pdf.

Cosimano, T.F., and M.T. Gapen, 2005, "Solving Ramsey Problems with Nonlinear Projection Methods," *Studies in Nonlinear Dynamics & Econometrics*, Vol. 9, spring.

Grimm, V., 2002, "Visual Debugging: A Way of Analyzing, Understanding, and Communicating Bottom-up Simulation Models in Ecology," *Natural Resource Modeling*, Vol. 15, No. 1, pp. 23–38, spring.

Huang, Y., X. Xiang, G. Madey, and S. Cabaniss, 2005, "Agent-based Scientific Simulation," *IEEE Computing in Science & Engineering*, Vol. 7, No. 1, pp. 22–29, Jan./Feb.

Macal, C.M., and M.J. North, 2005, "Validation of an Agent-based Model of Deregulated Electric Power Markets," in *Proceedings of the North American Association for Computational Social and Organizational Science (NAACSOS) Annual Conference,* Notre Dame, IN, June 26–28.

Matlab; available at http://www.mathworks.com/products/matlab/.

Repast; available at http://repast.sourceforge.net/.

Sargent, R.G., 1998, "Validation and Verification of Simulation Models," in *Proceedings of the 1998 Winter Simulation Conference,* pp. 121–130.

Xiang, X., R. Kennedy, and G. Madey, 2005, "Verification and Validation of Agent-based Scientific Simulation Models," in *Proceedings of Agent-directed Simulation Conference*, San Diego, CA, April.

HOW SIMPLE IS SIMPLE ENOUGH? MILITARY MODELING CASE STUDIES

T.E. MEYER,* M. KOEHLER, P. BARRY, and B. TIVNAN,
The MITRE Corporation, McLean, VA

ABSTRACT

All models are abstractions of the real world. Determining the appropriate level of abstraction is a balancing of the complexity of the system being modeled, the available data resolution provided by data sources and subject matter experts, the needs of decision makers, and the limitations of the computational and developmental resources. Results from algorithmically linear, physical, closed-system simulations can often be improved by using higher-resolution inputs and by modeling lower-order phenomena. It is not as obvious; however, that ever-increasing resolution will necessarily improve the results from modeling complex systems. Two military course-of-action (COA) development case studies are examined to determine what level of model resolution is sufficient to provide significant insight into COA development. We examine the appropriate level of fidelity for modeling force structures and behaviors as well as the appropriate level of detail for modeling the terrain and physical environment. Methods for evaluating and comparing the results of varying model resolutions are presented.

Keywords: Model fidelity, distillation, abstraction, brigade

INTRODUCTION

Because of their nature, it is very difficult to build models of combat that are able to provide complete answers to a military decision maker's questions. Most combat situations are open systems whose initial conditions are poorly known and for which the motivation of opposing forces is unknown. Further, because the (1) motivation of allies is not well known, (2) communication and interactions among entities are extremely complex, and (3) number of possible courses of action available to the participants could be huge, closed-form modeling is problematic. Typically, the military modeler's response to this problem has been to build increasingly complex models to represent as significant a portion of the combat situation as possible. As a result, these models have tended to be extremely expensive and to require an extensive amount of training, time, and human resources to set up and run. Given the uncertainties of the inputs, the results of even the most competently validated models are suspect at best. Several military organizations are addressing these issues by implementing models focused on specific questions and by establishing processes for executing the models a sufficient number of times to see the potential range of outcomes.

Both the Army G8 Laboratory and U.S. Marine Corps (USMC) Warfighting Laboratory are investigating the use of agent-based models (Axtell 2000) and high-performance computing (HPC) to provide the analytic capability to examine alternate tactical courses of action at the

* *Corresponding author address*: Theodore E. Meyer, The MITRE Corporation, 7515 Colshire Dr., Mail Stop H305, McLean, VA 22102; email: tedmeyer@mitre.org.

brigade level. Scenarios to address specific questions are developed relatively quickly in the agent-based modeling environment and tested in real time. The scenarios become the “sandbox” for testing alternate courses of action (COAs) in a fast-turnaround process. As is always the case, however, modelers, subject-matter experts, and analysts frequently identify aspects of the model or scenario data that can be improved. Because results from algorithmically linear, physical, closed-system simulations can often be improved by using higher-resolution inputs and by modeling lower-order phenomena, it has been assumed that similar improvements will enhance the verisimilitude of question-focused agent-based models. The purpose of this paper is to document some preliminary examinations of the effects of input data precision on brigade-level scenarios implemented in agent-based models.

ABSTRACTING BRIGADE-LEVEL COMBAT

Distillations

Everything should be made as simple as possible, but not simpler.

Any intelligent fool can make things bigger, more complex….

It takes a touch of genius and a lot of courage to move in the opposite direction.

~ Albert Einstein

Data farming (Brandstein and Horne 1998) is a process developed by Project Albert, a USMC program aimed at supporting military decision makers. It is the process of running models many times in an HPC environment and varying the initial conditions in order to find outliers, examine the potential range of outcomes, and test the model across its parameter space. Models used within the paradigm of data farming are referred to as “distillations” (i.e., abstractions). It is recognized that all models are distillations of the real world. It is only by the judicious implementation of specific aspects of a system that we can produce models that are helpful. The quotes by Einstein above capture the intent of modeling within the realm of data farming. Distillations should be complex enough to address the question, but no more complex than that. Distillations should be:

- Intuitive — the users must be able to understand the parameters and rules that define the model and how they relate to the system being modeled;
- Transparent — the users must be able to understand how the behaviors in the model emerged from a set of parameters and rules; and
- Transportable — the model must be portable to a data farming environment (Horne and Meyer 2004).

Although any model could be data farmed, distillations are intended to be a bottom-up reduction to the essence of a question. Typically, distillations are expected to be developed quickly — potentially in a matter of a few days or hours. Current distillation development applications use abstraction judiciously, thus representing a number of phenomena in a few paradigms. For example, various types of interchanges (such as food, resources, and positive or negative messages or propaganda) may be abstracted and proxied by weapons exchanges in some distillation modeling environments. Location or proximity in a model can be abstracted to represent relative aspects of other relational parameters. Modeled obstacles can represent walls, floors, borders, or sociological or psychological obstructions in non-geoterrain or combat interchanges. In short, using creative abstraction can keep the models computationally simple, allowing for a large number of model executions in a relatively short period of time. This guideline encourages distillation modelers to innovate and use imagination to define abstractions.

Military decision makers often must address questions with answers that are dependent on intangibles; examples are “How will the morale of my men affect this battle?” “How tired are they?” and “What does the enemy know about my positions?” These types of questions rarely have precisely defined initial conditions or a complete set of algorithms that describe the system being considered. As implied above, we have been using data farming with a wide variety of possible variable combinations to provide insight into these complex questions. Looking at the distribution of results over a large number of runs can provide insight that can be used to address these complex questions. The accomplishment of this data farming relies on two basic ideas:

1. Use HPC to execute models many times over varied initial conditions to gain an understanding of the possible outliers, trends, and distribution of results and

2. Develop models called distillations that are focused to specifically address the question.

Model Scenarios: IMEF and Army G8

Figure 1 is a screen shot of the Map Aware Non-uniform Automata (MANA) modeling environment (Roger et al. 2002) implementing a scenario near an airfield in Camp Pendleton, California. This model is an illustrative example of the type of models that are currently of interest to the USMC First Marine Expeditionary Force (IMEF) and the Army G8. The models often represent hundreds of units including tanks, unmanned aerial vehicles (UAVs), ground troops, artillery, aircraft, and unmanned vehicles, and all of the associated weapons, armor, communication, and command structure.

The scenario in Figure 1 takes place over a 12-kilometer grid and incorporates ground-based units such as infantry, field artillery, and armored vehicles from two opposing forces: red and blue. The pictured scenario includes more than 200 agents. The blue and red dots are agents that represent various entities within the blue and red forces. The blue and red lines represent the entities firing on the opposing force. The entities are displayed on a standard topographic map. Not visible but an essential part of the scenario are the topographic data that the agents within the scenario sense and respond to, both from a visibility and a mobility perspective. This scenario was used as a demonstration of MANA’s ability to handle scenarios of this scale when we began to design a larger scenario that covered a different, wider (35-kilometer-square) area; incorporated air support; and included more than 400 agent types. The analysis described in the subsequent sections used this more complicated scenario.

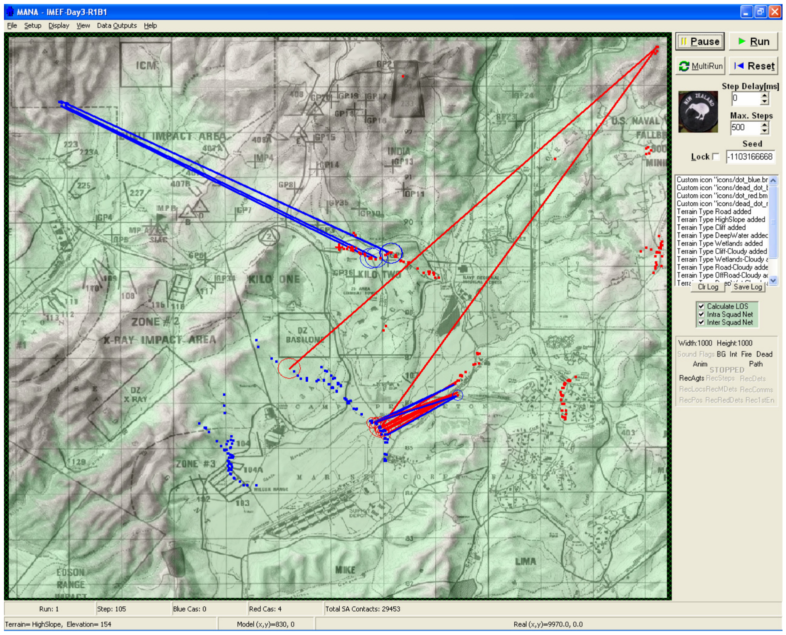

FIGURE 1 MANA modeling environment

Abstracted Model Parameters

Models can become more complex by adding additional detail to the algorithms, rules, and data inherent in them. One of the easiest ways to "improve" a model is to use readily available, higher-resolution data sources with input having more "real-world" accuracy. The question this paper examines is whether this added detail is useful in providing better and more insightful results.

For this study, two input data sources are examined at several levels of abstraction. The MANA simulation environment uses an elevation grid to determine line-of-sight. Weapons are modeled by defining a probability of kill (Pk) profile that determines effectiveness at various ranges. Although implementing "improvements" to these input data is straightforward, it was unclear whether there would be any significant improvement in the quality of the results produced. Therefore, as part of this study, the resolution of the digital terrain elevation data (DTED) was varied, as was the Pk. The specifics of this process are discussed below.

In the scenario, four red force 155-mm Howitzer artillery units play a vitally important role in the outcome. The units are well-protected and, depending on the specifics of the Howitzers' implementation in the model, can be devastating to the blue forces. Various public data sources can be used to acquire information about the effectiveness of Howitzer shells. These sources are questionable, though, and may not be based on anything but an untrained observer's comments. Classified sources can provide more "accurate" data, potentially creating a more valid representation of the weapon in the model.

However, before spending time and resources implementing a classified version of the model in order to execute runs using validated Pk profiles, it is appropriate to test the model by using various unclassified Pks to determine whether variations in the weapon's Pk profiles have a significant effect. Because of the significance of these weapons in the outcome of the battle, it was anticipated that variations in the Pk would be likely to have a significant impact.

Three variations of the Howitzer's Pk profiles were examined for this study, as shown in Figure 2. The base case uses a nominal profile acquired from unclassified sources. Pk A (on the left in Figure 2) represents a hypothetical "real" Pk profile with "fine" adjustments to the profile. The right side of Figure 2 represents Pk B, a "simple" version of the profile.

The model currently uses DTED Level 0 at a 1-kilometer resolution. Given the vagaries of sensor limitations, tree lines, unit positioning, and other terrain features, the importance of implementing high-resolution elevation data was an open question. For the purposes of this study, two abstractions of the scenario's elevation were tested: the DTED Level 0 data and a flat surface. With DTED Level 0 data, the line-of-sight is affected by elevation obstacles. With the flat surface implemented, only an agent's sensor range and the terrain type affect how far the agent can see.

Data Farming and Results

In order to examine the impact of the changes to the Howitzer Pk and elevation, the model was data farmed. The model was run 500 times — 100 times for each of these five variations:

1. Base Case — Base-case Howitzer profile and DTED Level 0;

2. Control — Same Howitzer Pk profile and elevation as in the Base Case, with only a change in the random seed;

3. Howitzer Pk A — The "refined" Howitzer blast radius Pk profile and DTED Level 0;

4. Howitzer Pk B — A simplified Howitzer blast radius Pk profile and DTED Level 0; and

5. Simple Elevation — Flat surface (no variation in surface height) and same Howitzer Pk profile as in the Base Case.

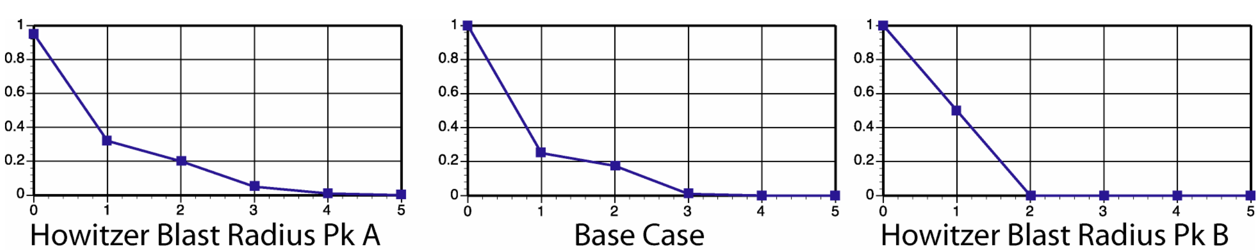

FIGURE 2 Probability of kill blast radius profiles (X axis, one unit is 40 m)

Data farming is a statistical sampling technique. Each execution of the model with a different random seed, referred to as a "replicate," is a sample of the almost infinite population of potential model outcomes for a particular set of parameter inputs. Varying input parameters produces "excursions" that expand the population and provide analysts with a comparative stratification of the outcomes. Data farming addresses the issue that a model may produce a range of results depending on random or designed variations in inputs. Only by sampling the input data space can the full scope of possible outcomes be understood. Methods using formal experimental design to effectively explore this excursion/replicate space have been developed (Lucas et al. 2002).

Each of the five variations enumerated can be considered an excursion. For each excursion, 100 replicates were run in order to acquire enough data to be able to statistically compare distributions.

For each of the 500 model runs, a set of more than 800 end-of-run measurements of effectiveness (MOEs) were generated by the model. These MOEs include red and blue force casualties for each agent group. Other outputs from the model can include data on agent position, casualty location, enemy contact and detection, and communication for every time-step. For the purposes of this study, the end-of-run MOE data were reduced to total blue casualties in order to examine the impact of small variations on the large-scale results.

Figure 3 represents the MOE results from these runs. Each plot represents a histogram of the distribution of the number of blue casualties over 100 runs of the model. The Base Case and Control represent the model scenario with no changes to the original input parameters. The difference between the Base Case and Control is a change in the set of random

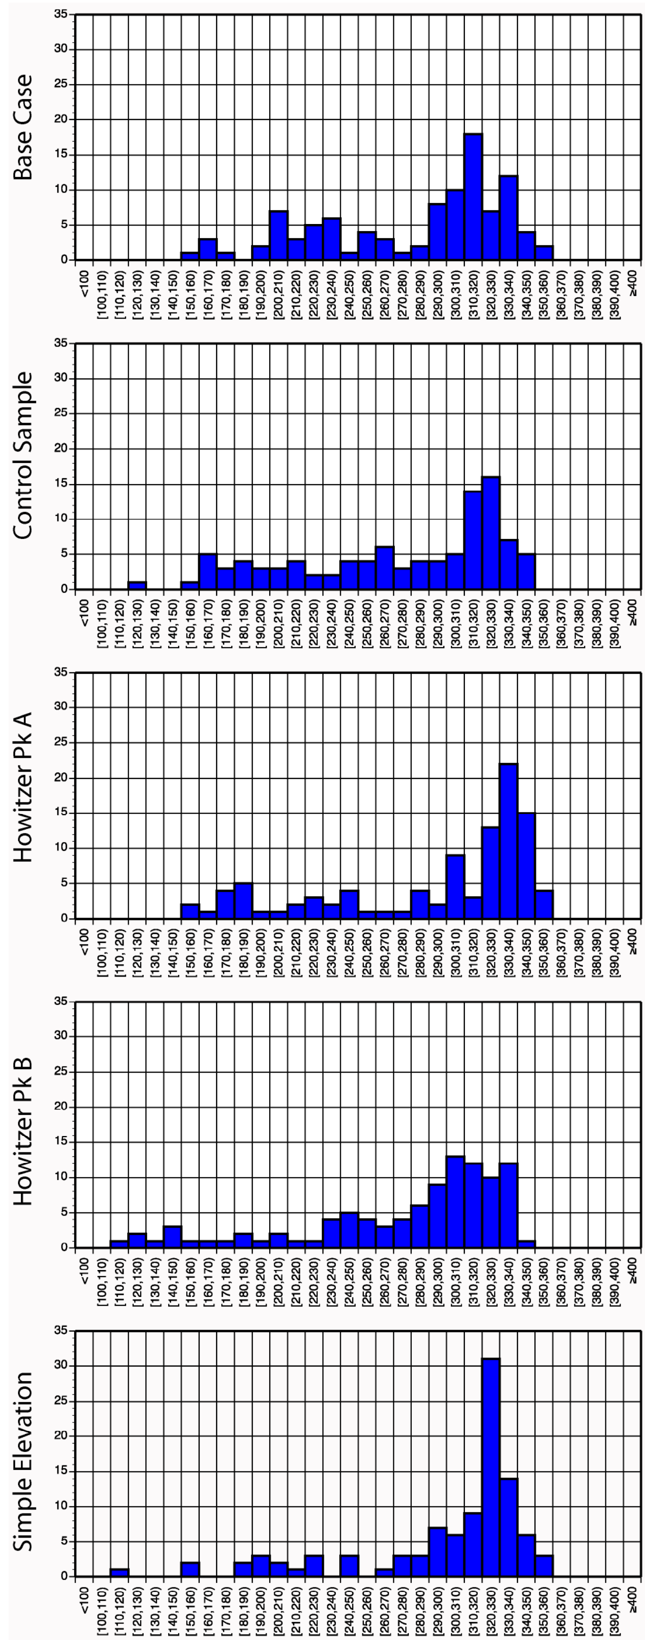

FIGURE 3 Blue casualty distributions

seeds used to provide small changes to initial starting positions and to adjudicate engagements. The authors' assumption was that the Control and Base Case should have MOE distributions that are statistically the same. The MOE distributions that resulted from the variations of Howitzer (Base case versus Howitzer Pk A versus Howitzer Pk B and Elevation (Base Case versus Simple Elevation) were examined in conjunction with the Control to discern whether these variations have any statistical impact on the model results.

Figures 4 and 5 represent quantile plots of the MOE distributions. Figure 4 shows the excursion distributions plotted against normal distributions. These plots indicate that the profiles of all of the distributions were very similar, but that there are visible, if not statistically significant, differences in the distributions.

The QQ plot on the left of Figure 4 shows the quantiles of Base Case and Control plotted against a normal quantile. The plot on the right includes the quantiles for the Base Case and the excursions. It is evident that the Control distribution is a reasonable match to the Base Case distribution, but that the excursions show some deviations from the Base Case. Figure 5 shows a QQ plot of the Control and excursions versus the Base Case. This plot indicates that the Control and Howitzer Pk B distributions adhere more closely to the Base Case and that the other two excursions have similarities in their profiles.

Table 1 provides the results of a statistical analysis of the distributions. Displayed are the Mann-Whitney nonparametric, independent, two-group "statistical comparisons." Mann-Whitney does not assume normal distributions and provides an indicator of the similarity of distributions. T-tests (which do assume normal distributions) gave similar results. In the table, significance values of less than 0.025 indicate that there is a difference in the distributions. It is interesting to note that the distributions fall into two groups (differentiated by coloring in the table): (1) Base Case, Control, and Howitzer Pk B group and (2) Howitzer Pk A and Simple Elevation group. This is also indicated (although not as obviously) in Figure 5.

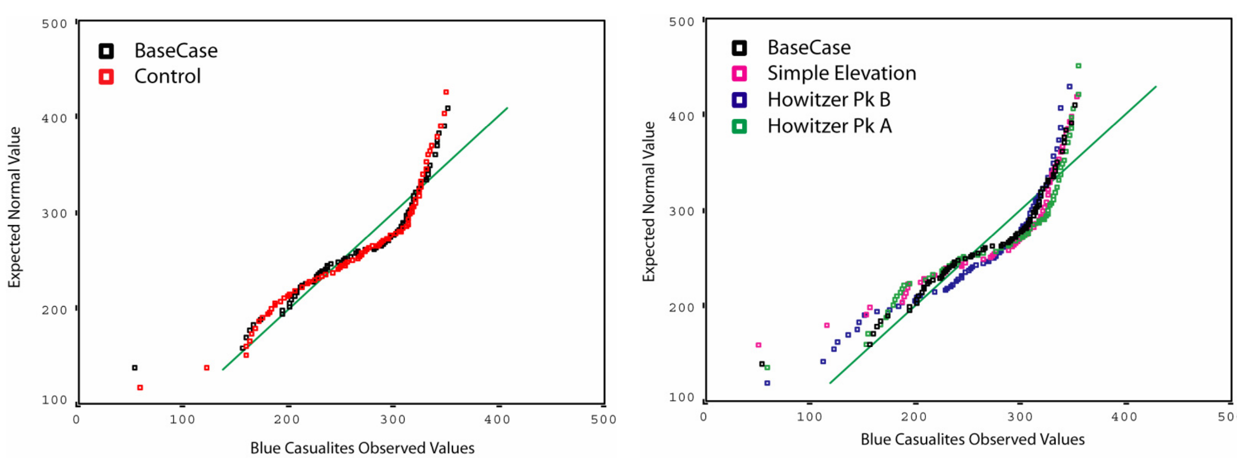

FIGURE 4 Normal QQ plot of Base Case and Control (left) and of Base Case and excursions (right)

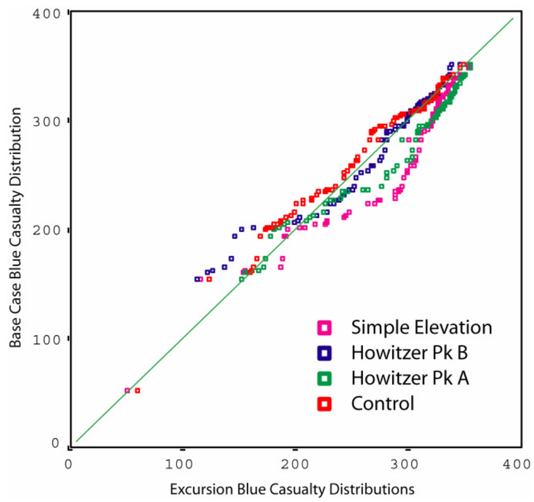

FIGURE 5 QQ plot of Base Case versus Control and excursions

TABLE 1 Statistical differences

Mann-Whitney Test	Base Case	Control	Howitz Pk A	Howitz Pk B	Simple Elevation
Base Case	-	0.396	0.005	0.495	0.002
Control	0.396	-	0	0.823	0
Howitz Pk A	0.005	0	-	0	0.489
Howitz Pk B	0.495	0.823	0	-	0
Simple Elevation	0.002	0	0.489	0	-

Why would the distribution fall into these groups and why is there a shared distribution between the Simple Elevation and the Howitzer Pk A group? After an examination of multiple model executions and their methods, the authors provide the following explanations.

Flattening elevation has a potentially large influence on both the red and blue forces' artillery effectiveness. Flattening the elevation provides all units with a 360° line of sight. Blue has significantly more artillery units that can take advantage of the flattened terrain. Figure 5 indicates that the Simple Elevation excursion is advantageous to the blue force.

At first glance, the profiles in Figure 2 would seem to show that the Pk B profile is a much larger variation from the Base Case profile than the Pk A profile. Pk A seems to consist of small adjustments to the Base Case profile. Pk B is a radical change in all but the zero radius point. Yet Pk A is statistically different than the Base Case, and Pk B is not. What is not obvious is that one "fine" adjustment to the Pk A is the adjustment of the zero radius point from a 100% kill probability to a 95% one. The Base Case and Pk B maintain a 100% probability at the zero radius point. The "minor" change to the zero radius point has more impact than a more radical change at higher radii. This change is also an advantage for the blue force and results in a distribution similar to that of Simple Elevation.

Is Pk A or Pk B more “real?” In the real world, if a Howitzer shell hits within a zero radius of a unit’s position, it is a direct hit and that unit is killed. This scenario and model, however, employ a grid cell that is 40 meters in length. A zero radius implies a hit for anything within that cell. A 95% Pk at zero radius can be considered an adjustment to the real-world data to account for the fact that within the model abstraction, hits within 40 meters are not necessarily deadly. For example, if an artillery shell directly hits a tank, the tank will very likely suffer damage that will make it combat-ineffective. However, if the artillery shell hits 30 meters away, the tank will experience little damage. In MANA, in the current scenario, both of the above hits are considered the same — causing similar effects. Here is an example of an incompatibility between levels of abstraction or simplicity. The actual Pk data indicate that at zero radius, a target should suffer from more severe damage; however, the abstraction of the terrain generates an absurd result with this Pk table.

As one creates models and struggles with determining the proper level of simplicity or abstraction, one must be consistent. Although it is very tempting to abstract model features that are less well understood and to increase the “realism” of better-understood features, one may be tempting fate. Increasing the realism of one area of a model while abstracting other areas may bias your results in significant ways. Worse yet, these biases are introduced by parts of the model that do not, on their face, appear to be problematic (i.e., “These are the actual data, how could that be the problem?”).

CONCLUSION

While all models are abstractions, the abstract models discussed above are used to glean insight into complex phenomena for which there are no closed-form analytic solutions. Further, it is frequently difficult, if not impossible, to check the “correctness” of results of a distillation. Unlike the modeling of physical phenomena, which produces testable results and usually has an underlying mathematical model, the modeling of intangibles and complex phenomena has no “correct” solution. However, the very act of (1) describing the distillation at the appropriate level of abstraction for the question addressed, (2) explicating the associated assumptions, and then (3) exploring the parameter space has been shown to be useful in providing insights to decision makers.

This paper has demonstrated that when a distillation is being developed, it is important that the abstractions be consistent. “Validated” real-world inputs may be incompatible with other model abstractions, such as terrain resolution. The model described in this paper is based on a grid with cells that are 40 meters across. Therefore, it makes no sense and, more importantly, contributes little to the value of the results to have an “accurate” real-world Pk profile that would model the degradation in Pk by the actual distance in meters from the impact. For this distillation, the Pk profile needs to be abstracted — its resolution and probabilities refactored to the scale and functional requirements of the model — before being incorporated into the model. Understanding the overall implications of the abstractions within the distillations is necessary and critical for their appropriate use. As was demonstrated, increasing the fidelity on some parameters (e.g., Pk) with scant regard for the other parameters (e.g., terrain resolution) contributes little and frequently can give a false sense of precision.

We have only begun to examine the question “How simple is simple enough?” in this study, and only for the model scenarios that were specifically analyzed. For the scenarios modeled, only a small set of input parameters — weapon Pk and elevation — over a small range of variation was examined. It is evident from this study that the level of simplicity (abstraction) of the model and its input data can affect the outcomes in statistically significant ways, especially if they are inconsistently “simple.”

The lesson learned from this exercise is as follows: *Validated “real-world” data may not increase the accuracy of a distillation model and may have a negative impact if they are incompatible with the model’s abstractions.*

REFERENCES

Axtell, R., 2000, *Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences,* Working Paper No. 17, Center on Social and Economic Dynamics, Nov.

Brandstein, A., and G. Horne, 1998, “Data Farming: A Meta-Technique for Research in the 21st Century,” *Maneuver Warfare Science 1998,* U.S. Marine Corps Combat Development Command, Quantico, VA.

Horne, G., and T. Meyer, 2004, “Data Farming: Discovering Surprise,” in *Proceedings of the 2004 Winter Simulation Conference,* Denver, CO.

Lucas, T., S. Sanchez, L. Brown, and W. Vinyard, 2002, “Better Designs for High-dimensional Explorations of Distillations,” *Maneuver Warfare Science 2002,* U.S. Marine Corps Project Albert, Quantico, VA.

Stephen, R.T., M.A. Anderson, and M.K. Lauren, 2002, *MANA Map Aware Non-uniform Automata Version 2.0 User’s Manual*.

X-MAS: SUPPORTING THE TEDIOUS WORK OF VALIDATION IN AGENT-BASED SIMULATION

Y.I.L. SUEMATSU* and K. SHIMOHARA, ATR International, Kyoto, Japan, and Kyoto University, Kyoto, Japan
K. TAKADAMA, ATR International, Kyoto, Japan, and Tokyo Institute of Technology, Tokyo, Japan
O. KATAI, Kyoto University, Kyoto, Japan

ABSTRACT

Validation is an essential issue in the growing field of agent-based simulation (ABS), as ABS has become a prominent paradigm in the study of social complex systems. However, the main difficulty faced in this validation process is the lack of techniques and tools to assure the reliability of models. Thus, the validation of models is still a tedious task. In our previous work, we proposed *cross-element validation*. This process consists of performing the validation *within* a model by comparing the simulation results of the model under several instances of some of its composite elements. Elements are, for instance, learning mechanisms or network iteration topologies. Even though it is relatively simple, this validation process requires performing several simulations of the model under the possible combinations that exist among a certain number of instances of some elements. In other words, it requires several implementations of the model to account for the above-mentioned combinations of elements. Therefore, tools to support this validation process are required. To support the cross-element validation process for ABS models, this paper presents cross-element validation for multi-agent-based simulation (X-MAS). This tool provides facilities for performing easy cross-element validation of ABS models and also facilitates the implementation of general-purpose ABS models. To illustrate the potential of X-MAS, a cross-element validation of a bargaining game model was performed by evaluating several learning mechanisms applied to the agents. The findings showed that simulation results can be strongly affected by even small variations in the elements. Therefore, cross-element validation should be performed before deep analysis of the implemented model.

Keywords: X-MAS, cross-element validation, agent-based modeling, bargaining

INTRODUCTION

Although agent-based simulation (ABS) is becoming a prominent paradigm in the study of complex social sciences, it still lacks the validation techniques and tools to assure the reliability of the models. One available technique is the so-called docking or alignment of computational models proposed by Axtell et al. (1996). This is a validation process where two models that deal with the same phenomenon are compared, in order to determine whether the results from one computational model match the results of another model. This process requires the replication of one model based on the other model's framework. However, the following

---

* *Corresponding author address*: Yutaka Inti Leon Suematsu, ATR International, Network Informatics Laboratories, 2-2-2 Hikaridai, Keihanna Science City, Kyoto 619-0288, Japan; e-mail: yutaka@atr.jp.

difficulties make it rather awkward: (1) models may be developed for different purposes, (2) common parts between models may be small in number, and (3) fair evaluation criteria may be difficult to define in order to guarantee the seeming equivalence of the models' simulation results.

In order to overcome the above-mentioned difficulties, our previous work proposed the concept of *cross-element validation* as a process that performs the validation *within* a model (Takadama et al. 2003), as opposed to docking, which consists of validation *between* models. Cross-element validation consists of detecting, analyzing, and comparing the model's macro-behavior under different variations of some of its composite elements. Elements are, for instance, how individual agents store knowledge, how to perform learning, or what iteration network topology is used. By understanding whether and how the elements' implementations affect the model's overall behavior, the reliability of the ABS model is expected to increase.

Cross-element validation must perform several simulations of one model with variations of the composite elements. The number of simulations exponentially increases as both the number of replacements of a certain element and the number of elements to be evaluated increase. Thus, the efficient performance of each simulation is required. Another difficulty is the modification of the model implementations, which is cumbersome. Because of these facts, even relatively simple cross-element validation process is a tedious task. Therefore, tools with the following three requirements are indispensable: (1) easy model implementation, (2) flexibility to simplify model element exchanges, and (3) the construction of efficient programs to accelerate the simulations.

Several ABS libraries and tools are available in the community for facilitating model implementation, such as Swarm (Swarm Development Group 2005), Repast (Repast 2005), and MASON (George Mason University 2005). Unfortunately, their use in the cross-element validation of models is quite difficult because they do not fulfill the above-mentioned requirements.

To support the cross-element validation process of ABS models, this paper presents cross-element validation for multi-agent-based simulation (X-MAS). This tool provides facilities for performing easy cross-element validation of ABS models by providing a framework wherein several variations of some elements involved in the model can be evaluated easily without reimplementing the model for the numerous combinations between elements. X-MAS has been developed as a general-purpose library to support diverse ABS models, ranging from social science to engineering.

To illustrate the potential of X-MAS, the cross-element validation of a bargaining game model was performed, comparing the results of the model with different learning mechanisms, such as evolution strategy (ES), learning classifier system (LCS), and reinforcement learning (RL), and also variations in discrete and continuous values in the representation of the knowledge.

This paper is organized as follows. The next section describes X-MAS and its features and gives a brief description of the composed layers. Then the bargaining model is described, including the simulation results and a brief discussion related to the experimental findings. Finally, a summary and future work are presented.

X-MAS

Features

The cross-element validation for X-MAS provides a rich framework that facilitates both the implementation and cross-element validation of ABS models. Additionally, X-MAS facilitates multi-intelligent agent implementation by providing a framework that embeds the commonly used learning mechanisms and knowledge representation schemes of agents in ABS models. X-MAS is a collection of generic object-oriented programming libraries. Because of performance, scalability, and portability issues, it was developed in Standard C++ with extensive use of STL and BOOST Libraries (2005). Cutting-edge techniques of meta-programming in C++ were employed to provide facilities for easily exchanging elements for their substitutes in the model. These characteristics make it possible to satisfy the flexibility required for the cross-element validation process that was enumerated in the previous section. X-MAS is portable, and the models can be compiled in several operating systems with no modifications. X-MAS only requires the use of compilers that support major features of the C++ Standard, including *partial template specialization*, which is supported in many modern compilers.[1]

X-MAS is considered a general-purpose library for multiagent systems. The core libraries are highly customizable, providing the scalability to build domain-specific libraries and user-customizable libraries. As shown in Figure 1, the simulation is developed in two levels. The left side represents the implementation of the model that will be executed in the back end. Then, visualization of the simulation results are shown and analyzed in an independent process, as shown in the right side of Figure 1. This structure will accelerate the simulation process.

Moreover, X-MAS facilitates the maintenance of the model implementation by keeping the core implementation of the model in one program regardless of the possible variations of the element to be evaluated.

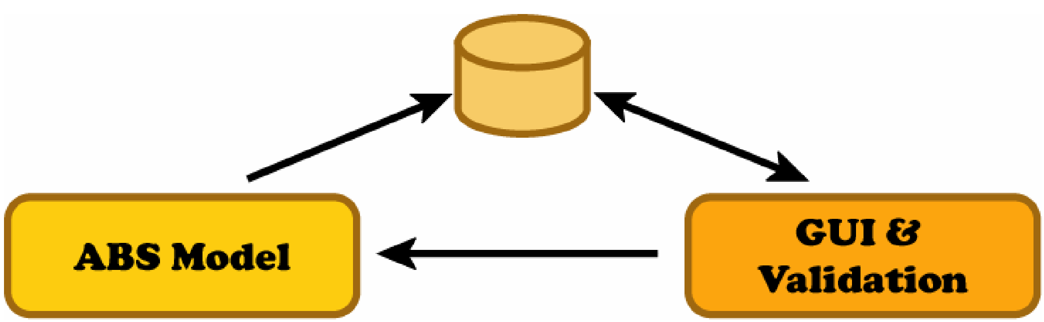

FIGURE 1 Simulation under X-MAS (The left side represents the implementation of the model that is executed in the back end. Then the simulation results are visualized and analyzed by another process, as shown in the right side. Additionally, some GUI interfaces for controlling the simulation and the exchange of elements for performing cross-element validation are available.)

---

[1] Most modern compilers support the major features of the C++ standard. However, even today, some compilers still do not support important minor features such as partial template specialization. X-MAS makes extensive use of the power of templates to fulfill the requirements of cross-element validation.

Modules

X-MAS consists of three main modules or libraries, as shown in Figure 2: (1) X-MAS core libraries, (2) domain-specific libraries, and (3) visualization and cross-element validation interfaces.

X-MAS Core Libraries

X-MAS provides a set of generic libraries and utilities for easy implementation of simulation models. Several algorithms commonly used in any model implementation are provided with a highly customizability. Some implementations, for instance, consist of the scheduling of agent interaction, selection algorithm, and simulation cycle control.

Domain-specific Libraries

The X-MAS framework allows the implementation of domain-specific libraries by customizing some of the already available libraries. X-MAS provides libraries for supporting intelligent agents by providing several learning mechanism algorithms such as RL, LCSs, genetic algorithms (GAs), and ESs. It also provides genotype-based, rule-based, and other algorithms with discrete and continuous values as knowledge representation schemes. Each algorithm provides a default setting and a highly customizable framework to include variations of the algorithm by allowing parts of the algorithm to be replaced with user's algorithms, such as different crossover algorithms applied to GAs.

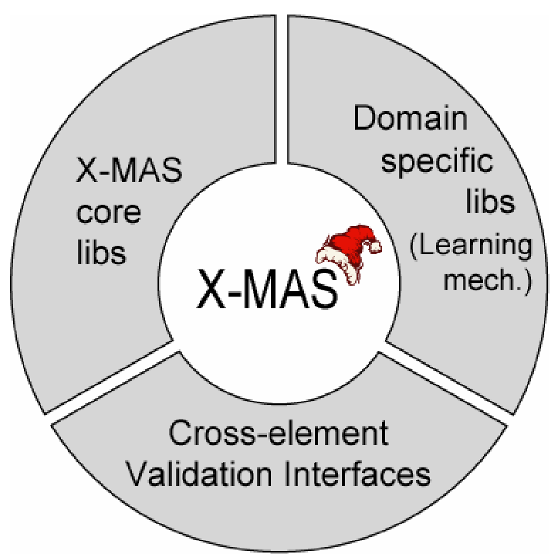

FIGURE 2 The three X-MAS modules

Visualization and Cross-element Validation Interfaces

The main purpose of X-MAS is to provide a set of interfaces to interact with the model. It includes some interfaces to visualize the simulation results of the model under all of the possible combinations between elements. For instance, as represented in Figure 3, it is assumed that the model consists of agents that have learning capabilities and representation capabilities for the knowledge. Three implementations for each involved element are considered. A1 to A3 and B1 to B3 represent implementations of the possible learning mechanisms and knowledge representation schemes, respectively. Therefore, this may require evaluating the model with the nine possible combinations between both elements. As shown in the left side of Figure 3, X-MAS allows the implementation of one model where the combination of the involved element to be evaluated can be easily selected, as opposed to the nine reimplementations of the model with traditional tools. Furthermore, X-MAS provides interfaces to plot the simulation results of the model under all possible combinations among elements, as presented in the right side of Figure 3. Finally, from these results, the equivalence and implications of the implementations of the model elements can be easily evaluated. In the right side of Figure 3, the red box represents equivalent simulation results in six implementations of the model. Therefore, it can be said that the model produces invalid (strange) results when using B3 as element B, as opposed to valid results for the other implementations. As a consequence, the model can be minimally validated under implementations B1 and B2 as element B and all three implementations of element A.

This module is still undergoing development. Because of the fact that there are no standard graphic libraries in the C++ standard, which is a major disadvantage, some interfaces with Python and Java are planned for further development.

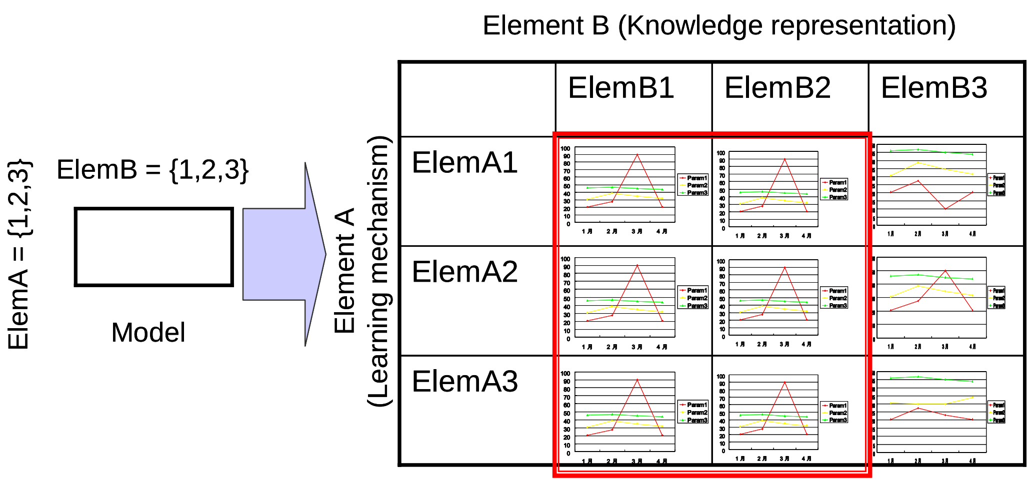

FIGURE 3 Image of the use of X-MAS

CASE STUDY

Bargaining Game

As a concrete example, the bargaining problem (Muthoo 2000) was employed, which addresses a situation where two or more players try to reach a mutually beneficial agreement in order to maximize their profits through negotiations. This problem was selected because it has been studied in the context of game theory (Osborne and Rubinstein 1994) for several years, and its results are well known. Therefore, simulation results can be evaluated by comparing them with the rational behavior of players.

The problem considered in this research is the one proposed by Rubinstein (1982). This model uses the following scenario. Two players, P_1 and P_2, have to reach an agreement on the division of a *pie*. For this purpose, they alternate offers, describing the possible division upon which they would like to settle. The player who receives the offer has to decide whether to accept it or not. If the offer is accepted, the negotiation process ends, and each player receives the share of the pie determined by the concluded contract (e.g., P_1 receives x and P_2 receives $1 - x$ at time t, where x is a value in the interval [0,1]). Otherwise, the receiving player makes a counter-offer, and all of the above steps are repeated until an agreement is reached, or the process is aborted when the limit number of offers is reached; in that case, both players receive a null payoff.

For experimentation, a finite-horizon model was employed, where the maximum number of steps in the game is fixed and known by both players as common information. In the case where the maximum number of steps is one (also known as the *ultimatum game*), the proposer (player P_1) makes the only offer, and the responder player (P_2) can either accept it or not. If P_2 refuses the offer, both players receive a null payoff. Since a rational player always takes actions that maximize her payoff, P_1 tries to keep most of the pie to herself by offering only a minimum share to P_2. Since there are no further steps to be played in the game, P_2 inevitably accepts the tiny offer, under the notion of “anything is better than nothing.”

By applying a backward induction reasoning to the situation above, it is possible to analyze situations where the maximum number of steps is greater than one. For the same reason as that of the *ultimatum game*, the player who can make the last offer in a finite game where payoffs are not discounted by time has a great advantage to obtain the larger share[2] of the pie by making a minimum offer (Stahl 1972).

Model

The implemented model was designed in the framework of the bargaining game as follows:

---

[2] In this paper, the terms *payoff* and *agent* are used instead of the terms *share* and *players* for their more general meaning in the bargaining game.

Model Structure

The basic structure of the agents was implemented using the following components as shown in Figure 4. Note that each agent has the same architecture.

< Memory >

Strategies memory. This stores a set of strategies that agents use during negotiation (in Figure 4, the number of strategies is *n*). Each strategy consists of a fixed number of paired offer (*O*)/threshold (*T*) values and the worth of the strategies (*w*). These strategies are similar to those used in Oliver (1996). The offer and threshold values are encoded by floating point numbers in the interval [0, 1], while the worth values are calculated as averages of acquired payoffs. In this model, agents independently store different strategies, which are initially generated at random.

Selected strategy memory. This stores the strategy selected to confront the strategy of an opponent agent. Figure 4 shows the situation where agent A_1 selects the *x*th strategy, while agent A_2 selects the *y*th strategy.

< Mechanism >

Learning mechanism. This modifies both offer and threshold values in order to generate good strategies that acquire a large payoff.

As a concrete negotiation process, agents proceed as follows. Defining $\{O,T\}_i^{A_{\{0,2\}}}$ as the *i*th offer or threshold value of agent A_1 or A_2, agent A_1 starts proposing the first offer $O_1^{A_1}$. Here, it is counted as one *step* when either agent makes an offer. Then A_2 accepts the offer if $O_1^{A_1} \geq T_1^{A_2}$; otherwise, it makes a counter-offer $O_2^{A_2}$ (i.e., the offer of A_2). This cycle is repeated

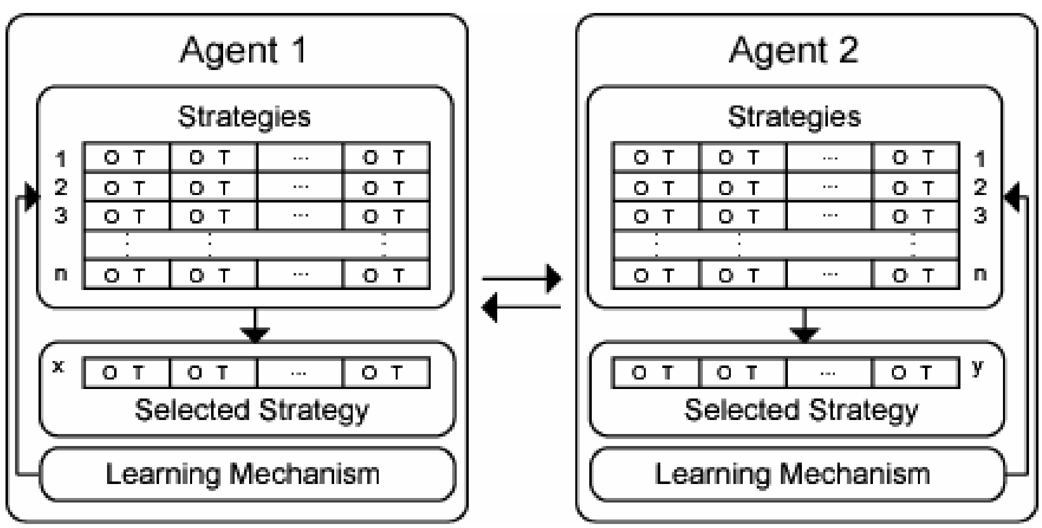

FIGURE 4 Bargaining model structure (Each agent consists of a memory, which is a set of strategies the agent uses during negotiation, and a learning mechanism to improve the strategies.)

until either agent accepts the offer of the other agent or the maximum number of steps is exceeded. To understand this situation, let's consider a simple example where the maximum number of steps is 10, as shown in Figure 5. Following this example, A_1 starts by offering 0.01 to A_2. However, A_2 cannot accept the first offer because it does not satisfy the inequality $O_1^{A_1}(0.01) \geq T_1^{A_2}(0.99)$. Then, A_2 counter-offers 0.01 to A_1. Since A_1 cannot accept the second offer from A_2 for the same reason, this process is repeated until A_1 accepts the 10th offer from A_2 where the offer satisfies the inequality $O_{10}^{A_2}(0.01) \geq T_{10}^{A_1}(0.01)$. In case the negotiation fails, which means that the maximum number of steps has been exceeded, both agents can no longer receive any payoff (i.e., they receive 0 payoff). Here, this is counted as one *confrontation* when the above negotiation process ends (satisfactory or unsatisfactory.)

Furthermore, the worth of each strategy is calculated by the average of payoffs acquired during a fixed number of confrontations (*CONFRONTATION*), where the strategies of the other agent are randomly selected in each confrontation. For example, the *x*th strategy of A_1 in Figure 4 confronts the randomly selected strategies of the other agent in the predefined number of confrontations, and then the worth of the *x*th strategy is calculated by the average of payoffs acquired during these confrontations. Since each agent has n number of strategies, the $(CONFRONTATION \times n \times 2)$ number of confrontations is required to calculate the worth of all strategies of both agents. Here, it is counted as one *iteration* when the worth of all strategies of both agents is calculated.

Elements for Cross-element Validation

The focus of this case study is to make some comparative studies to investigate the influence of different learning mechanisms and knowledge representation schemes (Takadama et al. 2003). For this purpose, each element was designed as follows.

Learning Mechanisms

When the learning mechanisms of agents are being implemented, several mechanisms can be considered. Among the many useful learning mechanisms, the following were employed: (1) ES (Back et al. 1991, 1993), (2) LCS (Goldberg 1989; Holland et al. 1986), and (3) RL (Sutton and Barton 1998).

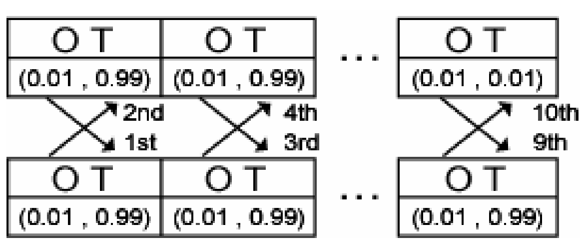

FIGURE 5 Negotiation process between two agents

Knowledge Representation Schemes

In the bargaining game, the representation of the agents' strategies must be considered, though there are no standard guidelines. From this fact, the following two types of knowledge representation capabilities were employed: (1) continuous real numbers (e.g., 0.01…) and (2) real numbers restricted to two decimal digits (e.g., 0.01; called discrete numbers in this paper). The reason why this knowledge representation was employed is because (1) social scientists may take the latter case for a concise representation and (2) a real number in offer and threshold values is critical in the bargaining game.

Simulation Results

Figure 6 shows the simulation results of the possible combinations among learning mechanisms (ES, LCS, and RL) with the two knowledge representation schemes (continuous and discrete). All figures indicate the results of the payoff. The vertical axis indicates the payoff, while the horizontal axis indicates the iteration number. In particular, Figure 6 shows the payoff of agent A1 in the lower lines and that of A2 in the upper lines. Figures 6(a) and 6(b) represent the results when using continuous values as knowledge representation schemes with ES and LCS, respectively, as learning mechanisms. Similarly, Figures 6(c), 6(d), and 6(e) represent the results when using discrete values as knowledge representation schemes with ES, LCS, and RL, respectively, as learning mechanisms. These results show that simulation results do not exhibit the same tendency when different learning mechanisms or knowledge representation schemes are applied to agents.

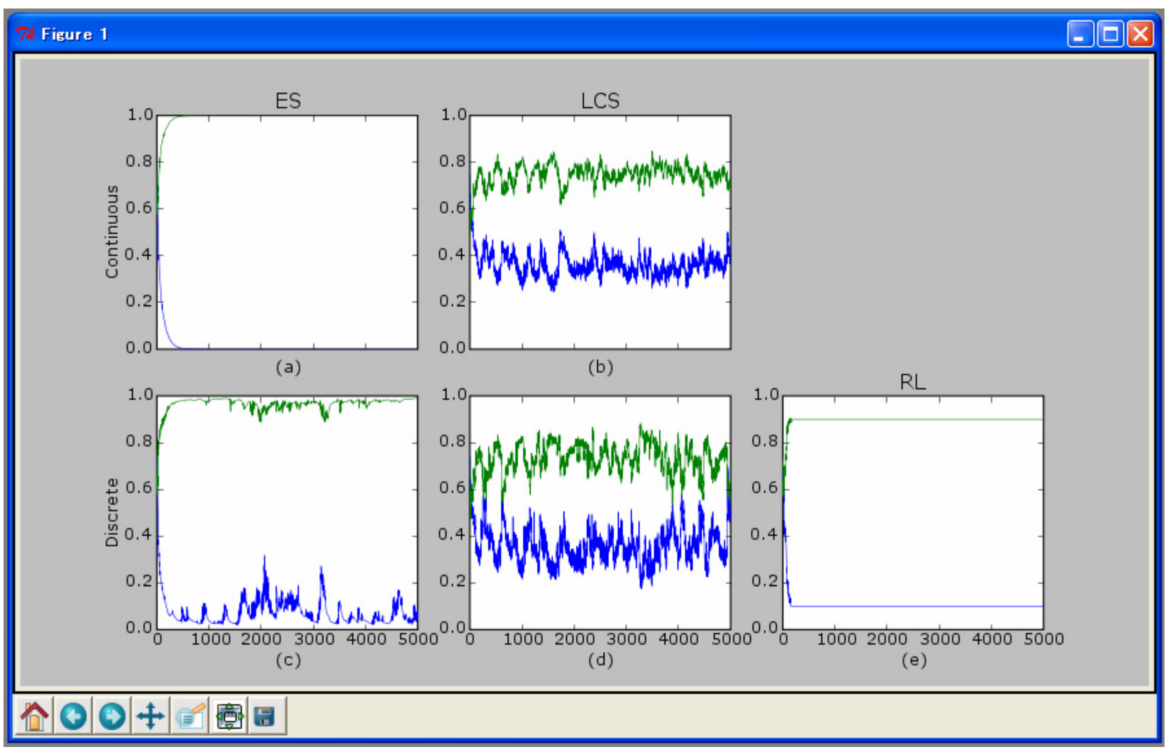

FIGURE 6 Simulation results

Discussion

Results of Cross-element Validation

Theoretical results from game theory prove that both rational agents A_1 and A_2 receive the minimum and maximum payoffs at the final negotiation process, respectively. This is because A_1 in our simulations has to accept any small offer proposed by A_2 at the 10th negotiation process; otherwise, A_1 cannot receive any payoff (i.e., it receives a null payoff). Therefore, it was expected that learning agents can acquire the maximum and minimum payoffs.

Analyzing Figures 6(a) and 6(b) shows that the payoff of ES-based agents finally converges at the mostly maximum or minimum value (i.e., 1 or 0), while that of LCS-based agents neither converges at a certain value nor becomes close to the maximum or minimum value. From Figures 6(a) and 6(c), it is observed that the results of ES-based agents using discrete knowledge representation degrade the results obtained when using continuous knowledge representation (note the rather wavy lines in Figure 6(c)). Finally, from Figures 6(c) and 6(e), it is observed that the payoff of ES-based agents was effected using two decimal digits, while RL-based agents converge at the mostly maximum or minimum value (i.e., 0.9 or 0.1).

These results show that ES-based agents with continuous knowledge representations and RL-based agents with discrete knowledge representations could produce results as expected by game theory. Therefore, both models are minimally validated.

From this analysis, it can be concluded that simulation results are sensitive to the learning mechanisms applied to agents. Also, even minor considerations in the knowledge representation, particularly discrete and continuous representations, may produce unexpected results.

As a result, it is strongly recommended that some cross-element validation of models be performed before deep analysis and interpretation of their simulation results.

X-MAS Compared with Other Tools

In order to help researchers in the field of social sciences simulate their models, several tools have been developed to reduce the difficulties of the programming process and enhance the understanding of the outcomes (e.g., Repast, Swarm, and Mason). However, performing cross-element validation will require the knowledge of some internal libraries to easily exchange elements in the model. In several cases, it may require reimplementation for all possible substitute elements in the model. The reason for this is that they are not designed for validation purposes but for easy program implementation. X-MAS, on the other hand, was designed to support the cross-element validation of ABS models and to facilitate program implementation.

X-MAS provides a framework for implementing generic models, and several variations of the model can be performed more easily. It is expected to be considered as a framework for the replication of models.

SUMMARY

Although ABS is becoming an essential tool in the study of complex social sciences, the validation of ABS models is still an important issue to be considered. Cross-element validation was proposed in our previous work. This process consists of performing the validation within a model by comparing the simulation results of the model under several instances of some of its composite elements. To support the cross-element validation process of ABS models, this paper presented the cross-element validation for multi-agent-based simulation (X-MAS). This tool provides facilities for simplifying the cross-element validation of ABS models. It also facilitates the implementation of general-purpose ABS models. The potential of X-MAS was tested by means of a bargaining game model, by evaluating several learning mechanisms applied to the agents. It showed that simulation results can be strongly affected by even small variations in the elements. In particular, arbitrary assumptions in the learning mechanism and knowledge representation schemes may produce unexpected results. Therefore, cross-element validation should be performed before deep analysis and interpretation of the implemented model.

Further research includes (1) implementating several GUIs for interaction with models and (2) performing cross-element validation of several models.

ACKNOWLEDGMENTS

This research was conducted as part of "Research on Human Communication" with funding from the National Institute of Information and Communications Technology (NICT), Japan. The research was also supported in part by the Sasagawa Scientific Research Grant from the Japan Science Society. In addition, it was supported in part by a Grant-in-Aid for Scientific Research (Young Scientists (B), 17700139) from the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan. We would like to express our thanks for all of the above support.

REFERENCES

Axtell, R., R. Axelrod, J.M. Epstein, and M. Cohen, 1996, "Aligning simulation models: A case study and results," *Computational and Mathematical Organization Theory (CMOT) 1*(1):123–141.

Back, T., F. Hoffmeister, and H. Schwefel, 1991, "A survey of evolution strategies," in *Proceedings of the Fourth International Conference on Genetic Algorithms*, pp. 2–9.

Back, T., F. Hoffmeister, and H. Schwefel, 1993, "Evolutionary programming and evolution strategies: Similarities and differences," in *Proceedings of the 2nd Annual Conference on Evolutionary Programming*, pp. 11–22.

Boost Libraries, 2005, home page; available at http://www.boost.org/.

George Mason University, 2005, MASON; available at http://cs.gmu.edu/~eclab/projects/mason/.

Goldberg, D., 1989, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.

Holland, J., K. Holyoak, R. Nisbett, and P. Thagard, 1986, *Induction*, The MIT Press.

Muthoo, A., 2000, "A Non-technical Introduction to Bargaining Theory," *World Economics 1*(2):145–166.

Oliver, J., 1996, *On Artificial Agents for Negotiation in Electronic Commerce*, Ph.D. thesis, University of Pennsylvania.

Repast, 2005, home page; available at http://repast.sourceforge.net/.

Osborne, M., and A. Rubinstein, 1994, *A Course in Game Theory*, The MIT Press.

Rubinstein, A., 1982, "Perfect Equilibrium in a Bargaining Model," *Econometrica 50*(1):97–109.

Stahl, I., 1972, *Bargaining Theory*, Economics Research Institute at the Stockholm School of Economics.

Sutton, R., and A. Barto, 1998, *Reinforcement Learning: An Introduction*, The MIT Press.

Swarm Development Group, 2005, home page; available at http://www.swarm.org.

Takadama, K., Y. Leon, N. Sugimoto, E. Nawa, and K. Shimohara, 2003, "Cross-element validation in multiagent-based simulation: Switching learning mechanisms in agents," *Journal of Artificial Societies and Social Simulation 6*(4).

AGENT-BASED MODELS AS QUANTITATIVE SOCIOLOGICAL METHODOLOGY: CALIBRATING SIMULATION MODELS TO DATA AND FINDING CONFIDENCE INTERVALS FOR MODEL PARAMETERS

S.P. WILCOX,* Northrop Grumman Corporation, Arlington, VA

ABSTRACT

A simulation model of neighborhood crime rates and its estimation using data illustrate how a simulation model can be employed as a supplement to prose-style sociological reasoning and how the simulation model can be used as an estimation methodology that replaces the traditional method of specifying regression models with mixtures of system-level and individual attributes as predictors. The core of the estimation framework is a generalization of the method of simulated moments (MSM) estimator of econometrics, which matches moments (i.e., expected values and variances) and a practical estimation methodology. A simulation meta-model giving the approximate relationship between model parameters and functions of the moments of model outputs, such as means and variances, is employed to calibrate the model to social data. The result is a simulation model-based replacement for the current paradigm for empirical sociological methodology that avoids reification and the ecological fallacy, which is applied to the estimation of a new model of neighborhood crime.

Keywords: Crime, quantitative sociological methodology, method of simulated moments, simulation model calibration, simulation meta-modeling

INTRODUCTION

To play off the title of a famous book on statistics (Christensen 2002), statistical methods in the social sciences often consist of “plane answers to complex questions.” To answer questions about complex social phenomena, such as neighborhood crime rates, sociologists often commit the ecological fallacy by positing the reality of abstractions and then employing their measures in multiple regression models, which fit a hyperplane to the behavioral characteristics of a complex system. Usually these models pertain to a closely related group of dependent variables. These models are justified through discursive, philosophical-style social theory that makes ontological and behavioral claims about system dynamics and the relationship between the individual levels of analysis, but the degree of logical rigor achieved does not necessarily justify the epistemological claims made for the regression model specification. From the perspective of social simulation modeling, however, system models are available, and the question is one of choosing the right model, assessing where a candidate simulation deviates from validation data sets, and finding a good set of model coefficients in an efficient manner.

Many have noted the discursive nature of sociological argument, either as a good thing (Sica 2004) or a shortcoming (Mahoney 2004). Computational social science is emerging as an alternative paradigm, but it is time to go beyond demonstrations of promise and develop

* *Corresponding author address*: Steven P. Wilcox, Northrop Grumman Information Technology, 2100 Washington Boulevard, Arlington, VA 22204; e-mail: steve.wilcox@ngc.com; spwilcox@acm.org.

agent-based modeling as an alternative methodology on an equal basis to the traditional paradigms. Since quantitative sociological methodology operates by performing statistical tests on model parameters that are estimated from data, it is natural to wish to do the same thing with agent-based models and other social simulations. The process of theoretical development being considered here is to encode discursive theoretical arguments into agent-based models, and then evaluate their implications by executing the models. The next step is then to perform statistical tests on various model parameters, in order to test the theories under consideration.

We approach the issue of estimating parameters in social simulation models as a problem of model calibration or validation, in which the model employs parameters that cannot be estimated by direct observation in the real world. In order to make the process of theoretical testing work well, however, methods of parametric hypothesis testing are needed for a broader range of models. This advance is achieved by employing simulation meta-modeling and response surface methodology as a method to achieve the matching of sample moments from simulation runs to empirical data, and by employing a jackknife estimator of the standard errors of the estimated parameters. The resulting methodology brings together elements of econometrics, operations research, and computational statistics in a new way that is computationally feasible and suitable for use as a quantitative social science methodology.

In this paper, we showcase the methodology from a recent paper on crime in the sociological literature by Browning, Feinberg, and Dietz, to be referred to as BFD (Browning et al. 2004). Here we attempt to employ an agent-based model as a methodological improvement over the theoretical arguments and methodological approach of the subject paper, in which it is desired to substantiate a view of crime as occurring within a context of a network of social exchange relationships that tend to impede the suppression of crime. We consider a new agent-based model of neighborhood crime and apply the proposed methodology to reanalyze the data correlation matrix presented.

SIMULATION MODEL ESTIMATION METHODOLOGIES

The question of how to estimate the parameters of a simulation model, calibrate it, and/or validate it is not new. The case at hand is the presence of parameters (or processes) that are not directly observed in the data or that otherwise do not directly correspond to the observables, and the context is complex, stochastic simulations that are not susceptible to analytic (calculus-based) approaches. First we take a quick overview of the approaches that have been taken, and then we discuss the genealogy of the proposed methodology.

The basic principle of simulation or parameter estimation from outputs is to match such things as means and variances of model outputs to data from a real system. The conventional economic methodology of deriving analytical results has been applied to agent-based economic models, such as the application to exchange rates by Alfarano et al. (2005). Axtell et al. (2002) performed a systematic search of an eight-dimensional parameter space in order to find a match to the archeological record of the Anasazi of Long House Valley, but this approach has exponential computational complexity. From the nonconvex optimization perspective, Gilli and Winker (2003) investigated a heuristic for matching the moments of an agent-based simulation model of exchange rates by refining a set of points in a three-dimensional parameter space that bracket the optimal solution. Many others have employed a variety of approaches to the optimization problem of finding optimal parameters for a process that is modeled by using simulation. We now consider why the matching of simulation moments is also an econometric idea, and how to make it computationally feasible in the case of complex stochastic simulations through the application of simulation meta-modeling.

The generalized method of moments (GMM) econometric estimation methodology (Hansen 1982) is a powerful methodology for estimating the coefficients for a wide variety of estimation problems. GMM is instructive in that it performs a match between features of a model and features of the data and then defines asymptotic standard errors for the resulting parameter estimates. Following the textbook by Mátyás (1999, Chapters 1 and 10), the basic approach is to define a continuous, $q \times 1$ vector function

$$f(x_t,\theta) = s(y_t,z_t) - \sigma(z_t;\theta)$$

of a parameter vector, θ, and data vectors $x_t = (y_t,z_t)$:$t = 1,...,T$, where x_t is divided into dependent variables y_t and independent variables z_t. We attempt to estimate θ_0, the true value of θ, by using the moment conditions

$$E\left[f(x_t,\theta)\right] = 0 \,.$$

A simple example of this is writing the mean and variance of X by using the notation of expected values and then subtracting their theoretical values, which are functions of θ. As we do not have these expected values, we employ instead the sample moments as a function of θ,

$$f_T(\theta) = T^{-1}\sum_{t=1}^{T} f(x_t,\theta) \,.$$

If there are q parameters to estimate, then the method of moments estimator is to solve the exactly identified system of equations, $f_T(\theta) = 0$, for $\hat{\theta}_{TM}$ in terms of the data. However, there are many possible moments from which to choose in this procedure. If there are fewer parameters to estimate than moment conditions, the problem is overidentified, and the GMM estimator defines a positive definite weighting matrix A_T, from which we obtain

$$Q_T(\theta) = f_T(\theta)' A_T \, f_T(\theta) \,,$$

a measure of the “distance” away from satisfying the moment conditions. Assuming, among other regularity conditions, that $f_T(\theta)$ is continuously differentiable, the minimization of $Q_T(\theta)$ over θ yields the GMM estimator, $\hat{\theta}_T$, which is obtained via solving the first-order conditions.

For several varieties of complex econometric models, GMM estimation is computationally infeasible due to such problems as high-dimensional multiple integrals being required to compute the expected value of the dependent variable, as might be present as part of $\sigma(z_t;\, \theta)$. The method of simulated moments (MSM), which originated with Pakes and Pollard (1989) and McFadden (1989), addresses this issue. MSM has also been suggested by Richiardi (2004) as an estimation methodology for agent-based computational economics. In MSM, one performs a simulation that generates estimates of $\sigma(z_t;\, \theta)$, such as the natural Monte Carlo estimator based on a sample of size R,

$$\hat{\sigma}_R(z_t;\theta) = R^{-1}\sum_{r=1}^{R} s\left[y_t^{(r)}(\theta), z_t\right],$$

where $y_t^{(r)}(\theta)$ is the r'th simulated value of y_t. Then the MSM estimator $\hat{\theta}_{MSM}^{R}$ is obtained by minimizing the criterion,

$$\left[\sum_{t=1}^{T} f_R(y_t, z_t;\theta)\right]' A\left[\sum_{t=1}^{T} f_R(y_t, z_t;\theta)\right],$$

where A is a positive definite weighting function and

$$f_R(y_t, z_t;\theta) = B(z_t)'\left[s(y_t, z_t) - \hat{\sigma}_R(z_t;\theta)\right].$$

This optimization is facilitated by the formation of a Monte Carlo estimate of the derivatives of $\hat{\sigma}_R(z_t;\theta)$ with respect to θ, which may be readily available from the simulation as analytical derivatives conditional on the various pseudorandom number values.

The MSM procedure does not quite fit the envisioned application on two counts. First, a correlation matrix is not a generalized moment, but a function of moments. Second, it is not suitable for the general social or agent-based simulation model in that the derivatives of $\hat{\sigma}_R(z_t;\theta)$ are not necessarily available. While the former is accommodated in the discussion by Gelman (1995) by a modification to the normal equations, we address the latter through the use of least-squares models in order to estimate the relationships between the simulation parameters and the correlations or other features of interest. Using these relationships, or simulation meta-model, we then find an approximate minimum of the criterion function. Here we apply least-squares again to find the optimal set of parameters to minimize the "distance" between the simulation's correlation matrix and the real-world one. This procedure thus creates a point of contact with the literature on response surface methodology (Kleijnen 1998; Myers and Montgomery 2002) and simulation optimization (Andradóttir 1998; Fu 2002).

In outline, the approach proposed for agent-based modeling is as follows. For a model with Q parameters, define a set of N features, such as moments, functions of moments, and other definable functions of the model outputs that will be the basis of distinguishing good from poor models. The expected value of these should be a smooth function of the parameters, and their sampling variance should go down with sampling size. Then, running the model in M batches, collect the sets of feature vectors, s_i, $i = 1,...,M$. Using ordinary least-squares (OLS) or weighted least-squares (WLS), the regression equation for all feature vectors, consolidating the simulation meta-models, is given in matrix-vector form as

$$s_i = \Theta b_i + \varepsilon_i,$$

where Θ is the $M \times (Q + 1)$ matrix of parameters by batch number, including a column of ones; each b_i is a column vector consisting of Q coefficients plus an intercept term; and s_i is the M-vector of generalized moments of type i. At this point we check the meta-model with respect to the regression assumptions, such as linearity and homoscedasticity. It is also possible to perform multivariate regression tests of significance showing whether any of the features are affected by a given parameter (see Johnson and Wichern 1992, Chapter 7).

In aid of finding the desired estimates, define the concatenations of column vectors

$$\hat{B} = [\, \hat{b}_1, \hat{b}_2, \hat{b}_3 \ldots \hat{b}_M \,] \text{ and } S = [\, \hat{s}_1, \hat{s}_2, \hat{s}_3 \ldots s_N \,] \;,$$

Then the predicted values from the regression model are given by the $M \times N$ matrix,

$$\hat{S} = \Theta \hat{B} \,.$$

We wish to find the values of θ that create the best fit to the actual feature data, s_0. $\hat{\theta}$ is obtained by minimizing

$$\left\| \theta' \hat{B} - s_0' \right\| = \left\| \hat{B}' \theta - s_0 \right\| \;.$$

This is a regression equation, which is estimated by using least-squares. Here, roles are reversed as the matrix of estimates from the first phase is transposed to become the set of predictor variables and as the target features form the dependent variable vector. We can think of this regression as finding the maximum likelihood estimator of θ conditional on the estimate $\hat{B}$ and assuming the accuracy of the meta-model for the relationship between the features and the parameters. From the perspective of inverse problem theory (Tarantola 2005), this regression is the solution to an inverse problem in which there is no prior information and the linear model is employed as an approximation to the parameter-data relationship. While inverse modeling is an established methodology in hydrology — as seen in Hill (1998) as well as Poeter and Hill (1998), for example — our setting has an additional complication in that the simulation model is stochastic.

Note that the usual MSM approach calculates standard errors for the coefficients on the basis of the availability of an accurate derivative of the moment function conditional on the various random number instantiations that occurred in the simulation. We replace this derivative with a least-squares estimate, whose sampling error affects the standard error of the estimates in addition to the consequences of any lack of fit in the meta-models.

As an alternative to an involved matrix formula for an asymptotic approximation to the standard error of the estimates, we consider utilizing the Quenouille-Tukey jackknife, whose use as a variance estimator is discussed in Efron and Stein (1981). The jackknife is a feasible and general-purpose method of conservatively estimating standard errors while also reducing estimation bias. To utilize the jackknife, form the pseudo-values

$$P_j = \hat{\theta} + (N-1)(\hat{\theta} - \hat{\theta}_{-j}) \qquad \text{for } j = 1, \ldots, N \,,$$

where $\hat{\theta}_{-j}$ is the estimator of choice computed with observation j removed. In this case, the "observations" that are removed are one of the several generalized moments or features being matched and its associated regression coefficients. The observation thus deleted embodies the sampling error for the original data as well as the estimation error of the simulation meta-model coefficients for that parameter. The jackknife estimators of the mean and variance of the estimate are the mean of the P_j and the customary formula for the sampling variance of a mean:

$$\hat{\hat{\theta}} = \frac{1}{N}\sum_{j=1}^{N} P_j$$

and

$$Var(\hat{\hat{\theta}}) = \frac{1}{N(N-1)}\sum_{j=1}^{N}\left(P_j - \hat{\hat{\theta}}\right)^2 .$$

For simplicity, these are the equations for a single component of the estimate vector, but a covariance matrix of the estimates can also be computed in the usual fashion from the P_j as well.

NEIGHBORHOOD CRIME MODELING

Theoretical Discussion

In BFD's review of the crime literature, it is clear that sociological theories of crime are based on hypotheses about social interactions in a neighborhood network. Not well-covered are economic theories of crime, such as rational choice theory (Becker 1968, and many others to follow) and the conception of peer effects as positive and negative externalities (Glaeser et al. 1996; Calvó-Armengol and Zenou 2004), which are also relevant.

Social disorganization theory (Shaw and McKay 1969; Kasarda and Janowitz 1974; Kornhauser 1978) views interpersonal social attachment as a good thing. According to Shaw and McKay, poverty, residential instability, and ethnic heterogeneity promote crime by inhibiting the formation of neighborly networks and attenuating community-level action against crime. According to Kasarda and Janowitz, extensive friendship and kinship bonds strengthen neighborhood attachment, and Kornhauser finds that weak bonds mediate the effect of disadvantage on the capacity for social control.

The cultural transmission model (Whyte 1937; Wilson 1996; Crane 1991) focuses on the legitimate social networks as bulwarks against a counterculture of crime. The criminal subculture emerges in opposition to mainstream culture, and strong networks in socially disadvantaged communities may facilitate its spread. Thus, there is a contagion of problem behaviors, for which gang culture provides social support.

BFD propose and empirically support a negotiated co-existence model, in which social networks are a source of general social capital for offenders, which tends to protect them. Thus the attitude of neighborly efficacy to fight crime tends to suppress criminal behavior but is offset to some degree by social capital. Thus, social disorganization theory is not quite right, but BFD wish to avoid attributing too much organizational capacity to the criminal networks as well.

The economic literature on peer effects in crime is intriguing as a contrast because it is inherently agent-oriented. The analyses by Glaeser et al. (1996) and Calvó-Armengol and Zenou (2004) highlight the importance of heterogeneity in agents' toleration for crime as a moderator of peer influence effects, which act as source of training and facilitation. The positive externalities due to the interactions between criminals contrast with their aggregate competition for resources. These competing phenomena help explain the variability of crime rates across time and space.

Critique of BFD's Model

BFD support the negotiated coexistence model by estimating a regression model by using data defined at the neighborhood level. In it, an interaction term between the level of the attitude of efficacy to fight crime and an attitudinal measure of network exchange shows up as a significant predictor of the crime rate. Individual-level attitude and household-level victimization data are aggregated into an area measure, by using hierarchical linear models to obtain empirical Bayes residuals as the dependent variable and main independent variables. This effectively partials out gender, age, race/ethnicity, education, employment status, marital status, years of residency, home ownership, and number of recent moves. The dependent variables are violent crime victimization (log odds) and the logarithm of the homicide rate. Control variables include measures of disadvantage, residential stability, population density, immigrant concentration, and the lagged homicide rate. Support is found for the negotiated coexistence model.

BFD's results may be critiqued in that using neighborhood-level data to support theories of agents in social networks succumbs to the ecological fallacy. Also, using attitudes as "independent variables" is questionable, since attitudes may be an accommodation to facts rather than their cause. Attitudes are clearly endogenous as a class, and there is always a question concerning the direction of causality. There is also the problem of completeness when reasoning in discursive ways. It is not always clear that the prose theory supports a certain sign of regression coefficient, since something may be left out of the reasoning.

Building an agent-based form of the theory of crime has promise for addressing the issues that arise in the consideration of BFD's analysis. By building the model at the agent level, the ecological fallacy can be avoided. By using an agent-based model to reason about the way in which different phenomena interact to produce an expected result, we avoid the problem of incomplete reasoning, although the problem may occur at a higher level in the form of the choice of models or perspective. Within the limits of causal reasoning, the endogeneity of attitudes can be addressed in an agent-based model by specifying path models with appropriate loops.

Agent-based Model Development

The model developed here and presented in Figure 1 incorporates a two-dimensional analogue of the circular social influence network considered by Glaeser et al. (1996), in which the features of nonhomogeneous occupational preferences and competition among criminals for scarce economic resources create a situation in which disparate equilibria are possible and the crime rate can vary significantly over time. In addition, it includes the interplay between the attitude of neighborly anti-crime efficacy and the behavior of being a criminal.

Since neighborhoods are geographical, this model represents the attitudes and behavior of residents on a two-dimensional 50×50 lattice but uses the toroidal topology to avoid edge effects. In a real neighborhood, one has more than the eight neighbors present in many lattice models. Here we use the Moore neighborhood of radius 4, giving a set of 80 neighbors with

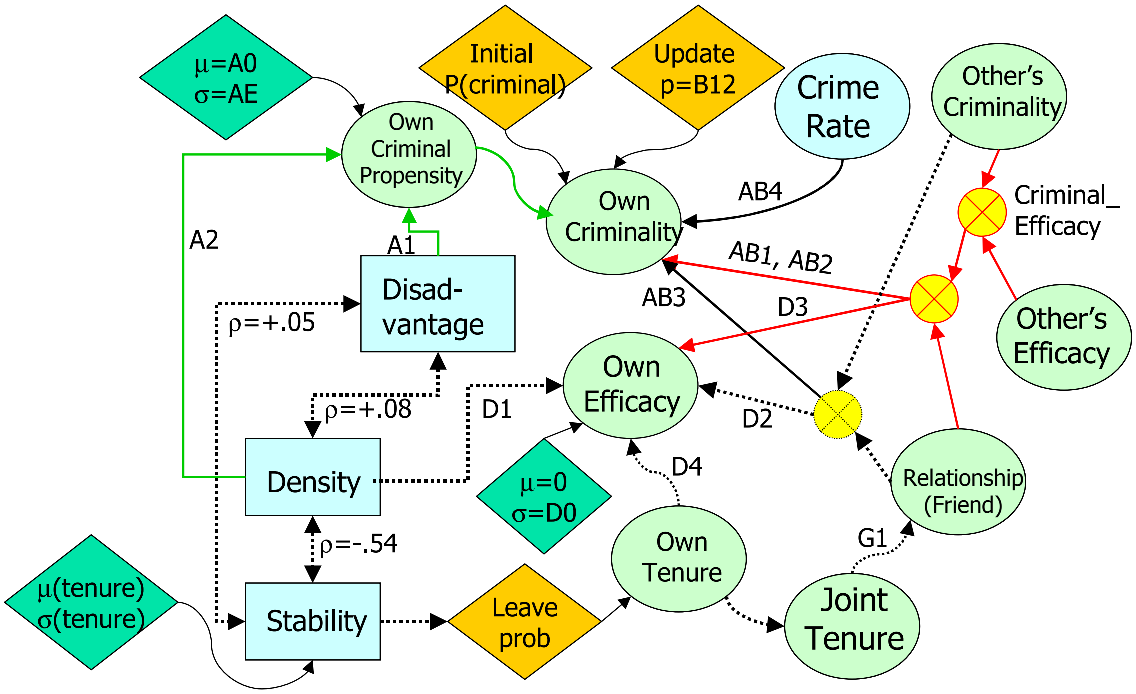

FIGURE 1 Path diagram of the agent-based model

whom an agent may form friendships and by whom an agent may be influenced. However, the crime rate's economic effect on occupational choice is taken over the entire set of 2,500 agents.

In this model, agents may be either criminals or law abiding, and they possess a level of "efficacy." Their underlying criminal propensities are heterogeneous within the neighborhood. Social influence is conveyed at three levels: friends, nonfriends in the Moore neighborhood, and the community as a whole. Agents may move out of the neighborhood and thus be immediately replaced. They also may change status between criminal and noncriminal and form friends. All friendships are two-way.

In the blue boxes, *Density*, *Stability*, and *Disadvantage* are unit-free exogenous variables that drive residents' *Criminal Propensities*, *Efficacy*, and *Probability of Leaving*, seen in parameters *A1*, *A2*, and *D1*. In the green diamonds are scale and location parameters related to these exogenous variables. *Tenure* is influenced by a yearly leaving probability that is derived from the stability parameter. *Criminality* is a choice that is initialized and then reviewed with probability *B12*. If it is to updated, the agent is a criminal according to a criminal propensity, which is heterogeneous between agents, as well as to the influences of friends and neighbors. In accordance with the peer effects literature, the global crime rate has a suppressing effect *(AB4)* on an individual's propensity to choose criminality, while friends who are criminals have a separate effect *(AB3)*. The value of efficacy used in the influence calculations depends on whether the person doing the influencing is a criminal, in which case the *Criminal_Efficacy* parameter is used. The effective efficacy of the other thus computed has two different effects *(AB1* and *AB2)* on the other's occupational choice, depending on whether the other is a friend or just a neighbor. The update of *Efficacy* is affected by the crime rate among the agent's friends according to parameter *D2*, and their *Effica*cy according to parameter *D3*. *Efficacy* is also influenced by tenure *(D4)* and density *(D1)*. Finally, friendships grow with joint tenure, according to probability *G1* each year.

Implementation Details

Each iteration of the simulation is nominally one year, during which each agent's status is stochastically updated in fixed order:

1. Moving out of the neighborhood,
2. Building friendships (which accumulate),
3. Deciding whether to reevaluate one's occupation, and
4. Updating attitude of efficacy.

First it is determined whether the agent will leave on the basis of the tenure and stability-related calculations of the probability of leaving detailed above. The probability of leaving does not depend on the crime rate. If the agent leaves, it is replaced with a new agent; otherwise, the agent is updated.

A new agent is given a criminal status and propensity according to the probability utilized in the initialization of criminal status, and it is given values of baseline efficacy, friendship probability, and friendship status according to the initialization procedure. The starting value of tenure is 1 year.

An agent who does not move out of the neighborhood is influenced by the *efficacy* of both friends and nonfriends in his Moore neighborhood, but according to separate coefficients. If the other is a criminal, the effective *efficacy* is set at a parameter; otherwise, the effective *efficacy* is the actual *efficacy*. For both classes of Moore neighbors, the influence of the other on one's own efficacy is on a per-agent basis, making influence of the two categories proportional to their numbers divided by the total number of Moore neighbors.

The latent probability of becoming a criminal is the $\exp(x)/[1 + \exp(x)]$ function of the sum of the agent's *criminal propensity* and products of parameters with the efficacies of friends and nonfriends in the agent's Moore neighborhood, as well as the crime rates among friends and globally. If a Bernoulli trial against the criminal status update probability is successful, criminal status is updated according to the latent probability.

Friendship cumulatively increases, and new friends are added from the Moore neighborhood according to the friendship creation probability. Tenure is incremented by one each iteration. Efficacy is based on the baseline efficacy for the agent calculated at initialization, plus the products of parameterized coefficients multiplied by the friend *crime rate*, the friend *efficac*y, and *tenure*.

One may comment on the model's complexity. The path model implied by theory is fairly complex yet incomplete, and things had to be added. The effect of time on relationship building is common sense, but not explicitly stated as theory. Cognitive consistency theories could be further exploited to suggest additional relationships between one's own attitudes and own behavior. The model is also too simple in that there is no distinction between the rates of victimization and the populations of criminals.

Model Parameterization

Table 1 presents the 19 parameters corresponding to the path model in Figure 1 to describe the simulation process. In the path diagram, the prefix containing the affected node is omitted, but it is included here for clarity.

DATA ANALYSIS

Empirical Data Reanalyzed

BFD report a correlation matrix of the data analyzed in their regression analyses, which form the feedstock for the demonstration of the proposed analytical method. The correlations reported by BFD and utilized in this paper are presented in Table 2. The correlations highlighted in orange are the exogenous variables. The presence of the 4-year lag of the crime rate gives us a reading on the level of consistency over time possessed by the phenomenon of crime, which is

TABLE 1 Agent-based model parameters

Number	Name	Type of Parameter	What It Is Multiplied by	What It Affects
1	pcriminal	Initial probability	-----	Initial criminal status
2	CriminalPropensity_A0	Constant	-----	*Criminal Propensity*
3	CriminalPropensity_AE	Variability	-----	*Criminal Propensity*
4	CriminalPropensity_A1	Coefficient	*Disadvantage*	*Criminal Propensity*
5	CriminalPropensity_A2	Coefficient	*Density*	*Criminal Propensity*
6	CriminalImpulse_AB1	Coefficient	*UnrelatedEfficacy*	*CriminalImpulse*
7	CriminalImpulse_AB2	Coefficient	*FriendEfficacy*	*CriminalImpulse*
8	CriminalImpulse_AB3	Coefficient	*FriendCrimeRate*	*CriminalImpulse*
9	CriminalImpulse_AB4	Coefficient	*CrimeRate*	*CriminalImpulse*
10	Efficacy_D0	Variability	-----	*Efficacy (initial)*
11	Efficacy_D1	Coefficient	*Density*	*Efficacy (initial)*
12	Efficacy_D2	Coefficient	*FriendCrimeRate*	*Efficacy*
13	Efficacy_D3	Coefficient	*FriendEfficacy*	*Efficacy*
14	Efficacy_D4	Coefficient	*Tenure*	*Efficacy*
15	Criminal_Efficacy	Constant	-----	*Effective efficacy*
16	Criminal_B12	Update probability	-----	*Criminality*
17	Friend_G1	Probability of Formation	-----	*Friendship*
18	Tenure_mean	Constant	-----	*Probability of Leaving*
19	Tenure_std	Variability	-----	*Probability of Leaving*

TABLE 2 Correlations Reported by Browning et al. (2004)

	Crime_Rate (Homicide)	*Disadvantage*	*Residential_stability*	*Density*	*Previous Crime_Rate (Homicide)*	*Collective_Efficacy*	*Network_Interaction*	*High_interaction*	*High_interaction*Collective_Efficacy*
Crime_Rate (Homicide)	1.00								
Disadvantage	0.76	1.00							
Residential_stability	0.04	0.05	1.00						
Density	-0.03	0.08	-0.54	1.00					
Previous Crime_Rate (Homicide)	0.81	0.77	-0.06	0.09	1.00				
Collective_Efficacy	-0.54	-0.56	0.38	-0.44	-0.60	1.00			
Network_Interaction	-0.14	-0.13	0.05	-0.18	-0.13	0.47	1.00		
High_interaction	-0.07	-0.02	0.08	-0.17	-0.07	0.36	0.75	1.00	
*High_interaction*Collective_Efficacy*	-0.30	-0.33	0.27	-0.26	-0.37	0.68	0.38	0.39	1.00

important with respect to the economic literature, as the level of instability has been the difficult part to explain. The means of the variables are not given. Many of the means and standard deviations are meaningless, since they are empirical Bayes residuals or attitude measures.

Since they are survey estimates and the outcome of empirical Bayes purification, many of the subject variables are subject to estimation error themselves. This creates the problem of unmodeled measurement error in the predictor variables, which causes bias in the regression estimates. For the present purposes, it would also be helpful to have estimates of the reliability of the predictor measures.

Simulation Model Data Collection

Ecological data are collected from the 2,500 persons in the simulated neighborhood. After initialization, the model is executed for 20 iterations prior to the collection of the lagged log odds crime rate, and then executed for 4 more iterations prior to the collection of the rest of the data. Residential stability is collected as the mean of tenure (in simulation iterations). Network interaction is the mean of the friendship status. Efficacy is also the mean of this variable before its adjustment for criminal status. The log odds of the crime rate are estimated by using an accommodation for the possibility of zero crime rates. The agent-based model was implemented in Matlab 5.3 (The Mathworks, Inc., 1999).

Correlations of interest were calculated from batches of 30 independent model runs and used as data for the later method of simulated moments analysis. To achieve this, variables were calculated to support the interaction effects in the regression analysis BFD report. From the network interaction variable, the 70th percentile was calculated in order to create the indicator for high network interaction and thus its product with collective efficacy.

The experimental design employed in data collection is to uniformly generate the simulation parameters within the upper and lower bounds determined by the experimenter. An initial set of test runs yielded a set of confidence bands, which were employed in the subsequent batch of data collection runs. A total of 117 batches were collected after the deletion of those with missing values due to taking the logarithm of zero.

Data Analysis

Statistical Procedures

Standard tools of multivariate analysis and regression model checking are employed to estimate and assess the meta-models of correlations in the simulated data. SAS version 9 (SAS Institute, Inc. 2002) was employed for the bulk of the post-simulation analysis, although the estimator was also implemented in Matlab (by using standard least-squares formulas taken from Judge et al. 1988). This being the first application of the application of meta-modeling to MSM, the difficulties encountered are instructive. Since correlations are the dependent variable in the meta-models, a transformation to correct heteroscedasticity was needed. Fisher (1915) suggests the

$$\tanh^{-1}(r) = \log_e\left(\frac{1+r}{1-r}\right)$$

transformation, but this was modified to

$$\log_e\left(\frac{1.02+r}{1.02-r}\right)$$

in order to avoid difficulties with the logarithm of zero, which was encountered in some of the regressions. Other difficulties arise when not all of the correlations are predicted equally as well from the model parameters during the first least-squares estimation. This results in heteroscedasticity in the second phase regression, which is addressed by using a vector of weights calculated as the inverse of the residual variance estimates in phase 1. For this exploratory analysis, the standard errors are computed by using weighted least-squares regression rather than using the jackknife variance estimate.

Model Critique

An advantage of the new methodology is its ability to assess whether model parameters affect the measurements being made. In this study, the first-stage regression analyses exhibited strong effects of some parameters on the sampled correlations, but not others, as determined by using multivariate tests of the parameters. High points of the significance parade include *Efficacy_D1*, which is the effect of *density* on *efficacy*, and *Efficacy_D3*, which is the effect of having criminals as friends on one's feeling of efficacy. The mean and standard deviation of neighborhood stability (average tenure) also get high marks. Table 3 presents multivariate tests of the effects of the simulation parameters on the correlation statistics collected. Parameters 6–8 of the simulation concern the effects of one's efficacy and criminal behavior on another's criminal behavior. A multivariate test that the correlations generated by the model were not related to these parameters rejected this null hypothesis with a statistically significant Wilks' Lambda p-value. However, it seems surprising that so many of the parameters appear to be immaterial when examined in this fashion, although one can always look to increasing sample size. The one hoped-for lack of statistical significance is that of *pcriminal*, which is an initialization constant for agents, but this was marginally significant.

TABLE 3 Multivariate tests of significance for model parameters

Hypothesized Zero Effects	Name	Wilks' Lambda	FValue	NumDF	DenDF	ProbF
1	*pcriminal*	0.536767	1.7	33	65	0.0343
2	*CriminalPropensity_A0*	0.651701	1.05	33	65	0.4198
3	*CriminalPropensity_AE*	0.625851	1.18	33	65	0.2826
4	*CriminalPropensity_A1*	0.763691	0.61	33	65	0.9391
5	*CriminalPropensity_A2*	0.736327	0.71	33	65	0.8626
6	*CriminalImpulse_AB1*	0.465699	2.26	33	65	0.0025
7	*CriminalImpulse_AB2*	0.370912	3.34	33	65	<.0001
8	*CriminalImpulse_AB3*	0.627128	1.17	33	65	0.2888
9	*CriminalImpulse_AB4*	0.495998	2	33	65	0.0086
10	*Efficacy_D0*	0.595334	1.34	33	65	0.1569
11	*Efficacy_D1*	0.234221	6.44	33	65	<.0001
12	*Efficacy_D2*	0.572457	1.47	33	65	0.0923
13	*Efficacy_D3*	0.387803	3.11	33	65	<.0001
14	*Efficacy_D4*	0.527811	1.76	33	65	0.0259
15	*Criminal_Efficacy*	0.599631	1.32	33	65	0.1719
16	*Criminal_B12*	0.46647	2.25	33	65	0.0026
17	*Friend_G1*	0.514365	1.86	33	65	0.0166
18	*Tenure_mean*	0.347084	3.71	33	65	<.0001
19	*Tenure_std*	0.409598	2.84	33	65	0.0002
6 to 9	*Influence_on_crime*	0.065114	1.95	132	261.32	<.0001
6 to 8	*Interpersonal_Influence*	0.119368	2.04	99	195.5	<.0001

The power of the meta-modeling approach is balanced by the need for model checking. Consider, for example, in Figure 2 (the residual plot versus the predicted values for w183), the transform of the correlation between *residential stability* and *network interaction*. There appears to be a curvilinear effect, which, however, is not the case universally, as seen in the residual plot for w117 in Figure 3 (the transformation of the correlation between *density* and *collective efficacy*). Here the plot is basically acceptable, except for some question about the narrowing of the residuals toward the right side boundary. The plots for the other variables show some combination of these issues as well. This indicates the need for an ad hoc approach to assuring that the statistical meta-model fits, rather than relying on an automated procedure.

Estimated Model Parameters

Table 4 shows the estimated model parameters from the method of simulated moments by using weighted least-squares. The parameter estimates and p-values show some disappointments and some surprises. The tolerance values are included as an indication of an identification problem.

While most parameters are not statistically significant based on the estimated t statistic, three stand out as having statistics of greater than 2.5 in magnitude. Parameter 7, *CriminalImpulse_AB2*, which is the effect of the efficacy of friends on the agent's criminal impulse, is positive, which is contrary to expectation. This might be explained by omitting the effect of behavior on attitudes in the path diagram, but it is not obvious how this explanation might apply in the current case.

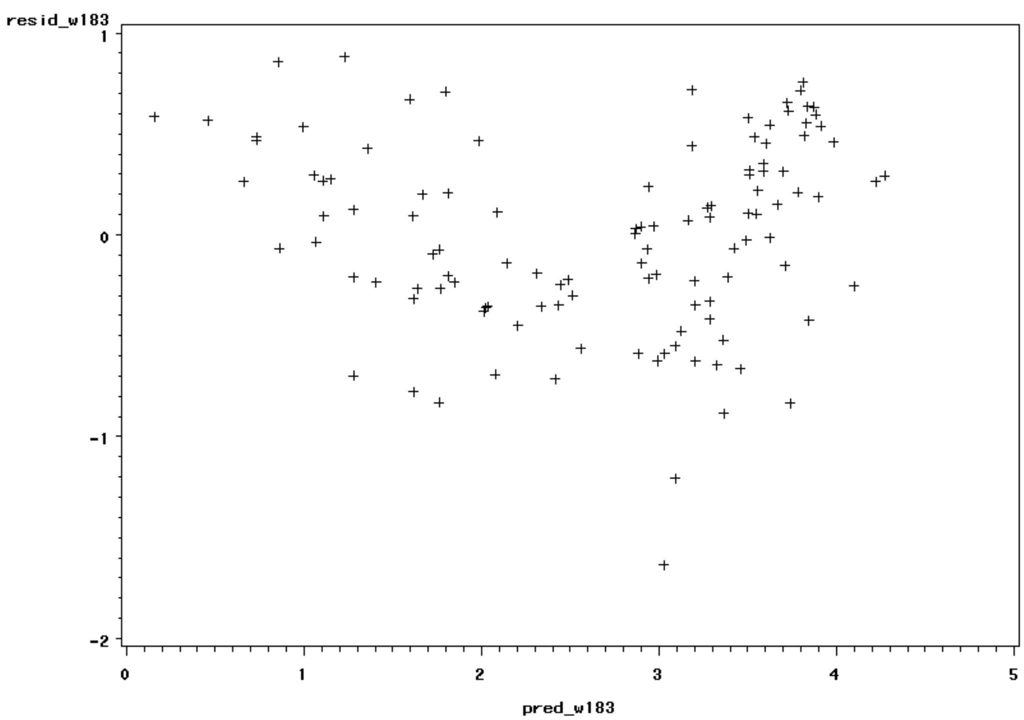

FIGURE 2 Residual analysis of the meta-model for w183

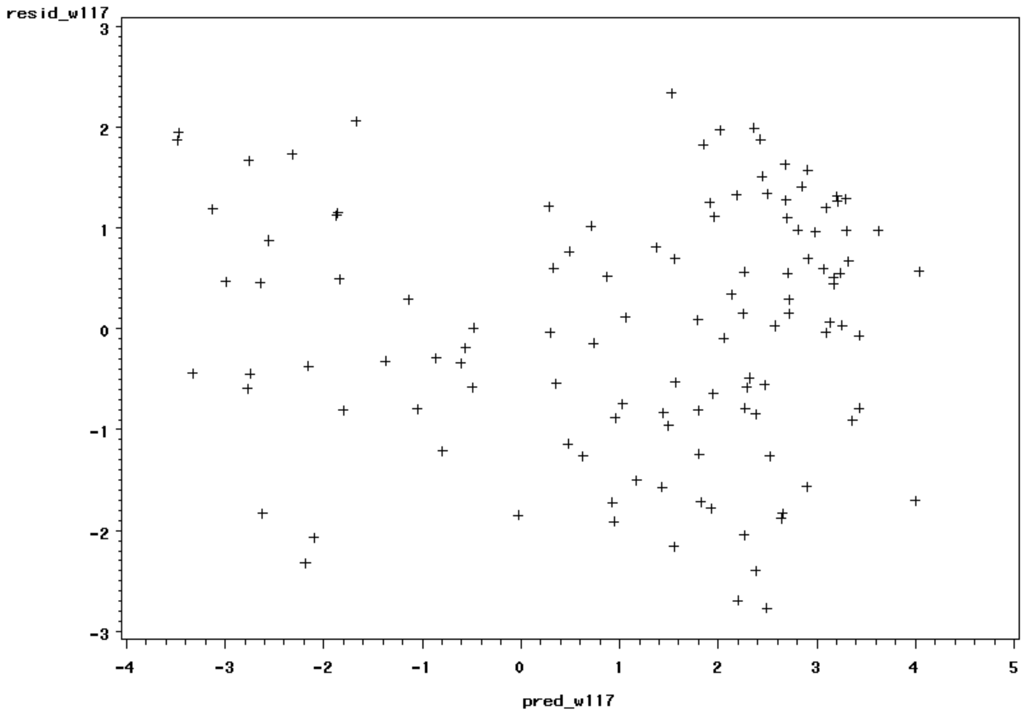

FIGURE 3 Residual analysis of the meta-model for w117

TABLE 4 Regression results

	Name	Sampled Minimum	Sampled Maximum	Parameter Estimate	Standard Error	t Value	Pr > |t|	Tolerance
1	*pcriminal*	-0.163	0.05	0.05817	0.2423	0.24	0.8137	0.04317
2	*CriminalPropensity_A0*	-2.85	1	-8.23298	5.40891	-1.52	0.1502	0.13064
3	*CriminalPropensity_AE*	-4.63	1	3.96582	3.855	1.03	0.321	0.18601
4	*CriminalPropensity_A1*	0.0018	3.34	5.76952	2.73867	2.11	0.0537	0.13844
5	*CriminalPropensity_A2*	-3.05	1.87	12.10151	7.95447	1.52	0.1504	0.07677
6	*CriminalImpulse_AB1*	-1.09	2.46	0.67283	5.25193	0.13	0.8999	0.07619
7	*CriminalImpulse_AB2*	-9.83	34.1	70.94446	23.66379	3	0.0096	0.06765
8	*CriminalImpulse_AB3*	-0.0804	2.28	6.61994	2.28734	2.89	0.0118	0.41749
9	*CriminalImpulse_AB4*	-22	14.8	-48.5363	36.60198	-1.33	0.206	0.12967
10	*Efficacy_D0*	-6.44	0.959	-4.32025	4.33556	-1	0.3359	0.04991
11	*Efficacy_D1*	-1.51	1.4	-0.76677	0.99723	-0.77	0.4547	0.1348
12	*Efficacy_D2*	-4.11	-0.0057	-2.6532	4.13226	-0.64	0.5312	0.12534
13	*Efficacy_D3*	-2.88	1.98	3.80379	3.3113	1.15	0.2699	0.03767
14	*Efficacy_D4*	-0.028	0.251	-0.4344	0.31996	-1.36	0.196	0.02597
15	*Criminal_Efficacy*	-20.7	-0.0151	29.16331	15.13305	1.93	0.0745	0.08524
16	*Criminal_B12*	-0.123	0.919	1.0151	0.77115	1.32	0.2092	0.14942
17	*Friend_G1*	-0.001	0.335	0.61942	0.22094	2.8	0.0141	0.08399
18	*Tenure_mean*	1.01	23.8	4.35497	12.99332	0.34	0.7425	0.05683
19	*Tenure_std*	-1.99	0.991	0.04896	1.7355	0.03	0.9779	0.03722

Parameter 8, the effect of the crime rate among the agent's friends on his own criminal impulse, also was estimated as being positive. However, the estimate is outside the range of the data, which was in part based on preliminary estimates based on the first 60 observations. Parameter 17, which is the rate of friend formation per year, was also outside the range of the data. Parameter 4, which is the effect of disadvantage on criminal propensity, was marginally statistically significant and positive. This is no surprise, but the estimate was also outside the range of the data.

The tolerance values indicate an approximate lack of full rank in the parameter estimates from the meta-models, which may indicate an identification problem with regard to estimating the original path diagram by using the available correlations. As an alternative, we may consider the stepwise weighted regression results (using the default settings) from SAS Proc Reg presented in Table 5.

As a subset of the original variables, the average criminal propensity, the effects of *urban density* and criminal friends on *criminal propensity*, the rate at which the criminal choice is updated, the probability rate of friendship formation, and neighborhood stability suffice to model the observed data as well as can be expected from the agent-based model developed here. Interestingly, it does not seem necessary to include the parameters associated with the particular interaction effect that was the centerpiece of the article by BFD. However, there is an issue with regard to the calibration of the time clock, as a mean neighborhood tenure of 35 years is too long, and a rate of friendship formation of 61% per year for the closest 80 neighbors seems high.

DISCUSSION

The application of simulation meta-modeling to estimate simulation parameters by using the MSM is feasible and scalable. With the methodology, it should be possible to extend agent-based models into the practice of quantitative sociological methodology by performing statistical tests of agent-based model parameters instead of regression parameters. The results of this approach are more accurate reasoning about the agents and activities reasoned about in substantive research. As befits a methodological pilot study, however, a number of critiques and research opportunities need to be addressed in further work. These issues are detailed below.

TABLE 5 Trimmed regression results

Number	Name	Parameter Estimate	Standard Error	T Value	Pr > |t|	Tolerance
2	*CriminalPropensity_A0*	-5.11923	2.55393	-2	0.0551	0.6836
5	*CriminalPropensity_A2*	10.84781	2.43872	4.45	0.0001	0.95291
8	*CriminalImpulse_AB3*	3.78883	1.70443	2.22	0.0348	0.87717
16	*Criminal_B12*	0.88839	0.36351	2.44	0.0213	0.78448
17	*Friend_G1*	0.61142	0.0881	6.94	<.0001	0.61627
18	*Tenure_mean*	35.42161	3.50078	10.12	<.0001	0.91326

There is much computational statistical work to be performed on the application of meta-models to MSM, including equations for the asymptotic variance matrix of the parameter estimates and simulation studies of the performance of the jackknife estimator and some alternatives. In this case study, a weighted least-squares estimate of the final parameters was employed as an approximation on the grounds that the meta-model is misspecified to some degree anyway. Another issue is the nonlinearity of some of the relationships between parameters and the generalized moments. While Fisher’s transformation helped, the one-size-fits-all approach had its limits in terms of addressing heteroscedasticity and nonlinearity. A difficulty with estimating nonlinear relationships is inverting them to determine the final parameter estimates, which is possible, but not as easy as solving a regression equation. A related difficulty is that the meta-model did not predict all the correlations equally well.

The agent-based model developed here had at its heart a path model of social influence. Any issue with such path models, such as identification, can be expected to present difficulties in this context as well. Further work with attitudes and behavior in relationship to crime would need to take care concerning model identification with respect to the underlying causal path model. The technique being explored here is not a substitute for collecting the right data and matching the right features of the data.

Employing an agent-based simulation as a replacement for theory increases the precision of one’s arguments, but at a price. With the simulation, the domain of modeling concern increases as one examines the arguments and aligns the theory. Because of the increased rigor of this process, the need for elaboration is made clear beyond what was apparent from the prose expression of the theory.

There are also additional substantive issues as well as issues with the simulation model that may be addressed. Calvó-Armengol and Zenou (2004) find that the equilibrium crime rates are sensitive to the geometry of the social network among criminals. Thus it is fair to ask how the social network assumptions made here affect the results. There is also a tendency for friends to be selected to match one’s choices, as can be seen in the case of adolescent sexual behavior (Billy and Udry 1985a,b). This may affect the friendship network insofar as influence effects are concerned. The current model has issues pertaining to the details of the simulation of friendship formation and the rate of leaving, which are a priority for model refinement.

In this paper, we have considered MSM by using weighted least-squares as a methodologically superior alternative to ecological regression models and prose sociological theory, and we employed a recent sociological journal article on neighborhood crime rates as a case study. Although in using this case study, a number of methodological issues and areas where hand statistical labor is required arose, the theoretical advantages of the methodology and its basic practicality are an important forward step in sociological methodology.

REFERENCES

Alfarano, S., F. Wagner, and T. Lux, 2005, “Estimation of Agent-based Models: The Case of an Asymmetric Herding Model,” to appear in *Computational Economics*; available at http://www2.warwick.ac.uk/fac/soc/wbs/research/wfri/prog2005/lux.pdf.

Andradóttir, S., 1998, "Simulation Optimization," Chapter 9 in J. Banks, ed., *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, New York, NY: John Wiley & Sons.

Axtell R.L., J.M. Epstein, J.S. Dean, G.J. Gumerman, A.C. Swedlund, J. Harburger, S. Chakravarty, R. Hammond, J. Parker, and M. Parker, 2002, "Population Growth and Collapse in a Multiagent Model of the Kayenta Anasazi in Long House Valley," *Proc. Natl. Acad. Sci. USA* **99** (Suppl. 3):7275–7279; available at http://www.pnas.org/cgi/content/abstract/99/suppl_3/7275.

Backus, G., and F. Gilbert, 1970, "Uniqueness in the Inversion of Inaccurate Gross Earth Data," *Philos. Trans. Royal Soc. London* **266**:123–192.

Becker, G.S., 1968, "Crime and Punishment: An Economic Approach," *Journal of Political Economy* **76**:169–217, March/April.

Billy, J.O., and J.R. Udry, 1985a, "Patterns of Adolescent Friendship and Effects on Sexual Behavior," *Social Psychology Quarterly* **48**(1):27–41.

Billy, J.O., and J.R. Udry, 1985b, "The Influence of Male and Female Best Friends on Adolescent Sexual Behavior," *Adolescence* **20**(77):21–32.

Browning, C.R., S.L. Feinberg, and R. Dietz, 2004, "The Paradox of Social Organization," *Social Forces* **83**(2):503–534, Dec.

Calvó-Armengol, A., and Y. Zenou, 2004, "Social Networks and Crime Decisions: The Role of Social Structure in Facilitating Delinquent Behavior," *International Economic Review* **45**(3):939–958.

Christensen, R., 2002, *Plane Answers to Complex Questions: The Theory of Linear Models*, 3rd Ed., New York, NY: Springer.

Crane, J., 1991, "The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping out and Teenage Childbearing," *American Journal of Sociology* **96**:1226–1259.

Efron, B., and C. Stein, 1981, "The Jacknife Estimator of Variance," *The Annals of Statistics* **9**(3):586–596.

Fisher, R.A., 1915, "Frequency Distribution of the Values of the Correlation Coefficient in Samples of an Indefnitely Large Population," *Biometrika* **10**:507–521.

Fu, M.C., 2002, "Optimization for Simulation: Theory vs. Practice," *INFORMS Journal on Computing* **14**:192–215.

Gelman, A., 1995, "Method of Moments Using Monte Carlo Simulation," *Journal of Computational and Graphical Statistics* **4**(1):36–54.

Gilli, M., and P. Winker, 2003, "A Global Optimization Heuristic for Estimating Agent Based Models," *Computational Statistics and Data Analysis* **42**(3):299–312.

Glaeser, E.L., and B. Sacerdote, 1999, "Why Is There More Crime in Cities?" *The Journal of Political Economy*, Vol. 107, No. 6, Part 2: Symposium on the Economic Analysis of Social Behavior in Honor of Gary S. Becker (Dec. 1999), pp. S225–S258.

Glaeser, E.L., B. Sacerdote, and J.A. Scheinkman, 1996, "Crime and Social Interactions," *The Quarterly Journal of Economics* **111**(2):507–548.

Hansen, L.P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* **50**:1029–1054.

Hill, M.C., 1998, *Methods and Guidelines for Effective Model Calibration*, Water-Resources Investigations Report 98-4005, Denver, CO: U.S. Geological Survey; available at http://wwwbrr.cr.usgs.gov/projects/GW_Subsurf/mchill/pubs/method/index.shtml.

Judge, G.G., R.C. Hill, W.E. Griffiths, H. Lütkepohl, and T.-C. Lee, 1988, *Introduction to the Theory and Practice of Econometrics*, 2nd Ed., New York, NY: John Wiley & Sons.

Kasarda, J.D., and M. Janowitz, 1974, "Community Attachment in Mass Society," *American Sociological Review* **39**:328–339.

Kleijnen, J.P.C., 1998, "Experimental Design for Sensitivity Analysis, Optimization, and Validation of Simulation Models," Chapter 6 in J. Banks, ed., *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice,* New York, NY: John Wiley & Sons.

Kornhauser, R.R., 1978, *Social Sources of Delinquency*, Chicago, IL: The University of Chicago Press.

Mahoney, J., 2004, "Revisiting General Theory in Historical Sociology," *Social Forces* **83**(2):459–489.

Matteo, R., 2004, "The Promises and Perils of Agent-based Computational Economics," Laboratorio Riccardo Revelli Centre for Employment Studies working paper; available at http://econwpa.wustl.edu:8089/eps/comp/papers/0401/0401001.pdf.

Mátyás, L., 1999, *Generalized Method of Moments Estimation*, Cambridge, UK: Cambridge University Press.

McFadden, D., 1989, "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica* **57**(5):995–1026.

McFadden, D., and P.A. Ruud, 1994, "Estimation by Simulation," *The Review of Economics and Statistics* **76**(4):591–608.

Myers, R.H., and D.C. Montgomery, 2002, *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, 2nd Ed., New York, NY: John Wiley & Sons.

Pakes, A. and D. Pollard, 1989, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica* **57**(5):1027–1057.

Pearson, K., 1893, "Asymptotic Frequency Curves," *Nature* **48**:615–626.

Poeter, E., and M.C. Hill, 1998, *Documentation of UCODE, a Computer Code for Universal Inverse Modeling*, Water-Resources Investigations Report 98-4080, Denver, CO: U.S. Geological Survey. Available at http://water.usgs.gov/software/code/ground_water/ucode/doc/ucode.pdf.

SAS Institute Inc., 2002, *The SAS System for Windows*, Cary, NC.

Shaw, C.R., and H.D. McKay, 1969, *Juvenile Delinquency and Urban Areas*, Chicago, IL: The University of Chicago Press.

Sica, A., 2004, "Why 'Unobservables' Cannot Save General Theory: A Reply to Mahoney," *Social Forces* **83**(2):491–501.

Stern, S., 1997, "Simulation-based Estimation," *Journal of Economic Literature* **35**:2006–2039, Dec.

Tarantola, A., 2005, *Inverse Problem Theory and Methods for Model Parameter Estimation*, Philadelphia, PA: Society for Industrial and Applied Mathematics.

The Mathworks, Inc., 1999, *Matlab 5.3*, *Student Version*, Natick, MA; available at http://www.mathworks.com/.

Whyte, W.F., 1937, *Street-corner Society*, Chicago, IL: The University of Chicago Press.

Wilson, W.J., 1996, *When Work Disappears*, New York, NY: Alfred A. Knopf.

DISCUSSION

Methods, Toolkits, and Techniques

(Verification and Validation, Thursday, October 13, 2005, 4:00–6:00 p.m.)

Chair and Discussant: *Roberta Davidson, Argonne National Laboratory*

Verification and Validation of Scientific and Economic Models

Michael North: I'd like to introduce Roberta Davidson, who's going to be the chair and discussant for the "Verification and Validation" session.

Roberta Davidson: This is our last session of the day, and we realize we're running a little late. We'll do our best to keep the rest of the talks on schedule. As Mike said, we're going to be talking about verification and validation, which is a hot topic in agent-based modeling. Our first presentation will be given by Ryan Kennedy from the University of Notre Dame.

Ryan Kennedy: I'm going to talk about verification and validation of scientific and economic models. First, I'll go over some of the basic concepts of what we mean by verification and validation, and then I'll move on to our research objectives and methods. We looked at this problem through two different case studies: an agent-based scientific model and an equation-based economic model. Finally, I'll present our conclusions and plans for future work.

[Presentation]

Kennedy: In terms of future work, we'd like to collect and evaluate more statistical data and perform more statistical tests. It would be great to compare our simulation results against some of your world data and then come up with more stringent and formalized verification and validation methods.

Davidson: Thank you. We'd like to open it up for questions now, but first, I have a question for you regarding the validation techniques that you've used for agent-based models. Do you find that's been received well by the users or the sponsors, or are they still questioning? Can you really validate an agent-based model? What kind of reception are you getting?

Kennedy: We're getting a little bit of both there. It's still rather new, so there's not much else for people to see. So far, they have been rather pleased with what we've shown them. We've been working on this for a while, so they know what to expect by now.

Charles Macal: I have a question regarding the validation bias that may be built into the process. Have you constructed any invalidation tests, particularly things tested or designed to find the special cases in which the validation actually breaks down as opposed to possibly being biased toward trying to find all the tests for which there's actually, even subconsciously, a positive comparison?

Kennedy: We haven't tried to do any of that kind of testing, but that would be something that we could definitely look into in the future.

Davidson: Are there any other questions for Ryan? Okay. Let's give Ryan a hand.

How Simple Is Simple Enough? Military Modeling Case Studies

Davidson: Our second speaker for "Verification and Validation" is Matthew Koehler. He comes to us from The MITRE Corporation.

Matthew Koehler: I'm going to go fast in the interest of time so feel free to ask me any questions afterward. Before I start, however, I did a great disservice to Steve [Bankes]. I failed to mention that OldMcData can also handle evolutionary programming and things like that to make a more dynamic experimental design. That's a cool feature that I should have mentioned before, but I waited to surprise you with it now.

Regarding this talk, there's a little bait-and-switch that'll happen with respect to the title. I'm just trying to continue the subterfuge I started in a previous discussion.

[Presentation]

Koehler: With that, I will take questions or simply let the next group come in.

Macal: In your model, you focus on space with respect to some level of resolution. Are there similar issues involved with respect to time, say time stepping versus whatever? Can you comment on that versus your thoughts and how it would be applicable to experience or existing models?

Koehler: Yes, we are absolutely struggling with time. You think something is 1–1000, 2–1000, so how hard can it be? Well, it turns out to be exceedingly difficult sometimes. In MOE, we're using 10-second time steps. Then, fire rates and sustainable fire rates and things like that get to be an incredible hassle. So, yes, time is just as important, if actually not more so, than these other abstractions. When you start saying that this agent is fire team, you gloss over it and hope no one notices. I don't know if I answered your question, except that I agree.

Macal: Well, yes, you answered the question. I wondered though how you were coming along with…

Koehler: How are we overcoming it?

Macal: Exactly. Do you want to put a positive spin on the answer, or just leave it at, well, it's impossible?

Koehler: Well, it's a chicken-and-the-egg thing. To actually start getting a handle on how these things need to be abstracted, you have to think about the size of the play box. If you're unlucky enough to be on a torus, how long will the battle take or how long will the simulation be in reality? From there, if it's going to be a 12-hour event, how many time steps is that going to be? You can ask, “What is a time step in the simulation, and how does that map into the reality of the rest of the system?”

It’s not impossible, but sponsors can be a little tricky. They like to throw out some ideas, have you go away, play around with them, come back, show them this wonderful work you’ve done, and then critique it. You then go back and tweak it and so forth. If they start changing, saying that 12 hours is not long enough, do 48, you have to go through and rescale every single number in your scenario. Thirty-five kilometers isn’t quite the right size because we didn’t get this other feature in there that we think is neat. So let’s make it 45.

So it’s not impossible, but if you can start out knowing those three things — the size, the duration, and how many time steps will be used, you can start getting a handle on it. If you have to do it piecemeal, however, it would be tricky. It’s hard to get sponsors to understand that up front and to force them to think that carefully about the scenario before they’ve seen it.

Davidson: Thanks. Let’s give Matthew a hand.

X-MAS: Supporting the Tedious Work of Validation in Agent-based Simulation

Davidson: Our third speaker is Yutaka Inti Leon Suematsu from Kyoto University and the ATR Network Informatics Laboratories.

Yutaka Inti Leon Suematsu: Good afternoon. I’m going to present this Christmas, which is a tool we are developing in Kyoto University. We plan to use this tool to support validation of agent-based models.

The structure of my presentation is as follows. First, I will give a brief introduction about the motivation of this project. Second, I will describe the cross-element validation, which is a validation process we proposed in our previous work. Third, I will give a description of the Christmas toolkit, along with a very simple case study. Finally, I will present our conclusions and future work.

[Presentation]

Davidson: Thank you. Are there any questions for Yutaka? I have a question that is based on reading your paper. It sounds as if Christmas could be applied to almost any agent-based model. Do you agree?

Suematsu: Yes. We are providing some generic libraries. It’s not such a framework that you have graphics and you can create some models easily. In this case, you have to implement everything so it’s a general-purpose library. You can use it in many kinds of multi-agent systems.

Davidson: Are there any other questions? I want to thank Yutaka for coming all the way from Japan for this.

Agent-based Models as Quantitative Sociological Methodology: Calibrating Simulation Models to Data and Finding Confidence Intervals for Model Parameters

Davidson: Our final speaker today is Steven Wilcox. He comes from Northrop Grumman.

Steven Wilcox: Good afternoon, almost evening. My talk is about traditional quantitative methodology in sociology and other social sciences, which you all are taking aim at with agent-based modeling, but the arrows have not quite stuck. There are many more traditional sociologists than agent-based models or modelers. The question is, "What needs to be done?"

First, I'll complain about the situation, then I'll present a possible way to fix the problem. I'll talk about a case study of urban crime in Chicago. I read a paper in sociology, and it gave this long theoretical, even philosophical, discourse on the previous literature. It had a long discussion about what was the right regression model. I was not sure whether they actually convinced me that the model was going to test their theories. But the point is not to complain; the point is to change the situation.

[Presentation]

Davidson: Thank you. Are there any questions for Steve?

Kostas Alexandridis: Given the stochastic character of the agent-based models that you described, traditional parametric estimation sometimes does not look at the probability of an event occurring. How would you deal with an issue where there might be some local convergence that the model picks up but is not a global maximum or the global optimum? How will the model predict, given the data? In agent models, a lot of agent models, how are the patterns good, given the process?

Wilcox: It sounds like you're talking about nonlinearity, and the issue of, well, I assumed it was linear, if you noticed, but is it really linear? And you have the issue of hidden nonlinearity, say, some pocket in the middle of the plane.

Part of the analysis process when you do statistics and regression is to take a look at your residual plots, plot the residual versus predicted value, and so forth and try to spot those nonlinearities and, of course, do something about them. This affects your answer. There's the problem of hidden nonlinearities, which I admit to, but I think the perfect is somewhat an enemy of the better here, and this is better than prose-based quantitative methodology. So I'll go with the better, I think.

Davidson: Thank you. Any more questions for Steve? All right. Let's give Steve a hand.

North: I'd like to thank Robbie Davidson for running the "Verification and Validation" session. I'd also like to thank all the session chairs for today and all of the speakers as well. Most important, though, I'd like to thank the audience for seeing through Day 1, which is nearing the 10-hour mark. The rest of the days in the conference will be a little easier on you. Tomorrow we'll be starting at 8:30 with a few awards, and then at 8:45 we have an invited distinguished speaker, Josh Epstein, who will be talking on "Generative Social Science Applications:

Applications of Agent-based Modeling." That's going to be an exciting presentation, so please try to arrive on time. With that, I think we're set for the day. See you tomorrow.

Friday, October 14, 2005

Social Simulation Applications

Invited speaker

Joshua Epstein

GENERATIVE SOCIAL SCIENCE: APPLICATIONS OF AGENT-BASED MODELING

J. EPSTEIN,* Senior Fellow
Economic Studies Program, The Brookings Institution, Washington, DC

ABSTRACT

Agent-based modeling and simulation have become central to complex adaptive systems research. These tools, initially used for artificial life simulation have been adapted to research in economics and now more extensively throughout the social sciences. Formerly, in the domain of computer science, agent-based models were also used to address a wide range of complex problems in government and industry. This paper discusses the notion of a generative explanation in the social sciences and the central role of agent-based computational modeling in generative social science. The presentation covers diverse applications drawn from such fields as epidemiology, civil violence, and archaeology.

Editors' Note: The full paper was not received in time for publication. The abstract is included to provide a frame of reference for the discussion that follows this session. Dr. Epstein's latest book, *Generative Social Science: Studies in Agent-based Computational Modeling* (Princeton, NJ: Princeton University Press) will be published in late 2006.

* *Corresponding author address*: Joshua Epstein, Center for Social and Economic Dynamics, The Brookings Institution, 1775 Massachusetts Ave., Washington, DC 20036; e-mail: jepstein@brookings.edu.

DISCUSSION

Social Simulation Applications

(Invited Speaker, Friday, October 14, 2005, 8:45–9:45 a.m.)

Chair and Discussant: *David Sallach, Argonne National Laboratory and The University of Chicago*

Generative Social Science: Applications of Agent-Based Modeling

David Sallach: Josh has been doing recent work in generative modeling and some not-so-recent work in generative modeling, which, as you know, is the theme of this conference and a major strategy for addressing complex processes. Josh is fortunate to have worked at, and been regarded as an exemplar at, Brookings.

We look to the early work of Schelling and Axelrod, but together with his coauthor, Rob Axtell, the Sugarscape model has played that kind of role. Since then, he's done a number of interesting applications, including in civil violence and some epidemiology work that he's going to report on today. One of the things that I especially appreciate is that Josh is concerned about the epistemological issues — what underlies the new methodological and scientific kinds of initiatives that agent simulation makes possible. We welcome the fact that those kinds of issues will be woven into the talk as well. We look forward to your talk.

Joshua Epstein: I'm very honored that the conference is called "Generative Social Processes." I think that's great. My talk is called "Generative Social Science: Applications of Agent-based Modeling." First, I will talk about agent models, just a little, because this audience doesn't need to hear much about that. Second, I'll talk about this idea of generative explanations, since I think that's what's at stake in a lot of the work that we're doing. Finally, I'll give some applications to epidemiology. I've been doing a lot of work in infectious diseases for the National Institutes of Health (NIH) and others, and I wanted to quickly show you a whole range of epidemic applications.

[Presentation]

Sallach: Thank you very much. Very, very interesting, very good work. We have time for discussion. I would like to start with a couple questions. I'd like to tie in your early comments about generative social science and explanations and so forth with other problems that go forward. A lot of smallpox and other such diseases do not have much social intermediation in the process. You began to refer to it when you said people might not be willing to go for shots or warehouses of vaccine might be…, well, let's not even talk about it. If you look at other kinds of diseases, like AIDS, cultural factors become extremely important, as do lifestyle factors and so forth. One of my questions has to do with the advances in substantive social theory, where I would assume that advances mean finding mathematical formulations for them. Also, I'd like to investigate what kinds of advances are necessary and what social agent simulation can do: what is its contribution to furthering substantive mathematical social theory?

Epstein: Okay.

Sallach: The second related question has to do with the construction of reference classes. If you think of problems like civil violence, there are so many forms and types of civil violence occurring in so many historical and economic circumstances and so forth. In some of your written work on generative social theory, you say that you're generating explananda, right?

Epstein: Right.

Sallach: You've shown us good examples of how you take one benchmark case and generalize from that, but you might also be framed in terms of how do you define the class of explananda that you want your effort to be relevant to. I also wonder whether there is a role for meta-theory, meta-rules that are themselves generated, generative. Can we generate some of the complexity that we would then like to be in that reference class?

Epstein: In the first one, I think the behavioral side of all the epidemic work is really all but absent, and I've made a big racket about that at NIH. We have a working group that tries to put behavior back into these things.

For example, in 1994 in one of the big provinces in India, 20 million people fled on the basis of a rumor that there was pneumonic plague. Actually, there was never even a confirmed case. So the idea that something is going to happen, such as pandemic flu, and that people are just going to go about their business is ludicrous.

The entire behavioral aspect of this is understudied, and there are epidemics of fear and bad information and silly prescriptions and the rest of it. Agent models are going to be very important in thinking about the behavioral dimension of epidemic behavior, so that's an important frontier.

Regarding the second question, it's complicated. When you build the civil violence model, I'm always careful to say that it is an idealization that tries to get into some core dynamics. I wouldn't allege to have outlined clearly the exact set of phenomena covered or definitely not covered. I'm not quite sure how to answer, except to say that the more data we have, the more clearly we will know the range of applicability of any of these models.

Sallach: Are there other questions?

Pam Sydelko: Emergency management is big right now, especially for any bioterrorism event. There is pressure, obviously, to use models, whatever they are. I agree that agent models are probably one of our best tools for use in emergency management prediction. How would you feel about giving a model to the local population for exploration, letting them play around with the models in their exercises and training scenarios?

Epstein: Yes.

Sydelko: I'm working on a project that is trying to do this, and there is this feeling that the local population doesn't have enough understanding to do these things, which I think is absolutely false, because they're actually the most knowledgeable about their local situations.

Epstein: Yes, you're probably right.

Sydelko: Can you comment on pushing these kinds of models as exploratory models into the local areas and letting them play around with and look at what-ifs?

Epstein: I think that's a great idea. Environments like Repast and others are perfect for that. The kind of toy, easy prototype, interactive, immersive experience, exploratorium aspect of it would be great. You could put these things in high schools, colleges, and the local town council.

Sydelko: Yes, but the pressure is to actually do it more in emergency management way, and actually put them into models ….

Epstein: Yes, but it's very top-down. A colleague of mine, Roz Lasker in New York, did a great survey on smallpox and other things. Her attitude was, "Let's find out what the communities would like to do and see if we can help them do that." A lot of people would like to isolate themselves and not go to a vaccination site. What's involved with that? We need to make sure that they will have enough to eat. Will they have all the details of how to implement a strategy, a bottom-up strategy, of epidemic containment?

But I think it's a great idea. I'd like to talk to you more about it. I think the technology is moving in a direction that would permit that. People could just log onto some site and play these games.

Reginald Tucker-Steely: I'm a doctoral student from the Harvard School of Public Health. It's great that these models work for epidemics, but how do we translate these types of models to chronic diseases where we see different vectors for diseases, such as the social determinants of health and poverty and things like that? Do you see that to be an easy transition? If so, what are some of the things that would be important in building those types of models?

Epstein: Great question. To be completely honest, I haven't done anything in that area myself; most of my work has been on infectious diseases. But I think that's a really great question. So I'm not positive. I don't actually know of modeling work in that area, but I'm happy to help poke around and see if we can come up with anything.

Li Yin: Will the scale of the model play a very important role in a simulation? As you mentioned in your presentation, you will have to build a large-scale model and a global model to simulate the real situation. But there will be complexity between the model and the accuracy of your simulation, so how do you balance them?

Epstein: That's an excellent question, and it's very much on our minds. Scaling these things up, of course, is an empirical question. Is there some curve? Does it all flatten out? Are complexity and scale beyond some point just superfluous? But we certainly don't know that, and our aim is to build these very big — hundreds of millions — agent models. I think that's a great question, and all I can say is that we're working like crazy to do it.

The cluster we're using now is, again, one of the trade-offs is, for example, in AnyLogic, there are all these very slick drag-and-drop prototyping tools — state charts. There's an active agent object that you just drag into a window, and then you can drag a variable down and add functions. There's huge overhead that comes with that ease of programming and so on.

So we've tried to strip out the native Java and ship it to the bigger machines. We're running now on a 64-node cluster, two CPUs per node, 5 gigs, Opteron nodes. There's overflow to the NCSA [National Center for Supercomputing Applications] grid that's been arranged and tested, so we can go. We can build very, very big models, and we'll see at what point the increasing scale is not returning any more insight. I don't think we're at that frontier yet. We just don't know.

When people say that we could do the same thing on the back of an envelope with differential equations, it's an empirical claim. We'll find out. If that's true, okay. We wasted a lot of time scaling up our agent models. My suspicion is that it's not true and that you'll get really novel things happening at very large scale that might not happen at smaller scales. We'll see. That's where we are.

Carl Johnston: I'm a doctoral student at George Mason with the Interdisciplinary Center for Economics Science under Vernon Smith. I was thinking that you're going broad in scale in one dimension, but there's also another dimension that I try to get people to think about, which is the supply side dimension. A lot of the problems that happen with intervention are, in fact, economic problems, for example, the vaccine problem and ownership of those vaccine-producing companies and so forth. It seems to me that there ought to be a way, in addition to everything else, to be able to model, for example, different ownership structures for vaccine-producing companies and how they respond to these needs. Those facets are just as significant in formulating a response as most biology. So I'm wondering if anybody has thought about modeling the supply side in addition to modeling more agents.

Epstein: Yes. It's an entirely different concept. Two people, Jenny Lanjouw, who was an economist at the University of California, Berkeley, and Michael Kremer, who is at Harvard and also Brookings, have looked at the vaccine, the pharmaceutical industry, and all of that. The main thing is that you think of this as a public good. We're talking about millions of people dying, so you'd like it to be a public good, not a commodity or a private sector thing.

The question is the government's role. Roche does this much better than the government, so you don't want to replicate their production line and their labs and their distribution network in the government. How do you give incentive to produce the adequate level of vaccine? One thought is to subsidize them to produce the appropriate level. I think it's one of these cases where markets are probably not quite up to it because once the thing happens, you can't rev up. Once the thing's gone, it's an irreversible mess at that point.

Xinrong Lei: My name is Xinrong Lei, from UIUC [University of Illinois at Chicago], a doctoral student. I benefited very much from your talk. If I want to implement this model and also consider a different structure of a relatively small scale and large scale because, as a small group, the structure might follow the small world phenomena, but for a larger one, the structure might be different. If I want to implement this model, would you have some suggestions?

And also, consider transportation. More than one type of network could do transportation. For example, the railroad connection may be different from the airplane companies' connections. If I want to implement more than one type of this kind of network connection in the model, what should I do? Do you have any suggestions?

Epstein: Well, you're certainly right that other modes would matter a lot, especially intercontinentally. It's probably mostly air that we care about, but inside the continental US, there's ground transportation, air transportation, trains, cars, and so forth.

There's an IBM model called the STEM [Spatiotemporal Epidemiological Modeler] that you should look at. I think you can just navigate to it from ibm.com [Editors' note: see http://www.alphaworks.ibm.com/tech/stem/download]. They had every census track and ZIP code, and they have all the ground transportation worked out reasonably for that.

For some of these bugs, there's an argument as to whether you should even model the air because it will spread quickly via the ground. If you just take out air, it might still spread all over the country in one week, and with air, it's one week and four days or something. So is it worth the computational effort to include all of that? I know it's easy to get lost.

You should also look at an article in *Scientific American* in March by Eubank and company about smallpox in Portland, which is their account of a model that was developed at Los Alamos, called EpiSims [Epidemiology System], which is an extremely detailed model. It's got about 1.6 million agents, but it's got every car, every stop sign, every corner, every bus, everything. [Editors' note: See *Scientific American*, "'EpiSims' Unleashes Virtual Plagues in Real Cities to See How Social Networks Spread Disease. That Knowledge Might Help Stop Epidemics," Chris L. Barret, Stephen G. Eubank, and James P. Smith, March 2005.]

If you want an example of a maximum resolution treatment, that would be a good place to start. I would then look at STEM for the nationwide version. Again, my concern about all those models is that you can't need every single puppy and fire hydrant in the model. It's just impossible. I think the temptation in the high-end computing side is to throw in way more than is needed. It's counter-indicated because it's going to clutter the thing so much you're not going to know what the central mechanisms are. I'm always thinking about what is the minimal model needed to give you some insight. The maximal model's easy. You throw everything in, which is a temptation to be avoided.

Joanna Bryson-Beth: Joanna Bryson-Beth. Your content is incredible, and I really enjoyed that part of the talk. I always wonder if I'm in the right sector when I hear something like this, but I'm going to ask a brutally technical question. You're talking about models you're literally thinking about 5 or 6 billion agents. You heard our talks yesterday where we were trying to get out things that model 20 or 30 agents. First of all, is there some kind of standard that government is using?

Epstein: Oh, no.

Bryson-Beth: Do people have access to these kinds of systems? I assume that you need not only the software, but also the platform, like the old-school parallel computing bank.

Epstein: Yes, they're big — right. Well, go ahead. I didn't mean to interrupt.

Bryson-Beth: It sounds like you're starting to answer the question. How do we get these machines? Should we be working on them? Are we directing our effort in the right places, or is there something else we should be looking at? Is the next generation out there? I've never done anything with 6 million agents.

Epstein: Yes, but I shouldn't attempt to answer that completely because I'm not positive what everybody's working on. All this stuff is happening at universities and research places. None of this is inside the government. It's funded by the government, but it's a research consortium of universities with a couple people from IBM and NCSA and other places.

Unidentified Speaker: You mentioned the name of one of the modeling systems. Are you doing all your modeling on one system with one cluster?

Epstein: No. Research Triangle Institute got the award to be the informatics core of this activity. Part of their funding was to put together a reasonably big cluster. Everybody in the project can use that cluster to run runs. We do big production runs on the cluster and so on.

Unidentified Speaker: Do you mean it's one project and then …

Epstein: No. We have routine access to the cluster, and it's got dynamic load-sharing and all sorts of other fancy things that I don't know about. To me, it's just transparent, you know. You log on, you write your code, you submit the job, and it gives you the answer.

The trick is, can you ultimately do all that in Java? I think that's the issue, whether you can do it in a high level using AnyLogic. They have this agent object that comes with tons and tons of functionality. We're saying, "Strip all that functionality out. We need a light agent object." If I'm writing, $y = x^2$, I don't want to call up the whole function library. Just let me hand code that. If I want to declare something as a variable that has all this baggage, state charts, all this fancy stuff, take that all out. It's great for prototyping, but for big modeling, it's not great.

I think there are many levels going on. We like to build these prototype models in systems Java-based environments. I guess Ascape is becoming a dying dialect, but Repast, AnyLogic, or a variety of others. For toy prototyping modeling, I think that's the way to go. But then when you say, "Okay, we want to scale that up to a zillion agents," do you have to port it to C++, or can we find a way to strip out, to stay with these environments, but modify them in ways that permit large-scale computing?

I'm trying to remember what the stats were. In our initial simple agent in AnyLogic, it was 200 kilobytes or something — a giant thing. All it had was five doubles for a name or something. It was this vast object. We need to have that down to 300 bytes or we just can't use it.

So, I don't know exactly. I think it's great to have a lot of things going on. I think it's neat to have toy prototyping languages that you could deploy at the local levels, so the Mayor of Cincinnati could run it. But I think at the moment, it's is tricky to scale up these right now.

I don't know if that answers the question, but it's not the government. It's just about buy, memory, and clock cycles.

Craig Stephan: Hopefully, one quick question. Craig Stephan, Ford Motor Company. That was a very intriguing talk. I certainly enjoyed it. I have a hypothetical question, more from the public policy standpoint than from the modeling standpoint. In the toy models at least, you basically assumed that the vaccine was 100% effective if you got it, and zero if you didn't. Are there intermediate things that one could do? Say, you might be able to reduce the transmissivity of the disease by 20% by some public policy measure, like handing out rubber gloves or something. How would that affect the spread of the epidemic?

Epstein: Yes, that's an excellent question. There's a whole bunch of work on drugs that don't affect me, so I get the bug. It doesn't affect my probability of dying, but it affects my capacity to transmit to others. In the area of AIDS, especially, there's a lot of work on a vaccine that would have that effect, and it makes a huge difference in the modeling. Changing the transmissibility of the bug is a huge deal.

Isolation is another huge deal. In the smallpox thing, if we isolate families of the sick, that takes a huge bite out of the epidemic. So there are self-protective measures; there are attenuated vaccines that can affect transmissibility, but not the pathogenicity of it. All those are very nice. There's literature on this. They had exactly that idea and talked a lot about the possibility of an AIDS vaccine that wouldn't save the person from HIV, but it would prevent transmissibility.

A lot of the antiviral drugs for flu are so-called neuraminidase inhibitors. Neuraminidase inhibits an enzyme that permits the replicated viruses to leave their host. It doesn't make me better, but it prevents me from transmitting it, so it's a real avenue that people look at.

Sallach: We'll take one more question.

Lars-Erik Cederman: Lars-Erik Cederman, ETH Zurich. This was a fascinating talk, but I want to go back to the beginning. You made a meta-theoretical point about generative social science, which I think is a very important one; I agree with you 100%. This is the way to, as I say, bring across the insights from computational modeling because computational modeling is so much more than just a technique.

The real challenge, of course, is to go beyond generative sufficiency, and I think that your talk went pretty far in convincing us that you can actually do that. I wanted to see if you could verbalize the strategies and how you would do that. Obviously, we could shoot for more than one macro pattern, to say, instead of having just one target to grow one particular regularity or pattern. We could try to do many, and thus narrow down the set of possible models that we are looking for. You could also look at the micro-level mechanisms and impose constraints on those, empirical, theoretical, or whatever. How would you sum up the best cocktail of heuristics and strategies to go beyond generative sufficiency?

Epstein: Well, this is a very central and good question, and I talk about it at some length in a paper called, "Remarks on the Foundations of Agent-based Social Science" that will be in the Princeton book, but is a Santa Fe working paper and a Brookings working paper, called "Remarks on the Foundations," and it will be a chapter in this computational econ book.

But I think of it a little bit differently, Lars. Let's imagine there's something you're trying to explain: the firm size distribution or power law distribution of something. You have come up with some micro-specification for the agents that does suffice to generate that in some reasonable sense. How do you know that that's the right micro-specification? That is the question that normally arises. You'd like to generate many competing micro-specifications, all of which produce that distribution, and then, as in any other science, you have to figure out which of these micro-specifications is the most plausible empirically. That may mean that you design laboratory experiments with humans to figure out is this one involves a cognitive load that's just not plausible, or if this one involves memory that isn't plausible, or if this one involves some other predisposition that I don't buy. I think that's really the way you'd have to do it.

It's hard to get any micro-mechanisms that generate the things, but let's assume this embarrassment of riches, where there are competing ones. It seems to me you're in a position of any other science, where you're trying to design laboratory or other empirical tests that would permit you to adjudicate between those competitors. I don't know how else to answer it.

I think we should have an argument at some point about validation. I hate that word. I think it should be banished. I'm not sure you'd ever be in a position to say it's generative necessity. I mean, you have generative sufficiency and there are other considerations. The plausibility at the micro level and parsimony and other things that go into the selection of any theory in any science …

Sallach: Well, we've obviously over-run our time, and pleasantly so. We want to thank you very much for your presentation and discussion.

Epstein: I'm honored. Thank you.

Public Policy

SIMULATING WATER USAGE DURING UNCERTAIN TIMES IN THE SOUTHWESTERN UNITED STATES: AN ABM OF STRATEGIES AND POPULATION LEVEL ACTIONS

W.A. GRIFFIN* and S.K. SCHMIDT, Center for Social Dynamics and Complexity,
Arizona State University, Tempe, AZ

ABSTRACT

Many of the nation's, and indeed the world's, most rapidly growing urban areas are in arid environments and face a future of greater water uncertainty. Arid cities therefore will benefit from a clearer articulation of the effects of climate change on urban water demand and supply and on community response to growing uncertainty. The Decision Center for a Desert City at Arizona State University is one of several new centers funded by the National Science Foundation to investigate human decision-making under climatic uncertainty. To address the uncertainty faced by water consumers, policy makers, and scientists, we are developing an agent-based model of water use (DesertWater) that integrates census data with municipality-supplied data on water use and implements plausible agent decision rules about water consumption, conservation, and media influence. We present our current version of the model and discuss our rational for the embedded decision rules.

Keywords: Water, ABM, agent-based model, uncertainty, policy, modeling

INTRODUCTION

Like an oasis, the Phoenix area — a complex of metropolitan cities — has emerged out of a desolate desert to become the fifth largest urban area in the United States. Having grown from a modest 300,000 in 1950 to 3.2 million in 2005, the population is expected to exceed 6 million by 2025 (Jacobs and Holway 2004). Not surprisingly, this influx of people is a continuing catalyst for new construction; residential areas, educational facilities, hospitals, retail centers, and other businesses are being developed to satisfy the evolving needs of the population. While Arizona's economy reaps the benefit of this expansion, it is questionable whether Arizona's ecology can sustain this rapid development.

The Phoenix transformation from saguaros and sand to concrete and cars is deceiving. Although metropolitan in appearance, Phoenix *is* a desert: it receives only 180 mm of annual precipitation and has typical summer temperatures of 115°F (Baker et al. 2004). As a result, the threat of a water shortage is omnipresent among today's residents of Phoenix, as it was with the earliest Sonoran dwellers — from the prehistoric Hohokam, who constructed 1,000 miles of irrigation canals, to the Euro-American farmers, who converted the dryland river valley into an agricultural paradise at the end of the nineteenth century (Gober 2005).

* *Corresponding author address*: William A. Griffin, Center for Social Dynamics and Complexity, YHE 2502, Arizona State University, Tempe, AZ 85287; e-mail: william.griffin@asu.edu.

Since its inception as a city, Phoenix, like most urban areas, has explored various options for water acquisition and management. These options collapse into three basic strategies: (1) seek more water, (2) conserve the available supply, or (3) implement some policy that involves strategies 1 and 2. Pros and cons exist for each strategy. Increasing supply is costly but will ensure a reservoir in drought conditions. Conservation works in theory, but the necessary amounts and strategies for implementation are not known. And an appropriate ratio of the two strategies may exist, but the proposed cost-benefit returns are purely suggestive and hypothetical. Given the lack of empirical data on all three strategies, debating the optimal strategy remains a scientific, policy, and political sport. There is simply not enough information about current social behavior, future climatic and hydrological change, or population growth and shifts to inform researchers about the best method to ensure the future water supply at a reasonable price.

Increasingly, it is recognized that even the best science will not significantly reduce uncertainty about global climate warming and the climate cycles that cause droughts, floods, hurricanes, and tornados. Society must learn to make better decisions in the face of uncertainty. The Decision Center for a Desert City (DCDC) at Arizona State University is one of several new centers funded by the National Science Foundation (NSF) to investigate human decision-making under climatic uncertainty. In 2004, the DCDC was founded to focus on water management decisions in the urbanizing desert of central Arizona. Under the charge of an NSF grant, the DCDC coordinates a program of interdisciplinary research and community outreach to improve water-management decisions in central Arizona. To that end, the DCDC studies the behavioral processes of individuals, examines how water managers make decisions, and then applies sophisticated models of decision science to water-allocation problems.

Decision-making, Incomplete Information, and Agent-based Modeling

The DCDC's central mission is to enhance and improve water management decision-making. Agent-based modeling is at the core of the decision-making tools being used at the center. In fact, a water-use agent-based model (ABM), named DesertWater, was immediately developed within this large multidisciplinary center. Why? Because an ABM can simulate processes in which decisions are decentralized and made by individuals and groups with different perceptions of uncertainty and attitudes toward risk. Our models quantify behavioral processes and then examine the reciprocal relationship between individual micro-social processes generated by explicit decision rules and group ontologies (Griffin 2003; Griffin et al. 2004; Schmidt et al. 2005). Groups, acting on the aggregate effect of individual rules, emerge as discrete entities that influence resource use and policy implementation and that, most noteworthy, by their actions, iteratively modify subsequent agent-level decisions. This reciprocal relationship between agent-level decisions and collective use of resources has been successfully modeled for other commodities. For example, North and colleagues (North 2001; Macal and North 2002) have examined the dynamics of electricity and natural gas consumption in competitive resource markets.

Modeling Water-use Decisions in the Southwest

Herein we present an overview of the rationale and algorithmic structure of the new, Repast-based, water-use ABM being developed at the DCDC. In its current form, aside from it having the more traditional aspects of any good ABM (e.g., a landscape populated with families composed of individuals [agents], with each agent having separate water-use preferences), we have developed several unique features within each agent of DesertWater that we hope provide a realistic representation and simulation of intra- and inter-familial water use.

Unique Model Features

First, aside from empirically based sociodemographic attributes (e.g., race, age, sex, income, and education obtained from Census data), agents are assigned values representing three relevant decision-making characteristics: (1) receptivity, (2) sensitivity, and (3) hierarchy. *Receptivity* refers to the ability to acquire or perceive information about either the relevant characteristic in the current scenario (e.g., price of water, media information) or the amount of water use by others. *Sensitivity* refers to the amount of change in water use that occurs in response to information obtained from other agents (via receptivity). *Hierarchy* reflects the intra-family influence that an agent has on other familial agents. For example, parents tend to have a higher rank than adolescents within the family (but not always), and if a parent decides to reduce water use, this change in behavior modifies the behavior of other family members.

Second, each agent is assigned a vision (i.e., sphere of perception of others) that extends from near neighbors (about 80%) to other agents far beyond its immediate geospatial location. This ability to perceive and retrieve information about another agent's water use is one of the factors that determines if, and by how much, personal water use is modified.

Third, the choice of which agents get monitored by other agents is based on tag matching (i.e., degree of homophily). Tags represent sociodemographic information (e.g., education, sex) that agents use to determine whether or not to attend to, and receive information from, other agents (Holland 1995).

Current Implementation

At this initial stage of model development, we are cross-referencing sociodemographic data with municipality-supplied water-use data. The data range from single-family households to office buildings; we have between 300,000 and 400,000 monthly water-use records for each year from 1995 to 2003. At this juncture of development, we are focusing on single-family households because the available data (e.g., usage) on this group are the most detailed. This provides an empirical basis for rule construction and expected consumer variation in response to price changes and media campaigns. Our objective is to construct agent interaction and information exchange rules that modify water use as a function of (1) water price fluctuations, (2) media information, and (3) perceptions of water shortage. Agents receive information about water through contacts with other agents, general perturbations (e.g., changes in water costs as indicated on water bills), or simulated media campaigns that encourage water conservation.

Consumer Information, Media Campaigns, and Population Penetration

Numerous U.S. states and several nations have instigated water conservation methods (e.g., see www.saws.org/conservation/, www.ec.gc.ca/water/, www.watercare.net/). Strategies to institute these measures generally fall into two categories: (1) the rough-and-ready Draconian (e.g., turn your water off or else) or (2) the Platonic, which emphasizes the cooperative tendencies of an informed public when adequate and truthful information is provided (Gilg and Barr 2005). Our water-use model is built on the latter style, along with the assumption that an informed public, when given a rationale with justification, will reduce consumption if members perceive that the problem is severe and observe that other people are also conserving water use. In effect, our model is built on this two-tier system: (1) media exposure and (2) near neighbor (agent) behavior. In its current form, the model integrates these two by using a differential media campaign (e.g., those least responsive are targeted more [e.g., Gilg and Barr 2005]), and each agent observes and responds inversely to the adjacency of other agents. The specific aggregate (i.e., population-level) response behavior is determined by one of two distributions: diminishing returns (i.e., each subsequent exposure unit of media has proportionally less impact) or a sigmoidal distribution with thick tails. The latter curve is based on the notion that consumer response will follow a contagion model; specifically, that some people are and will remain immune, and that among the susceptible others, the rates at which the new conservation behaviors move through the population will follow well-established epidemic trajectories (i.e., initially slow entry, then rapid explosion until the individuals that will eventually modify their behavior actually do).

Model Components and Overview

Census data are used to populate the households. Age determines initial water use for each family member and is calculated as a percentage of the initial water use seed. The first member of the family is an adult, and all tag characteristics (sex, race, education, and income) are derived from estimates of the likelihood as generated from the census data. Member *receptivity* (normal [truncated] distribution with a mean of 0.5, range of 1.0, standard deviation of 0.3) and *sensitivity* (normal [truncated] distribution with a mean of 2, range of 4, standard deviation of 2) are assigned randomly to each household member. The distribution (i.e., likelihood of being present) of tag characteristics for family members is also derived from census data. When the number of persons in the family is greater than three, there is a 15% chance that a grandparent is present. After populating the landscape and generating initial water-use data, the model generates weekly consumption estimates based on the assumed influence of the media and price. The process of generating these estimates is described next.

Observation and Adjustment by Comparison

Each week, 10% of the agents are randomly selected. These individuals then compare their water use to that of a pool of other similar agents, and from this comparison, they either reduce or increase their own water use accordingly. More specifically, after randomly selecting a household member (of the 10%), 200 random individuals who match on at least two tags with the selected individual are pooled such that 80% of them are within the same census tract, and an additional 15% are drawn from another, noncontiguous, tract. The final 5% are randomly selected from the remaining population. From this pool of similar individuals, 8–13% are selected, and their collective water use is averaged. From this value, an *influence score* is generated as the product of the percentage difference between the individual's water use and the pool's average water use and the *receptivity* of the individual (i.e., influence score = percentage difference × receptivity of the member). This influence score is then multiplied by the individual's *sensitivity* score to produce a new water use value. If the member who adjusts is either an adolescent or an adult, a percentage (30–75%) of this modification is distributed to all members of the family. A higher percentage is applied if the adult being modified is the head of the household.

Media — An Example

Although the emphasis has traditionally been on price manipulation to modify water consumption, it is probably more economically prudent to instill long-term behavior changes by using education. Within this perspective, we think of education as being two pronged. First, there is the formalized method of teaching conservation methods to young children. Second, there is the advocacy of reducing water use across the age range via satiation. We focus on this latter aspect of education. As currently implemented, we can inundate a selected population (i.e., chosen by receptivity) by using a myriad of media outlets, including television, radio, billboards, print, and mail. All outlets have equal weighting (i.e., influence) in the current model; the affect parameter of each outlet can be easily modified. We expect to implement a differential weighting scheme as we acquire either a theoretical rationale or empirical evidence.

OUTPUT

For data analyses and transfer, the output is in a comma-delimited file containing tick count (i.e., week), STFID (state federal ID [plot location]), (*x*, *y*) coordinates of the agent location, present water use of the agent, and present price of water for each agent every 52 ticks. In addition, we have graphic output of the (1) agents displayed in the (*x*, *y*) coordinate system, with color coding according to household water use (clicking on the agent provides household composition information including the current states of the individuals); (2) total water use of the whole population at any given point in time; (3) average water use according to age (Figure 1), and usage histogram (Figure 2); and (4) average water use according to water provider area (Figure 3). Figures 1, 2, and 3 provide visual feedback to the user at each iteration, and in response to any manipulation (i.e., change in media exposure) during a run.

FUTURE DIRECTIONS

In its current form, DesertWater provides some plausible scenarios for modifying Phoenix water use via a media conservation campaign. There are, however, several components that will need to be considered in future evolutions of the model. To remain consistent with the three arenas (i.e., science, policy, and politics) involved in determining the best strategy for water management, we approach our concerns and future intentions as specific to each.

Science

From a scientific perspective, the immediate concerns for improving the model focus on incorporating factors that will improve the ecological validity of DesertWater. For example, Phoenix is a major producer of citrus. Acres of orange groves and grapefruit trees are housed within the metropolitan and surrounding area. As a result, approximately 58% of all water use in

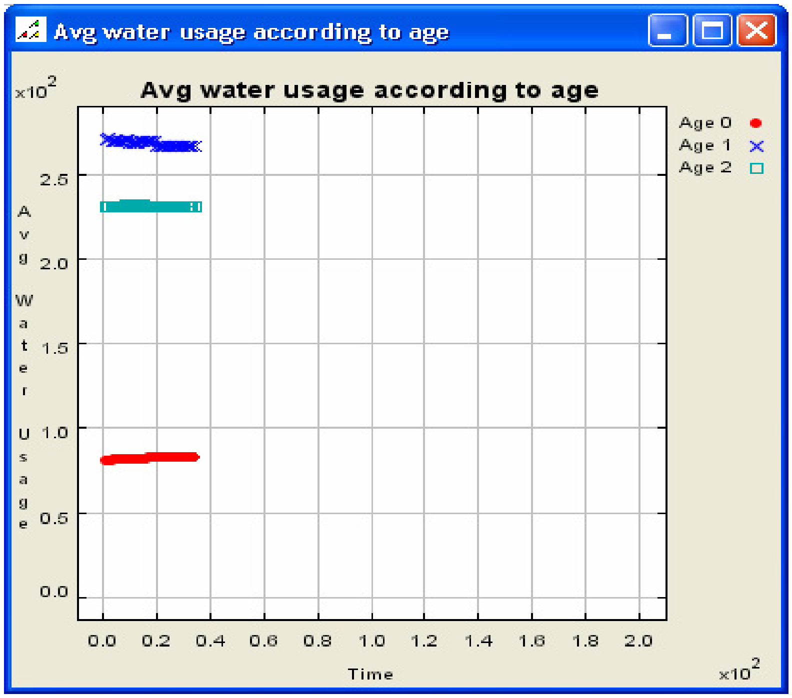

FIGURE 1 Water usage by age

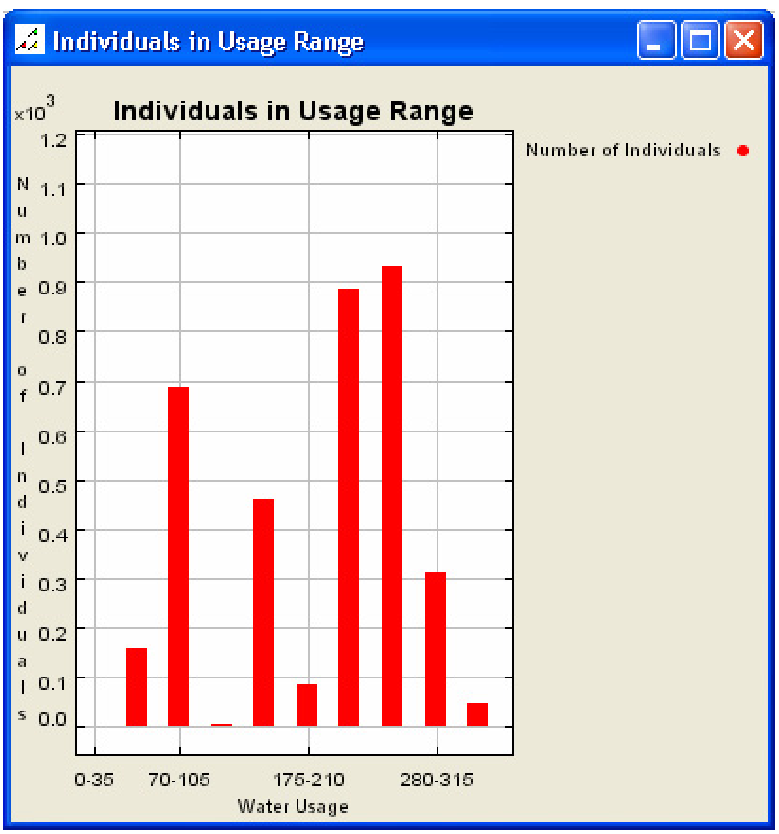

FIGURE 2 Usage histogram

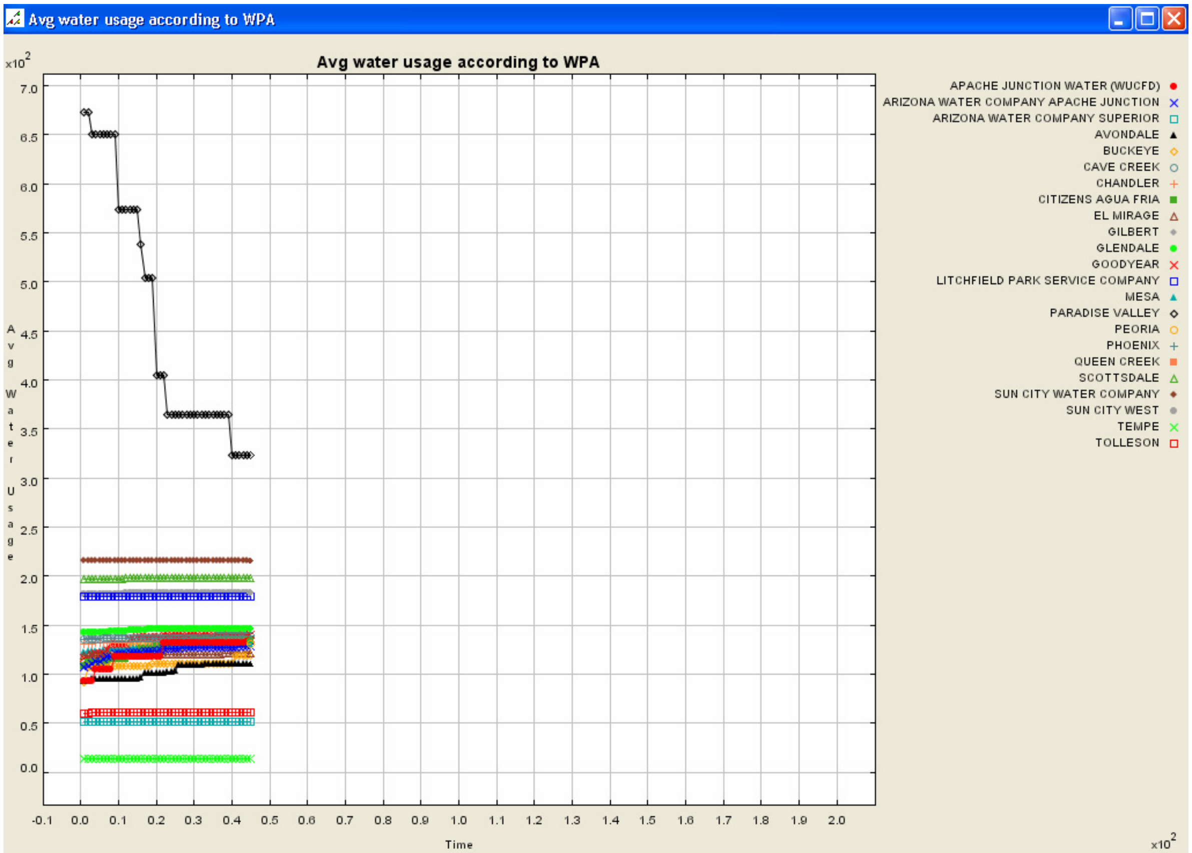

FIGURE 3 Water usage by provider

the valley can be ascribed to agricultural draws (ADWR 1999). Clearly, this is an aspect that needs to be included in the model.

Of the remaining 42%, industrial and commercial businesses account for a small proportion of water consumption, leaving a residential majority. While the current model already incorporates single-family households, it is limited in its application of family households for two reasons. (1) A large proportion of seasonal residents, college students, and low-income families live in multi-unit complexes (e.g., apartments, condominiums, and dormitories). (2) Baker and colleagues (2004) estimate that more than 70% of all residential water consumption is allocated for outdoor use (e.g., swimming pools, plants, and lawn care). Each of these factors has important implications with regard to capturing the true dynamics of the population and water-use landscape. Again, the incorporation of these two factors will be crucial for improving the validity of the model.

A final aspect is the trajectory of urban growth for the Phoenix metropolitan area. For the majority of U.S. metropolitan areas, geographic barriers limit spatial development. As a result, new construction typically requires the reallocation of land use. The Phoenix valley, however, is an exception to this constraint. Because Phoenix is housed within the desert and thus surrounded by vast open areas of desert, spatial growth is far from bounded. While the majority of new construction sites are converting agricultural lands to residential use, new areas of the desert are being transformed into housing developments in an outward direction, estimated at almost one-half mile per year (Gober and Burns 2002). This expansion has resulted in making the Phoenix region the largest contiguous metropolitan area in the United States (Melnick 1995).

One by-product of this increased construction is the nighttime attenuation of cooling by the re-radiating structures (Baker et al. 2002). This phenomenon, also known as the “urban heat island effect” has resulted in an increase of 0.1°C per year in Phoenix’s average minimum temperature over the past 50 years. In addition, Baker and colleagues also note that the number of “misery hours per day” (hours in which the temperature is above 38°C) in the valley has doubled since 1948. Not surprisingly, higher temperatures increase the need for cooling and irrigation — variables that are directly related to increased water demands (Larson et al. 2005). Given the climatic and ecological impact of urban sprawl in the Phoenix valley and the accompanying modification of water consumption, future iterations of the model will include scenarios describing how continued growth may modify water use.

Policy

Media

With respect to policy implementation, DesertWater currently maps the fluctuation of household water use as modified by exposure to water conservation media. The use of such a campaign is not novel; most local, national, and international campaigns typically include media messages that encourage residents to use water wisely. However, very few of these campaigns consider the varying likelihood of infiltration and response to those messages. By assigning our agents with sensitivity and receptivity thresholds, we have taken the first step toward modeling the complexity involved in the effectiveness of standard conservation campaigns.

Currently, the implementation of a differential response to media occurs randomly across the population of DesertWater. Yet, existing literature suggests that a conservation response is not random, but rather that there are four distinct patterns of response to media-driven conservation attempts: (1) consistent and high-frequency conservationists, (2) consistent and low-frequency conservationists, (3) individuals who practice conservation only if little or no personal sacrifice is required, and (4) individuals who engage in no conservation practices (Gilg and Barr 2005). In addition, this literature suggests that females, home owners, the well-educated, and the politically liberal are most likely to fall into groups 1 and 2, and that their male, renter, minimally educated, and politically conservative counterparts fall into group 4. Finally, Gilg and Barr note that high-income individuals are the most likely to cluster into group 3. For this particular study, older individuals were also more likely to cluster into groups 1 and 2, but this finding has not been replicated elsewhere; in fact, the opposite is typically found (Schultz et al. 1994).

With the exception of the uncertainty about age, there appear to be consistent patterns related to the demographic characteristics of individuals who engage in conservation behaviors. Given these findings, it appears that water conservation attempts by the media need to target individuals accordingly. In standard media campaigns, those in groups 1 and 2 may attend to messages to conserve water, but given the likelihood that they have already minimized use, it is not likely that the messages will result in further decreases for these individuals but rather will give them a psychological affirmation that they are doing their part. To get an added increase in conservation behavior from these individuals, it may be more efficient to use "reinforcement media." This type of media would focus on acknowledging how helpful these individuals have been, then subsequently encourage them to do just a little more. Reinforcement has been shown to be effective in perpetuating and increasing various behaviors at the individual and population level (e.g., Franzini et al. 1991), and it is likely a better method for appealing to individuals already engaged in the desired behavior.

Because individuals clustering into group 3 tend to have high incomes, it is hard to ascertain if their unwillingness to sacrifice comfort for water conservation is a result of the hurried lifestyle often seen in high-income families (e.g., working 10+ hours per day), the fact that they can afford higher water bills, or some combination of both. Clearly, standard water campaigns will not appeal to these individuals if they believe that they are entitled to use more water for either of the reasons stated above. Instead, messages targeting these families might be more effective if they recognized and empathized with their busy schedules, hard work, and stress *before* asking them to sacrifice the comfort of taking a longer shower or the convenience of running a half-empty dishwasher.

Unlike the individuals in group 3, conservation for the individuals in group 4 is not related to sacrifice. It is related to trust. These individuals consistently report disbelief in the media's call for action on conservation issues (Gilg and Barr 2005), citing exaggeration of the event (e.g., drought) or government attempts to deceive the population as the primary reason for their skepticism. Like those targeted to the other three groups, it is unlikely that standard media campaigns will affect the behavior of these individuals. To increase the likelihood of conservation actions by this group will first require addressing their bias toward media disbelief *before* asking them to join in the effort — the assumption being that very few people will act on something they think is false, especially when the action requires energy.

Given that four cluster patterns have been identified in the literature and that the demographic characteristics of each cluster are known, selecting distributions of agents that are skewed toward each of these media types will be the next likely step toward improving the model. In fact, by building scenarios that iterate the competing methods of standard media campaigns and differential media campaigns, we can tract the modification of water use in each method. Doing so should help us provide a case for the utility of our model in informing future policy decisions related to water conservation attempts via the media.

Education

While the media may constitute one way to encourage conservation, an additional, and perhaps superseding method, is the education of children. Intuitively, changing an already-existing behavior is much harder than shaping a new behavior. If young children receive the consistent message that using water wisely is the expected norm, they will be more likely to adopt water conservation behaviors and continue using them throughout adulthood. This effect has been used in other campaigns, ranging from encouraging seat-belt use to discouraging tobacco use, and it has been relatively effective (Roberts and Fanurik 1986; Jason and Pokorny 2002).

In addition to shaping children's water use behavior, educational methods may also vicariously modify adult behavior. Instead of relying on individual volition to comply in media campaigns, a family contagion effect may be created via education: children will bring the information home and act as the "policers" of household water use. Parents will, no doubt, vary in their degree of willingness to comply with the rule-oriented requests and reminders of their children. However, a child's personal request is harder to ignore than a billboard or television commercial and, as such, is a more probable means of effectively altering existing behavior.

As a result, we ultimately intend to build scenarios into DesertWater that include a distribution of agents receiving "education." Once implemented, we will be able to track the ability of these agents to modify water use in the short-term and long-term scenarios of the model. On an immediate basis, we hope to show the possibly of a contagion effect modifying overall household water use. Again, developing scenarios of competing methods — including education versus no education — will help establish the utility of our model for policy decision-making.

Politics

Policy and politics typically go hand-in-hand; ideally, the relationship is reciprocal, with each arena informing and influencing the other. With water, however, and specifically water in Arizona, this becomes a difficult task. While the policy makers desire to implement Platonic methods of "encouragement without enforcement," the politicians at the state level have serious doubts whether these methods can actually work or work in time (Arizona's Colorado River supply may soon be decreased by as much as 30%; see Larsen et al. working paper for details). As a result, the duel between Arizona water policy and politics has become a joust between nice and necessity. With the possibility of drought conditions and a decrease in an already limited supply, the politicians' concerns are legitimate. In order to ensure federal government assistance in the event of a water emergency, Arizona must show that it is trying to take the necessary measures to ensure that there is a long-term water supply for the Phoenix metropolitan area.

Water use behavior in the valley clearly needs to change, yet existing Platonic methods have not produced the desired result, and few politicians want to experience citizen backlash by imposing more Draconian methods of enforcement. Ultimately, water bans and the prevention of "water-unfriendly" landscaping may have to be imposed. In the interim, however, politicians continue to search for a less aversive yet equally effective strategy of decreasing water use in the Phoenix valley. Moreover, Arizona's water is regulated by individual water providers instead of the government, thereby complicating the policy and pricing structure of this commodity. These water providers designate the amount of available supply and price of water distributed in their area. Collectively, delegates from each area work together to set the ceiling price for water in Arizona. On one hand, the business approach to water management is beneficial to the residents, because it keeps the monthly water bill at a reasonable rate. On the other hand, it limits the ability of politicians to step in to offer either price incentives or disincentives in an attempt to modify water use.

Although the current state of water pricing is not a malleable topic, we have already incorporated this feature into the model. The basis for our decision was simple: people less responsive to other means of conservation attempts may be more or exclusively responsive to financial reinforcement or constraints. While we will primarily continue to focus on the things that do have the ability to be altered (i.e., media, education), we think it would be remiss on our part to not consider the possibility that altering the price may provide the "tipping point" needed to obtain a large-scale change in water-use behavior. We think that this scenario will be important if other methods fail, and the government is faced with making decisions about regulating the water business in Arizona.

ACKNOWLEDGMENTS

This work is based on work supported by the National Science Foundation under Grant No. SES-0345945, Decision Center for a Desert City (DCDC). We appreciate the assistance of our colleagues at the DCDC, especially Pat Gober. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

ADWR (Arizona Department of Water Resources). 1999. *Third Management Plan: 2000–2010 Phoenix Active Management Area.* Phoenix, AZ.

Baker, L.A., A.J. Brazel, N. Selavor, C. Martin, N. McIntyre, F.R. Steiner, A. Nelson, and A. Musacchio. 2002. "Urbanization and warming of Phoenix (Arizona, USA): Impacts, feedbacks and mitigation." *Urban Ecosystems 6*, 183–203.

Environment Canada, undated, *Freshwater Website;* available at http://www.ec.gc.ca/water/.

Franzini, L.R., L.M. Sideman, K.E. Dexter, and J.P. Elder. 1990. "Promoting AIDS risk reduction via behavioral training." *AIDS Education and Prevention 2*(4), 212–321.

Gilg, A.W., and S. Barr. 2005. "Encouraging environmental action by exhortation: Evidence from a case study in Devon." *Journal of Environmental Planning and Management 48*(4), 593–618.

Gober, P. 2005. *Metropolitan Phoenix: Place Making and Community Building in the Desert.* Philadelphia, PA: University of Pennsylvania Press.

Gober, P., and E.K. Burns. 2002. "The size and shape of Phoenix's urban fringe." *Journal of Planning Education and Research 21*(4), 379–390.

Griffin, W.A. 2003. "Agent-based models and computer simulation as a methodological aid for the study of micro-social behavior." Presented at Society for Research in Child Development, biennial meeting, Tampa, FL, April.

Griffin, W.A., S.K. Schmidt, and O. Kornienko. 2004. "An agent-based model of interaction between intimates." In C.M. Macal, D. Sallach, and M.J. North (Eds.), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*. ANL/DIS-05-6. Co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 7–9.

Holland, J.H. 1995. *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Addison Wesley.

Jacobs, K.L., and J.M. Holway. 2004. "Managing for sustainability in an arid climate: Lessons learned from 20 years of groundwater management in Arizona, USA." *Hydrogeology Journal 12*(1), 52–65.

Jason, L.A., and S.B. Pokorny (Eds.). 2002. *Preventing Youth Access to Tobacco.* Binghamton, NY: Haworth Press.

Larson, E.K., N.B. Grimm, P. Gober, and C.L. Redman. 2005. "The paradoxical ecology and management of water in the Phoenix, USA metropolitan area." Working paper. School of Life Sciences, International Institute for Sustainability, and Decision Center for a Desert City, Arizona State University.

Macal, C.M., and M.J. North. 2002. "Simulating energy markets and infrastructure interdependencies with agent-based models." In C. Macal and D. Sallach (Eds.), *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*. ANL/DIS-03-1. Co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 11–12.

Melnick, R. 1995. "Growth and the future of greater Phoenix: Quantity, quality, or both?" Morrison Institute for Public Policy, School of Public Affairs, Arizona State University.

North, M.J. 2001. "Toward strength and stability: Agent-based modeling of infrastructure markets." *Social Science Computer Review,* 307–323.

Roberts, M.C., and D. Fanurik. 1986. "Rewarding elementary schoolchildren for their use of safety belts." *Health Psychology 5*(3), 185–196.

San Antonio Water System, 2005, *Conservation;* available at http://www.saws.org/conservation/.

Schmidt, S.K., W.A. Griffin, L.D. Hanish, C.L. Martin, R.A. Fabes, H. Barcelo, and P. Greenwood. 2004. "PlayMate: New data, new rules, and model validity." In C.M. Macal, D. Sallach, and M.J. North (Eds.), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*. ANL/DIS-05-6. Co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 7–9.

Schultz, P., S. Oskamp, and T. Mainieri. 1995. "Who recycles and when? A review of personal and situational factors." *Journal of Environmental Psychology 15*, 105–121.

WaterCare, undated, Web site on South Australia's water issues; available at http://www.watercare.net/.

USING AGENT-BASED MODELING TO BETTER UNDERSTAND LOCAL HEROIN DEALING ORGANIZATIONS AND DRUG MARKETS

L.D. HOFFER,* Washington University School of Medicine, St. Louis, MO

ABSTRACT

This project applies agent-based modeling (ABM) to better understand the operation, organization, and structure of a local heroin dealing network and illegal drug market. The project will develop a series of computer simulations of a local heroin dealing organization and market using an eighteen-month ethnographic dataset. The ethnographic research, collected in Denver, Colorado during the 1990s, is a historic account of how a group of heroin users and dealers operated within the Larimer-area heroin market. The narrative addresses the local transformation from an open-air market to a closed/private market; the dealers' conversion from a street-based freelance operation to a private and professional organization; how dealers managed ongoing relationships with customers; how dealers made profits within different levels of the organization; and how new drug dealers initiated sales enterprises. Simulations will use a protocol linking findings from the dataset to the ABM agent and environment architecture. Simple prototypes will incorporate features of the findings, including: a typology of customers; drug purchasing and consumption patterns; social and economic exchange interactions; the dealers' credit system; wholesale and retail heroin prices; and profit data. Local environmental and historic considerations that influenced outcomes will be integrated to accurately represent agent behaviors in the organization and market.

Editors' Note: The full paper was not received in time for publication. The abstract is included to provide a frame of reference for the discussion that follows this session.

* *Corresponding author address*: Lee D. Hoffer, Washington University School of Medicine, 660 South Euclid Avenue, St. Louis, MO 63110; e-mail: hofferl@epi.wustl.edu.

MULTI-AGENT MODEL PROTOTYPE FOR CHILD VEHICLE SAFETY INJURY PREVENTION

Z. KOBTI,* S. RAHAMAN, and R.D. KENT, School of Computer Science,
and A.W. SNOWDON and T. DUNLOP, Faculty of Nursing,
University of Windsor, Ontario, Canada

ABSTRACT

Injury prevention is a growing concern in the health care sector mainly as a result of the mounting health care costs for full patient recovery. In particular, the safety of children's car seats is important to parents, legislative regulators, and automotive manufacturers. Deploying the child restraint mechanism; selecting the correct type of car seat on the basis of the age, height, and weight of the child; and correctly using the seat present a number of challenges for the driver. Driving frequency, distance, and traffic, along with the driver's experience and history, are only some of the factors that can contribute to an accident. We base our model parameters on measures colleted from recent surveys. We build a multi-agent simulation in Repast that mimics the driver's behavior that leads to the selection of a child restraint and the injury outcome after an accident. In this initial iteration of development, we report on the design phase and initial prototype. We define the level of agency and autonomy involved in the system by identifying the roles and attributes of individuals within a social context in a closed yet dynamic environment. We formulate a number of key index indicators, such as driver experience, correctness of child restraint use, accident probability, and the individual injury level following an accident. The prototype is initialized with known survey data and configured with probes to measure the injury outcome and various demographics. The prototype sets the stage for the next iteration, where we will deploy learning mechanisms to enable agents to harness the safety knowledge available in the system and employ it to decrease injury.

Keywords: Agent-based modeling system, vehicle safety, child seats, injury prevention

INTRODUCTION

Injuries due to road crashes are the leading cause of death worldwide for persons aged 0 to 44 years. The World Health Organization (WHO) has investigated the impact of road crashes globally and projects that they will be the third leading cause of death worldwide by the year 2020 (WHO 2004). Although developed countries fare better than undeveloped countries, these injuries remain a significant health issue in countries such as Canada and the United States. In Canada, injuries account for $12 billion of health care spending annually. Globally, 1.2 million deaths are attributed to road crashes — approximately 3,200 deaths per day. One of the major issues in injury prevention research relative to road crashes is the complexity of the factors that influence vehicle safety, particularly for vulnerable populations, such as children.

* *Corresponding author address*: Ziad Kobti, School of Computer Science, University of Windsor, 401 Sunset Avenue, Windsor, Ontario, Canada N9B-3P4; e-mail: kobti@uwindsor.ca.

Traditional approaches to injury prevention research identify multiple factors or variables that influence crash outcomes, often using cross-sectional statistical approaches. However, the multi-dimensional nature of injury outcomes from vehicle collisions requires new and innovative approaches to data analysis in order to fully understand the multi-dimensional nature of injury prevention.

The fundamental characteristics of vehicle safety and injury prevention research encompass an understanding of human behavior and the environment, as well as a number of associated social and economical components. Health researchers teamed with computer scientists from the University of Windsor to embark on developing a virtual model to replicate such an environment. Critical aspects of the model involve human intelligence exhibited in software agents with the ability to act rationally through the use of knowledge and information and in pursuit of a goal. Agents differ in important characteristics by reacting to environmental and socioeconomic variations, as well as by deliberately employing knowledge, reasoning, and even learning to plan to achieve their goal. The collaboration of agents is another critical property in social systems. The dynamic environment provides agents with a nonlinear change over time, presenting agents with a constant need to adapt and learn from their own as well as from others' experiences. The changes are often the result of feedback that the agents receive as a result of ongoing activities. Individual agents tend to organize into groups or hierarchies and form a structured organization, often influencing the underlying system's evolution. Investigators probing these systems often can identify emerging properties arising from the interactions and actions of the individual agents. These attributes together characterize a complex system. A multi-agent simulation consequently acts as a tool in modeling real-world complex systems, in which an agent provides the observer with a first-person's eye view of the world as it unfolds in the simulated world (Kobti 2003).

The purpose of this pilot study is to employ agent modeling as a strategy for examining the interactions of the multiple dimensions of injury outcomes associated with vehicle crashes. The first section presents an overview of the literature related to children vehicle safety, followed by a characterization of agent-based modeling and social systems. The next section provides a detailed methodology for the simulation prototype and finally reports the results of initial testing and model validation strategies.

LITERATURE REVIEW

Research that examines children's safety in vehicles has grown substantially in the past few years. One of the most consistent and alarming findings in the research is that motor vehicle accidents continue to be the leading cause of death and serious injury among children under the age of 14 years (Zaza et al. 2001). In Canada, these collisions result in hundreds of child fatalities annually. Despite the availability of effective safety restraint devices, an additional 15,000 young Canadians are severely injured annually as a result of roadway accidents, (Transport Canada 2003). When child safety restraints are used correctly, the risk of death and/or serious injury can be reduced by as much as 74% (Biagioli 2002; Weber 2002). Many countries are moving toward legislation for child safety seats in response to this growing body of evidence, but legislating increased use alone may not be adequate to protect children from trauma due to vehicle collisions.

Safe Seat Transitions

Continued deaths and injuries to children are occurring partly as a result of vehicle safety practices not being accurately matched to the occupants' body structure. The majority of safety experts agree that height, weight, and age (for infants) are the key determinants in child restraint choice and transition. All child restraint devices and children under the age of 12 should be in the rear seats of vehicles. Initially, infants must ride in a rear-facing seat — until they weigh 22 pounds and until they reach the age of at least one year (O'Day 2001). The second transition requires parents to move their child from the infant seat to the forward-facing seat. Children remain facing forward until they weigh 40 lb and are at least 40 inches tall. The next transition occurs for children who are between 40 and 57 inches tall and weigh between 40 and 80 lb. At this height and weight, children are most safely restrained in a booster seat (Apster et al. 2003; Lee et al. 2003). Unfortunately, with the exception of two Canadian provinces, children over 40 lb can be legally secured by a vehicle seat belt system. Lap and shoulder belts have always been developed and tested for adults; however, they do not offer adequate protection for children weighing less than 80 lb. Most parents do not know that a seat belt offers less than optimal protection for their school-aged child (Rivara et al. 2001). Without knowledge of the importance of safe booster seat use, parents prematurely graduate their children to vehicle seat belts, completely unaware of the risks or falsely believing that they made a correct choice (Simpson et al. 2002). One study suggests that two-thirds of children were inappropriately restrained by seat belts (Ramsey et al. 2000). Similarly, a recent Canadian survey reported that only 28% of school-aged children are restrained in booster seats (Safe Kids Canada 2004). These results call attention to numerous missed opportunities for controlling unnecessary deaths and injuries to children in vehicles.

Consequences of Inappropriate Transition

The risks associated with premature graduation to seat belts by young children have been established in the literature (Winston et al. 2000). Certainly, the most extreme outcome of not using a booster seat is fatalities following collisions. However, inadequately restrained children are also more likely to sustain serious injuries to the head, neck, spine, and abdomen (Winston et al. 2000). Injuries occur because children's small torso and underdeveloped pelvic bones offer a less-than-optimal fit for both shoulder and lap belts (Slatter and Vargish 1998).

In the event of a crash, a shoulder harness positioned somewhere other than over the sternum allows for excessive movement of the head and upper body, commonly associated with negative spinal outcomes, skull fractures, and severe brain injuries (Winston et al. 2000). Similarly, ill-fitting lap belts place undue pressure on the abdomen. Findings from a crash surveillance study (Winston et al. 2000) revealed that only children prematurely restrained in seat belts incurred abdominal injuries. Other lap belt injuries frequently reported in the literature involve the lumbar spine and intestinal tract (Lane 1994). Preventing early seat belt transitioning is an important means of decreasing the prevalence of child deaths and injuries in vehicles. Figure 1 illustrates the percentage of premature seat belt use among children of varying ages as collected from field survey data. The trend reveals that as a child's age increases, drivers are more likely to use a seat belt to restrain the child instead of the appropriate child restraining device.

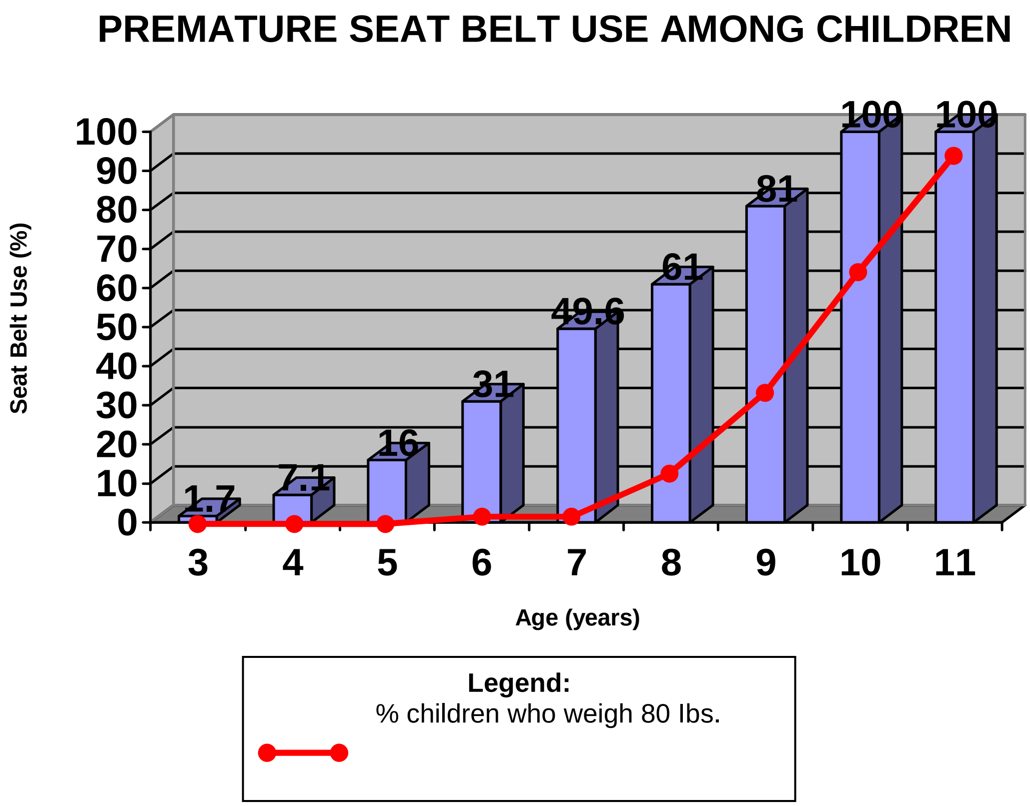

FIGURE 1 Premature seat belt use among children by age

Issues of Use and Misuse

While the majority of parents attempt to use vehicle safety systems to protect their children, misuse and nonuse continue to be significant factors in injury outcome for children traveling in vehicles. Biagioli (2002) reported that more than 80% of safety seats are misused. Correct use requires that the child safety restraint be (1) appropriate for the child's height and weight, (2) accurately installed and positioned in the vehicle, and (3) used every time the child is transported in the vehicle. Progressive patterns of child growth and development pose particular challenges for many parents. Normal physical and cognitive changes throughout phases of the child's life cycle require parents to learn to install and use a variety of vehicle restraints for each stage of their child's growth and development. For example, infants quadruple their weight in the first two years, and children "normally" gain weight steadily, at 4 to 6 lb per year, until adolescence (Wong 1999). There is little doubt that installation also contributes to high rates of misuse of child restraint systems because of the wide variety of safety seat models, seat belt systems, and seating configurations in vehicle interiors. Some of the most common difficulties with restraint installation are ensuring tightness of harness straps and safety belts and remembering to use locking clips or tether straps (Kohn et al. 2000; Lane et al. 2002).

However, the challenges that parents face in making vehicle safety decisions also depend on the accuracy and availability of the information that supports the effective use of restraint systems. The accuracy of information on the correct installation of safety seats varies, depending on the manufacturer's instructions, the directions of sales personnel, and the advice of family or friends. It is estimated that only 50% of parents actually read the product manual on how to install the safety seat properly. For those who do, the instruction manual's vocabulary often exceeds the parent's comprehension level (ability to readily understand the information and follow the directions for proper utilization) (Block et al. 1998; Decina and Knoebel 1997; Gaines et al. 1996; Wegner and Girasek 2003). In a recent study of 107 manuals from 11 different manufacturers, a grade 10 reading level was required, on average, in order to fully comprehend the information presented (Huggins 2003).

Nonuse is another significant issue. The rationales that parents describe for choosing not to use a child safety restraint for their toddlers or preschool children include the child's fussiness and discomfort, the inconvenience of using the device, and the need for using a restraint device for a younger child (Decina and Knoebel 1997).

Intervention Research

A systematic review of interventions designed to increase use of child safety seats was conducted in 2001 (Zaza et al. 2001). The review focused on the effectiveness of five interventions aimed at increasing child safety seat use. The success of each intervention was evaluated in terms of changes in the use of child safety seats or injury rates. At the time of the review in 1998, more than 3,500 citations were screened, and 72 met the inclusion criteria for the reviews. The results of this review identified "strong evidence of effectiveness for child safety seat legislation and distribution plus education programs" (Zaza et al. 2001, page 31). Education-only programs directed toward parents, young children, health care professionals, or law enforcement personnel (Zaza et al. 2001, page 31) were found to be less effective than communitywide information/enhanced enforcement campaigns and incentive-plus-education programs.

In 1999, Rivara identified the need for intervention research to investigate methods to increase the use of booster seats. It was suggested that different strategies might be needed for this older population than what had been previously used for younger children. Knowledge of the benefits and purpose of booster seats alone is not sufficient to promote increased use (Simpson et al. 2002). Parental perception of risk, awareness/knowledge, and parenting style were identified as issues affecting use when comparing parents of children in booster seats with those who had put their children into seat belts. In this study, media campaigns, improved laws, education for parents, and extending the use of child restraints to older children were among the strategies suggested by parents in focus groups for increased use of booster seats.

Ebel et al. (2003) conducted a prospective, nonrandomized, controlled community intervention trial to evaluate the effectiveness of a multi-faceted community booster seat campaign designed to increase observed booster seat use among child passengers in motor vehicles. By 15 months after the start of the campaign, booster seat use did increase significantly in the intervention communities relative to the control communities. The success of this intervention was in part due to the multi-faceted approach that was undertaken.

Agent-based Modeling Systems

Multi-agent-based interoperation is a new paradigm distinguished by features such as requests that are specified in terms of “what” and not “how”; agents that can take an active role, monitoring conditions in their environment and reacting accordingly; and agents that may be seen as holding beliefs about the world (Huhns and Singh 1997).

Agent-based modeling is gaining popularity because of its versatility for encapsulating and abstracting complex system models. Furthermore, its application has evolved beyond the discipline of computational and computer science, and it has been adopted in a variety of disciplines ranging from administration and economics to health care and policy making.

In this study, an agent-based simulation is developed to analysis the child safety issues in North American societies. Various statistical analyses and surveys have been adopted into the model in order to initialize it and provide a realistic perspective on various agent functions. Child safety issues depend on a variety of social, mental, and emotional factors, and traditional surveys and statistical analyses, because of their multi-dimensional nature, fall short in discovering emerging aspects of these issues under dynamic conditions. For instance, the perspective that parents have and choices that parents make about their children’s safety depend on their knowledge about this issue, personal way of thinking, ethnic background, cultural influences, financial position, and enforced laws, along with observed and learned experiences.

Repast (Recursive Porous Agent Simulation Toolkit) is an open source tool for developing agent-based simulations created at The University of Chicago and maintained by Argonne National Laboratory (ANL). It is especially suitable for implementing a simulation model that involves living social entities or beings. Repast provides an excellent framework and specifications for developing agent-based simulations. The framework’s support is not limited to core simulation strategies (like event scheduling and random number generation), intuitive user interface, and graphic generation and visualization tools; it also provides built-in support for genetic algorithms, neural networks, and geographic information systems (GISs). Since it is open source and distributed under the GNU general public license (GPL), the user can customize and extend it virtually in any way he or she chooses. Currently, Repast is available in three flavors: Repast J for the Java platform, Repast .Net for Microsoft .Net framework, and Repast Py for Python scripting. In this prototype, we used Repast J 3.0 in order to fully exploit the flexibility of Repast and the portability of Java (Xu 2004).

MODEL DESIGN AND IMPLEMENTATION

An agent in the current model is designed as a household consisting of individual members, including their gender, age, and kinship status, as well as information about the household’s income level and other relevant ethnic and social characteristics. Household members can be adult individuals, including parents, and children. Associated with each household is a set of vehicles. A driver is a designated individual who satisfies the rules, such as age and license issue. The event of driving is abstracted by the class Trip. In the AutoSimModel, the creation of the agents and initiation of the events in the model along with various probes may be implemented. The class diagram, representative of the basic classes in the object-oriented Repast model, is illustrated in Figure 2.

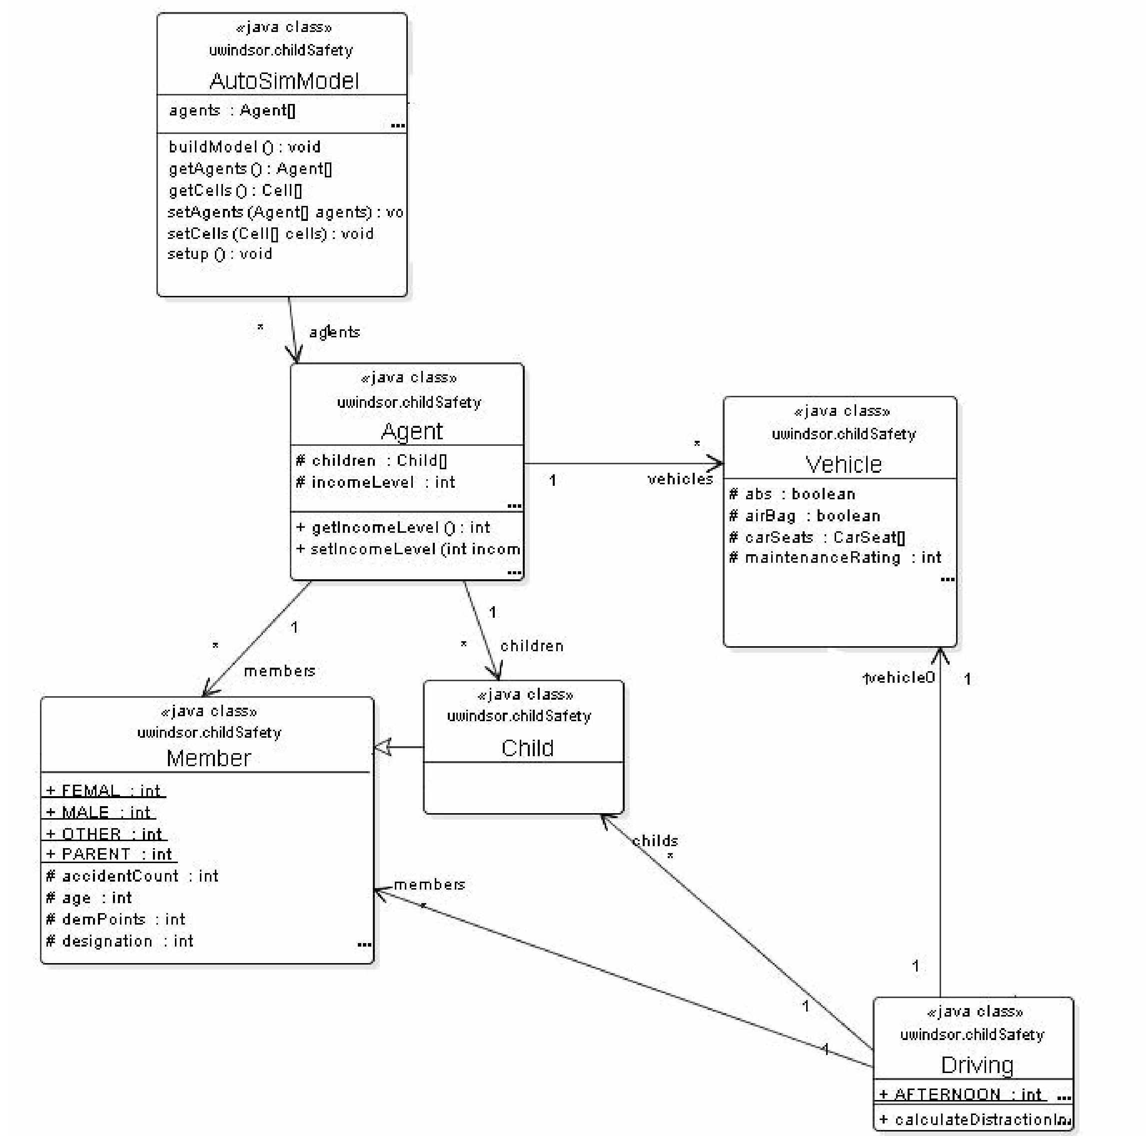

FIGURE 2 Class diagram: Java package uwindsor.childSafety

Indices and Functions

To initialize the model data and functional behavior, we rely on known measures and outcomes for common procedures. Specifically, we use empirical indices to identify the states of each event. A number of indices that have been formulated are described below.

The probability of a driver to become involved in an accident mostly depends on his/her personal driving experience, the condition of the vehicle, the traffic conditions at the time of driving, the trip distance and duration, and some other factors that may cause distraction while driving. In order to estimate the probability of an accident, we consider three indices: driver's experience index (DEI), distraction index (DI), and vehicle safety index (VSI). Each of these indices assumes a value in the closed range 1 to 10, with the possibility of a value of 0 indicating the absence of influence from such index. These indices are employed in the probability of accident index ($P_{Accident}$) by using the following rule:

$$P_{Accident} = (DEI \times W_{DEI} + DI \times W_{DI} + VSI \times W_{VSI})/100, \tag{1}$$

where W_{DEI}, W_{DI}, and W_{VSI} correspond to the weight percentage for each of the DEI, DI, and VSI indices, respectively. The weight is an estimated measure for the effect that a given index has on the overall circumstances for the likelihood of an accident. If the weights are not known, they are assumed to be divided equally for each of the indices. Note that the value of $P_{Accident}$ is bounded to assume a value in the closed range [1–10]. In the current prototype, we kept all the weights equal. Consequently, Equation 1 reduces as follows:

$$P_{Accident} = (DEI + DI + VSI)/3. \tag{2}$$

Table 1 summarizes all the factors currently considered in calculating each of these indices.

The model assumes that a higher value for $P_{Accident}$ indicates increased chances for an accident to take place, and a lower value reduces the likelihood of an accident (an indicator of safer driving conditions). In Equation 2, $P_{Accident}$ is directly proportional to the individual values and the chances for an accident. Hence, the following observations are implied:

1. A higher DEI value indicates a risky driver, and a lower value indicates a safer driver.

2. A higher DI means a higher distraction level, and a lower DI implies the opposite.

3. A higher VSI means the vehicle is more prone to an accident, and a lower VSI means the opposite.

TABLE 1 Indices used in the simulation

Index	Factors Considered	Comment
Driver's experience index (DEI)	Years of experience, age , sex, training, ethnicity, birthplace, years living in Canada, education level	Can assume a value of 1 to 10, with lower values indicating safer driver
Distraction index (DI)	Passenger's average age, time of day, traffic condition, road rage, cell phone usage	Can assume a value of 1 to 10, with lower values indicating lower level of distraction
Vehicle safety index (VSI)	Model, year, mileage	Can assume a value of 1 to 10, with lower values indicating safer car
Probability of accident index	Driver's experience index, distraction index, vehicle safety index	Can assume a value of 1 to 10, with lower values indicating lower chance of an accident

The value of DEI (Table 2) is calculated by using the following rule (Equation 3), again scaled to the range 1 to 10:

$$\begin{aligned} DEI = {} & W_{yearsOfExperience} \times yearsOfExperience + W_{age} \times age \\ & + W_{sex} \times sex + W_{training} \times training + W_{ethnicity} \times ethnicity \\ & + W_{educationLevel} \times educationLevel + W_{birthPlace} \times birthPlace \\ & + W_{yearsLivingInCanada} \times yearsLivingInCanada. \end{aligned} \tag{3}$$

In the current implementation of Equation 3, the weights are as shown in Equation 4:

$$\begin{aligned} DEI = {} & 0.25 \times yearsOfExperience + 0.25 \times age \\ & + 0.05 \times sex + 0.10 \times training + 0.05 \times ethnicity \\ & + 0.10 \times educationLevel + 0.10 \times birthPlace \\ & + 0.10 \times yearsLivingInCanada. \end{aligned} \tag{4}$$

The DI encapsulates the major relevant factors that may contribute to driver distraction while driving, consequently leading to reduced driving performance and an increased likelihood of an accident. Table 3 shows the calculation details for this method. Each of the factors again assumes a value 1 to 10, with 1 being the lowest level of distraction and 10 being the highest.

Equation 5 details the computation of DI; the equation is simplified since the weights are equal in this case.

$$\begin{aligned} DI = {} & (AgeIndex + RoadRageIndex + TrafficIndex \\ & + TimeOfDayIndex + CellPhoneIndex)/5. \end{aligned} \tag{5}$$

Some vehicle models have improved safety over others. The model, age, and mileage of the car are the main indicators that reflect the capability of the vehicle to sustain accidents. The VSI represents the overall contribution of the vehicle choice in an accident scene. Each of the factors in Table 4 assumes a value of 1 to 10, with 10 being the worst and 1 being the best choice for a vehicle in terms of safety.

Equation 6 shows the computation of VSI:

$$\begin{aligned} VSI = {} & (W_{type} \times TypeIndex + W_{age} \times AgeIndex \\ & + W_{mileage} \times MileageIndex)/100. \end{aligned} \tag{6}$$

The vehicle selection function is used at the very beginning of the simulation in order to initialize each agent (or household) with a set of vehicles. Its rules are based on Table 5 to assign a vehicle to a household depending on its income level and family size.

In order for the model to establish a frame of reference for the correct selection for the child restrain, we identify in Figure 3 the proper selection procedure for seat usage depending on the child’s age and height. In the model, we are now able to compare the driver’s actual decision to the correct one.

TABLE 2 Driver experience index calculation table

Factor (Variable Name)	Description/Rule	Weight (%)
Years of experience	YearExp < 5 ⇒ 10.0 YearExp < 10 ⇒ 9.0 YearExp < 15 ⇒ 7.0 YearExp > 15 ⇒ 4.0	25
Age	Age < 18 or Age > 80 ⇒ 10.0 Age < 20 or Age > 75 ⇒ 9.0 Age < 25 or Age > 70 ⇒ 7.0 Age < 30 or Age > 60 ⇒ 6.0 Age < 35 or Age > 50 ⇒ 5.0 Age < 40 or Age > 45 ⇒ 4.0 Age ≤ 45 or Age ≥ 40 ⇒ 3.0	25
Sex	Sex = male ⇒ 9.0 Sex = female ⇒ 7.0	05
Training	Country of driver's training: Canada ⇒ 5.0 USA ⇒ 8.0 Others ⇒ 9.0	10
Ethnicity	Caucasian ⇒ 6.0 Native ⇒ 6.0 Asian ⇒ 7.0 African ⇒ 7.0 European ⇒ 7.0 Others ⇒ 9.0	05
Education level	Below high school ⇒ 9.0 High school graduate ⇒ 7.0 College graduate or higher ⇒ 5.0	10
Birthplace	Canada ⇒ 6.0 USA ⇒ 8.0 Other ⇒ 9.0	10
Years living in Canada	< 5 years ⇒ 8.0 < 10 years ⇒ 7.0 < 15 years ⇒ 4.0 < 20 years ⇒ 2.0	10

TABLE 3 Distraction index calculation table

Factor (Variable Name)	Description/Rule	Weight (%)
Passenger's age group (average age)	Average age < 15 ⇒ 10.0	20
	Average age < 20 ⇒ 9.0	
	Average age < 25 ⇒ 8.0	
	Average age > 30 ⇒ 7.0	
	Average age > 35 ⇒ 6.0	
	Average age > 40 ⇒ 5.0	
	Average age > 45 ⇒ 4.0	
	Average age > 50 ⇒ 3.0	
	Average age > 55 ⇒ 1.0	
Road rage	Involved ⇒ 10.0	20
	Not involved ⇒ 00.0	
Traffic conditions	Rush hour ⇒ 9.0	20
	Busy ⇒ 8.0	
	Moderate ⇒ 7.0	
	Light ⇒ 6.0	
	Very light ⇒ 5.0	
Time of day	Morning ⇒ 9.0	20
	Noon ⇒ 7.0	
	Afternoon ⇒ 8.0	
	Evening ⇒ 6.0	
	Night ⇒ 4.0	
Cell phone usage	Using ⇒ 9.0	20
	Not using ⇒ 0.0	

The current simulation is specially designed to work with a comparable demographic pattern as experienced in the Auto 21 survey literature. When initializing the agents, it uses the known statistics, such as the age group distribution of the adult population, average size of households, and average number of children per family. Some other statistics are also used from other sources detailed in the model for each specific parameter. Initially, the model distributes the vehicle objects among the households by using the vehicle selection function (Table 5). Indices like DEI and VSI are also calculated at this time. The DI is calculated at simulation time. The sequence of operations to initialize the model is depicted in Figure 4.

Each time-step in the simulation is equivalent to a day in the life of an agent. The simulation starts by scheduling each driver's driving assignments; it is, of course, possible for a driver not to drive on a given day. While selecting a household for driving, some statistics were used from the survey literature (e.g., 66% of parents drive their children on an everyday basis). The DI and probability of accident index ($P_{Accident}$) are calculated at this point. Depending on the $P_{Accident}$ values, some of the trips become involved in an accident. After an accident occurs,

TABLE 4 Vehicle safety index calculation table

Factor (Variable Name)	Description/Rule	Weight (Total 100%)
Model/type	SUV ⇒ 9.0	W_{type} = 20
	Mini van ⇒ 7.0	
	Van ⇒ 6.0	
	Truck ⇒ 6.0	
	Coupe ⇒ 6.0	
	Sedan ⇒ 4.0	
	Station wagon ⇒ 5.0	
Age = current year – model year	Less than 1 year old ⇒ 2.0	W_{age} = 40
	Less than 2 years old ⇒ 3.0	
	Less than 4 years old ⇒ 4.0	
	Less than 6 years old ⇒ 5.0	
	Less than 8 years old ⇒ 6.0	
	Less than 10 years old ⇒ 7.0	
	Less than 12 years old ⇒ 8.0	
	Less than 14 years old ⇒ 10.0	
Mileage	Less than 50,000 km ⇒ 2	$W_{mileage}$ = 40
	Less than 100,000 km ⇒ 3	
	Less than 150,000 km ⇒ 5	
	Less than 200,000 km ⇒ 6	
	Less than 220,000 km ⇒ 7	
	Less than 250,000 km ⇒ 8	
	Less than 300,000 km ⇒ 9	
	More than 300,000 km ⇒ 10	

the health indices of the involved persons are consequently updated, along with the vehicle damage. The output graphs and charts probing the model are updated every 50 steps. The main operations and the life cycle of an agent are depicted in Figures 5 and 6, respectively.

The simulation is executed for 365 time-steps (days), and the results summarizing demographic distribution, correctness of child car seat usage, and average health over time are presented. The demographic distribution histogram shown in Figure 7 is generated right after the simulation is initialized, and it is scheduled for an update every 365 ticks (1 year). The height of the bars represents the frequency of each age group. Most likely drivers are in the age range of 31 to 36 years.

Figure 8 shows the correct child seat usage for each of the child age groups: infant (1–12 months), toddler (13–48 months), and walker (49–144 months). This graph is updated every 50 ticks (days). If there is no legislation in effect, drivers do not learn the correct child seat usage, and subsequently this graph remains almost the same over time in the current simulation.

TABLE 5 Vehicle selection function based on family size and household income[a]

Family Size (No. of Members)	Household Income ($ per Annum) Below 20 K	Below 30 K	Below 40 K	Below 60 K	Below 80 K	Below 100 K	Over 100 K
1	0 or 1 sedan 9-12	1 or 2 sedan/van 7-10	1 or 2 any type 5-8	1 or 2 any type 4-7	1 or 2 any type 2-5	1 or 2 any type 0-3	1 or 2 any type 0-3
2	0 or 1 sedan 9-12	1 or 2 sedan/van 7-10	1 or 2 sedan/van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3
3	0 or 1 sedan 9-12	1 or 2 sedan/van 7-10	2 or 3 sedan /van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3
4	0 or 1 sedan 9-12	1 or 2 sedan/van 7-10	2 or 3 sedan /van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3
5	0 or 1 sedan /van 9-12	1 or 2 sedan/van 7-10	2 or 3 sedan /van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3
6	0 or 1 sedan /van 9-12	1 or 2 sedan/van 7-10	2 or 3 sedan /van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3
7 or more	0 or 1 sedan /van 9-12	1 or 2 sedan/van 7-10	2 or 3 sedan /van 5-8	2 or 3 any type 4-7	2 or 3 any type 2-5	2 or 3 any type 0-3	2 or 3 any type 0-3

[a] In each cell, the first line is the number of vehicles the family may have. The second line is the type of vehicles the household may own, and the third line indicates how many model years old the vehicle may be.

Figure 9 shows the correct and incorrect seat usage comparison among different weight groups of school-aged children. The figure shows a snapshot of the graph at the very beginning of the simulation. This graph shows no change over time if there is no legislation or change in driver knowledge about child safety. Should there exist any such characteristics of learning or enforcement, the outcome would change. We would expect an increase in the correct usage frequency and a decrease in the incorrect usage.

The overall measure of the effectiveness of proper seat usage and increase in safety is revealed in Figure 10, which shows the average health of adults and children over time. The health index is based on a normalized value from 0 to 10, where the higher the number means the healthier the individual. This value is normalized to reflect the level of injury sustained by individuals after an accident. A recovery period can be implemented to enable improvement in

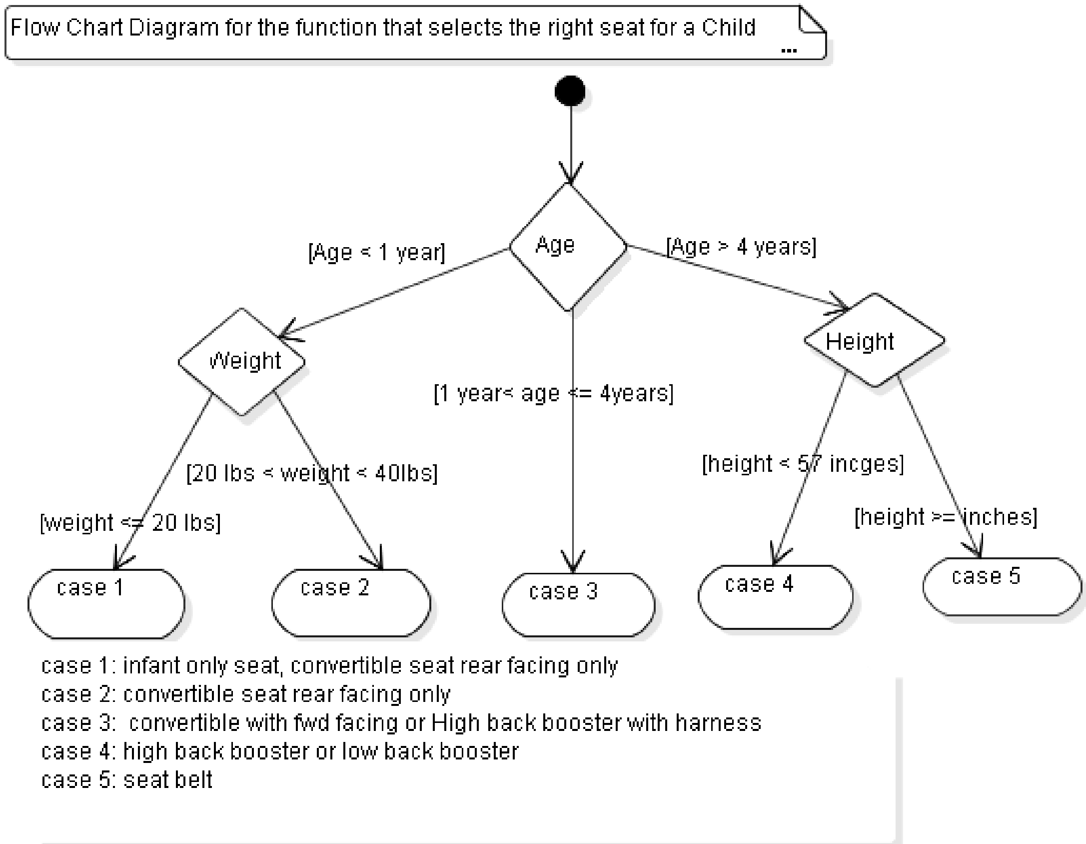

FIGURE 3 Flow chart for the method that assigns a child correctly to a seat on the basis of his/her age, weight, and height

health over time (i.e., healing). Some injuries, however, are terminal, which is indicated by lower values, and death is when an individual reaches a health index of 0; subsequently, the dead individual is removed from the system.

CONCLUSIONS AND FUTURE WORK

This paper presents a new approach for studying vehicle injury prevention for the general population and for children in particular. The first phase of the study is to build a prototype capable of modeling agent behavior, including vehicle selection, driver assignment, child seat usage, and accident generation. Many parameters were used to initialize the model; they were based on published studies and field surveys. Different measures were then implemented to probe the outcome of the system. Some histograms can reveal insights about the population being modeled and thereby enable us to understand the characteristics of the population and various demographic aspects, such as age and gender distributions of drivers. Other graphs can measure the levels of correctness of child seat use, indicative of the knowledge of the drivers in the modeled system. Furthermore, we can examine the overall population health, even breaking it down to adult and child health separately, in order to see the effect of learning, prevention, or legislation on the driver behavior and overall injury levels in the population over time.

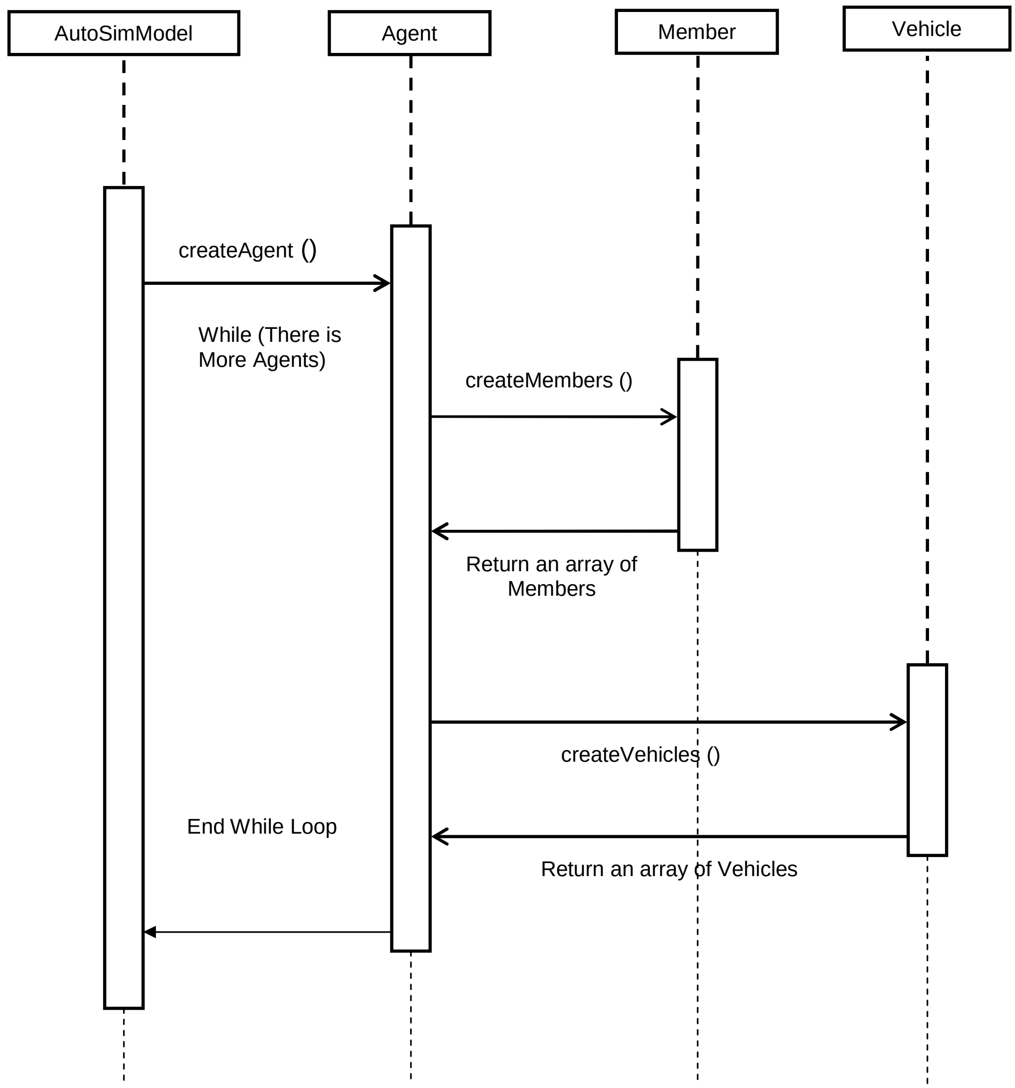

FIGURE 4 Model initialization sequence diagram

In the current prototype, we enable the agents to behave in a probabilistic, nearly deterministic manner. The next phase of work will involve the addition of the ability of agents to learn and adapt to various injections and feedbacks in the model. For instance, given an enforceable legislation, a driver who receives a fine for incorrect use of a seat for a child would make the driver learn from his/her mistake and make the correction in the future, thereby reducing the injury level in the case of an accident. Other parameters can also be added into the system, and additional measures can be probed in order to validate the system.

ACKNOWLEDGMENTS

This work was partially supported by a grant from Auto21, supported by the government of Canada, and contains the input of various industry experts.

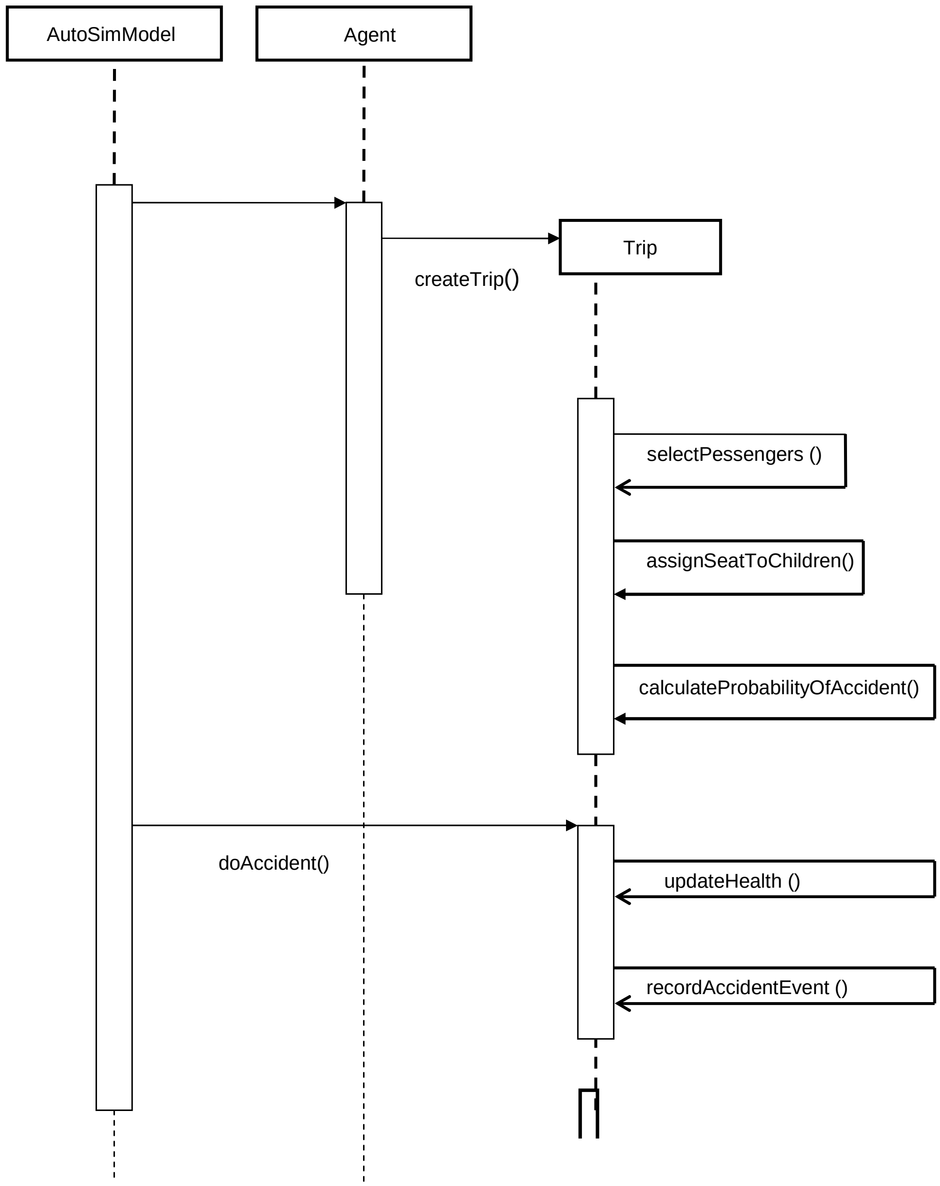

FIGURE 5 Summary of agent actions in a single time-step

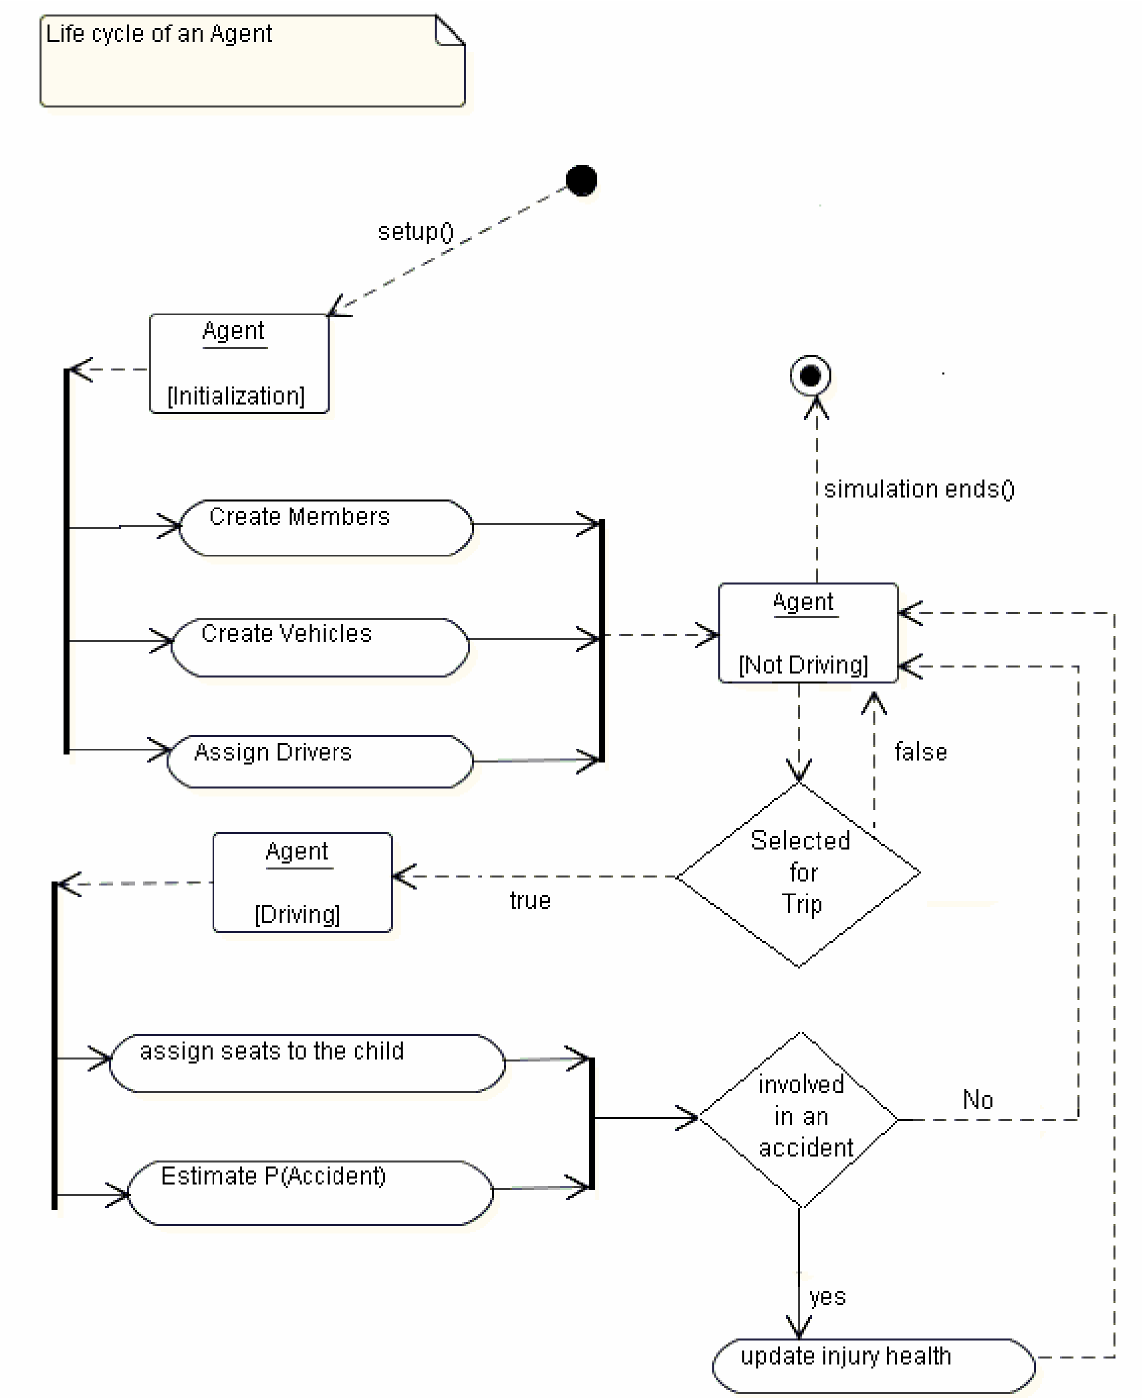

FIGURE 6 Life cycle of an agent through the entire run of the simulation

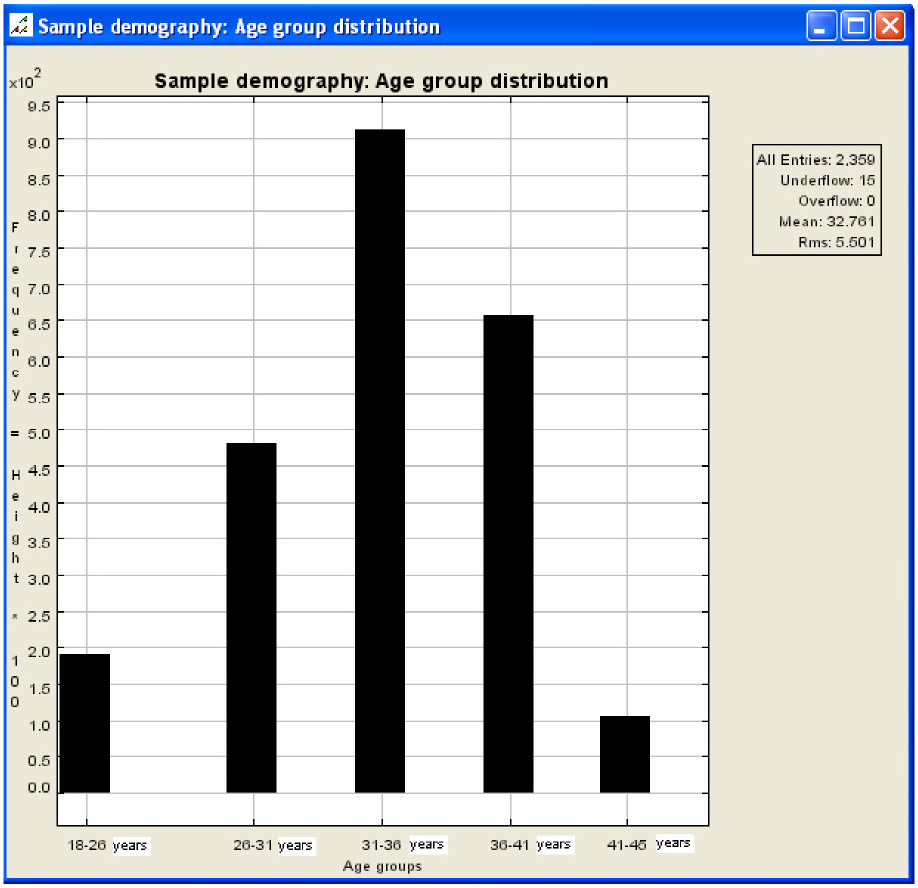

FIGURE 7 Age distribution of the adult population (Age groups are on the x axis, and frequencies are on the y axis, where frequency = height × 100.)

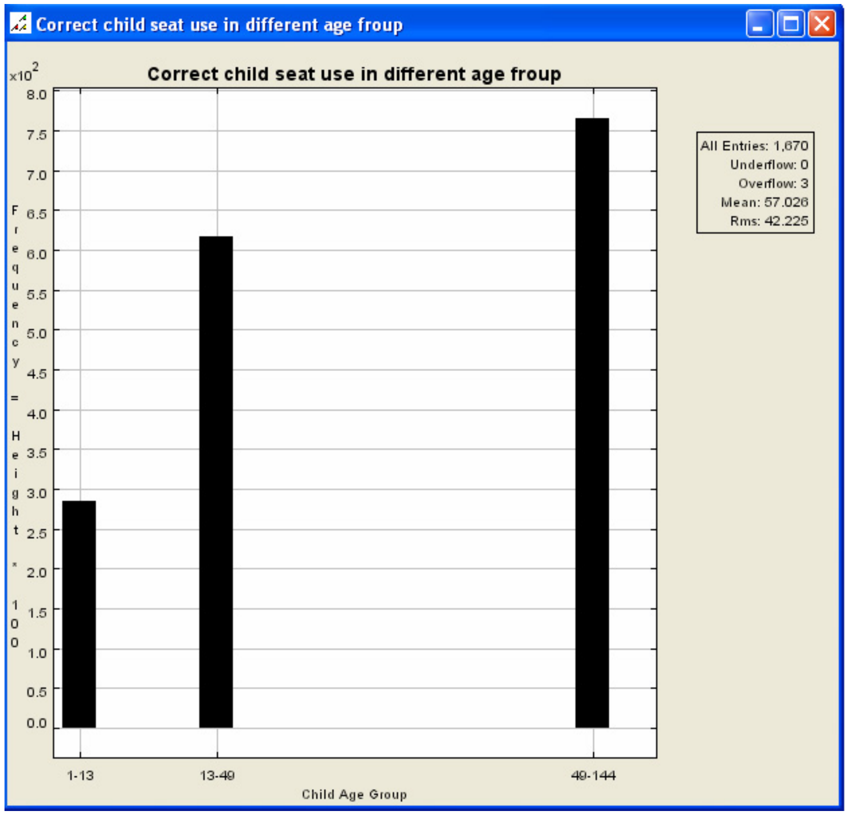

FIGURE 8 Correct seat usage distribution among all children (Among 1,900 children, approximately 290 infants, 620 toddlers, and 760 school-aged children use the correct seat; the rest use incorrect seats.)

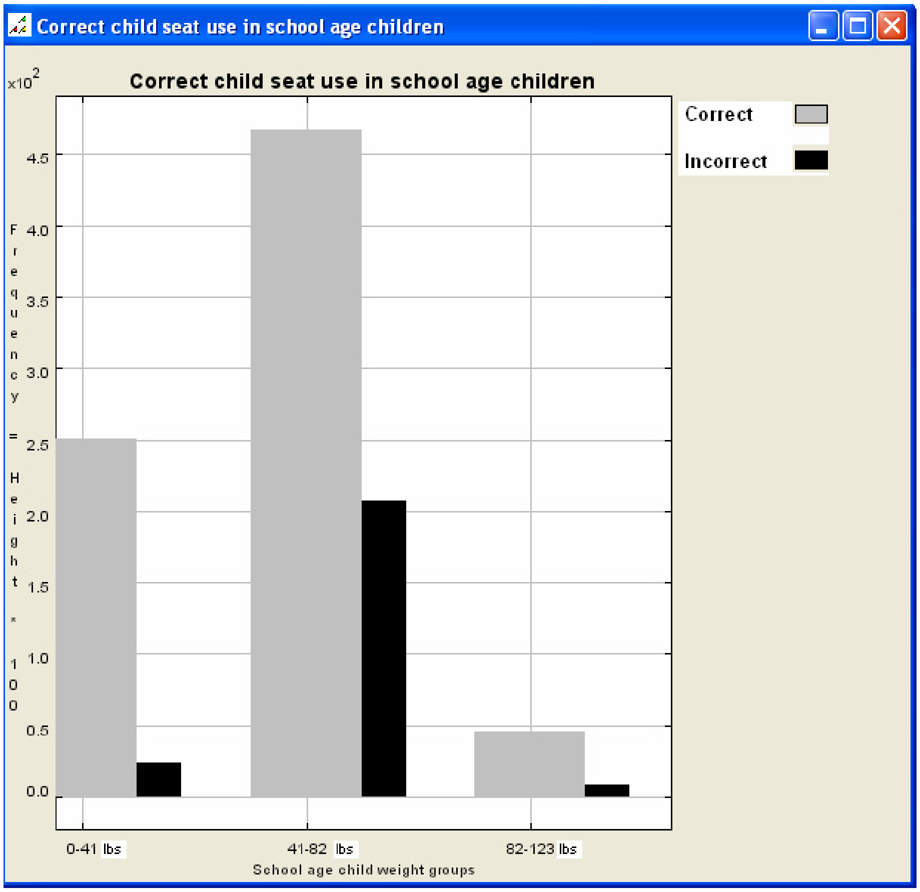

FIGURE 9 Correct vs. incorrect seat usage among different weight groups of school-aged children

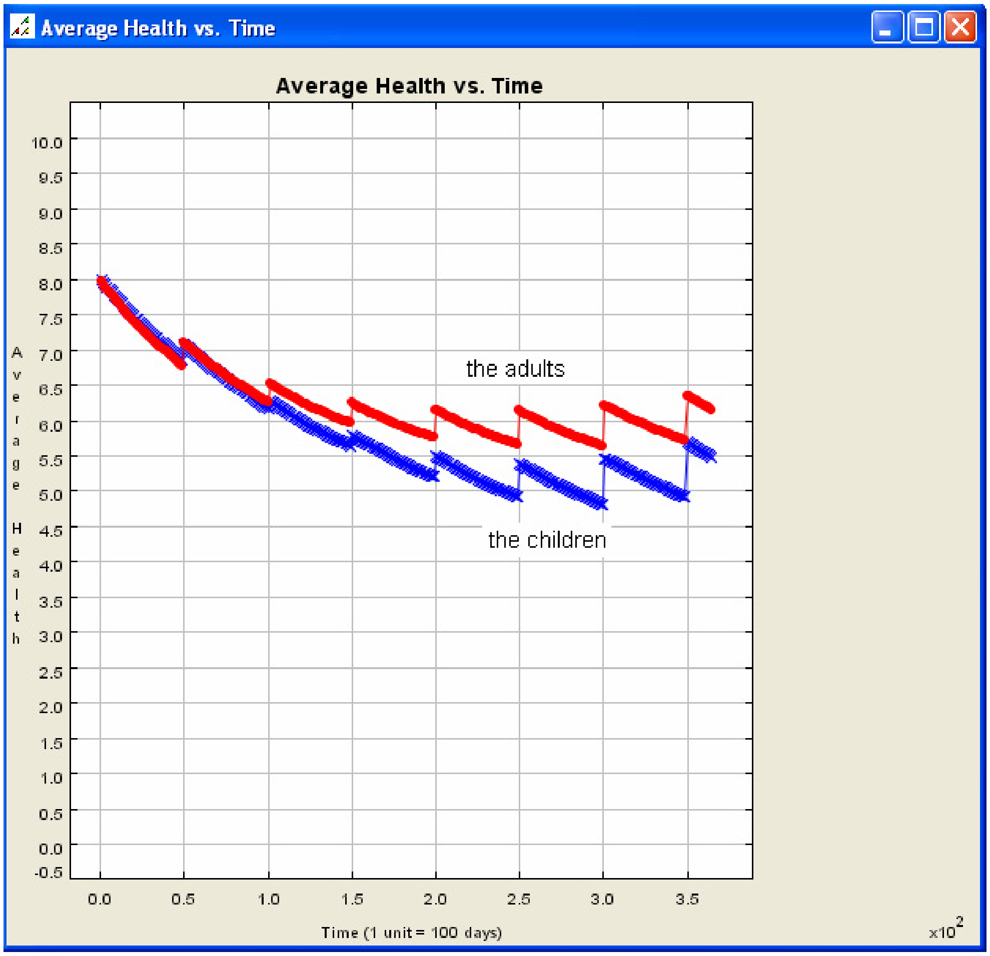

FIGURE 10 Average health index for adults and children over time

REFERENCES

Apsler, R., S.W. Formica, A.F. Rosenthal, and K. Robinson, 2003. Increases in booster seat use among children of low income families and variation with age. *Injury Prevention* 9, 322–225.

Biagioli, F., 2002. Proper use of child safety seats. *American Family Physician* 65(10), 2085–2090.

Block, D., T. Hanson, and A. Keane, 1998. Child safety seat misuse: Home visiting assessment and intervention. *Public Health Nursing* 15(4), 250–256.

Decina, L., and K. Knoebel, 1997. Child safety seat misuse patterns in four states. *Accident* 29(1), 125–132.

Ebel, B.E., T.D. Koepsell, E.E. Bennett, and F.P. Rivara, 2003. Use of child booster seats in motor vehicles following a community campaign: A controlled trial. *Journal of the American Medical Association* 289(7), 879–884.

Gaines, S., B. Layne, and M. DeForest, 1996. Promoting automobile safety for young children. *The American Journal of Maternal/Child Nursing* 21, 148–151.

Huggins, C., 2003. Child car seat instructions hard to understand. *MEDLINEplus*.

Huhns, M.N., and M.P. Singh, editors, 1997. *Readings in Agents*. California: Morgan Kaufmann Publishers, Inc.

Kobti, Z., R.G. Reynolds, and T. Kohler, 2003. A multi-agent simulation using cultural algorithms: the effect of culture on the resilience of social systems, in *Proceedings of the IEEE Conference on Evolutionary Computation, CEC '03,* Vol. 3, pp. 1988–1995, Dec. 8–12.

Kohn, M., K. Chausmer, and M.H. Flood, 2000. Anticipatory guidance about child safety misuse: Lessons from safety seat "check-ups." *Archives of Pediatrics and Adolescent Medicine* 154(6), 606.

Lane, J.C., 1994. The seat belt syndrome in children. *Accident Analysis, and Prevention* 26(6), 813–820.

Lee, J.W., K. Fitzgerald, and B.E. Ebel, 2003. Lessons for increasing awareness and use of booster seats in a Latino community. *Injury Prevention* 9, 268–269.

O'Day, K., 2001. Breakpoint's booster seat and burn awareness issue. Loyola University Health System, *Breakpoint Newsletter*, Feb. 4(2).

Ramsey, A., E. Simpson, and F. Rivara, 2000. Booster seat use and reasons for nonuse. *Pediatrics* 106(2).

Rivara, F.P., E. Bennett, B. Crispin, K. Kruger, B. Ebel, and A. Sarewitz, 2001. Booster seats for child passengers: Lessons for increasing their use. *Injury Prevention* 7, 210–213.

Safe Kids Canada, 2004. *Booster Seat Use in Canada: A National Challenge,* June; available at http://www.safekidscanada.ca/ENGLISH/IP_PROFESSIONALS/SafeKidsWeek2004/documents/EnglishReport_BoosterSeats.pdf. Accessed May 2005.

Simpson, E.M., E.K. Moll, N. Kassam-Adams, G.J. Miller, and F. Winston, 2002. Barriers to booster seat use and strategies to increase their use. *Pediatrics* 110(4), 729–736.

Statter, M., and T. Vargish, 1998. The spectrum of lap belt injuries sustained by two cousins in the same motor vehicle crash. *The Journal of Trauma* 45(4), 835–837.

Transport Canada, 2004. *Canadian Motor Vehicle Traffic Collision Statistics: 2003;* available at http://www.tc.gc.ca/roadsafety/tp/tp3322/2003/menu.htm. Accessed May 2005.

Weber, K., 2002. Crash protection for child passengers: A review of best practice. *UMTRI Research Review* 31(3), 1–27.

Wegner, M.V., and D.C. Girasek, 2003. How readable are child safety seat installation instructions. *Pediatrics* 111(3), 588–591.

Winston, F.K., D.R. Durbin, M.J. Kallan, and E.K. Moll, 2000. The danger of premature graduation to seat belts for young children. *Pediatrics* 105(6), 1179–1183.

Wong, D. 1999. *Whaley and Wong's Nursing Care of Infants and Children* (6th ed.). St. Louis, Missouri: Mosby.

Xu, J., Y. Gao, and G. Madey, 2004. A docking experiment: Swarm and Repast for social network modeling, in *Proceedings of the Simulation Conference*, Vol. 1, Dec. 5–8.

Zaza, S., D.A. Sleet, R.S. Thompson, D.M. Sosin, J.C. Bolen, and Task Force on Community Preventive Services, 2001. Reviews of evidence regarding interventions to increase use of child safety seats. *American Journal of Preventive Medicine* 21(4S), 31–43.

MARKET PENETRATION OF MORE SUSTAINABLE VEHICLES: THE HEV CASE

J.L. SULLIVAN,* C. STEPHAN, B. GOODMAN, and M. EVERSON,
Research and Innovation Center, Ford Motor Company, Dearborn, MI

ABSTRACT

Modern light-duty vehicles are desirable conveyances that are safe, secure, affordable, and comfortable and provide owners both utility and the freedom of access to destinations, both near and far. However, these vehicles consume considerable fossil energy and generate about a sixth of U.S. carbon dioxide emissions. With mounting evidence on climate change and recent record-high gasoline prices (an energy security issue), some combination of government policies, market incentives, or both are likely to become operative and thus lead to reduced petroleum consumption and concomitant carbon dioxide emissions. Automotive manufacturers are certain to be affected by such changes and hence need to develop more sustainable vehicles. But because the transition to more sustainable transportation is likely to be uncertain and risky, understanding the evolution of new advanced-technology product offerings to the auto marketplace would be of great benefit to vehicle manufacturers. We believe that agent-based modeling can be very helpful in this regard. This bottom-up approach, ultimately to be applied to the hybrid electric vehicle market, permits modelers to estimate the extent, rate, and robustness of the market penetration of more sustainable, advanced automotive product systems as a function of consumer sensitivity to fuel prices, consumer/agent product preferences, vehicle market sensitivity to fuel prices, and government actions (carbon taxes, fuel taxes, CAFÉ) motivated by petroleum consumption and carbon emissions. A detailed description of the model is presented, and preliminary results are also discussed.

Keywords: Agent-based modeling, auto market evolution

INTRODUCTION

Recent record-high gasoline prices and accumulating evidence on global average temperature increases and ice-cap shrinkage are reinforcing concerns about energy security and climate change and, as such, portend change in the economic and political spheres. While market forces for gasoline are likely to reduce consumption, the geopolitical and environmental dimensions of these issues are prone to spawn new energy policy initiatives that will reinforce the effects of current energy market trends. In such an environment, consumer use of vehicles is apt to change in a number of ways, including increased car pooling, reduced frequency of vehicle use, reduced miles driven, and ultimately changes in the kinds of vehicles demanded. Understanding potential shifts in consumer preferences for vehicles and their timing, magnitude, robustness, and rate is clearly of value to an auto manufacturer, particularly as it pertains to the introduction of more sustainable advanced vehicle technologies, like hybrid electric vehicles (HEVs), into the auto market.

---

* *Corresponding author address:* John Sullivan, Physical and Environmental Science Department, Research and Innovation Center, Ford Motor Company, 2101 Village Road, Dearborn, MI, 48124; email: jsulliv8@ford.com.

A number of business modeling approaches facilitate better understanding and characterization of the marketplace. These include optimization (of an organization's objective function) and discrete event simulations. Unfortunately, these approaches assume consumer behavior is fixed and do not allow for consumer adaptation, an important consideration in any evolving process, such as a marketplace. System dynamics (SD) modeling, a "stock and flow" approach, has also been widely and successfully applied to a number of business-related questions (Sterman 2000). Some have even been applied to the diffusion of new technology into the marketplace (Bass 1969). SD does capture the influence of feedback, a critical process to capture in the characterization of any complex system. When used with Monte Carlo methods, a wide range of scenarios (changes in the course of the economy, fuel price changes, etc.) can be explored for their impact on product sales and profitability. Nevertheless, because SD considers consumers in large blocks and not as individual entities, important consumer feedbacks get missed. For example, low-income agents with large vehicle miles traveled (VMTs) would have to reduce miles driven when faced with higher gasoline prices, while others would not. We feel a better approach to characterizing marketplace transitions is agent-based modeling (ABM). In ABM, agents interact on the basis of assumptions about their characteristics, objectives, interests, and behavioral and decision rules. Hence, the method is aptly suited to represent the interactions of a diverse and heterogeneous population of consumers, manufacturers, regulators, and fuel producers and to follow how they play out in an evolving marketplace.

Agents are virtual decision-making entities in software. What makes agents so interesting is that they can demonstrate emergent behavior (i.e., aggregate or macro behaviors that are surprising and cannot be readily explained on the basis of individual behaviors). ABM has been used for a wide range of applications. It has led to insights into human and social behavior (Axelrod 1999), including interpretation of organizational behavior (Carley 2000). It has also been applied to long-term policy analysis (Lempert et al. 2003), especially under conditions of deep uncertainty. For transportation, ABM is becoming increasingly employed. There has been considerable application of ABM to traffic modeling simulations, particularly traffic congestion. For a sampling of that work, see a special issue of Transportation Research (2002).

Of most interest to us are applications of ABM to market-related questions. Macal and North (2002) used ABM to simulate the interactions, co-evolution, and response to market conditions and physical disruptions of the interdependent natural gas and electricity infrastructures. The behavior of their agents, representing the two infrastructures, included agent objectives, bidding and pricing strategies, and ability to learn and adapt in an evolving market. In a separate paper, North et al. (2002) detail their approach to simulating decentralized electricity markets. This ABM employs many diverse and heterogeneous agents representing generating companies, demand aggregating companies, transmission companies, consumers, system operators, and government regulators, each having its own business strategy, objectives, and decision rules. After demonstrating good correlation between the bidding behaviors of their model to those of a live six-player market simulation game, they presented some preliminary results to illustrate the model's capabilities. For example, utility price volatility was much lower for pay-as-bid pricing for power than for pay-as-MCP (market clearing price) approach. The former was found to benefit consumers, but producers lost money. Agent-based simulations have also been applied to the financial markets. For example, LeBaron (2000) found that in a population of investor agents, as investing strategies become homogeneous across the agent population, the risk of market crashes increased considerably, especially near sharp price declines.

One of Ford's interests in ABM is its use as a tool to anticipate shifts and evolving trends in the automotive sector, particularly to the introduction of more sustainable automotive products into the marketplace. There is clearly a business case for developing such a tool, not only to anticipate market shifts defensively but also to identify new business opportunities. Recently, Stephan and Sullivan (2004, 2005) have applied ABM to the evolution of the hydrogen infrastructure. It was found that the evolutionary process is a classic "chicken and egg" problem dependent on the initial hydrogen filling station density and distribution relative to area population concentrations (urban and other areas), fuel subsidies, vehicle tax incentives, and driver agent comfort with a new technology. In the model, when it came time for drivers to buy a new car, the perception of fuel availability was important, as well as colleague and neighbor buying practices. The model was a two-agent model, including vehicle owner and fuel station owner agents. In a U.S. Department of Energy (DOE) research project, this work is now being extended to include more agent classes.

The purpose of this paper is twofold: to (1) to introduce our model recently developed at Ford to characterize the evolution of the HEV product system into the U.S. automotive market and (2) present some preliminary results. A detailed description is given of the agents, their attributes, and decision rules. For an exogenous increase in fuel price, preliminary results are presented for fuel sales, fuel revenues, VMT, and vehicle sales. For this set of conditions, who buys which vehicle is also discussed.

THE MODEL

The purpose of this model is to explore the influence of market factors on the evolution of more sustainable, advanced-technology vehicles into the U.S. auto market. A current example is HEVs. Of particular interest is the influence of exogenous shocks to the system, particularly their influence on vehicle sales, fuel sales, carbon emissions, fossil fuel consumption, and miles driven. The model also monitors who buys which vehicles during transitional periods following shocks. The model is an agent based simulation comprising four agents: vehicle owners, fuel suppliers, government, and original equipment (vehicle) manufacturers (OEMs). A depiction of their interactions is shown in Figure 1. Lists of agent characteristics that figure into their decision rules are given in Table 1. In principle, many of these characteristic can change endogenously during a simulation, although at this time, only a few do. While the vehicles provided to the market by the OEMs have numerous characteristics as listed in the table, they are, nevertheless, a data type and not agents. At this time, we consider three OEMs, each producing three vehicles, and one fuel producer. One central dealer of used cars, assumed here to be a special instance of the OEM class that does not make cars, sells all the used cars. Fuel availability is assumed ubiquitous. Notice that the lists in the table for government appear to be more like actions than characteristics. In the context of our model, we assume the relevant government characteristics are to monitor and plan.

The model has been written in objected-oriented Fortran and developed in the Visual Fortran Design Studio on a desktop computer. A more complete description of the model will appear in forthcoming publications, but a brief description is appropriate here. All owner agents are assigned an income on the basis of U.S. income distribution. On the basis of transportation statistics, a certain fraction of that income, depending on magnitude, is devoted to transportation. It is divided into three categories: variable cost (fuel expenditures), fixed cost (vehicle

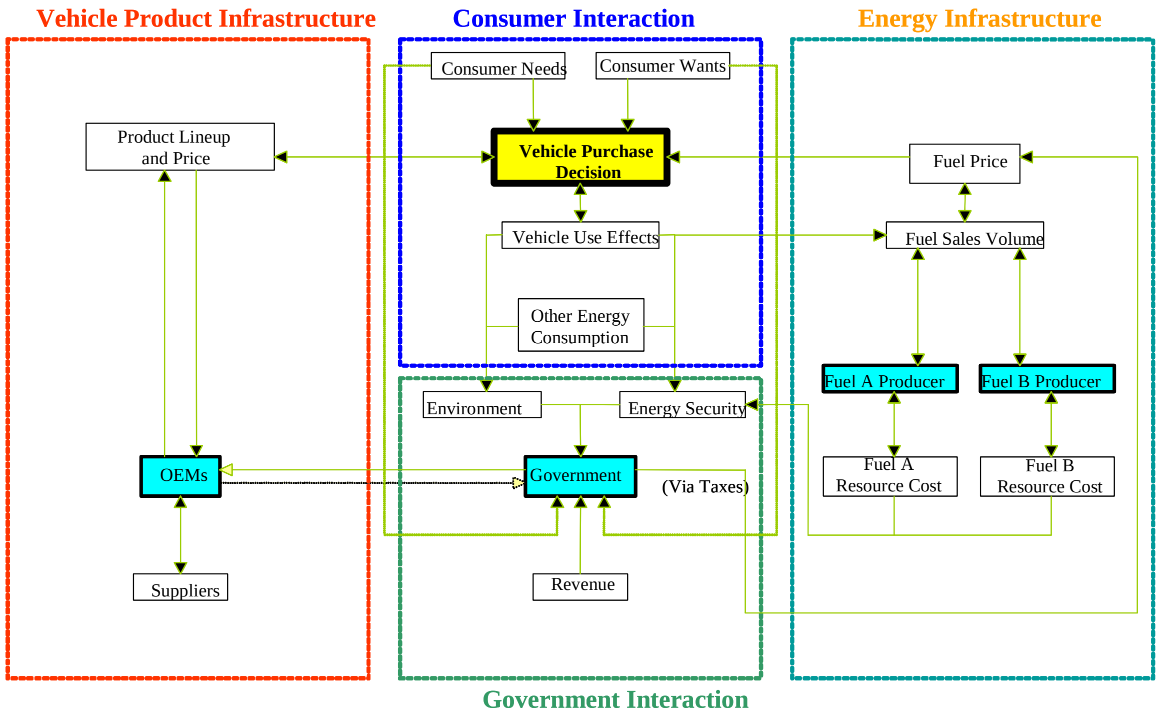

FIGURE 1 Agents and their interactions

depreciation), and other expenses (insurance, maintenance and repair, etc.). Owners live in neighborhoods commensurate with their income (low, middle, and upper) located in either urban or suburban regions or one of two townships, all found on a 100 × 100-mile grid (Figure 2). Agents live and work on the grid and drive for commuting, errand, and leisure purposes. Each OEM sells three different vehicles, covering a comparatively wide range of vehicle sizes, performances (acceleration times from 0–60 miles per hour [mph]), fuel economies, and prices. When it is time to get another car, owners can buy either new or used cars, depending on their new/used preference characteristic that remains unchanged. Some "owner" agents have incomes so low that they cannot own a vehicle. Instead, if these agents live in the urban region, they take the bus; otherwise, they carpool, though this option has yet to be implemented in the model at this time. The actions of these agents, who represent approximately 10% of the owner agent population, must be considered, as there are fossil fuel consumption and carbon dioxide (CO_2) emission burdens associated with them.

For purposes of compact notation, we henceforth denote owner actions as OAi, where the "i" indicates an action based on a particular set of decision rules. In this report, owners respond directly and indirectly to a fuel price increase (i.e., they change the amount of fuel purchased or perhaps eventually buy more-fuel-efficient vehicles). A listing of various owner actions and those of the other agents is given in Table 2.

Other agent actions include those of government, fuel producers, and OEMs. For this initial exercise of our model, we assume the government agent monitors fossil fuel and carbon flows only. It does not act, although it could establish fossil fuel consumption and carbon

TABLE 1 Characteristics (factors) on which agent behaviors are based

Owners	OEMs	Government	Fuel Suppliers	Models
Income	Number of vehicle offerings	Compute fleet fossil fuel use	Number of fuels provided	OEM_ID
Transportation budget	Sales performance	Compute fleet CO_2	Revenues	Model_ID
Vehicle Age_b4_sell	Reputation		WTW[a] fossil burden per gallon	Size
Mileage_b4_sell	Profit		WTW CO_2 per gallon	Performance
Commute distance	Fleet CAFÉ		Feedstock cost	Vehicle segment
Errand distance	Market share		Set fuel price	City fuel economy
Leisure trip distance	Development time to introduce a new vehicle			Highway fuel economy
Fuel demand elasticity				Mileage
Home address				Age
Work address				Price
Preferred car size				OEM profit per unit
Preferred car performance				Share of segment
Response time to fuel price increase				
Delay time in buying a new car because of fuel price increase				
New or used vehicle				

[a] WTW denotes well to wheels.

emission targets and/or impose a number of policy options, including levying carbon taxes, fuel taxes, gas guzzler taxes, or new CAFÉ regulations. Such actions could be triggered exogenously or endogenously — an exercise to be explored in a future application of the model.

At this time, the only fuel producer action, FAi, considered here is passing fuel price increases due to supply shocks, dollar for dollar, through to the consumer. In future exercises of the model, the influence of alternative fuels will also be examined. For example, could ethanol blended with gasoline offset the effect of rising crude oil prices on fuel prices? Or could ethanol have a significant influence on consumer vehicle purchases if appreciable carbon taxes were in effect? In addition, future applications of the model will address fuel demand shocks on the auto market.

TABLE 2 Partial list of agent actions

Label	Purpose of Action	Details of Action
OA1	Purchase gasoline	Purchase gasoline to meet mileage needs. If gasoline price changes, adjust gasoline purchase by changing errand mileage as per agent's fuel demand elasticity, ranging between –0.4 and –0.1.
OA2	Purchase car	If present car's age and mileage exceed those of owner's target values, buy another car. Depending on owner income, survey available used or new vehicles. Select two vehicles that are the closest to budget neutrality for operating and ownership cost targets, one over and the other under target, and choose the one that most closely meets or exceeds owner's vehicle size and performance desires.
GA1	Affect change in fossil fuel consumption	On the basis of fossil fuel consumption and carbon dioxide emissions when compared to national targets, if observed levels exceed targets, do nothing.
FA1	Set fuel prices	If a supply shock occurs, pass price increase through to consumers, dollar for dollar.
VSA1	Set vehicle prices	If market share of vehicles shifts in the used or new car market, adjust prices to bring market share back into balance.

Vehicle manufacturers sell vehicles at market prices. They monitor vehicle sales, market shares, and change prices as appropriate. Generally, they seek to maximize market share and profits. If market trends suggest a market shift in consumer preferences, OEMs can elect to introduce new vehicle lines, though a time lag before introduction is incurred as a result of design and manufacturing considerations. The lag varies from one OEM to another. Actions taken by the OEMs are denoted VSAi (i.e., vehicle sales actions). Each vehicle has two owner cycles: one as a new car and another as a used car. After the first cycle, the vehicle goes to a used car lot, where, depending on its mileage, it is priced at 40% to 60% of the new car price. When an owner buys a used car, his or her current vehicle, assumed to be a used car completing its second owner cycle, goes to scrap.

RESULTS AND DISCUSSION

The intention of this article is to demonstrate the potential of the model to address the market penetration of more sustainable, advanced-vehicle technologies. Toward that end, we explore the effect of an exogenous shock to fuel prices on key market indicators, including vehicle sales, fuel sales, miles driven, fuel revenues, and who buys which vehicles. Because the model has yet to complete all verification steps and be calibrated and validated, actual application of the model to characterizing the market penetration of HEVs or any other advanced vehicle technologies must await discussion in future reports.

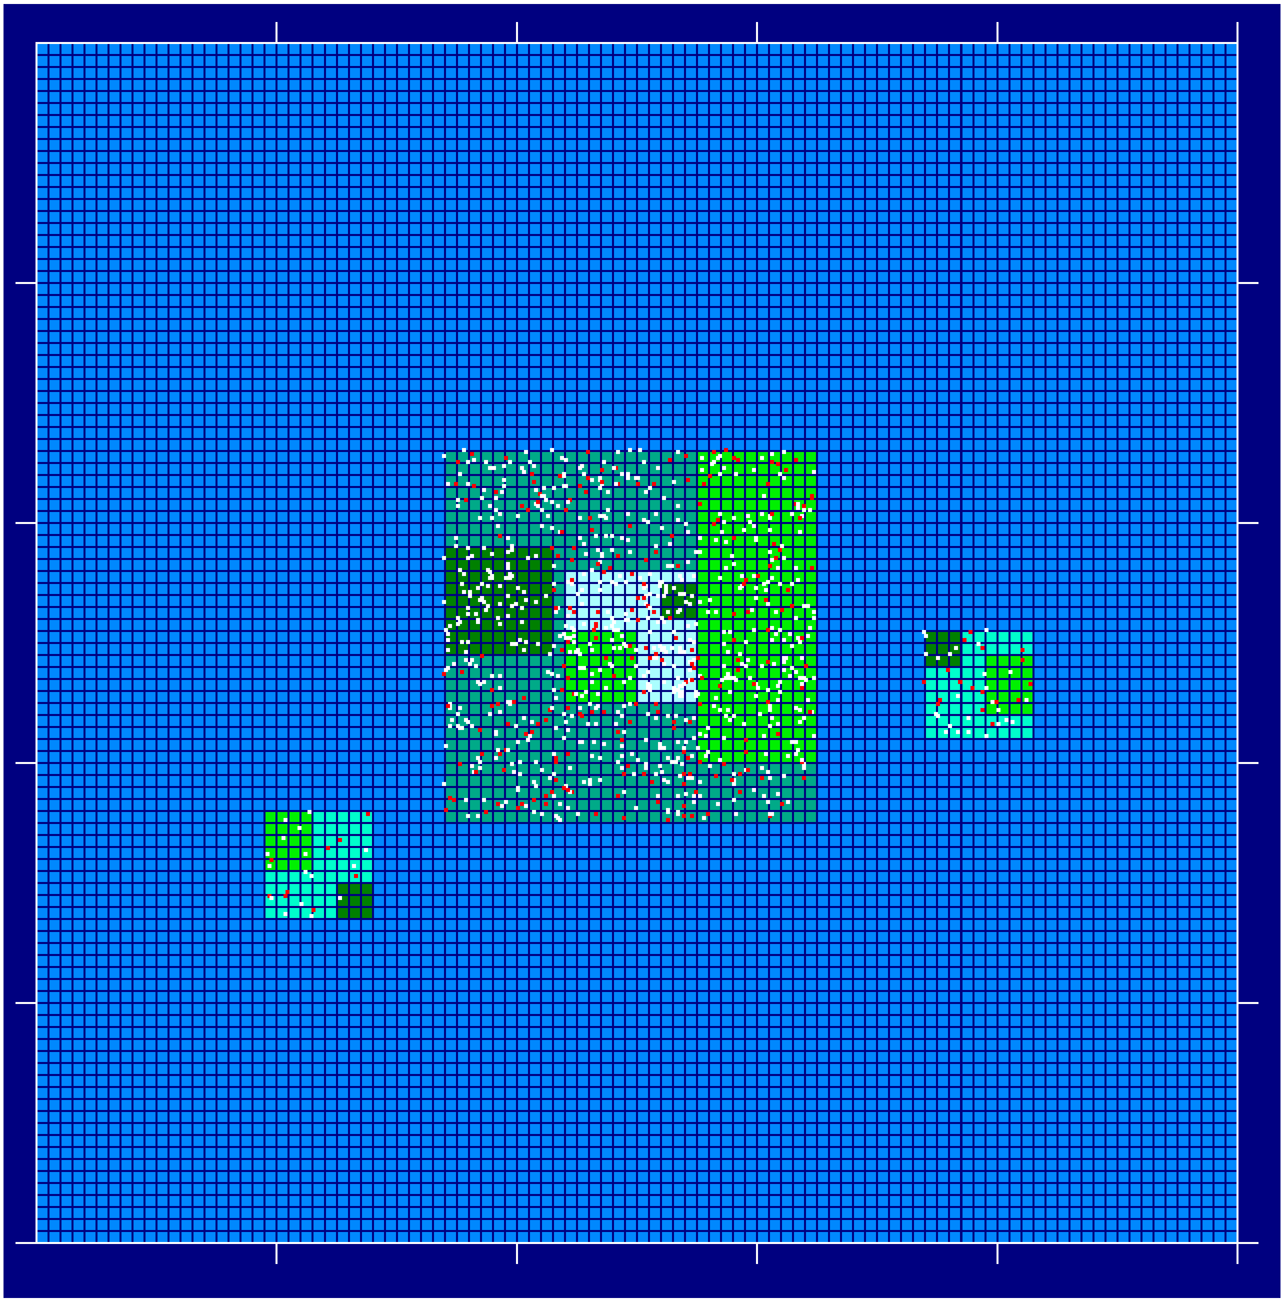

FIGURE 2 HEV world showing urban (center), suburban (ring region), and satellite towns (Incomes of areas are light green for lower levels, teal and light gray tones for middle levels, and dark green for upper levels. In this case, red agents, approximately 240 of them, are those that recently purchased a more-fuel-efficient vehicle.)

Because it is difficult to anticipate exactly the economic and market conditions that lead to a stable market in the absence of any exogenous shocks, we found it necessary to allow the model to run for 100 cycles as a conditioning step to reach equilibrium before applying a fuel price increase. Each cycle represents one month in the model world. For our results, we apply decision rules OA1, OA2, and FP1; OEM and used car dealers do not respond. As seen in Figure 3, in the absence of a price shock, the model shows no consistent change or drift in fuel revenues, although the small-amplitude long period variation observed in the figure is typical. The reason for this periodicity is unknown. Fuel sales in total gallons are not seen in the figure, as they are is superposed by miles driven. With no change in fuel price, miles driven must remain constant, as per the fuel demand elasticity assumption employed here.

Figure 4 shows the aggregate response (fuel revenues, fuel sales in gallons, and miles driven) of about 900 vehicle owners (out of 1,000 consumer agents; some cannot afford cars) to a $0.50 increase in gasoline price from an initial value of $2.00 per gallon. There are a number of features in the figure that merit comment. First, notice that because of the fuel price increase, fuel revenues rise abruptly and unchecked for the first month. This is due to a 1-month time lag before consumers (vehicle owners) begin to respond. After the first month, fuel revenues drop off steeply for several months. The reason for this can be seen in gasoline sales trace. The price increase in fuel has caused vehicle owners to use less fuel by driving fewer errand miles. per each owner's short-term demand elasticity for fuel, assumed here to range between –0.4 to –0.1. Depending on the owner agent, it takes from 1 to 3 months to respond to the fuel price increase and change driving behavior. Thus, the dropoff in fuel sales (gallons) and miles driven leads to a reduction in fuel revenues. Changes in total vehicle miles driven occur only in the first few months after the shock and remain constant thereafter.

Following the steep decline in fuel revenues and sales, a more gradual decline in fuel sales and revenues ensues before the decline finally levels out more or less at around 40 months. The gradual decline is due to owners shifting to more-fuel-efficient vehicles. When the time comes for their next vehicle purchase, they consider fuel costs, the time elapsed since the fuel price increase, and their vehicle preferences. An inspection of Figures 5 and 6 shows a shift to the most-fuel-efficient cars (vehicles 8 and 9), which cost less to own (lower price) and operate (better fuel economy) than the fuel-inefficient vehicles. Despite its comparatively good fuel economy, not many owners buy vehicle 7, almost certainly because of its higher price. Hence, both the growth of vehicles 8 and 9 in the on-road population and the more or less stagnant growth of vehicle 7 appear to be economically driven. Figures 5 and 6 also show that the vehicle model distribution of the on-road fleet has decreased for vehicles 4 through 7 and stayed about the same for vehicles 1 through 3, although the used vehicle fraction of these two groups has decreased for the former and increased for the latter. It appears that, despite higher gas prices, some used car buyers can afford to satisfy their preferences and purchase the more expensive used cars (models 1–3) that cost more to operate. A more detailed discussion of which agents buy what cars awaits a future report. However, some insight into this question is evident in Figure 2. There it is seen that about 240 owners (red dots) changed to more-fuel-efficient, less expensive vehicles in the course of the simulation. None of them live in the higher-income neighborhoods, suggesting that economic considerations for the next vehicle purchase of some owners are not forcing them to compromise on their vehicle size and performance preferences.

Incidentally, because of stochastic elements of the model that affect owner income, types of miles driven, vehicle preferences, and fuel demand elasticity, the dependencies shown in Figures 3–6 vary somewhat from run to run. Nevertheless, the results presented here are typical.

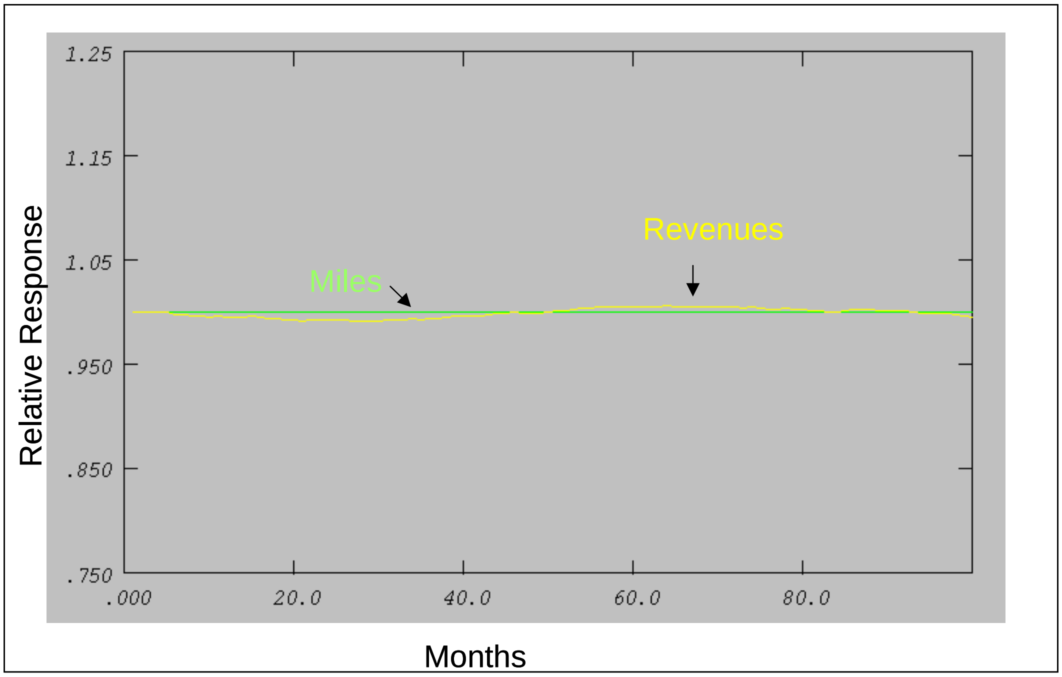

FIGURE 3 Miles driven and fuel revenue ($) responses of agents to a regime with no fuel price increase relative to values just before a fuel price increase (Vehicle makers and dealers do not respond.)

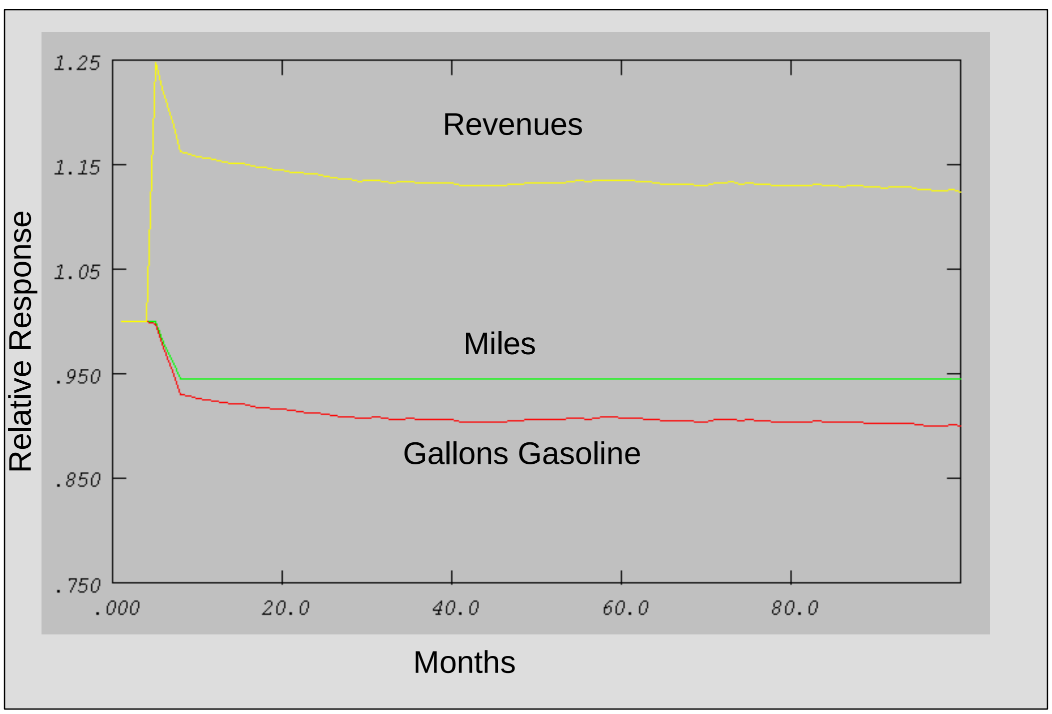

FIGURE 4 Miles driven, fuel revenue ($), and fuel sale (gallons) responses of agents to a fuel price increase relative to values just before a fuel price increase (Vehicle makers and dealers do not respond.)

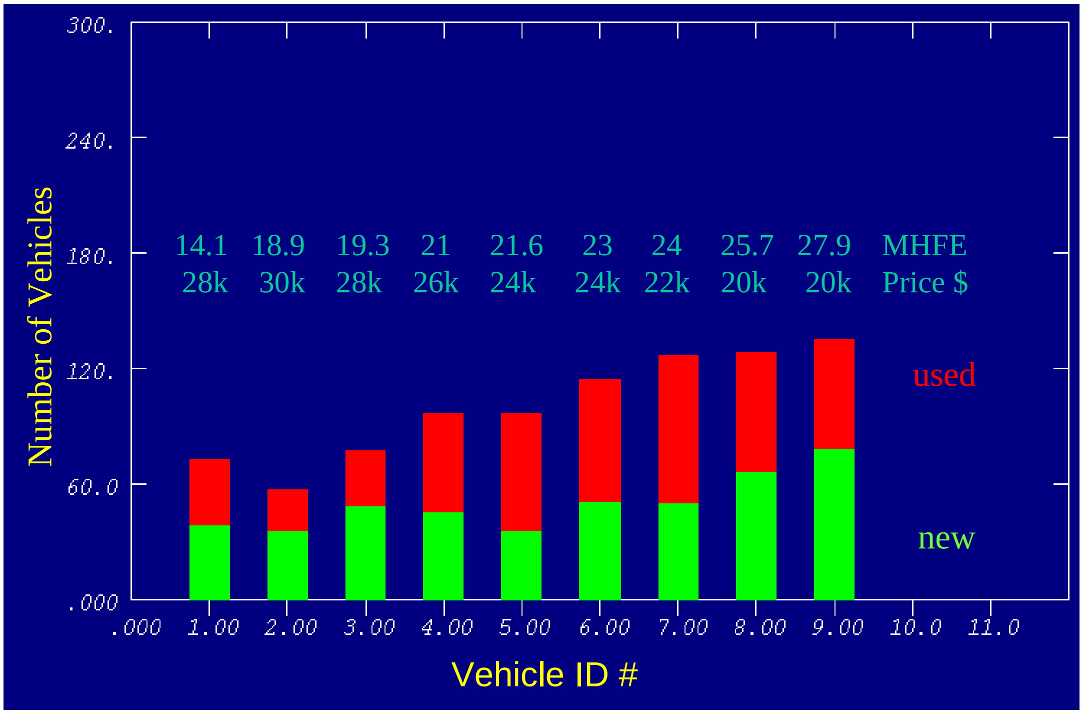

FIGURE 5 Initial fleet vehicle distribution over the nine available models after a conditioning cycle (Vehicle fuel economy and new vehicle prices above.)

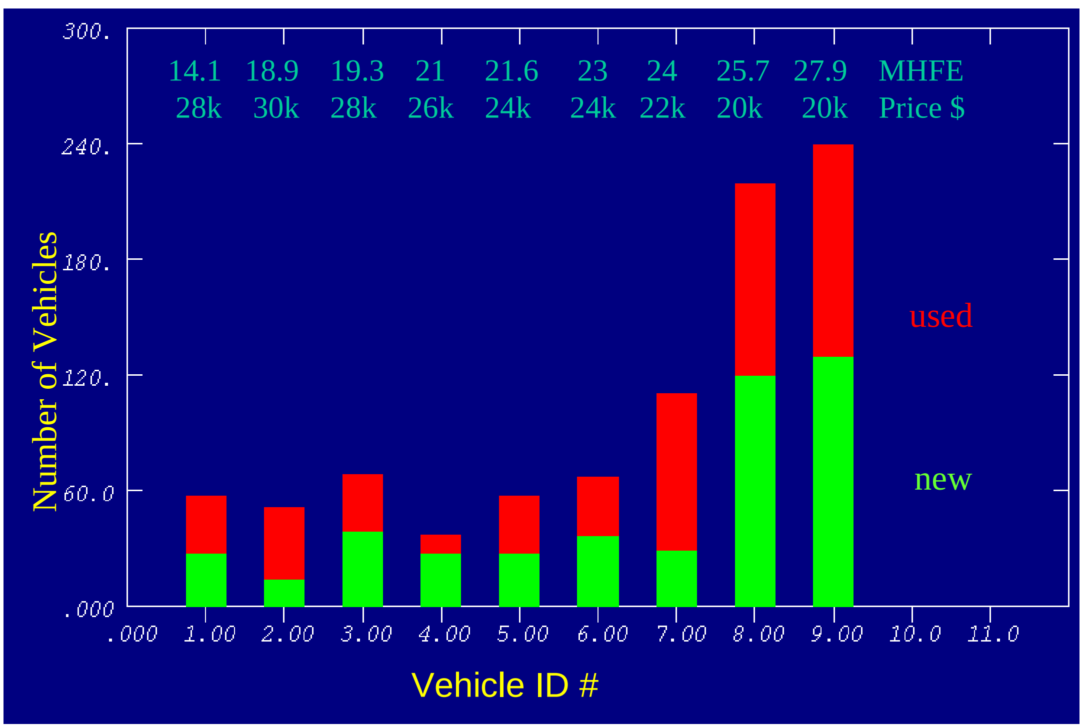

FIGURE 6 Final vehicle fleet distribution 100 months after an exogenous fuel price shock to the system (Vehicle fuel economy and new vehicle prices above.)

As a verification step, the observed short-term gasoline demand elasticity is –0.24, which is in good accord with the average of the randomly distributed elasticity values (that range from –0.4 to –0.1) assumed herein for owner agents. This value is computed from gasoline sales (gallons) just after VMT levels out (see Figure 4), about 4 months after the fuel price shock. But when long-term fuel price elasticity is considered, which includes both changes in vehicle purchasing preferences and miles driven, we get a revised elasticity of –0.44. This value, computed from fuel sales at the end of the simulation, emerges from our model and is directionally consistent with long-term elasticity values reported in the literature (Graham 2002). Graham and Glaister (2002) report short-term demand elasticity (up to 1 year) in a range of about –0.3 and long-term values (2 to 10 years) in a range of –0.6 to –0.8. Their long-term values, developed from an extensive review of the literature, are attributed to continued changes in driving behavior and changes in vehicle fleet fuel efficiency and vehicle population numbers. Admittedly, a more detailed comparison of behaviors of our owners with those of actual vehicle owners needs to be done.

CONCLUSIONS

A model to characterize the evolution of markets for advanced vehicle technology has been described, and preliminary results have been presented. The model, to the extent it has been tested, behaves as expected. That is, if fuel prices increase, owner agents buy less fuel and drive less, and some buy more-fuel-efficient vehicles. Although the features represent endogenous behavior in the model world, they are nevertheless reminiscent of the real world. But before model results can be taken as truly representative of the real world and thus used to guide product development strategy, more model verification, calibration, and validation are needed.

It is clear that more complicated interactions and factors need to be added to the model; these include leasing, brand loyalty, and consumer switching between purchasing new and used vehicles versus leasing.

REFERENCES

Axelrod, R., and M.D. Cohen, 1999, *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, The Free Press, NY.

Bass, F., 1969, "A New Product Growth Model for Consumer Durable," *Management Science* 15:215–227.

Carley, K., 2000, "Computational Social Science: Agents, Interaction, and Dynamics," in D. Sallach and T. Wolsko (Eds.), *Proceedings of the Agent 2000 Conference on Simulation of Social Agents: Architectures and Institutions,* ANL/DIS/TM-60, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 6–7, pp. 3–29.

Graham, D.J., and S. Glaister, 2002, "The Demand for Automotive Fuel: A Survey of Elasticities," *J. Transport Econ. and Policy* 36:1–26.

Lempert, R.J., S.W. Popper, and S.C. Bankes, 2003, *Shaping the Next One Hundred Years: New Methods for Quantitative Long Term Policy Analysis*, Rand Pardee Center, Santa Monica, CA.

LeBaron, B., 2000, “Financial Market Efficiency in a Coevolutionary Environment,” in *Agents 2000, Simulation of Social Agents: Architectures and Institutions,* p. 33.

Macal, C., and M.J. North, 2002, “Simulating Energy Markets and Infrastructure Interdependencies with Agent-based Models,” in C. Macal and D. Sallach (Eds.), *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution,* ANL/DIS-03-1, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 11–12, p. 195.

North, M., C. Macal, R. Cirillo, G. Conzelmann, V. Koritarov, P. Thimmapuram, and T. Veselka, 2002, “Multiagent Modeling of Electricity Markets,” in C. Macal and D. Sallach (Eds.), *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution,* ANL/DIS-03-1, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 11–12, p. 215.

Stephan, C., and J. Sullivan, 2004, “Growth of the Hydrogen Transportation Infrastructure,” in C.M. Macal, D. Sallach, and M.J. North (Eds.), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 7–9, p. 731.

Stephan, C., and J. Sullivan, 2005, “An Agent-based Hydrogen Vehicle/Infrastructure Model,” in *Conference on Evolutionary Computation,* pp. 1774–1779, Portland, OR, June 19–23.

Sterman, J.D., 2000, *Business Dynamics: Systems Thinking and Modeling a Complex World,* Irwin/McGraw Hill.

Transportation Research, 2002, “Part C: Emerging Technologies,” Vol. 10, Issues 5–6, pp. 325–527, Oct.–Dec.

DISCUSSION

Social Simulation Applications

(Public Policy, Friday, October 14, 2005, 10:00 a.m.–noon)

Chair and Discussant: *Pam Sydelko, Argonne National Laboratory*

Simulating Water Usage during Uncertain Times in the Southwestern United States: An ABM of Strategies and Population Level Actions

Pam Sydelko: It's great that Dr. Epstein's talk set us up for using agent-based modeling for policy. In the years I've been doing these agent workshops, there has been an increase in the number of applications in policy, which is not surprising. What is surprising, though, is the increase in the amount of funding for policy-related work, which I think is encouraging. That means federal agencies and other agencies are starting to see the importance and learning how valuable this kind of modeling could be for policy.

I'd like to introduce Bill Griffin from Arizona State University. We're going to be looking at very important policy applications for water use in the Southwest.

William Griffin: Those of you who have heard me talk before know I haven't ventured into studying water. I don't know a lot about water, and I know less about policy. So Shana Schmidt will talk about that in the second half of this presentation. My area is modeling micro-social interactions in people. Several years ago, the National Science Foundation had a call out for large infrastructures looking at decision making under uncertainty. Arizona State, which has a long history in ecology and land management, decided they would submit a grant. I got a call one day asking for a modeler, specifically for modeling water. I said that I didn't know anything about water, but I could do something about people. They thought that was close enough. And so I joined 37 other scientists, and we submitted this proposal. We were one of three locations that ended up getting a center. We call ours the Decision Center for Desert City — Arizona State, Columbia, and Carnegie Mellon.

[Presentation]

I'm going to turn this over to Shana because she's much more versed in the policy implications of this.

[Presentation Continues]

Shana Schmidt: These people probably work 10+ hours a day, and the last thing that you're thinking about after getting home after a 12- or 14-hour day is not going to be about turning the faucet off while you're brushing your teeth.

Unidentified Speaker: Did you find there was any correlation between the people in groups of one, two, and three?

Schmidt: Sometimes, but not always. We found differences in both. In the water data that we have, we're not able to draw any conclusions about it, officially, but it goes both ways.

Unidentified Speaker: …[Unintelligible]

Schmidt: They are for this particular study. Again, we haven't done this with our own data, so hopefully we'll get that done within the next few months.

[Presentation Continues]

Sydelko: Thank you. We have a couple minutes for some questions.

Craig Stephan: Craig Stephan, Ford Motor Company. I think you mentioned that agriculture used 49% of the water, and yet you didn't seem to address their usage at all in your plan.

Griffin: We're implementing the model in stages. One of the things that we decided as a group was to go after homeowners first because there are some intricacies about agricultural use that I'm not familiar with. We have some individuals on our group that are, but water use is very different from what I understand. Its pricing structure is very different, and that's another very different political arena to play in.

If you stay in the urban water-use group, we have, for example, single-family dwellings, commercial-use properties, multi-family dwellings, and those types of things; those are currently more trackable from a modeling point of view. That involves less political tampering. Once you move to agriculture, water use becomes very different. It's priced very different, and, by the way, in Arizona, there are dozens of contracts among the federal government and the Indian tribes and the State of Arizona. Some of them have been in court for over 60 years, and water is tied to all of those. One of the cases goes back to 1890, which allowed building of the Hoover Dam. We thought for the first year or two that we should just model households. That's the simplest way to explain it. It gets very complicated.

Sydelko: We have time for one more question.

Kostas Alexandridis: Kostas Alexandridis, from Purdue. I was wondering if you're using the long form, the public-use micro-data sample from the census. Are you using the proxy versus the number of planning facilities and number of kitchens and bathrooms in the household?

Schmidt: We have it, but we haven't implemented it into the model yet. So, yes, to all of those things. The renter/owner will go in first as the binary variable. Then it will go into number of bathrooms and how many we had, down to the census block level. Of course, we know how many households are in that area and how many family members are typically in that area. We also have data on the kind of structure of the homes as well. So, yes, it's all going to go in long term.

Griffin: We're laying the GIS information on top. Getting the census data is not a problem. Getting the name from a water consumption bill becomes a problem because of the confidentiality. Some WPAs and their adjacent WPAs will give us some information; others will not give anything. So we don't know the price of water and how much was consumed. Or we know how much was consumed, but we don’t know the price. Or we know the price, but we don’t know how much was consumed.

Sydelko: Okay, thank you very much.

Using Agent-based Modeling to Better Understand Local Heroin Dealing Organizations and Drug Markets

Sydelko: Our next speaker is Lee Hoffer from the Washington University School of Medicine. He has an interesting application in looking at the implications of heroin dealing.

Lee Hoffer: I’d like to start by thanking the organizers of Agent 2005 for giving me an opportunity to talk to you today. I’m rather new to the area of agent-based modeling, and I’d like to discuss an experimental research project in which we are developing a laboratory of agent-based simulations of a local heroin-dealing organization and market using an ethnographic research data set, which I’ll talk about in a moment.

However, there are two additional aims that I’d like to mention. The intent of the project is to develop simulations so we can learn more about the structure and function of organizations and drug markets. Another purpose is to inform policy.

[Presentation]

Sydelko: That’s great. We have time for one question.

Steven Bankes: It strikes me that using ethnography to specify micro-level rules is potentially very important. I’m imagining that in some future destination of the field, you’ll be a prototype for people who are interested in this area and who are ethnographers who go and write models to try and understand what they know. In the near term, we’re most likely to have people who are modelers who aspire to do a modeling area, who want to use ethnographers or ethnography to inform their work, so I’m curious. In your opinion, do you think that what you’ve done is going to be viable for them? Can they say that they want to build a model and then go read some ethnographies or talk to some ethnographers to use as a source? Or is that going to tend to be incomplete, and they’ll find out that what’s required is a reverse relationship where, after having done some modeling, they realize they need to know how people respond to advertising or some other crucial part of their model? They might then ask an ethnographer to do additional research to obtain that data for them. If you’re writing a grant proposal, would you include that second phase?

Hoffer: I think you raise an important point, and we can use epidemiology as a great model for your question. Epidemiology in HIV prevention has been very interested in the sharing of syringes. That’s an enormous risk behavior. In this very large NIDA-funded study, it was multiple sites, 16 or 17 cities. In their survey, they asked drug users if they ever shared syringes, and they answered that they hadn’t. I mean, the rate was pretty low in most cities.

They funded a number of studies to try to find out what was going on, and they used ethnographers to do the local stuff and look at the exchange behaviors. What they learned — well, what *we* found, actually because I did one of those studies in Denver — was that drug users don't talk about drug sharing. They don't admit to sharing. What they talk about is giving their syringes to someone after they're done with them.

I think that, as a result, they've integrated ethnography and are trying to understand how their data can be interpreted because epidemiology — survey data — are not … well that data can't tell you much about a process, a social process. I think the best studies have combined methods. That's how I see it with agent-based modeling. If you want to create a model, you need to talk to, or interview, subject experts and whatnot, but you also need to think about putting someone out there who is trained in using these methods. There're a lot of methods that can be brought to bear here where anthropologists create decision trees, and there's a whole protocol for analyzing that data. I think all those methods can be used in advance of going back and saying, "Hey, we're missing something here."

So I'd like to see them work together with agent-based modelers. Mike Agar, for example, is also working on a model that's doing things that are similar to my work.

Multi-agent Model Prototype for Child Vehicle Safety Injury Prevention

Sydelko: Our next speaker is Anne Snowdon, who will speak to us about child safety, vehicle safety, and prevention.

Anne Snowdon: My clinical background is critical illness in children, and, of course, one of the reasons children arrive in intensive care units is because of major trauma as the result of road crashes. Today I'd like to share with you the magnitude of that problem and tell you that it is growing; I'm speaking of the methodology, of course, of agent modeling. This is a very preliminary prototype of trialing that methodology to provide greater insights than we have to date on the issue of injury prevention. I will share with you some of our early observations and conclusions and give you a sense of where we think we are going from here.

[Presentation]

Sydelko: Thank you, Anne. We have time for a question.

Robert Reynolds: You focus on the individual drivers or agents, but public transportation is also an option. Certainly, as gas prices rise, it's becoming a more viable option for people. Are you looking at incorporating public transportation into your model? You also mentioned distractions, but I didn't really see anything about the environment. Certainly, there are cycles in the environment and those issues. How are those incorporated into your model?

Snowdon: We haven't begun to do either of those, but the public transportation issue is an interesting one because by and large the rate of child deaths and injuries is quite low in public vehicles. Canada doesn't have the number of large cities with subway systems, for example, as other countries. If you don't live in Montreal or Toronto, you're not riding public transportation systems because our communities are so far apart. They tend to be less viable.

Many people ask me why we aren't doing this with school buses. We do have school bus accidents. We just don't see the injury data and the severity we see in the family vehicles. So it's more of an issue of priority, not one that we wish to ignore, but certainly one that we'll look to building in later development.

As for environment, you bring up a great question. This area is probably the one I'm most interested in getting some input on because the vehicle/driver interface is a very complex environment. But it's embedded in the road infrastructure, weather conditions, traffic, etc., environment as well. For example, there's some research that talks about the number of accidents increasing dramatically at intersections, specifically at left-turning lanes. We need to be building in some of that data we have in the Transport Canada data or in the provincial data around those environments and how they influence. The complexity's huge, and as I start to have a better confidence with understanding agent modeling, in general, I hope to accomplish that. I'm grateful for your thoughts and suggestions.

Sydelko: Thank you very much, Anne.

Market Penetration of More Sustainable Vehicles: The HEV Case

Sydelko: Our next speaker is John Sullivan from Ford Research. Keeping in the vehicle theme, we're now going to be looking at sustainable vehicles.

John Sullivan: Thank you. It's a pleasure to be here today to talk about an area of application of agent modeling that pertains to the automotive market. I'd like to talk about market penetration of more sustainable vehicles. Our ultimate aim is to characterize the penetration of the hybrid electric vehicle systems into the U.S. vehicle market. The authors are myself, Craig Stephan, who was here in this forum last year and talked about our hydrogen infrastructure model and characterization of the evolution of that infrastructure, and two members from our Business Modeling Group at Ford Research.

By way of a quick agenda, first of all, I'd like to give you a notion of the company's motivation in doing this work, make the connection with complex systems, and describe how we think it could help us in addressing some of these challenges. Second of all, I'll briefly describe the model and present some preliminary results. Finally, I'll give a few concluding remarks.

[Presentation]

Sydelko: Thank you, John. I have time for one question.

Kostas Alexandridis: I find this presentation really good for emerging a major issue that appears to be in every discussion of sustainability. I'm working with change, and at every stakeholder meeting and in every community I've seen an apparent difference between sustainability goals and economic development. I think the most successful sustainable goals were coupled with the worst scenario for the business profit sector. Do you have any thoughts on that?

Sullivan: Well, first of all, if you look at the definition of sustainable development and the so-called three legs of the stool that it all rests on, which are economics, social equity, *and* environment, you'll see that solutions that lead to failure economically are just unacceptable. What we're looking for then are solutions that manage to do all of those things as best as possible. Clearly, there'd be some trade-offs, but no one is going to sign up for prescriptions for market failure. Do I understand your question?

Alexandridis: Yes …

Charles Macal: Since we're a little behind schedule, we're going to change the schedule slightly and have lunch. Zhian Li has graciously agreed to give his presentation entitled, "An Agent-based Model for Simulation of the West Nile Virus," as the first presentation in the 3:30 p.m. session, which is National Security and Emergency Management. There are probably some implications there, too. So thank you very much, Zhian, and for all the people in that session as well, for changing their schedule.

An Agent-based Model for Simulation of the West Nile Virus

[Editors' Note: This paper was moved to the session, National Security and Emergency Management, Friday, October 14, 2006, 3:30–5:30.)

Parallel Applications Session I —

Economics and Environmental Policy

MARGINS AND TRANSACTION TAXES IN AN INTRADAY CONTINUOUS DOUBLE-AUCTION FUTURES MARKET

L.J. USSHER,* Queens College of the City University of New York, Flushing, NY

ABSTRACT

Futures market quotes and transaction prices are derived endogenously in an artificial market with bilateral trading in a continuous double auction. The market participants — speculators, hedgers, and scalpers — have different strategies and reasons for trading futures. Risk-neutral speculators with heterogeneous expectations leverage themselves and try to maximize profit, while being constrained by margin requirements and trading costs. Hedgers provide a fundamental price anchor, and scalpers act as market makers. Emphasis is placed on exchange rules and regulations that govern trading rather than agent learning. The futures exchange imposes real-time gross settlement, margin requirements, and a one-way transaction fee or tax on speculators. Despite a lack of individual rationality and well-behaved demand functions, our model creates a bid-ask spread that, although turbulent, converges to the exogenous cost of trading for speculators and a mid-price that strongly detects the black box Walrasian equilibrium price. In a market with only speculators and hedgers, prices appear to have a lower level of kurtosis, or volatility, when the market is less leveraged, or, in other words, margin requirements are high. The raising of transaction taxes in such a market only serves to reduce trading and increase price volatility. By adding market makers to the model, trading volumes are maintained even in high tax regimes, making the market price more resilient and reducing price volatility.

Keywords: Margins, transaction tax, continuous double auction, futures market, agent-based model, scalpers

INTRODUCTION

Market microstructure emphasizes market design and the mechanics of trading. This paper simulates trading on a futures exchange. Unlike most papers in this genre, this paper ignores the role of information, learning, and rationality, instead investigating a market structure with diverse agents bound by trading rules or traditions. This is similar to the Gode and Sunder (1993) model of zero intelligence agents, where the budget constraint is critical to allocational efficiency. Despite its relative simplicity, this preliminary study may provide insights for market design. Our project is to analyze the presence of liquidity, efficiency, and stability at the aggregate level, without imposing exorbitant assumptions on micro behavior, such as rationality, or the Arrow and Debreu (1954) restrictions.

Most economic theories rest on the premise that aggregate relationships are stable over extended intervals of time. The creation of microfoundations to underlie these aggregate economic stylized facts has relied on maintaining the falsehood that aggregates behave the same

* *Corresponding author address*: Leanne J. Ussher, Queens College, City University of New York (CUNY), Flushing, NY 11367; e-mail: leanne.ussher@qc.cuny.edu.

way as their component parts and that therefore the behavior of the aggregate is attributed to that of some fictional representative agent (Martel 1996).

In the next section, we begin by describing futures trading on the floor of the exchange. There is a discussion of what is meant by liquidity and of how margin or tax policy might impact trader activity and market liquidity. Section 3 explains the model. In opposition to most papers in this genre (Arthur et al. 1997; Farmer and Joshi 2002; LeBaron 2000, 2002), this paper simplifies the trading process by removing information and learning from agent behavior. Speculator expectations are given at the outset and do not change. A continuous double-auction (CDA) market is implemented with real-time gross settlement (RTGS). Trading is derived from simplified but recognizable agent trading rules, which may involve backward-bending demand functions, leveraged trading, short selling, and asynchronous trading. In Section 4, we experiment with some preliminary simulations and suggest ways that notions of liquidity and institutional rules, such as margin requirements, transaction taxes, and RTGS, might be evaluated. Preliminary results on the impact that margins and taxes may have on price volatility are presented. We find that despite the potential for individual instability, market stability is a common trait.

FUTURES MARKET TRADING

Open Outcry

In an *open-outcry* futures market, as described by Silber (1984), all *bids* and *offers* must be announced publicly to the pit through the outcry of buy or sell orders. In particular, no prearranged trades are permitted on futures exchanges. Strict priority is kept, where the highest bid price and the lowest offer take precedence, and this is known as the *inside spread*. Lower bidders must keep silent when a higher bid is called out, and higher offers are silenced when a lower offer is announced, although simultaneous offers and simultaneous bids at the same price can occur. To increase the probability of execution, a trader can raise his bid or lower his offer, and then other traders must remain silent. This rule is designed to insure best execution, in the sense that sales occur at the highest bid price and purchases occur at the lowest offering, and all bids or offers do not live longer than the moment needed to make a transaction.

Scalpers, also known as *locals* because of their exchange membership, are floor traders who trade on their own account and have low transaction costs and more flexible margin requirements than speculators. Like *dealers*, in bond or foreign exchange markets, scalpers regularly quote a bid price at which to buy and an ask price at which to sell, *making a market* and thereby offering to complete orders quickly, typically at a price close to the last price, for those anxious to trade. By inserting this spread between the buy and sell, the scalper thereby receives a profit for providing the service of *immediacy,* which is just one dimension of *liquidity*. Scalpers may also provide *depth* commensurate with the quantity they are willing to buy or sell. While scalpers typically provide liquidity, it is important to note that they can also "consume" liquidity when they liquidate or *offset* positions, by selling at the bid price or buying at the ask price. This reduction in liquidity may cause temporary instability (Schwartz 1988).

An ordinary trader (nonscalper) can either tender his own ask or bid quote that competes with the scalper, called a *limit order,* or accept the price currently quoted in the market, called a *market order*. When a market participant accepts the market bid, he is said to *hit* the bid. When he accepts the market ask, he is said to *lift* the ask. The following example highlighting the choices of a nonscalper who wants to buy contracts is taken from Silber (1984, page 940). A commercial hedger can instruct his *broker* (on the floor) to buy 50 contracts *at the market*, in which case the broker lifts the asks of others in the pit. Alternatively, the commercial hedger can try to buy more cheaply by instructing the floor broker to bid for 50 contracts at the prevailing bid price in the pit. In the first case, the *market order* uses the immediate execution services provided by the offerers in the pit (from scalpers or whomever) consuming liquidity. In the second case, the bid represented by the floor broker can be used by others to sell into, thereby providing liquidity.

Our study is partly to consider how effectively a financial market with asynchronous trading operates without intermediaries, such as scalpers. Often the mismatch between buyers and sellers that typically exists at any given instant is resolved by some agents who are willing to play the role of *market maker* and provide *liquidity*.

Bid-Ask Spread and Liquidity

The academic literature on market microstructure recognizes that the arrival of random traders to buy or sell is asynchronous, and market activities are temporally discrete. This literature treats such moment-to-moment aggregate exchange behavior as an important descriptive aspect of markets (Garman 1976) and has led to many interesting questions, such as: How are market structure and the trading process related to the price process or the valuation of securities? What sort of trading arrangements maximize efficiency? How is information impounded in prices?

There is rarely a single price in microstructure analyses, and the research into the CDA and the various prices derived from this — either quoted, averaged, or actually traded on — are components of the *bid-ask spread*. The size of the spread is an important dimension of *liquidity*. Modeling the spread is an extremely complex matter, given that the markets are composed of numerous *limit traders* (which include dealers and ordinary traders) embedded in a dynamic, interactive environment. Such a system may best be modeled with an agent-based methodology.

The analytical bid-ask spread literature (Stoll 1978; Ho and Stoll 1981) explains the demand for immediacy from the asynchronous arrival of random traders to buy or sell. It is often assumed that dealers participate in every trade, known as a *quote-driven market.* The behavior of the market maker or dealer is typically described as a trader who inserts a spread between the buy and sell and thereby receives a profit for providing the service of *immediacy* in what might otherwise be a fragmented market. This view of the market maker, as a provider of predictable immediacy, was first formalized by Demsetz (1968) and then elaborated on by Garman (1976) and many others. It is generally accepted that the bid-ask spread is representative of the risks faced by the dealer as a result of inventory control and asymmetric information. When scalpers provide for market orders, they profit from impatient traders but lose to traders more informed. It is usually concluded that with competition, the spread is reduced to the dealer's trading costs. This theory has formalized the idea of dealers as being providers of liquidity and controllers of the size of the spread. Inventory control costs are assumed to be reasonably constant over time, while risks of asymmetric information are not (Engle and Lange 1997, page 4). The size of the premium charged by *immediacy providers* to cover these expected costs determines the size of the spread and thereby the extent of illiquidity in the market.

Liquidity is defined in many different ways. If the bid-ask spread reflects the price at which immediacy can be obtained by ordinary investors trading via market orders, then a market is commonly thought of as perfectly liquid if trades can be executed with no cost (O'Hara 1997; Engle and Lange 1997). By this definition, a narrower spread means a more liquid market. This simplified characterization and measure of liquidity has recently gained popularity (see Flemming 2003), although many other definitions have long been debated.

Liquidity is usually said to have four dimensions: *immediacy*, *width*, *depth*, and *resiliency*. *Immediacy* refers to how quickly trades can be arranged at a given cost. *Width* refers to the cost of doing a trade for a given size. *Depth* is the size of a trade for a given cost. *Resiliency* refers to how quickly prices revert to former levels after they change in response to large order flow imbalances (see Harris 2003, pp. 398–405).

Liquidity is often described as being supported by a particular group of traders. Market makers are considered the primary providers often endowed with the responsibility of balancing order flow — choosing prices that equate supply with demand. As a key participant in the price discovery process, the market maker acts as a matchmaker, bringing public buyers and sellers together.[1]

Schwartz (1988) argues that too much emphasis has been made of market makers and their spread. While they may be needed in illiquid markets, they are not a necessity for liquidity. Schwartz emphasizes the resiliency dimension of liquidity and argues that more attention should be paid to the manner in which ordinary traders supply immediacy to each other and compete to reduce market spreads with the scalpers (Cohen et al. 1979, page 814). Schwartz (1988) also warns that for market makers to stabilize a market, they must commit capital or inventory risk, and this may become substantial. Injecting liquidity into a system to stabilize prices might also be just as quickly withdrawn at a later date if shortages are incurred or if the market makers seek to rebalance their portfolios.

Alternatively, Bernstein (1987) and Black (1986) emphasize that *noise traders*, with their diverse opinions, help provide liquidity or *resiliency* to a market. Those who trade on noise allow others to trade on information. It is the noise traders who provide *depth*, *breadth*, and *resiliency* to a market. At the same time, however, noise traders add volatility to prices and push prices into overvaluation or undervaluation, attracting information traders who push prices back to fundamentals. Hence, noise trading actually puts noise into prices, and prices are less efficient. "What's needed for a liquid market causes prices to be less efficient" (Black 1986, page 32). Bernstein argues that this process leads to a curious paradox: "…depth, breadth, and resiliency, in other words are not ends in themselves, but a means to induce information traders to trade. Efficient prices are possible only with noise traders creating inefficiencies by their buying and selling" (Bernstein 1987, page 56). This is similar to the analogy of annealing: the market needs to be heated up and made more liquid in order for the efficient price to be found. It is not true, however, that liquidity is not an end in itself. With the segmentation of market roles into different agents, there are some (such as the managers of markets [e.g., a central bank, a stock or

1 See Stoll (1985) and Schwartz (1988) for further discussion and references on alternative views of dealers and scalpers.

futures exchange, or an investment bank managing a line of corporate bonds]) who are only concerned with making their market liquid and who leave the price level or efficiency goal up to the informed speculators.

Harris (2003, pages 402–403) has a different view from Bernstein (1987) and from Black (1986). Along more traditional lines Harris argues that liquidity is present when prices return to their “fundamental value”; hence, information is key in his description. He argues that it is the *value trader* who promotes resiliency. Value traders are the informed traders who collect as much information about fundamental values as is economically sensible. Value traders supply liquidity, under his notion, when prices differ substantially from their estimates of value, and they trade in quite large sums that may be held over extended time periods. Harris argues that uninformed traders can have a negative impact on prices because dealers are passive traders and do not have an opinion about fundamentals, and they are unable to distinguish between informed and uninformed traders.

However, Harris (2003, page 394) also argues that liquidity is best described as the object of a bilateral search (i.e., in which buyers search for sellers, and sellers search for buyers). Liquidity is easiest to find when many people on both sides of the market are looking for it at the same time. This reiterates Berstein’s and Black’s analysis that it is noise traders that make this search easier.

Value traders contribute to mean reversion and *scalpers* provide immediacy, but this may not necessarily be efficient. All these arguments depend on which traders have cash, or leveragability, on hand and are ready to modify their investment exposure at the cheapest possible price, through limit orders, thereby offering liquidity. Following from Schwartz’s contra-side orders, it would seem easier to not allocate liquidity to a group of traders (value, noise, intermediary) but rather to state that market orders remove liquidity and limit orders provide liquidity. Included in the limit order category is the dealer bid and ask. These different approaches may be considered in our model.

Margins and Transaction Costs

In a futures market, transactions are promises rather than actual transfers of assets. Each promise to buy or sell a commodity at the future spot date is backed with collateral, which can be held as cash or treasury bills with the exchange (or broker). If it is held as treasury bills, then it can earn a rate of return. A minimum margin (collateral) requirement is specified by the exchange to guarantee the fulfillment of each contract an agent holds, whether long or short. The margin requirement is typically quoted as an absolute value per contract (e.g., for a contract of 5,000 bushels of July wheat on the Chicago Board of Trade [CBOT], the *initial margin* requirement is $1,800 per contract). This amount is usually changed by the exchange during the contract’s life; it is increased as the contract approaches maturity or when price volatility increases. A competitive exchange tries to minimize the margin requirement so that it just covers anticipated overnight price changes. For example, if price changes are thought to have even a small chance of moving 10%, then the exchange would like to make sure that traders have margin holdings of at least 10% of the contract value, ensuring contract fulfillment. In addition to margin requirements, a percentage transaction fee is often imposed on round-trip transactions. Brokers, exchanges, or the government can institute this as a tax.

Typically thought of as liquidity augmenting, policies to reduce margin requirements and transaction costs are advocated because they increase the amount of trading in a market, which is often thought to reflect liquidity and reduce price volatility. In opposition, there have been a number of economists who argue that excessive trading can increase volatility. They wish to remove noise trading by raising margin requirements (Shiller 2000; Schlesinger 2000) or imposing a transaction tax (Tobin 1974; Pollin et al. 2001) to reduce *excessive speculation* and price volatility in foreign exchange, equity, and futures markets. There are many debates on whether such policies would increase or decrease extreme price volatility (fat tails in the price distribution). Critics argue that taxes would only increase volatility and cannot stop large price movements from occurring (Davidson 1997). Insights might be garnered by an agent-based modeling approach to these policy debates.

AGENT-BASED MODEL OF A FUTURES MARKET

Model Environment

We present a model of speculators, scalpers, and hedgers in a futures trading pit with *open-outcry* and a CDA trading mechanism. In this simplified model, all trader expectations, although heterogeneous, remain constant in order to place focus on the trading mechanism and the impact of trader budget constraints. This is a partial equilibrium model with two markets: a speculative futures market for grain and a residual money market. The price of money is normalized to 1, and agents operate on their budget constraint, which is a function of their wealth, transaction costs, and futures-contract margin requirements. There is no restriction on short selling. RTGS is implemented such that traders settle with each other and the exchange at their time of trade, rather than waiting until the end of the day.

Margins are implemented in this paper in a simplified manner, although they are still relevant to modern market design. First, we mark-to-market trader positions by using RTGS. Thus, instead of using the close-of-day *settlement price* to calculate margin calls, settlement is adjusted continuously throughout the day, and the settlement price used to calculate margin calls is the average of the bid and ask price, or mid-price. This means that profits and losses transfer hands between the exchange and the traders continuously, removing the risk of accumulated losses and trader default. This payment transfer is called the *variation margin*.

Second, the model analytically simplifies the margin calculation by making the *initial margin* and *maintenance margin* the same and specifying the margin requirement as a fixed percentage of the contract value rather than an absolute dollar value per contract. By using a margin requirement that changes with the percentage change in prices, we get closer to the essence of what the exchange considers in setting the margin.

Given these two margin features, our model offers considerable price and quantity feedback opportunities. Although for an individual, high-risk, speculative trader, marking to market is a cautionary act and reduces counterparty risk, it can also result in a volatile market price the higher the settlement frequency is (Farmer et al. 2004). If traders are on their budget constraints, then they will liquidate some of their position when prices move against them in order to stay within their margin requirements, and this creates backward-bending demand functions, as introduced in the next subsection.

Each type of trader has his own rules for trading. *Speculators* are risk neutral and differ only in their expectation of what the futures price should be and in their wealth. Expectations of the next-period futures-contract price stay constant during the trading period. Being risk neutral, speculators are typically at the corner solutions of their budget constraint, maximizing their futures position (long or short) at every chance they get to trade. There is a one-way fee imposed by the exchange, charged as a percentage of each transaction at the point of sale or purchase. The contract size is perfectly divisible, and prices are always non-negative. Speculators are required to hold a minimum amount of cash in their margin account, which is a percentage of the futures contract value. To safeguard contract fulfillment, the exchange carries out RTGS with *variation margins* imposed on every transaction.

Scalpers are members of the exchange and operate on the floor of the exchange without paying a trading fee. They do not have an opinion on the fundamental price and instead try to buy as low and sell as high as they can. They want to maximize the turnover of buys and sells while minimizing their inventory holding. Scalpers prefer to place limit orders (quotes) and to buy at their bid quote and sell at their ask quote. Scalper activity assists in balancing order flow over the long run, which does promote price efficiency, but it could create price instability in our bilateral CDA, either when they offer liquidity to so-called noise traders or when they are forced to liquidate their own inventory holdings with market orders.

Hedgers play only a limited role in setting up the fundamental demand and supply of contracts in the market. There are only two representative hedgers — one going long and the other going short — the difference being the net hedge. They only place market orders to fill their desired contract positions. The quantity of contracts desired is exogenous to the model and does not change. Once their futures position is attained, they stop trading, and together they leave a net excess demand or supply for the rest of the traders in the market to sort out.

Within the CDA, speculators and scalpers (if included) are selected for a sequence of bilateral trading through random nonreplacement in each round, so that each trader has an equal chance of trading. The hedgers are placed last in this sequence, which represents one round. The intraday period of futures trading has several rounds of quoting or transacting, at the bid or ask price. Trades and transaction prices are registered at each time t.

Speculator's Demand Function

In our model with leveraged speculation, κ represents the limit on how much larger a speculator's futures position — price multiplied by the number of contracts ($p_t x_t$) — can be than a trader's wealth m_t. For example, if $\kappa = 4$, then a trader can have up to 4 times his wealth dedicated to a long or short futures position. In other words, the margin requirement is 25%, $1/\kappa = 0.25$. The collateral kept in the margin account by speculator i is held as either Treasury bills or money, represented here as m_t^i. Money held must be greater than the margin requirement, $m_t^i \geq p_t x_t^i / \kappa,$ for the current futures position at all times (to the extent that trading allows). There will be several transaction prices throughout the day, which represent a trade at either a quoted bid p_t^b or a quoted ask p_t^a. If there is not enough collateral in the margin account to meet the margin requirement, then speculator i will have to liquidate his position with an offset purchase or sale at his next turn to trade.

The futures position x_t at price p_t is taken on by the speculator as a contract at time t to sell or buy x units of the underlying commodity at price p_t on the spot or maturity date of the futures contract. Since our speculator does not intend to make delivery on this contract, the purpose of holding this position is to flip the position and profit on price changes. On the basis of price expectations $p^{i,\theta}$ about the next transaction price p_{t+1}, speculator i will decide to go either long or short in futures. If the expected short-term gain does not compensate the cost of trading over the next period:

$$(p^{\theta}{}_t - p_t)\, x_t \leq \varpi\, p_t |(x_t - x_{t-1})| \,,$$

then the speculator will hold his current position instead of trading. The trader is myopic, and upon opening a position, there is no consideration of costs incurred for reversing the position.

Each speculator is risk neutral and simply maximizes expected wealth π from t to $t + 1$:

$$\pi^{e}_{t+1} = \left(p^{\theta} - p_t\right)x_t + m_t \ .$$

The speculator's demand curve is derived in the appendix via linear programming. In summary, speculator i's demand for futures in each period t is a slightly simplified version from Ussher (2004):

$$x^{i}_{t}\left(p_t; x^{i}_{t-1}, m_{t-1}, p_{t-1}, p^{i,\theta}, \kappa, \varpi,\right) ,$$

where:

p_t = Intraday futures market transaction price at time t,

x^{i}_{t-1} = Previous contract position,

m^{i}_{t-1} = Previous cash position in margin account following last transaction,

$p^{i,\theta}_{t}$ = Price expectation p^{θ} of the next futures price p_{t+1} at time t.

$1/\kappa$ = Margin requirement as a percentage of futures position value, and

ϖ = Percentage transaction tax on a one-way trade (paid each way).

A futures demand curve is usually represented as a smooth downward-sloping line from the top of quadrant II to the bottom of quadrant I in the two-dimensional R^2 space in Figure 1. Our model produces a nonlinear demand function as a result of the inherent corner solutions from the risk-neutral speculators' wealth constraints and the regulatory setting of margin limits $1/\kappa$, transaction costs, and RTGS.

Each risk-neutral speculator maximizes the next period's expected wealth by holding money as collateral and buying or selling futures (going long or short in futures). The decision to

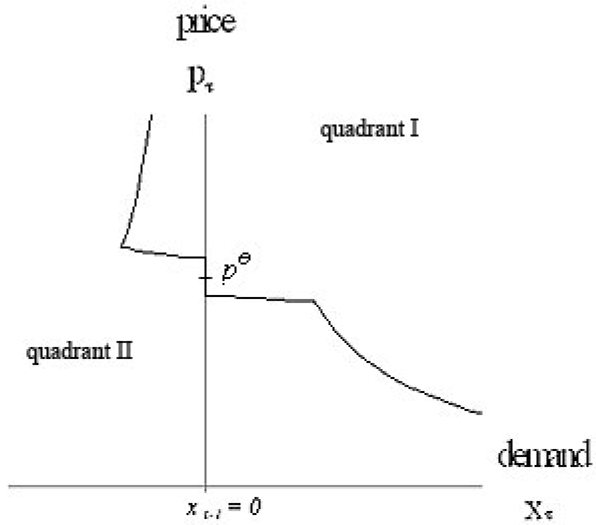

FIGURE 1 A speculator's demand for futures x_t as a function of p_t, with a past zero position x_{t-1} and price expectations of p^θ

buy or sell futures depends on whether the speculator expects prices to rise or fall, respectively. There is no restriction or disincentive to *short selling* (i.e., selling commodities that one doesn't own). A trader will trade only when price expectations p^θ are far enough away from the actual prices p_t to pay for the one-way transaction costs. Figure 1 has a zero contract position held over from last period. If a speculator currently has a futures position, then margin calls can lead to forced liquidation of the position when prices move against expectations. The possibility of a backward-bending demand function, as in Figure 2, is a result of the collateral px, which underlies demand for x, being priced in the same market.

The speculator will sell (buy) futures if he expects the price to fall (rise) when the slope of the demand function is positive. The demand function has a negative slope when purchasing power is declining from higher futures prices or when collateral is devalued and the speculator must liquidate part of his position to maintain the margin requirement.

At each t, the variation margin is calculated and net wealth is adjusted. The mid-price p^m is the average of the bid quote p^b and ask quote p^a:

$$p^m = (p^a + p^b)/2 \ .$$

By using the mid-price, the profit or loss is calculated with price changes and paid from the losing agent to the winning agent via the exchange, equivalent to:

$$\left(p_t^m - p_{t-1}^m\right)x_{t-1}^i \ .$$

Each speculator estimates his net wealth at each t, given prices (p^a, p^b, p^m), which determines his decision on how many futures contracts to buy or sell to maximize expected

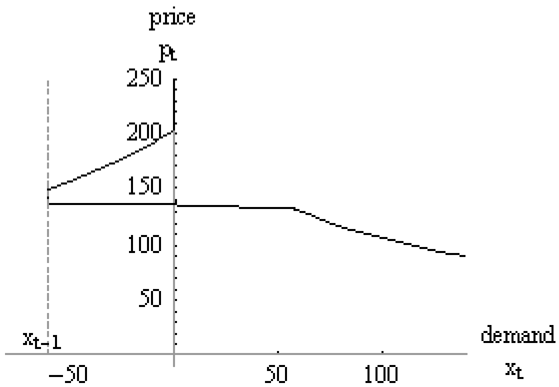

Speculator is short futures $x_{t-1} = -60$

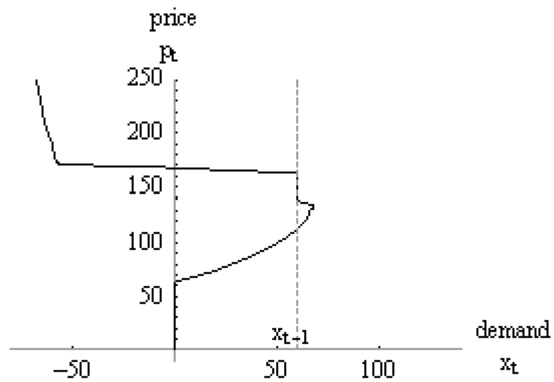

Speculator is long futures $x_{t-1} = +60$

FIGURE 2 A speculator's demand curve with either a short or long starting position: m_{t-1} = 5,000, p^{θ} = 150, and κ = 2 for each graph

wealth, while at the same time meeting his margin requirement, which is κ multiplied by net wealth. The mid-price is used in accounting for net wealth every period, as long as a position is held.[2]

Bidding and Trading Process

Central to our model is the auction process that simulates the open-outcry on the floor of an exchange, leading to transactions and thus transaction prices. It is a tâtonnement mechanism where both bid and ask prices adjust and where out-of-equilibrium trades take place when an agent agrees to sell contracts to another agent who is bidding for them, or when another agent decides to buy contracts from the agent who is asking for them. This process of quoting and trading is repeated many times, giving each market participant the chance to quote and trade several times and fill his orders. No new information is brought into this process; expectations remain constant.

The competitive bidding algorithm presented here is drawn from several sources. The manner in which speculators compete and how their price expectations interact with the bid-ask spread during the bidding process comes from Chan, LeBaron, Lo, and Poggio (CLLP; see Chan et al. 1998) and Yang (2002). An important modification to their model, apart from keeping expectations constant, is our distinction of risk-neutral speculators with collateral constraints and transaction costs. In addition, we have drawn on another algorithm derived from Silber (1984),

2 On the initial purchase of a market order the trader must pay a *variation margin* of $\left(p_t^m - p_{t-1}^m\right)\left(x_t^i - x_{t-1}^i\right)$. Important in this calculation of variation margin is that we keep the distinction between those that profit by buying at the bid or selling at the ask, versus those who are considered impatient and sell at the ask or buy at the bid. When a contract is bought and $\left(x_t^i - x_{t-1}^i\right) > 0$, if it is bought at the bid with a *limit order*, then the variation margin is positive $\left(p_t^m - p_t\right) > 0$. If, however, it is bought at the ask with a *market order*, then the variation margin is negative $\left(p_t^m - p_t\right) < 0$. This results in a transfer of wealth from the trader who is willing to pay for *immediacy* to the trader who gets paid for providing liquidity and *making the market*. The maximization of expected wealth by the speculator takes into account only the expected change in the trade price $\left(p^{\theta} - p_t\right)$, without anticipating whether the transaction is by market order or limit order.

emphasizing inventory control and noncompetitive bidding by our scalpers. Hedgers act similarly to speculators but only place market orders and hence do not compete in the bid-ask spread. These agents are used to represent fundamental supply and demand.

This asynchronous bilateral bidding process allows two or three traders to participate at any one time, offering or bettering limit order quotations or carrying out market order trades. Agents take turns entering into the inter-dealer market to quote price and quantity, to transact, or to exit. A round is completed when all agents have participated once, with the hedgers coming last. This is repeated for a different random sequence of scalpers and speculators for more than 50 rounds. The repetition or trading rounds represent competition within the price mechanism and help the convergence to equilibrium of market demand and supply.

Auction Algorithm for a Speculator

Half of the bid-ask spread is often thought of as a measure of the cost of executing a market order (the difference between the mid-point price and the payment price). We shall represent this price difference by the lowercase letter *s*. The size of this spread is actually endogenous to the bilateral trading process. In our model, speculator *i*'s reserve price is his expected price $p^{i,\theta}$ plus the one-way transaction tax ϖp_t.

At times when there is no bid or ask, a speculator will announce his own noncompetitive limit order on the basis of expectations $(1 \pm S\ \varpi)\ p^{i,\theta}$. In this case, *S* is a percentage of the transaction fee. If *S* is greater than 100%, then the new limit order will guarantee that a new hit or bid occurs with a demand different from zero.

We present the speculator algorithm with three traders: agent *k* has the best bid to date, agent *j* has the best ask to date, and agent *i* is the new entrant who makes a trade choice under the following four scenarios. Agents *j* and *k* are offering the best ask and bid quote to date, respectively, and are scalpers or speculators. Agent *i* represents a speculator who enters the market and witnesses the current bid-ask spread. Speculators attempt to profit by positioning themselves in each period to maximize short-run profit over every single period *t*.

- Scenario 1 (Figure 3a). The ask, $p_t^{j,a}$, and bid, $p_t^{k,b}$, currently exist with nonzero offers, at time *t*.

    1. If $p^{i,\theta} > p_t^{j,a}$, speculator *i* will post a market order and buy at this ask price — lift the ask quote.

    2. If $p^{i,\theta} < p_t^{k,b}$, speculator *i* will post a market order and sell at this bid price — hit the bid quote.

    3. If $p_t^{k,b} \le p^{i,\theta} \le p_t^{j,a}$ and $< \left(p_t^{k,b} + p_t^{j,a}\right)/2$, speculator *i* will post a sell limit order at a price of $(1 + S\ \varpi)\ p^{i,\theta}$ and thus quote his own ask, replacing agent *j*.

    4. If $p_t^{k,b} \le p^{i,\theta} \le p_t^{j,a}$ and $\ge \left(p_t^{k,b} + p_t^{j,a}\right)/2$, speculator *i* will post a buy limit order at a price of $(1 + S\ \varpi)\ p^{i,\theta}$ and thus quote his own bid, replacing agent *k*.

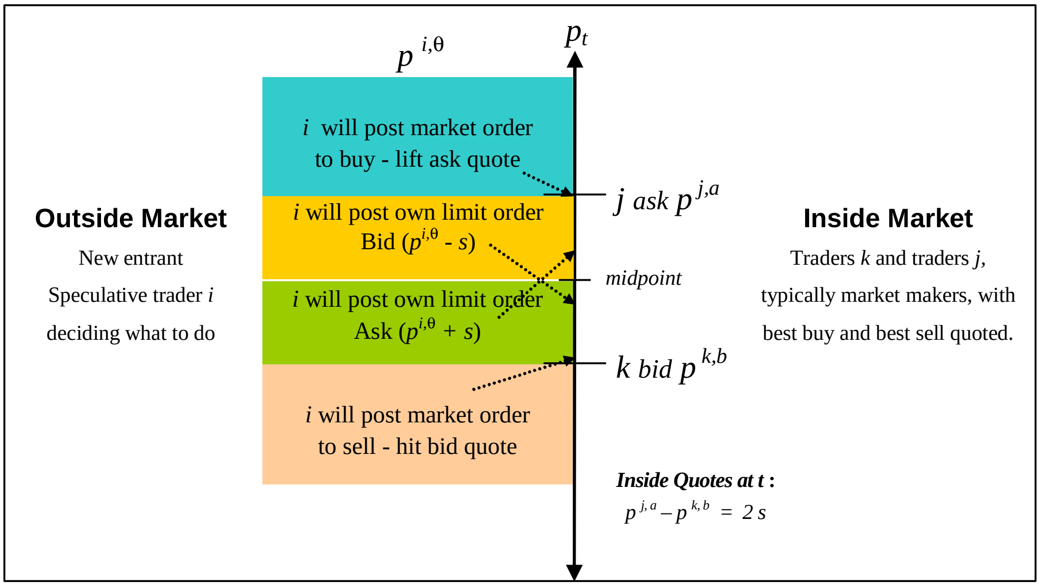

FIGURE 3a Scenario 1, in which both competitive quotes — bid and ask — exist in the marketplace prior to new entrant

- Scenario 2 (Figure 3b). Only the best ask, $p_t^{j,a}$, exists; that is, at $p_t^{k,b}$, the demand to go long must be zero as $\left(x_t^k - x_{t-1}^k\right) \leq 0$.

  1. If $p^{i,\theta} > p_t^{j,a}$, speculator i will post a market order and buy at this ask price.

  2. If $p^{i,\theta} \leq p_t^{j,a}$, speculator i will post a buy limit order $p_t^{i,b}$ at a price of (1 – S ϖ) $p^{i,\theta}$, but only if excess demand at this price is $\left(x_t^i - x_{t-1}^i\right) > 0$.

- Scenario 3 (Figure 3b). Only the best bid, $p_t^{k,b}$, exists; that is, at $p_t^{j,a}$, demand to go short is zero as $\left(x_t^j - x_{t-1}^j\right) \geq 0$.

  1. If $p^{i,\theta} < p_t^{k,b}$, speculator i will post a market order and sell at this bid price.

  2. If $p^{i,\theta} \geq p_t^{k,b}$, speculator i will post a sell limit order $p_t^{i,a}$ at a price of (1 + S ϖ) $p^{i,\theta}$, but only if excess demand at this price is $\left(x_t^i - x_{t-1}^i\right) < 0$.

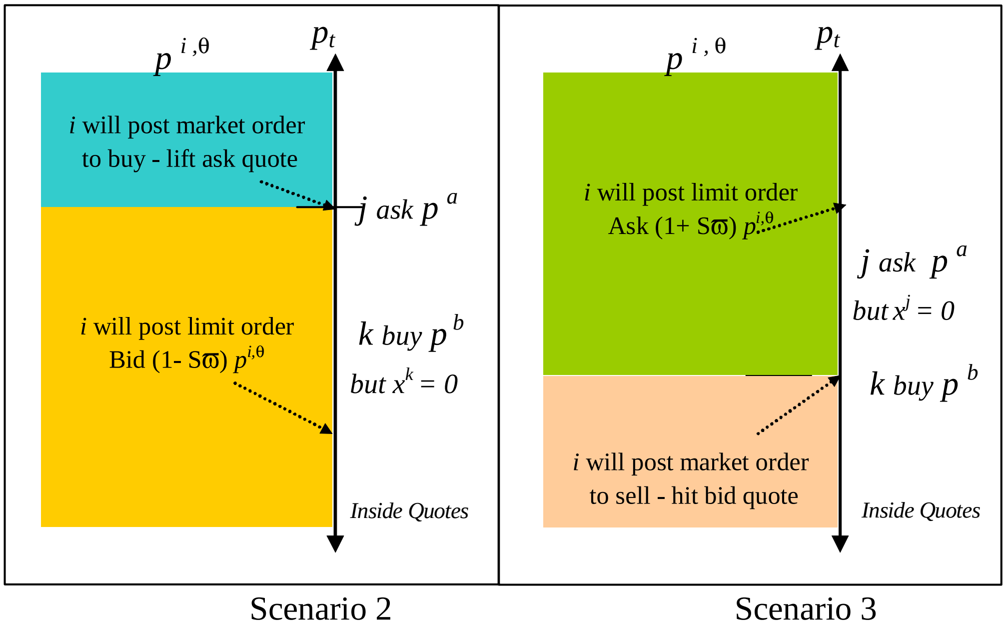

FIGURE 3b Scenario 2, in which an ask but no bid exists prior to new entrant, and Scenario 3, in which a bid but no ask exists prior to new entrant (0 = θ, τ = *t*)

- Scenario 4. No bid or ask effectively exists; that is, at $p_t^{j,a}$, $\left(x_t^j - x_{t-1}^j\right) \geq 0$, and at $p_t^{k,b}$, $\left(x_t^k - x_{t-1}^k\right) \leq 0$.

    1. The new entrant speculator will post a buy and/or a sell limit order at $(1 - S\,\varpi)\,p^{i,\theta}$ and/or $(1 + S\,\varpi)\,p^{i,\theta}$, respectively, as long as his bid is quoted for a buy of greater-than-zero contracts and his ask for a sell of greater-than-zero contracts. If this is not the case, then the current bid-ask remains, even though both traders have zero demand, and entrant *i* exits to join the queue to trade again later.

In our model, under Scenario 2 or Scenario 3, the speculator tendering the best bid (ask) might have had prices move against him; for example, if he was long (short) and prices fell (rose). He may remain offering a bid (ask) price to buy (sell), but at a quantity of zero. Now he wants to offset his position and sell (buy) so that excess demand is less (greater) than zero.

Scenario 2: $\left(x_t^k\left[p_t^{k,b}\right] - x_{t-1}^k\right) \leq 0$, where x_t^k is a function of $p_t^{k,b}$

Scenario 3: $\left(x_t^j\left(p_t^{j,a}\right) - x_{t-1}^j\right) \geq 0$, where x_t^j is a function of $p_t^{j,a}$.

Effectively under Scenario 2 (Scenario 3), agent *k* (agent *j*) falls silent and will eventually be replaced by a new entrant, as long as the new entrant has $p^{i,\theta} < p_t^{j,a}$ (has $p^{i,\theta} > p_t^{k,b}$) and as long as $[x_t^i(1 + S\varpi)\,p_t^{i,\theta} - x_{t-1}^i] > 0$ (as long as $[x_t^i(1 - S\varpi)\,p_t^{i,\theta} - x_{t-1}^i] < 0$); otherwise, agent *k* (agent *j*) will remain. Only when agent *k* (agent *j*) is replaced and exits the market will he be given the chance to satisfy margin requirements by liquidating his position with a market order, in turn, in the random trading round.

This model considerably changes the CLLP rules, which emphasize the manner in which price formation feeds back into the market by agents updating their expectations, to one where price formation feeds back into the market via quantity constraints, margin requirements, and inventory control. This model allows for leveraged trading and short selling and makes the method of settlement a central variable of the model.

Auction Algorithm for a Scalper

In addition to speculators, *scalpers* also participate in the futures open-outcry. Scalpers will try to charge as high a price as possible when selling and as low a price as possible when buying, while still competing with other traders to make a sale or purchase. Only the highest bid and lowest ask are heard in the trading pit. All other noncompetitive quotes must remain silent. Since speculators must compete on price, only speculators are able to narrow the inside-market bid-ask spread. Scalpers balance market order flow by using the *interdealer market* to offset their own inventory excesses. Taking a loss in order to liquidate an unbalanced inventory position forces other interdealer scalpers to also liquidate, and this dries up liquidity in the market until prices are modified.

The scalper algorithm is a simplified version of one stated in Smidt (1985). The objective is to buy at the bid and sell at the ask, maximizing a profit equal to the turnaround of inventory multiplied by the spread, while minimizing inventory risk with a very simplified control mechanism. There is a maximum net inventory ceiling K for each scalper. Netting out the long and short trades by a single agent consolidates the inventory x_t. Scalper inventory must be:

$-K \le x_t^n \le K$ for scalper n.

In actual markets, K is often as small as one contract and could be different for different scalpers. In our model, all scalpers have the same $K = 10$. When a scalper enters the trading floor from the random sequence, if his inventory is less than his maximum limit K, he always has the right to replace any agent in the *inter-dealer* (inside spread) market by simply matching the agent's quoted bid and ask. This is in contrast to speculators who must offer a better price to replace the agents in the inter-dealer market. If, however, the scalper's inventory is on his limit, then the scalper will place a market order to offload all inventory, if possible. The scalper algorithm is one of simple inventory control:

- New entrant scalper n

    1. If $-K < x_t^n < K$, replace the current market makers and quote both bid and ask at the current quotations $p_t^{k,b}$ and $p_t^{j,a}$, and for quantity $K - x_t^n$, buy, and for quantity $-K - x_t^n$, sell.

    2. If short and $x_t^n \le -K$, hit the market bid for a maximum $-x_t^n$ and post no quotes.

    3. If long and $x_t^n \ge K$, lift the market ask for a maximum $-x_t^n$, and post no quotes.

The dealer inventory control model outlined here, where a scalper will choose to make a market order rather than change his limit order prices, is in contrast to most accepted inventory control models such as Garman (1976) and Amihud and Mendelson (1980). These authors present dealers as changing their bid and ask to induce an imbalance of incoming orders, in order to reduce inventory. Hasbrouck (2003) questions this latter model and claims that as a general rule, most empirical analyses of inventory control refute this method of changing the quote for inventory control. He argues that a dealer who would pursue the hypothesized mechanism would be signaling to the world at large his desire to buy or sell. This would put him at a competitive disadvantage (Ibid 2003, p. 78). Our simplified mechanism does not touch on information signaling, yet it does avoid this specific criticism.

Auction Algorithm for a Representative Hedger

Hedgers are only concerned about filling their expected sales or purchases at the spot date via market orders in futures. They always come last in each round of the random sequence of speculators and scalpers.

- Hedger scenario

    1. The future purchaser of the commodity at spot, agent *q*, will lift the ask, $p_t^{j,a}$, for the maximum ask quote quantity, until the market buy order is filled, $x_t^q = x^{q*}$.

    2. The future seller of the commodity at spot, agent *r*, will hit the bid, $p_t^{k,b}$, for the maximum bid quote quantity, until the market sell order is filled, $x_t^r = x^{r*}$.

Since speculators and scalpers do not usually offer large size contract lots, it may take several rounds for our hedgers to finalize their purchases or sales. The hedgers contribute so-called *fundamentals* to our speculative market.

Trading Sequence

The setup for trading begins with a random ordering of 60 speculative agents and, when included, 10 scalpers. The two representative hedgers come last in this sequence, which, once completed, is called a trading round. Speculators have equal endowments and heterogeneous expectations taken from a symmetric distribution with a mean p^θ of 150. Speculators come together, along with hedgers and scalpers, in bilateral trades to create a CDA.

Two randomly selected traders begin with market quotes set at $p_0^b = 100 : p_0^a = 110$. A new entrant, randomly selected from the remaining traders (not a hedger), enters the floor to either accept or better the prices quoted. If a bid or ask is accepted, a trade is done and a transaction price p_1 occurs for, say, a market order by the new entrant. If, instead, the entrant replaces a bid or ask or both, then a new set of quotations $p_1^b : p_1^a$ (bid : ask) is created, with no transaction price. A sequence of quotes, and transaction prices, is generated during the trading round, with only transaction prices and volumes registered. Repeating the round, drawing a new random sequence of speculators and scalpers each time, creates an interday trading session.[3] This trading sequence is summarized here:

1. Speculators are initialized with initial wealth and random price expectations with a mean of 150. Two randomly selected speculators or scalpers begin with initial quotes of $p_0^b = 100 : p_0^a = 110$ and their respective buy and sell quantities (which may be zero), given their expectations.

2. The random sequence of speculators and scalpers to enter the market with nonreplacement is determined, with hedgers coming last.

3. With one or two agents quoting a bid-ask spread, the new entrant can either submit a new bid or ask, accept the existing bid or ask, or hold (pass).

4. A transaction occurs when the existing bid or ask orders are accepted and the transaction price is recorded accordingly. The transaction is the minimum of the quantities proposed for exchange by each bilateral trader.

5. At each point, mid-point prices are used to calculate speculator budget constraints in real time. On the basis of the past transaction price, each agent's wealth is updated, taking account of all margin calls (profits and losses).

6. Steps 3 through 5 are repeated for n times, where n = number of traders (one round).

7. Steps 2 through 6 are repeated for N times, where N = number of rounds.

8. The final market price is recorded as the last transaction price for this trading session.

The CDA bilateral search and trade algorithm is similar to a repetitive annealing process where the market is heated up through turbulent trading (when margins are low). This might be representative of a hot or liquid market, and this is warranted in order for the equilibrium point to be found. If traders become satisfied with the price and reduce their trading, then the market cools and converges to its fixed price or the efficient market price. But once the market cools, it becomes brittle, and a single trader can disrupt the price with a new quote $(1 \pm S\,\varpi)\,p^{i,\theta}$, causing a credit crunch and trading volume increases. The market heats up again, and the process is repeated.

SIMULATION OF INTRADAY TRADING

In creating a market that consists of highly speculative individual agents who are inherently unstable because of their leveraged positions and settlement constraints, we wish to discover how robust and stable our market is as a whole, given the regulatory framework of

---

3 In this paper, we stop at this point. But if one session was considered to be one period of constant expectations, in between the updating of expectations, then such trading sessions, when strung together, could be seen as a day of trading.

transaction taxes and margin requirements. With the imposition of no changes in agent expectations, we focus our analysis on the impact of margin calls and trading volatility on price formation. This is quite separate from the volatility that comes from expectations and information issues. We will consider how efficient trading is in converging to a stable equilibrium price that equates aggregate supply and demand. In addition, we will measure the presence of extreme price movements by looking at the kurtosis of our price distribution. This measure of volatility is most relevant to exchange governance that tries to maintain fair and orderly markets.

Simulation for 60 Speculators, Two Hedgers, and No Scalpers

We begin by simulating a CDA with just speculators and hedgers to consider how speculators alone can effectively replace formal market intermediaries, as suggested by Schwartz and Economides (1995). We use 60 speculators with the same wealth and randomly designated expectations drawn from three different normal distributions. All have a population mean of 150 and a standard deviation, σ, of either 1, 2, or 5. In this paper, we have used only one realization for each simulation,[4] where the 60 agents together have a sample mean of 150.4, a standard deviation of 2.7, and a kurtosis measure of 4.7. The following parameter trials included a tax of either $\varpi = \{0.1\%, 0.5\%\}$ on each one-way transaction and a margin requirement of either $1/\kappa = \{100\%, 33\%, 25\%\}$. All plots below are by transaction dates; hence, competitive quote changes do not show up when the bid-ask spread is narrowed through new limit orders. Only when a market order or purchase occurs are the bid-ask, mid-point, and transaction prices recorded at time *t*. Figures 8 and 9 (which appear later) actually show whether the trade took place at the bid or the ask.

Figure 4 shows a market with only speculators. The lines connect prices over time: bid, ask, and p^*; this is the Walrasian equilibrium price solution at time $t = 0$, given expectations, wealth, and the net hedge. While the bid-ask spread is usually converging through competitive quotes, the large swings outward by the bid-ask spread are due to the exogenous prices $(1 \pm S\,\varpi)$ $p^{i,\theta}$, limit orders that new entrants get to yell out when one of the inter-spread dealers is offering a zero-quantity bid or ask that they want. In these simulations, $S = 8$.[5] The only difference between the two graphs in Figure 4 is the decline in the margin requirement from 100% to 33.3%.

In Figure 4**,** as with all our simulations, the CDA bid-ask spread quickly detected the Walrasian equilibrium price of 151 despite starting with bid-ask quotations of 100:110.[6] There was a narrowing of the spread as speculators competed with each other, but it never got smaller than the transaction costs, ϖp_t, per unit *x*. When the aggregate desired demand across all agents in the market as a function of the ask approximated the aggregate desired supply as a function of

---

4 A full Monte Carlo of our model must be completed before conclusive results can be obtained. The results shown here are only preliminary, but they nevertheless show the potential of this type of analysis.

5 We also used $S = 1.5$, which resulted in the same average price, except that the market was very slow to equate demand and supply.

6 Given the relative symmetry in our wealth-weighted market expectations and given the small net hedge in the market ($x^{r*} - x^{q*} = 5$), this may be expected, but it can only be tested with further investigations across different initial conditions.

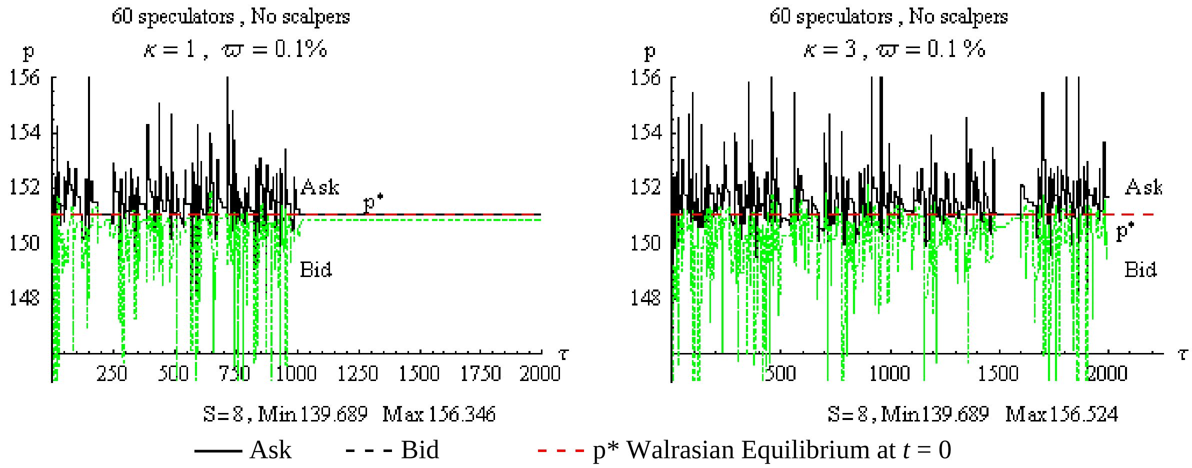

FIGURE 4 CDA bid-ask quotations at each transaction (margin requirements are 100% versus 33%)

the bid, then the bid-ask spread and mid-price were more likely not to change, as in Figure 4 for $\kappa = 1$ or 100% collateral. During these periods, the order flow is kept balanced in aggregate, as most speculators are satisfied and stop trading and as only very small trades occurred between two active traders post-$t = 1{,}000$, as seen by their trading positions in Figure 5. These two traders kept trading longer than their peers because their expectations were very close to $p^*(1 \pm \varpi)$ within the bid-ask spread. This also occurs when taxes are high (see Figure 7, which appears later).

The leverage positions, or contract value relative to wealth, of our 60 speculators are shown individually in Figure 5 for the two simulations with either 100% or 33% margin requirements. Traders who hold a contract position must be either long or short, which will show up as either positive or negative on the vertical. There is always a net zero-sum of contracts in a derivatives market. In our simulation with no leverage and $\kappa = 1$, contracts as a proportion of wealth appear to be relatively steady. In this market, trades are small, and most agents remain in the same or a similar position, even prior to $t = 1{,}000$. This is quite different when the margin requirement is reduced to 33%, $\kappa = 3$, for both Figures 4 and 5. Fewer agents are satisfied with their position as a larger number of traders remain below their leveraged limit and as more trading takes place.[7]

In both simulations, those agents with the more extreme price expectations will spend most of their trading time on or close to their limit — more so than those with "more accurate" price expectations (i.e., expectations that are closer to p^*). Although promoting leveraged trading by reducing the margin requirement stimulated trade activity, it did nothing to the average or standard deviation of the mid-price (both simulations approximated a mean of 151 and standard deviation of 1).

---

7 For trading points $t = 100$ to $t = 1{,}000$, the average trade when $\kappa = 1$ was $0.08x$ contracts. This compares to an average trade volume of $6x$ contracts when the margin is reduced to 33%, $\kappa = 3$.

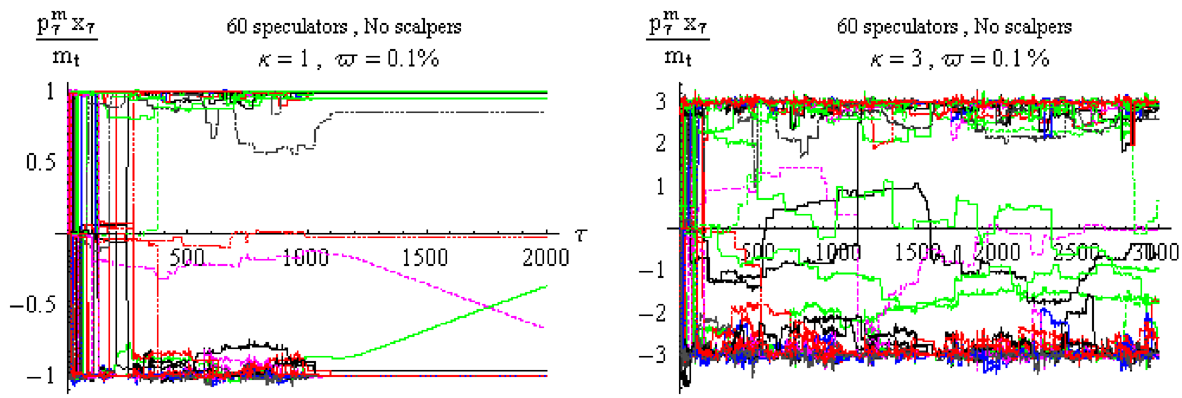

FIGURE 5 Leverage position by each speculator over time (long is positively leveraged, and short is negatively leveraged; margin requirements are 100% versus 33%)

Within our simple model, there is some support for Friedman's (1953) Darwinian suggestion that speculation is efficient because the noise speculators die out and fundamental speculators prosper. This argument is often used in policy circles for the reduction of margin requirements to leverage those traders with information. This assumes that there is only one right price, which is the case in our model, represented by p*. Our model confirms, as shown in Figure 6, that those traders with expectations furthest away from p* do lose money trading. Despite all trader expectations being drawn from a population with a mean of 150, the aggregate income for the group of 20 speculators with a population standard deviation of $\sigma = 5$ is shown to decline much faster in the lower margin environment of $\kappa = 3$ than $\kappa = 1$. The agents with the least noise (smallest dispersion) of expectations around the population mean have greater capital gains because they are more likely, as a group, to be paid for providing immediacy (placing limit orders rather than market orders), and this compensates their cost of trading.

We experimented with lowering the margin requirement from 100% to 25%, raising the transaction costs of trading from 0.1% to 0.5%, and comparing a market without and with scalpers. This was done for a group of 60 speculators and two hedgers with the same wealth and expectations for each simulation. Table 1 shows results for single trace runs of each scenario. In each column, the kurtosis of the mid-price is given first, then the median of the bid-ask spread as a percentage of the mid-price is presented in brackets. We chose mid-price kurtosis since this is representative of the price volatility relevant to the exchange in setting the margin requirement (see Ussher 2004) and since both kurtosis and the bid-ask spread may be considered as measures of liquidity in terms of price resiliency or cost of transacting, respectively.

This evidence, although anecdotal prior to a proper Monte Carlo analysis, suggests that in a market with no scalpers, lowering the margin to 25%, or $\kappa = 4$, may lower the median bid-ask spread and increase the price kurtosis, the more so when taxes are high. The time series of this simulation for prices and leverage is presented in Figure 7. By increasing the costs of transacting, we again reduce trading activity, despite the low margin requirement.[8] Trading is quite orderly, as shown by the leverage time series, with bursts of activity when prices readjust. In the price

8 Under a regime of $\kappa = 4 : \varpi = 0.1\%$, the average trade was 8 contracts, whereas for $\kappa = 4 : \varpi = 0.5\%$, it was 3.6 contracts.

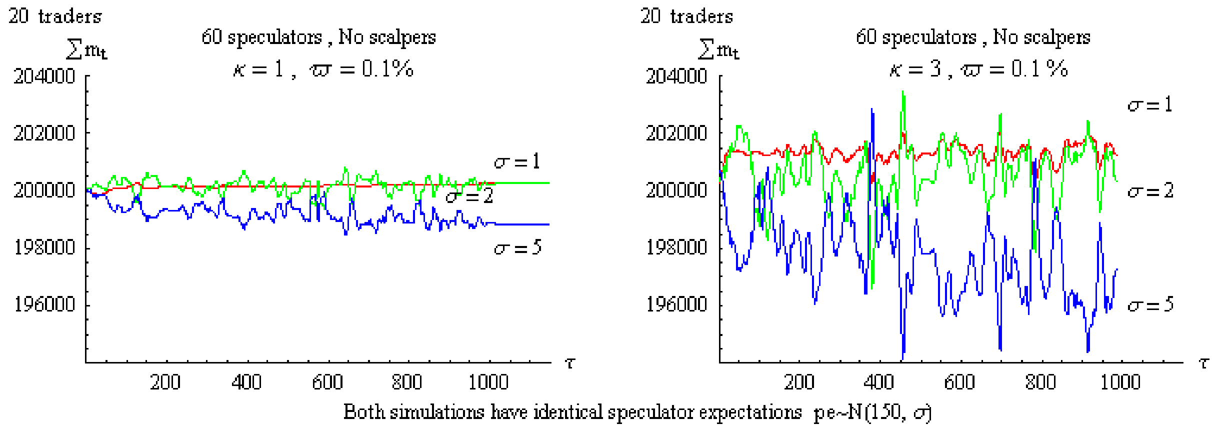

FIGURE 6 Aggregated wealth of each group of 20 speculators, smoothed with a 15-period moving average (margin requirements are 100% versus 33%)

TABLE 1 Measures of *liquidity* from four simulations across different tax rates, margin requirements, and scalpers or no scalpers[a]

	κ			
	No Scalpers		10 Scalpers	
ϖ	1	4	1	4
0.1%	11	20	11	6
	(0.5)	(0.2)	(0.4)	(0.9)
0.5%	23	35	4	5
	(0.5)	(0.5)	(0.5)	(0.8)

[a] Values are kurtosis of mid-price and, in brackets, the median of the bid-ask spread as a percentage of the mid-price. All statistics drop the first 100 price realizations and are from t = 100 to t = 1,000.

series graph, we see that the bid-ask spread is still ϖp and that this tends to flatline more often with the higher tax. This explains the higher price kurtosis value of 35 versus 20 for the same margin requirement, as kurtosis is a measure of the peakness of the price distribution. With the higher tax, a larger number of speculators remain below their leverage limit than when the tax is 0.1%. The CDA tâtonnement price process still detects the Walrasian price, but the mid-price is not as closely matched post-t = 1,500.

The market without scalpers appears to suggest that transaction taxes will increase the level of kurtosis in a market. This may support Davidson’s (1997) claim that a Tobin tax will not

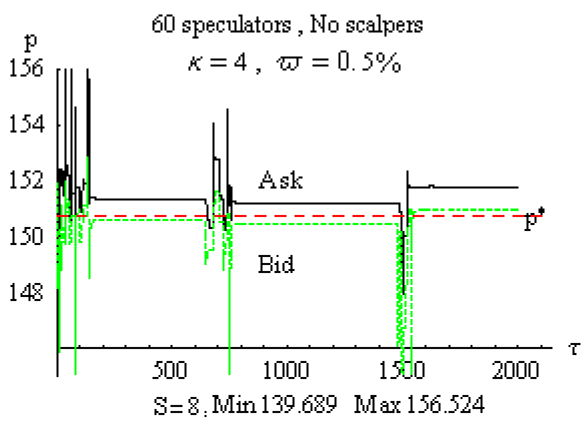

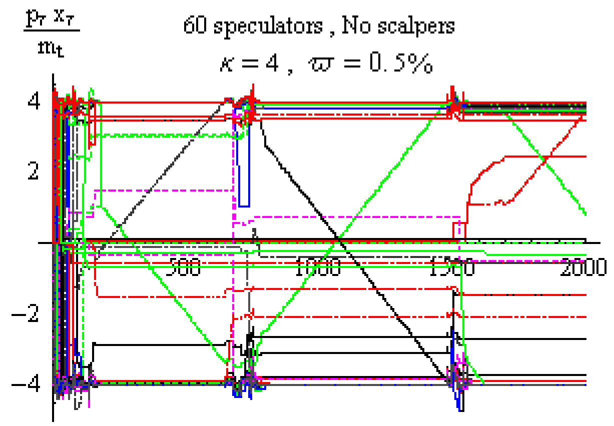

FIGURE 7 Prices and leverage positions for speculators (margin requirements are κ = 4 and tax ϖ = 0.5%)

reduce price volatility but will only reduce market liquidity. In our highly leveraged speculative market, it is interesting to see just what causes volatility, as it has nothing to do with changing price expectations. We shall consider in detail the increase of price volatility from $t = 640$ to $t = 780$ in Figure 8 for the same high-tax, high-leverage market. The size of the black and red dots are representative of a log transformation of the trade size. A black dot is a market order by a speculator to either buy at the ask because he expects prices to rise, or to sell at the bid because he expects prices to fall. A red dot represents a liquidation of a position in order to meet margin requirements and pay for losses following an adverse price change. A liquidation trade is usually on the backward-bending part of the demand function. Since most traders have already taken up their position on the basis of expectations, a lot of the trades that take place are red dots. Prior to $t = 650$, transactions were randomly distributed between the bid and ask, and trade size averaged around 1.2 contracts, with the spread equivalent to the transaction cost, 0.5% of the price.

In studying the above price destabilizations, we have found that buys usually follow buys and sells follow sells. Following what Hasbrouck (2003, page 13) noted for stock market data, trades at the bid tend to maintain trades at the bid, and trades at the ask maintain trades at the ask. In our model, this has nothing to do with expectations formation or trend following behavior; rather, it is due to collateral constraints causing credit crunches and the forced search for immediacy through market orders due to margin calls. A sudden downward bid is not brought back up but rather stays for a time at that low level. The trades at the low bid are followed by more price transactions at that low bid, despite expectations having not changed.

While a Tobin-like tax appears to add price volatility to our speculative market without scalpers, it may be a possible to use this policy to stabilize markets with scalpers, as seen in Table 1. In these simulations, all scalpers have the same inventory limit of K = 10 each. Scalpers are excluded from holding margin or paying transaction costs, as they are presumed to be local exchange members and do not go through brokers. This allows them to provide immediacy even when taxes are high. Figure 9 is a section of time series for two simulations with low versus high tax regimes — 0.1% versus 0.5% on a one-way trade. The grey trades are market orders done by scalpers. Scalpers will place a market order only when their inventory has reached its limit of K = 10; at all other times, scalpers provide limit orders at the bid and ask (all limit orders are the counter trade to market orders).

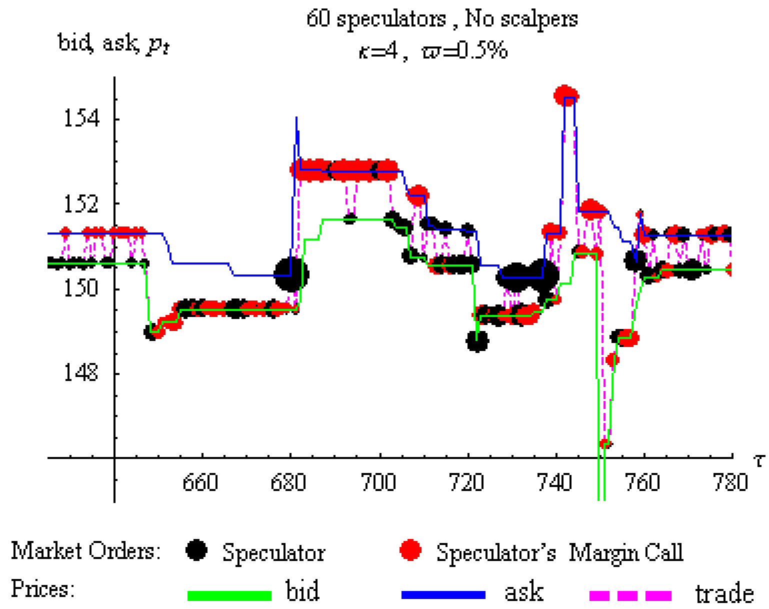

FIGURE 8 Prices and trade size by speculators and hedgers (trades are market orders only).

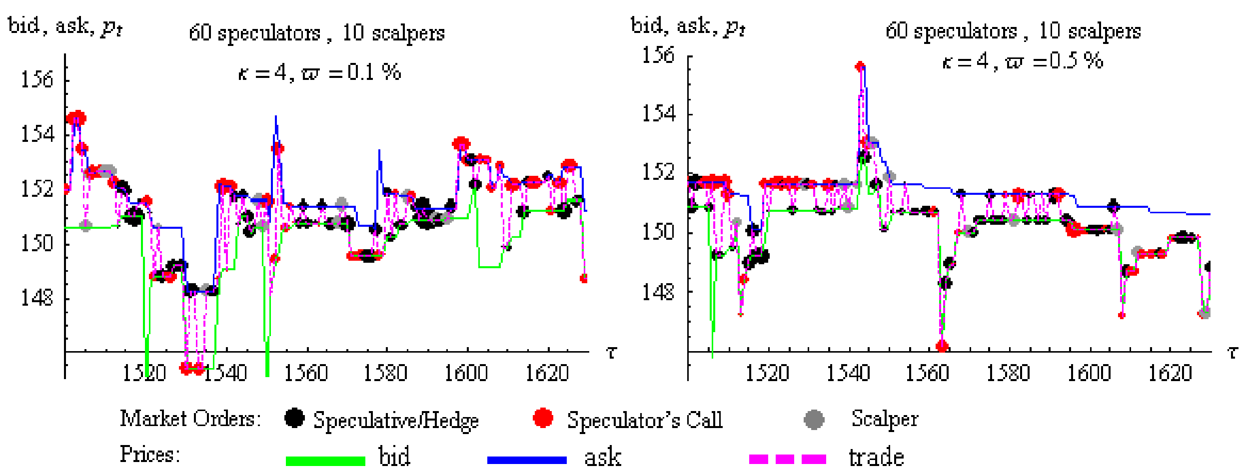

FIGURE 9 Prices and trade size by speculators, hedgers, and scalpers (who make market orders only)

An interesting result from introducing scalpers is that both market prices and activity are less sensitive to changes in the tax rate. Unlike the case with no scalpers, an increase in transaction taxes from 0.1% to 0.5% did not reduce trade activity and trade size in our market with scalpers. It also did not increase the average bid-ask spread, nor the level of kurtosis. While the lowest spread in each simulation with scalpers reflects the low or high tax rate ϖp, the average and median bid-ask spread in the two simulations are very similar.[9]

As specified by the model, scalpers do not narrow or determine the size of the spread, given their overly simplistic trading rules. Instead, myopic speculators determine the narrowness to which the bid-ask spread converges on the basis of their one-way cost of trading. Scalpers, who have no opinion about fundamental prices, will provide liquidity not only to those traders who stabilize markets but also to those traders who destabilize markets. Scalpers appear to not only maintain the bid-ask spread but also indirectly widen it when a *noise* or uninformed trader trades. Scalpers are ready to accommodate such prices. Scalpers may widen the spread even more by reversing their own excessive inventory position from accommodating the "uninformed trader," lifting or hitting the opposite side of the market, creating a zero limit order, and leading to a widening of the spread again. This will mean that the mid-price may be mean-reverting, but the spread initially widens on both sides before narrowing.

However, the larger bid-ask spread in scalper markets does not seem to indicate less liquidity. Instead, price resiliency (low kurtosis) is improved, even in markets with higher transaction costs. Scalper markets might be considered to be more liquid if one considered volume as an indicator of market liquidity or the proportion of speculators below their leverage limit and desiring to trade, as in Figure 10. By exempting market makers from a Tobin tax, this policy might still be successful in removing noise traders, but not at the cost of liquidity and less price resiliency.

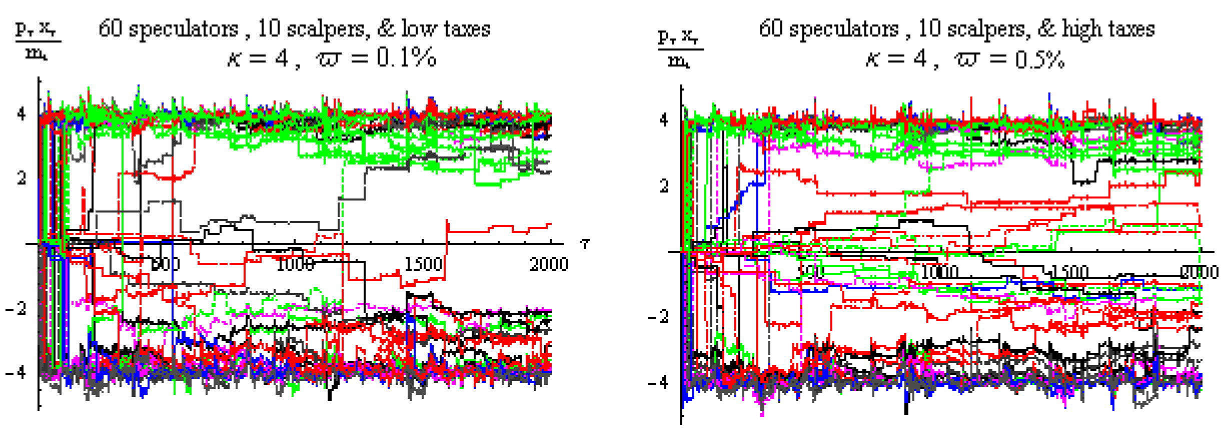

FIGURE 10 Leverage positions for speculators in a market with margin requirements of κ = 4 and scalpers (comparison of 0.1% versus 0.5% transaction tax)

---

9 For the simulation with 10 scalpers and $\kappa = 4$: $\varpi = 0.1\%$, the mean and median bid-ask spreads as a percentage of price were 0.92 and 1.3, with a minimum of 0.11. For the simulation with 10 scalpers and $\kappa = 4$: $\varpi = 0.5\%$, the mean and median bid-ask spreads as a percentage of price were 0.8 and 1.2, with a minimum of 0.5.

CONCLUSION

This attempt to unpack the Walrasian black box of a speculative futures market has shown that even with the inclusion of leveraged trading, short selling, RTGS, and transaction costs, our market is ultimately stable and reverts to the Walrasian equilibrium price in the long run. A market of speculators with diverse expectations and no market makers will produce a competitive bid-ask spread that fluctuates and often narrows to the cost of transacting ϖp. Adding market makers, or scalpers, however, does create greater price resiliency and may allow for policies that raise transaction costs without adding to market price volatility.

Margins and transaction taxes directly affect the distribution of market orders to limit orders for a fixed distribution of expectations. Without scalpers, lowering the margin requirement increases the sensitivity of demand to price changes and increases the degree of trading activity in the market.

When no scalpers are present, speculators with expectations that are closer to the long run price p^*, especially those within the tax threshold, gain from trading as a result of their ability to play the role of market maker and earn a spread from the noisier traders. The greater the leverage that is allowed in the market, the more impoverished the noise traders become. This follows Friedman's (1953) Darwinian process that low margins quickly sort out the "smart" traders from the "noisy" ones. A larger transaction tax can increase the peakness and fat tails of the price series, making it difficult for exchanges to use the observed probability of prices to set margin requirements. An increase in the tax threshold increases the number of speculators who compete to offer limit orders, which does tend to stabilize prices despite the higher kurtosis. While prices may be more stable in this market, they are less resilient (higher kurtosis).

When scalpers are included in the trading mix, the bid-ask spread is wider, order flow is turbulent, and trading volume is much greater. Despite the larger spread, this may be characterized as a more liquid market, and mid-prices are dramatically more resilient. Changing transaction taxes has less impact on both trade activity and price volatility. In this market, raising taxes can accomplish the goal of impoverishing traders with expectations far away from p^* without adding to extreme price movements or being detrimental to liquidity.

APPENDIX

The risk-neutral speculator maximizes next period's expected wealth (1). The first four of our boundary constraints represents the limit on a speculator's investment by the margin requirement when one is short in futures (2) and (3) versus the extent to which futures can be bought long (4) and (5). We have two each of these restrictions to take into account the one-way tax on both buys and sells $\varpi\, p_t|(x_t - x_{t\text{-}1})|$ for speculator i. If the transaction tax is positive, then this boundary constraint will be slack. This dual tax restriction also impacts the budget constraint (6) and (7). The bankruptcy conditions (8) through (12) stop money wealth from going below zero.

For speculator *i*:

Maximize:

$$\pi^{e}_{t+1} = \left(p^{\theta} - p_t\right)x_t + m_t \tag{1}$$

Subject to:

$$p_t x_t \geq -\kappa\left[\left(p_t - p_t^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} - \varpi\; p_t\left(x_t - x_{t-1}\right)\right] \tag{2}$$

$$p_t x_t \geq -\kappa\left[\left(p_t - p_t^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} + \varpi\; p_t\left(x_t - x_{t-1}\right)\right] \tag{3}$$

$$p_t x_t \geq \kappa\left[\left(p_t - p_t^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} - \varpi\; p_t\left(x_t - x_{t-1}\right)\right] \tag{4}$$

$$p_t x_t \geq \kappa\left[\left(p_t - p_t^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} + \varpi\; p_t\left(x_t - x_{t-1}\right)\right] \tag{5}$$

$$m_t \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} - \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{6}$$

$$m_t \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} + \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{7}$$

$$0 \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} - \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{8}$$

$$0 \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} - \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} - \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{9}$$

$$0 \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} + \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} + \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{10}$$

$$0 \leq \left(p_t^m - p_{t-1}^m\right)x_{t-1} - \left(p_t^m - p_t\right)\left(x_t - x_{t-1}\right) + m_{t-1} + \varpi\; p_t\left(x_t - x_{t-1}\right) \tag{11}$$

$$m_t \geq 0 \tag{12}$$

ACKNOWLEDGMENTS

The author would like to thank Duncan K. Foley, Rob Axtell, and Salih Neftci for helpful comments and suggestions. All mistakes are my own.

REFERENCES

Amihud, Y., and H. Mendelson, 1980, “Dealership Markets: Market-making with Inventory,” *Journal of Financial Economics* 8:31–53.

Arrow, K., and G. Debreu, 1954, “Existence of an Equilibrium for a Competitive Economy,” *Econometrica* 22(3):205–290, July.

Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer, and P. Taylor, 1997, “Asset Pricing under Endogenous Expectations in an Artificial Stock Market,” in A.W. Brian, S.N. Durlauf, and D.A. Lane, eds., *The Economy as an Evolving Complex System II: Proceedings* (Santa Fe Institute Studies in the Sciences of Complexity, Vol. 27), Addison-Wesley, Redwood City.

Bernstein, P.L., 1987, "Liquidity, Stock Markets, and Market Makers," *Financial Management*, Vol. 16, summer.

Black, F., 1986, "Noise," *The Journal of Finance* 41(3):529–543, July, Papers and Proceedings of the Forty-Fourth Annual Meeting of the America Finance Association, New York, NY, Dec. 20–30, 1985.

Chan, N., B. LeBaron, A. Lo, and T. Poggio, 1998, "Information Dissemination and Aggregation in Asset Markets with Simple Intelligent Traders," working paper, Massachusetts Institute of Technology.

Cohen, K.J., S.F. Maier, R.A. Schwartz, and D.K. Withcomb, 1979, "Market Makers and the Market Spread: A Review of Recent Literature," *Journal of Financial and Quantitative Analysis* 14(4):813–835.

Davidson, P., 1997, "Are Grains of Sand in the Wheels of International Finance Sufficient to Do the Job When Boulders Are Often Required?," *Economic Journal* 107(442):671–686.

Demsetz, H., 1968, "The Cost of Transacting," *Quarterly Journal of Economics* 82:33–53.

Engle, R,. and J. Lange, 1997, "Measuring, Forecasting and Explaining Time Varying Liquidity in the Stock Market," Discussion Paper No. 97-21R, University of California at San Diego, CA, Nov.; available at ftp://weber.ucsd.edu/pub/econlib/dpapers/ucsd9712r.pdfd.

Flemming, M.J., 2003, "Measuring Treasury Market Liquidity," *FRBNY Economic Policy Review*, pp. 83–108, Sept.

Farmer, J.D., J. Geanakoplos, S. Thurner, and D.J. Watts, 2004, "Optimal Leverage, Margin Calls, and the Emergence of Bank-hedge Fund Crises," mimeo, Santa Fe Institute.

Farmer, J.D., and S. Joshi, 2002, "Price Dynamics of Common Trading Strategies," *Journal of Economic Behavior and Organization* 49(2):149–171, Oct.

Friedman, M., 1953, "The Case for Flexible Exchange Rates," *Essays in Positive Economics*, The University of Chicago Press, Chicago, IL.

Garman, M., 1976, "Market Microstructure," *Journal of Financial Economics* 3:262–267.

Gode, D.K., and S. Sunder, 1993, "Allocative Efficiency of Markets with Zero Intelligence (ZI) Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy* 101(1):119–137.

Harris, L., 2003, *Trading and Exchanges: Market Microstructure for Practitioners*, Oxford University Press, New York, NY.

Hasbrouck, J., 2004, *Empirical Market Microstructure: Economic and Statistical Perspectives on the Dynamics of Trade in Securities Markets*, teaching notes for B40.3392, fall 2003; available at http://www.stern.nye.edu/~jhasbrou. Accessed May 3, 2004.

LeBaron, B., 2000, “Agent Based Computational Finance: Suggested Readings and Early Research,” *Journal of Economic Dynamics and Control* 24(5–6):679–702.

LeBaron, B., 2002, “Calibrating an Agent-based Financial Market,” working paper.

Martel, R.J., 1996, “Heterogeneity, Aggregation and a Meaningful Macroeconomics,” pp. 127–143 in D. Colander, *Beyond Microfoundations: Post Walrasian Macroeconomics*, Cambridge University Press, Cambridge, England.

O’Hara, M., 1997, *Market Microstructure Theory,* Second Edition, Blackwell, Malden, MA.

Pollin, R., D. Baker, and M. Schaberg, 2001, “Securities Transaction Taxes for U.S. Financial Markets,” Working Paper 20, Political Economy Research Institute, University of Massachusetts, Amherst, MA.

Schiller, R., 2000, “Margin Calls: Should the Fed Step in?” *Wall Street Journal*, April 10; available at http://cowles.econ.yale.edu/news/shiller/rjs_00-04-10_wsj_margin.htm. Accessed Dec. 3, 2003.

Schlesinger, T., 2000, “Dealing with Asset Bubbles: The Fed’s Changing Story and the Historical Record,” *Financial Markets Center;* available at http://www.fmcenter.org/fmc_superpage.asp?ID=223. Accessed Dec. 3, 2003.

Schwartz, R., 1988, “A Proposal to Stabilize Stock Prices,” *Journal of Portfolio Management* 15:1, fall.

Schwartz, R.A., and N. Economides, 1995, “Electronic Call Market Trading,” *Journal of Portfolio Management* 21(3):10–18, spring.

Silber, W.L., 1984, “Marketmaker Behavior in an Auction Market: An Analysis of Scalpers in Futures Markets,” *The Journal of Finance* 34(4):937–953, Sept.

Stoll, H.R., 1978, “The Supply of Dealer Services in Securities Markets,” *Journal of Finance* 33:1133–1151.

Stoll, H.R., 1985, “Alternative View of Market Making,” in Y. Amihud, T. Ho, and R. Schwartz (eds.), *Market Making and the Changing Structure of the Securities Industry,* Lexington Health, Lexington, MA.

Stoll, H.R., and T.S.Y. Ho, 1981, “Optimal Dealer Pricing under Transactions and Return Uncertainty,” *Journal of Financial Economics* 9:47–73.

Tobin, J., 1974, “The New Economics One Decade Older,” *The Janeway Lectures on Historical Economics*, Princeton University Press, Princeton, NJ.

Ussher, L.J., 2004, “An Agent Based Model of a Speculative Futures Market,” submitted in partial fulfillment of the requirements of the Ph.D. in economics, Graduate Faculty, New School University.

Yang, J., 2002, “The Efficiency of an Artificial Stock Market with Heterogeneous Intelligent Agents,” *Evolutionary Computation in Economics and Finance*, S.H. Chen (ed.), Springer-Verlag, pp. 85–106.

ASSESSING EX-URBAN RESIDENTIAL MARKETS: AN AGENT-BASED MODEL

L. YIN,* School of Architecture and Planning,
State University of New York at Buffalo, Buffalo, NY

ABSTRACT

Many of the traditional models used by urban planners are not well-tailored to the ex-urban residential market. Recent literature recommends that people's preferences on natural amenities, density (neighbor avoidance/large lot), and accessibility play important roles in shaping the ex-urban residential market. In this research, I explore the use of an agent-based approach to investigate and examine how ex-urban residential location patterns may result from the behaviors of decentralized and heterogeneous individual households that reflect their preferences influenced by these three drivers of ex-urban development. Three agent-based models are constructed to detect the dynamic ex-urban sprawl influenced by the three drivers, one at a time. Simulation results suggest that the agent-based models built in this research have a potential to represent the ex-urban residential market at a reasonably high level of accuracy.

Keywords: Exurban development, agent-based model

INTRODUCTION

Computer models of urban growth have a long history. Yet many of the traditional approaches used by urban modelers are not well-adapted to ex-urban environments. Most of these models assumed that households have similar location preferences — close to work (Alonso 1964) — while recent literature recommends that people's preferences for natural amenities, density (neighbor avoidance/large lot), and accessibility play important roles in shaping the ex-urban residential market (Davis et al. 1994; Nelson 1992; Riebsame et al. 1992; Irwin 1998). Agent-based modeling (ABM) can be used to represent the behaviors of heterogeneous homeowners and the evolution of every individual parcel at a relatively high level of complexity by using a process-based approach. Furthermore, ABM provides a means to assess temporal, decentralized, and autonomous ex-urban residential development decision making at the household level and link these decisions to aggregate ex-urban land use changes (Parker et al. 2003).

In this research, I explore the use of an agent-based approach to examine ex-urban residential location and to investigate how ex-urban development patterns may emerge from the behaviors of decentralized and heterogeneous individual households, reflecting their preferences influenced by the three drivers of ex-urban development. Three agent-based models are

* *Corresponding author address*: Li Yin, Department of Urban and Regional Planning, School of Architecture and Planning, State University of New York at Buffalo, 116 Hayes Hall, 3435 Main Street, Buffalo, NY 14214; email: liyin@ap.buffalo.edu.

constructed to detect the dynamic ex-urban sprawl influenced by different location preferences on accessibility, amenities, and density (neighbor avoidance, large lot development).

BUILDING AN AGENT-BASED MODEL

In this research, I formulate two types of agents to represent two types of households searching for ex-urban residential locations: commuters and second-home owners. Commuters are the traditional type of households in conventional urban models who value a short commuting distance to work most. Second-home owners or amenity-seekers reflect nontraditional types of households found in ex-urban areas whose location choices are strongly influenced by density level and site amenities (e.g., the presence of open space or a stream, large lots).

I build three theoretical models by using Repast, an ABM platform based on Java, and I explore household land conversion rules according to three types of location preferences (Table 1). Beginning with a model including only one type of household with a preference for urban accessibility only, I add the second type of household — second-home owners – favoring amenities and density in the second and third model, respectively. I examine the effects of accessibility, amenities, and density one at a time.

TABLE 1 Locational preferences by different types of households

Locational Preference Priority	Second-home Owners	Commuters
First	Proximity to natural amenities (public land, lakes, or streams)	Accessibility (proximity to jobs or highways)
Second	Density (quiet environment or large lot)	Density (high level of development in the neighborhood)
Third	Accessibility (proximity to roads or shopping)	Proximity to natural amenities (open space, lakes, or streams)

Agent-based models are built in three phases:

- Model I Assess the effects of accessibility on location.
- Model II Assess the effects of amenities on location.
- Model III Assess the dynamic effects of density/lot size preferences on location. Second-home owners are assumed to be space sensitive; they like a large lot. Commuters are space-neutral.

The first and the second model are designed to detect the static effects of accessibility and natural amenities, respectively. The third model simulates dynamic effects of density with respect to location. These three models are built on an abstract grid. It is a two-dimensional array of regular spaces represented as a mosaic of grid cells. An ASCII file is created and imported to Arcview to create the abstract grid. It has 150×150 cells, with a resolution of 100×100 meters per grid cell, as shown in Figure 1. This grid is the initial state of the development (i.e., the development at time step 0). All three models will be simulated on the basis of on this abstract grid. Road network, two rural places, public owned lands, a lake, and some streams are drawn randomly and added to the grid. This grid also sets up the basis for calculating accessibility, justify amenities and density variables. One household moves in to the environment/the abstract grid at each time step. Each time step or iteration in a Repast model can be any time period it takes for the next development activity taking place.

The geographic information system (GIS) plays a role in data compiling, processing, and spatial database building. The multi-agent-based modeling tool Repast simulates the temporal and spatial land conversion from one state (undeveloped) to another (developed) according to a set of predefined transitional rules based on households' preferences for accessibility, amenities, and density.

RESULTS AND DISCUSSION

In each of the Model I, II, and III runs, households' preferences and behaviors are adjusted in accordance with the purpose of the model. Commuters and second-home owners enter the environment (the abstract lattice) and interact with it. One household takes up one site or cell in each time step (iteration) depending on the bid it offered according to its location preferences and the result of the bid.

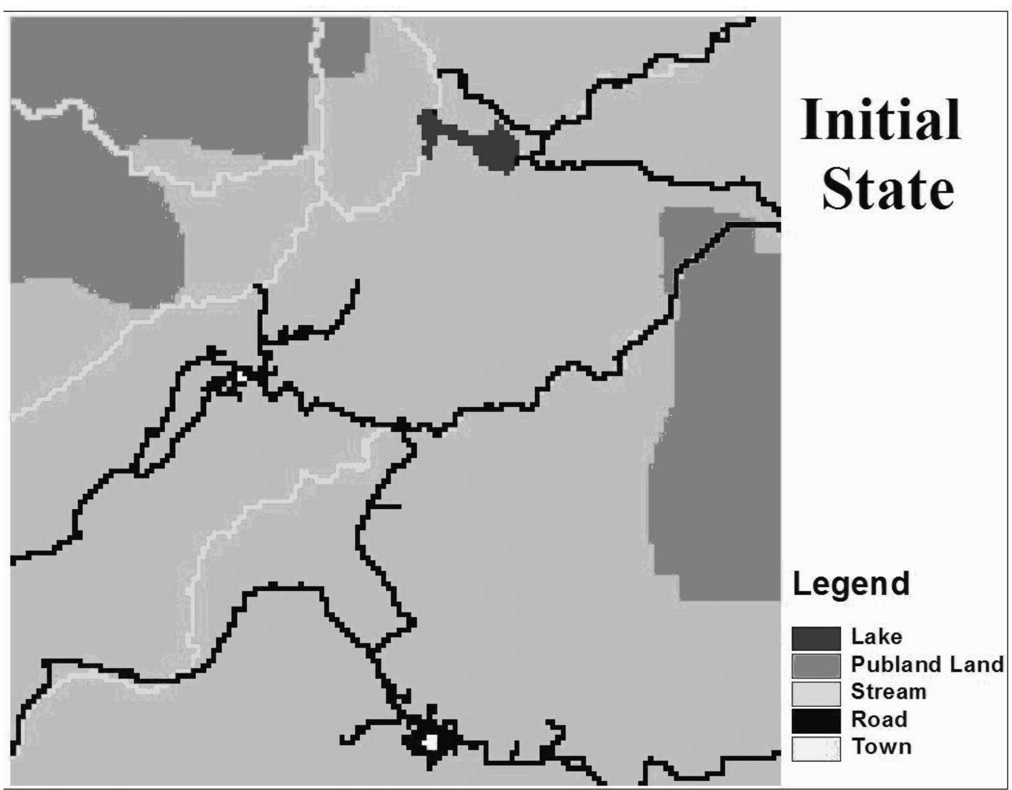

FIGURE 1 Initial state

Model I

Model I results show that build-out first occurs in the areas around the rural places/employment centers where accessibility is considered the highest and then extends to the areas along the transportation corridors. Since commuters favor accessibility factors, all the development gets pulled toward the areas with a high level of accessibility.

Model II

In Model II, because second-home owners chose to develop places in close proximity to natural amenities and commuters chose sites with good accessibility, clusters emerged along transportation corridors and around job centers, as well as in the areas with rich natural amenities (i.e., lakes and public land) or with easy access to both road and natural amenities.

Model III

Model III results show that at the early stage when there are a large number of empty places, second-home owners can find various sites or cells that satisfy their needs to a great extent. They tend to offer higher bids and win their bids more often. Therefore, development patterns are scattered as a result of second-home owners' bid triumph. After some time, development is seen in two extremes: clustering on cities and roads, and dispersion with some tendency to be close to roads. However, there are still more cells developed by commuters than by second-home owners.

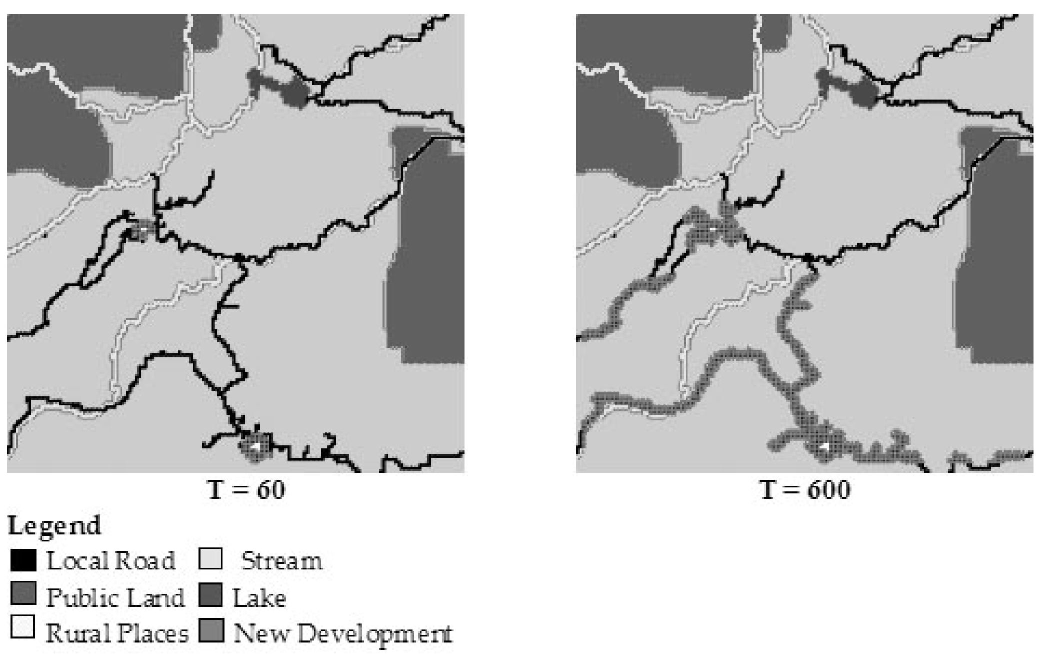

FIGURE 2 Model I

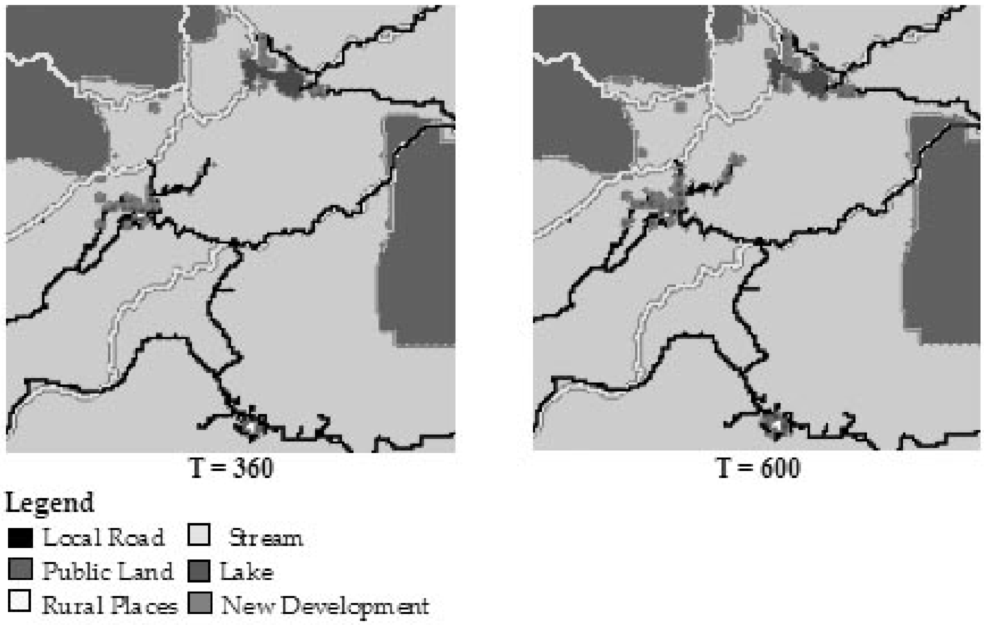

FIGURE 3 Model II

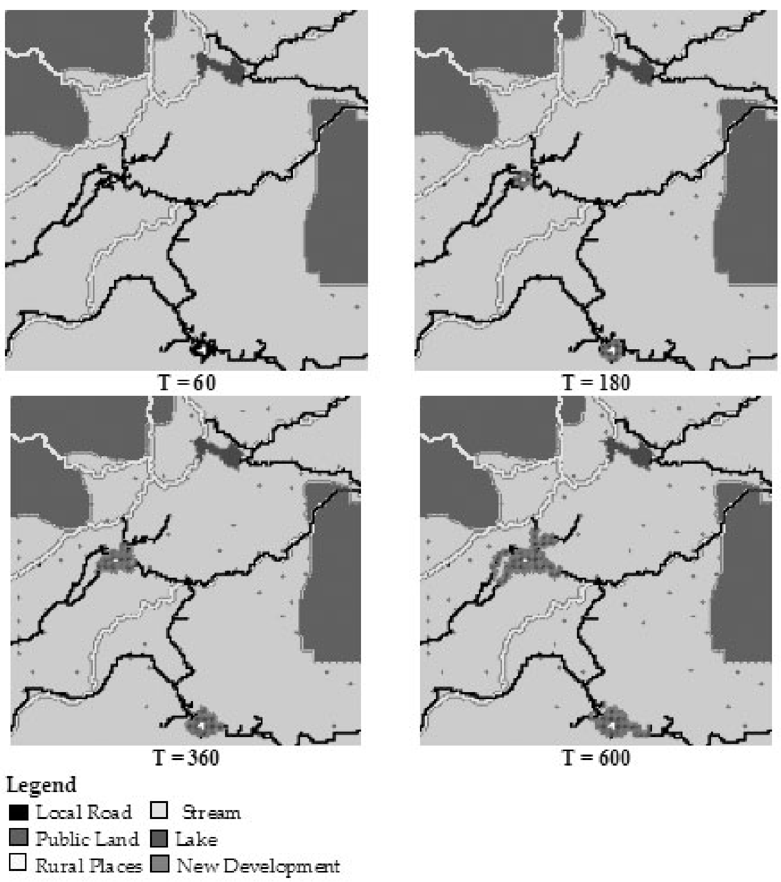

FIGURE 4 Model III

When development gets more and more densified and when accessibility clusters are stretching out, second-home owners find fewer and fewer cells that suit their needs, while commuters still find many sites that highly satisfy their needs. Some time after scattered development is pushed into the areas away from roads because of lack of space for second-home owners, large clusters emerge around areas with good accessibility. This is because commuters win bids more often now.

SUMMARY AND CONCLUSIONS

Model I illustrates how clusters emerged around the urban places/employment centers and along the transportation corridors because of households' preferences for accessibility. However, when compared with the actual development pattern, it concentrates development too tightly around rural places and transportation networks. In Model II, growth goes into areas with either rich amenities or easy access to highways and jobs, reflecting the attractiveness of sites with the presence of public lands or water bodies to second-home owners or of sites in close proximity to existing urban services and accessibility to major highways to commuters. The mix of preferences for amenities and accessibility from different types of households makes the model a more accurate predictor of ex-urban development than models based exclusively on accessibility for only one household type. However, in comparison with the actual development pattern, development tends to be too tightly around natural amenities.

Model III illustrates two extremes of development density patterns: cluster and dispersion resulting from different household locational preferences. Commuters' preferences for the limited available areas with close proximity to work and transportation networks fuel higher-density development in lands surrounding cities and highways; behavior preferences of second-home owners for large lots, spacious and isolated spaces, and neighbor avoidance push developments into the wilderness, which may become seeds for later development. Model III indicates that patterns of ex-urban growth density are influenced by factors such as spacing of lots and distance from infrastructure. Households interact with each other in ex-urban locational decision making. Second-home owners skip over properties close to other developed sites and obtain bigger isolated lots further out. This creates pressure for low-density development and a persistent dispersion pattern, and a significant and disproportionate reduction in the average density of development at the aggregate level. Model III also demonstrates development phasing effects at which ex-urban development shifts from a land market dominated by second-home owners to commuters. Yet the switch occurs only gradually after second-home owners reach the density threshold.

The agent-based models built in this research help researchers understand how ex-urban residential location patterns may result from decentralized and heterogeneous individual households' preferences for natural amenities, low density, and accessibility. Simulation results suggest that these models have the potential to represent the ex-urban residential market at a reasonably high level of accuracy.

REFERENCES

Alonso, W., 1964, *Location and Land Use: Toward a General Theory of Land Rent*, Harvard University Press, Cambridge, MA.

Davis, J.S., A.C. Nelson, and K.J. Dueker, 1994, "The New 'Burbs: The Exurbs and Their Implications for Planning Policy," *Journal of American Planning Association* 60:45–59.

Irwin, E., 1998, *Do Spatial Externalities Matter? The Role of Endogenous Interactions in the Evolution of Land Use Pattern*, unpublished Ph.D. dissertation, University of Maryland.

Nelson, A.C., 1992, "Characterizing Exurbia," *Journal of Planning Literature* 6(4):350–368.

Parker, D.C., S.M. Manson, M.A. Janssen, M.J. Hoffmann, and P. Deadman, 2003, "Multi-agent Systems for the Simulation of Land-Use and Land-Cover Change: A Review," *Annals of the Association of American Geographers* 93(2).

Riebsame, W.E., H. Gosnell, and D.M. Theobald, 1996, "Land Use and Landscape Change in the Colorado Mountains I: Theory, Scale, and Pattern," *Mountain Research and Development* 16(4):395–405.

ENVIRONMENTAL RACE TO THE BOTTOM: MIXED AGENT-BASED MODEL

A. ALEXEEV* and D.H. GOOD, School of Public and Environmental Affairs,
Indiana University, Bloomington, IN

ABSTRACT

Whether environmental policy is best determined by the federal government or whether it is best set by the states is far from settled. On the one hand, federal policy creates uniformity and eliminates potentially wasteful competition or spillovers between states. On the other hand, states can adapt environmental policies to reflect their own circumstances. However, states might also use this flexibility to pursue goals other than protecting the environment and goals that might even be environmentally harmful. This is the so-called race to the bottom. The purpose of this work is to develop an agent-based modeling approach that can describe state-level policy-setting behavior. This model forms a policy testing platform in which alternative schemes can be examined to limit the race-to-the-bottom process. At a minimum, this model illustrates that states might adopt bifurcated strategies, which may explain the lack of empirical support for the phenomenon, and that these strategies might develop from stochastic events rather than measurable state differences.

Keywords: Race to the bottom, agent-based, environmental policy

INTRODUCTION

Decentralization of the U.S. environmental policy remains a quite controversial issue of public policy design. While there are many presumed behaviors of states as they pursue economic growth and environmental policies, the empirical evidence is quite mixed. Proponents of strong federal authority argue that if a state is given authority to set its own environmental standards, the state will set lower standards to attract economic activity. This interstate competitive effect is known as the "race to the bottom." On the other hand, their opponents are convinced that decentralization has its advantages. States may use their flexibility to either set policies that are even stricter than federal standards or to adopt alternative methods for achieving the same environmental outcomes that the federal standards would achieve but at a lower social cost by taking advantage of local circumstances.

The number of works quantitatively evaluating the process of decentralization, and consequently the existence or absence of the race to the bottom, is limited in most cases, because systematic, reliable, and timely available empirical data from the states are rarely available. In situations like that, agent-based modeling can be a powerful tool for the study. Agent-based models provide the possibility to flexibly change and control the variables of interest: the biophysical world where agents operate, agents' attributes, and rules of cooperation, punishment, learning, etc. This flexibility and adaptability make it possible to study "what if" questions. For

* *Corresponding author address*: Alexander Alexeev, School of Public and Environmental Affairs, Indiana University, 1315 E. Tenth Street, Bloomington, IN 47405; email: aalexeev@indiana.edu.

example, what if agent preferences or given rules of cooperation change? How, then, may this affect the outcomes, efficiency, and effectiveness of the environmental policy? That broad control in the computer experiment is usually restricted, if not impossible, for most empirical studies because the "what if" question seems to have limited applicability to and consistency with regard to empirical data.

In the present work, we attempt to define the main features of an agent-based modeling approach that may provide us with an insight into state-level environmental policy setting and help us understand how economic development within the states and competition between the states may affect their environmental regulations and vice versa. To define the main attributes of the agents — the rules of their behavior and interaction — we consider a simple model of the state's economy producing one good by using one input and generating hazardous waste. In this approach, we assume the stringency of the environmental policy is defined by the environmental tax, which is endogenous with respect to agents. Finally, we discuss some difficulties in establishing the rules of agents' behavior and possible situations that may lead to complex dynamics.

PRIOR LITERATURE

Federalism and federal-state relations within environmental policy are extensively discussed by many scholars in the field (e.g., Sussman et al. 2002; Scheberle 2004; Braden and Proost 1997). Decentralization of federal power with respect to environmental policy may follow different processes. So-called "delegated" programs offer to the states a partial preemption of primacy in setting environmental quality. Standards established by a state must be at least as strict as the federal standards. Under the Clean Air Act, states are delegated with the right to design old source performance standards while new source performance remains under federal jurisdiction. Under the Clean Water Act, states may also choose to accept primary control over both new and old sources (and some states do) in setting water quality standards. Under the Resource Conservation and Recovery Act (RCRA), states are given broad control over the tax structure and permissible practices for the generation, treatment, disposal, and transportation of hazardous wastes. The federal government does not provide complete funding to the state running its own programs. Sussman et al. (2002) notes the distinct differences in states in choosing primacy as their motivation. These beneficial factors include the following: (1) the state earns flexibility in setting the environmental standards and may provide a higher level of environmental protection; (2) states may set priorities for their programs to take into consideration local circumstances; (3) states can establish their own in-state programs when there is no similar federal program; (4) states can consider their own administrative capacity in setting goals; (5) states may take into account local political and technological concerns; (6) states may establish the respective agency of a different type and size that answers the state's needs. This also benefits federal government, providing a possibility to save some funds for other social programs together with achieving a positive environmental outcome.

But the added flexibility brings the concern that states may set lower environmental standards or postpone the implementation of federal standards (see, for example, Engel 1997 and Revesz 2001). Scholars, advocating stronger federal authority, argue that interstate competition in industrial development causes states to relax their environmental standards in favor of their business communities (e.g., Duerksen 1983). This effect leads to a situation where each state may export the cost of its more lenient clean air standards: more pollution in the air to other states. At the same time, "the economic growth attracted by lower standards occurs exclusively within each state's boarders" Potoski (2001). The literature on these issues follows purely theoretical and purely empirical tracks.

Oates and Schwab (1988, 1996) developed a set of theoretical models. In their models, interstate competition results in states establishing environmental standards at the socially effective level. At the other extreme, Engel (1997) characterizes the problem as a one-shot prisoner's dilemma game. She identifies the crucial difference between the two models. Under her model, the *two* noncooperating states always produce the race-to-the-bottom outcome, while Oates and Schwab's model assumes a number of competitors that is large enough so that the market equilibrium remains unchanged when one of the states attracts capital by lowering its environmental standards. The large number of competitors (states) use the competitive equilibrium assumption that no state is large enough to distort the market. Engel suggests that even with more than two states, some might collude to relax standards and have a nontrivial impact on attracting capital, creating the race-to-the-bottom outcome.

Relatively few papers take an agent-based modeling approach for anything remotely close to this subject. Teitelbaum (1998) uses agents to model the ways in which firms adapt to changed regulatory environments. Teitelbaum showed that the government controls pollution more effectively when firms are given time to prepare for the onset of pollution regulations rather than being surprised by them and that the effects of pollution controls can vary widely across firm types. Batroszchuk and Nakamori (2002) used agent-based modeling together with empirical data to testify on the existence of the environmental Kuznets curve for carbon dioxide emissions for four European countries.

There are a number of purely empirical studies that attempt to test the race-to-the-bottom hypothesis. Potoski (2001) builds a regression model with the dependent variable being the number of criteria pollutants for which a state exceeds national ambient air quality standards (NAAQSs). Only the variables characterizing community-based action appeared to be significant. Another paper (Bond et al. 2004) gives similar results, which show that the level of environmental ambient standards for both air and water are defined by the level of democracy and environmental preferences of agents (citizens, groups, communities) acting in the state. The economic variables, like polluting industry strength, appeared to be insignificant, providing no evidence for the race to the bottom. List and Gering (2000) found no evidence for the race to the bottom in environmental quality from their study of abatement expenditures of the states. Dinan et al. (1999) had the same result for drinking water. In summary, most empirical works do not provide convincing evidence that the environmental race to the bottom exists. However, it is not clear if any of these studies appropriately control for simultaneity and/or selectivity. Policies may be strictest where the pollution problem is the most severe.

The race to the bottom is closely related to the "pollution haven" hypothesis. Several recent studies found empirical evidence for the existence of a pollution haven effect: pollutants are exported to states with less stringent environmental regulations, and firms in polluting industries are more likely to locate to those states. But even more studies do not support this result. Kahn and Yoshino (2004) analyzed pollution intensity and distribution of the bilateral manufacturing trade. The panel data analysis of whether richer or poorer nations specialize in exporting dirty goods done for 1980–1997 in 128 countries for 34 manufacturing industries supports the pollution haven hypothesis. Contrasting with this, one of the more influential studies (Ederington et al. 2004) examined whether trade liberalization affects environmental quality in the United States and found no such relation. Smarzynska and Wei (2004) studied investments flows to 25 developing countries. The study was done on an individual firm level, and it also provided no support to the pollution haven hypothesis. The paper by Copeland and Taylor (2004) includes an extensive literature review in which the authors expressed their vision of the problem and theoretical analysis of the topic. In another paper, Taylor (2004) develops a theoretical model of the pollution haven hypothesis by dividing the hypothesis into a series of logical steps, linking assumptions on exogenous country characteristics to predictions on trade flows and pollution levels. As was mentioned above, the application of the game theory to the pollution problem has some serious technical and conceptual limitations: in reality, many parties participate in setting environmental standards; the setting of standards is a dynamic process that includes different kinds of dynamic interactions between players in both vertical and horizontal planes, etc. (see, for example, Brander 1985).

In summary, one may conclude that the race to the bottom is one of the possible behavioral responses of the state to the decentralization of the environmental policy. There is still no solid empirical proof that states do or do not reveal this type of behavior in setting the environmental standards. Most empirical works conclude that the data being used do not provide evidence of the race to the bottom. Together, they suggest that agent-based models have the potential to contribute a lot to this subject.

ELEMENTS OF THE AGENT-BASED MODEL

In approaching this issue, we assume that the model will consist of n agents representing states (or jurisdictions). We assume that each agent (for example, agent i), produces two aggregated commodities: a clean good x_{ci} and a dirty good x_{di}. Production of the dirty good also implies the production of waste x_{pi}. We also assume that there is only one mobile input to production: labor L_i. The production functions are increasing, concave, and homogeneous of the first order. The use of only one input has a number of modeling advantages for describing economic behavior. First of all, we don't need to set a set of strong assumptions about the mobility of the capital, the comparative advantage of each agent, capital or labor abundance, etc. (see Oats and Schwab 1988) putting limitations on the number of agents. Second, it offers us the ability to describe behavior within and among states without the difficulty of establishing internal and external markets so that we can avoid the necessity of solving for a general equilibrium at each time period. Third, as mentioned above, there is some empirical evidence that the effect (price) of labor on waste generation has a higher magnitude then environmental regulations. With regard to the last point, we assume, consequently, that there is no incentive for an agent to transfer its capital to another agent, which is reasonable unless we don't model the location decision process of the firms.

The environmental regulations are set endogenously by the individual agents (state governments) each time period. We assume that the stringency of the regulations is defined by an environmental tax rate τ_i for the disposal of wastes within the state's boundaries. These tax rates may be set differently by each state. The states can also allocate resources L_{ai} for the abatement of hazardous waste generation. The abatement function $x_{pi} = g(L_{ai}, x_{di})$ increases in x_{di} and declines in L_{ai}. The higher the amount of the dirty good x_{di} that is produced, the more waste x_{pi} is generated. The practicality of this assumption is driven by our desire to include only one kind of agent in the model. If we included other agents (firms, for example), they would decide on the levels of effort to apply to abatement, possibly on the basis of command-and-control-type regulations. The view of firm decision making in this model is that it is very simplistic and nonstrategic. Once their production activities are included in the states' aggregate production, firms simply do what the states want them to do. As a last dimension to state-level behavior, they decide what to do with the wastes that are generated. They may dispose of them within state boundaries, making them subject to the waste disposal tax. Or, alternatively, they may decide to ship them to another state for disposal. This makes them subject to the other state's disposal tax plus any shipping and handling fees. The diagram describing the production-waste generation process for two agents is shown in Figure 1.

Ceteris paribus, there are two main pieces of information that affect the agent's decision to ship waste out of the state that are made by the agent: the total shipment cost and difference in taxes. From a rational point of view, one may consider the net benefit, which is equal to the difference between revenue due to taxes colleted for the waste kept by the agent and the total cost of the shipment to the other agent. These two pieces of information can be summarized by the "vision" φ_{ki} that agent k has of agent *k*. A higher fraction of wastes are sent to states with high vision than with low vision. The waste from state *k* sent to state *i* is described by $x_{pki} = \varphi_{ki} x_{pi}$.

To consider a simple situation, suppose that there are no transportation or handling charges. Under these circumstances, states send waste to whichever state has the lowest tax rate. This is the presumed behavior behind the race to the bottom. Under these circumstances, and with only two agents, if the environmental tax set by the agent-sender *j* is higher than the tax of the agent-acceptor *I*, then $0 < \varphi_{ji} \leq 1$, and zero otherwise. This situation may well be formally described by using the Heaviside function $\varphi_{ji} = \vartheta(\tau_j - \tau_i)$. In general, the cost of shipment is

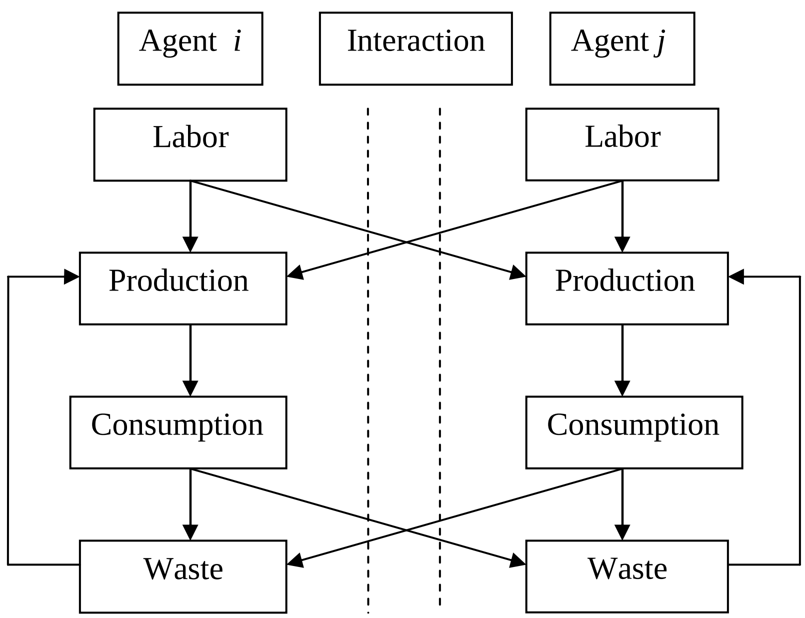

FIGURE 1 Diagram of the production-waste generation process

proportional to the spatial distance d_{ji} between the agents, and we use a logistic allocation function to reflect these two components in the "vision" function:

$$\varphi_{ji} = \vartheta(\tau_j - \tau_i)\frac{\exp(\tau_j - \tau_i - \gamma d_{ji})}{\sum_{all\ k}\exp(\tau_j - \tau_k - \gamma d_{jk})}, \qquad (1)$$

where γ is a marginal transportation cost set to be equal across all agents.

At each period, we assume that each agent maximizes a discounted sum of its utility function $U(X - D)$ over the time horizon. The utility function consists of two parts. The first part relates to the total income for the agent: $x_{ck} + p_d x_{dk} + \tau_k x_{pkr} - (\tau_i + \gamma)x_{pko}$, where x_{ck} has a unit price); τ_k and τ_i are the waste taxes set by the agent k and agent i, respectively; x_{pkr} is the amount of waste received from all other agents; and x_{pko} is the amount of waste shipped to agent i. We assume that agent k will likely ship to only the agent considering the tradeoff between the lower tax and minimal transportation cost. The second part of $U(X - D)$ represents the damage function $D = D(x_{pk} + x_{pkr} - x_{pko})$; x_{pk} is the amount of waste generated by the agent k. It is reasonable that the damage function is an increasing function of the total amount of wastes that are stored in the state, though it is likely that it will increase at a decreasing rate. What is ultimately important in this model is the way that the net amount of wastes disposed of in the state enter into the calculation about total utility.

There some potential problems that may appear to lead the behavior of the agent to the race to the bottom. Assume that environmental tax τ reduces the amount of waste by $(1 - \tau^*)$ as follows (see, for example, Lempert et al. 2003):

$$x_{p,t+1} = \delta(1 - \tau^*)x_{pt} \ . \qquad (2)$$

This transformation describes a regime with a unique trajectory while τ^* is a constant. In reality, some U.S. states have no taxes at all, while others use step-functions (CCH 2002). For the sake of simplicity, we assume that the tax rate may be expressed in the following form:

$$\tau_t^* = \begin{cases} \tau_{1t} x_{pt} & if \quad x_{pt} \le x_{p0t} \\ \tau_{2t} & if \quad x_{pt} > x_{p0t} \end{cases} . \qquad (3)$$

Note that τ_1 and τ_2 measured in different units are a tax ceiling. Also, in reality, τ_t does not change much over time. In the case of the United States, the change of the disposal tax schedule happens once in 5 to 10 years. Hence, while modeling, we can treat it as a constant during up to 10 time cycles. On the other hand, an environmental policy decision-maker ought to have (and we assign him with) the possibility to intervene in the policy at any time cycle. By substituting Equation 3 into Equation 2, one gets the nonlinear recurrent expression:

$$x_{p,t+1} = \delta(1 - \tau_{1t} x_{pt})x_{pt} \qquad (4)$$

for the amount of waste, which itself may produce unexpected behavior.

The main issue in modeling the race to the bottom is how to classify the phenomena. The difficulty here is that the modeler predefines rules according to which agents behave, and then they simply reflect that behavior. The stringency of the environmental regulations is an endogenous parameter with respect to the agent, and each agent itself should define and implement the rule of how to adjust it. But even this simple-enough system may reveal complex, dynamic behavior. Under a certain combination of parameters, one may expect additional complexity because of the interaction between the agents. For example, the use of Equation 3 for pollution calculations may get the system switched to the chaotic regime, producing bifurcations. The same regulations applied to different agents may produce different effects. Recall that most empirical works do not find convincing evidence that the effect takes place because the respective regression coefficients were found to be not significant. The reason may be that the measured value lies in such a zone.

DISCUSSION

The race to the bottom is one of the possible responses of the state to the decentralization of environmental policy. There is no still solid empirical proof that states do or do not reveal this type of behavior in setting environmental standards. Most empirical works conclude that the data being used do not provide evidence of the race to the bottom. On the other hand, it is a presumed part of environmental policy making that this behavior exists. The models that give rise to this kind of behavior (e.g., prisoner's-dilemma-type models) do not reflect the reasonable dynamic context of the problem.

Recall that the main difficulty here is that the decision to relax the environmental regulations should be endogenous for each agent. The decision to relax policy is to be made by an agent itself. It should be clear that if the modeler establishes the respective rule for the agent, the agent will switch to the race to the bottom, following exactly that rule. In this situation, the inverse approach may be reasonable: initially the system of the agents starts its development with agents interacting with each other. Then we randomly choose an agent and relax its environmental standard to study how the dynamics of both the whole system and the agent change. Another approach proposed above is to allow the agent to maximize its utility adaptively to choose the optimal behavior with respect to the environmental tax.

Finally, we discussed some difficulties that are likely to be encountered in the modeling approach used for the race-to-the-bottom phenomenon and possible nonlinear dynamic behavior that complicates the process.

REFERENCES

Bartoszczuk, T., and Y. Nakamori, 2002, "Environmental Kuznets Curve for Some Countries Regression and Agent-based Approach," in C.A. Brebbia and J.F. Martin-Duque, eds., *Air Pollution X* , WIT Press, pp. 93–102.

Bond, C.A., and Y. Hossein Farzin, 2004, *Freedom from Pollution? The State, the People, and the Environmental Kuznets Curve;* available at http://repositories.cdlib.org/are/arewp.

Braden, J., and S. Proost, 1997, *The Economic Theory of Environmental Policy in a Federal System*, Cheltenham, UK: Edward Elgar.

Brander, J., 1985, "Economic Policy Formation in a Federal State: A Game Theoretic Approach," in "Royal Commission on the Economic Union and Development Prospects for Canada," Vol. 63, in R. Simeon, ed., *Intergovernmental Relations,* Toronto, Ontario, Canada: University of Toronto Press, pp. 33–66.

CCH, 2003, *State Tax Guide*, Commerce Clearing House, Inc.

Copeland, B., and S. Taylor, 2004, "Trade, Growth, and the Environment," *Journal of Economic Literature* 42(13):7–71.

Dinan, T.M., M.L. Cropper, and P.R. Portney, 1999, "Environmental Federalism: Welfare Losses from Uniform National Drinking Water Standards," in A. Panagariya et al., eds., *Environmental and Public Economics: Essays in Honor of Wallace E. Oates*, Cheltenham, U.K.: Edward Elgar, pp. 13–31.

Duerksen, C., 1983, *Environmental Regulation of Industrial Plant Sitting: How to Make it Work Better,* Washington, DC: Conservation Foundation.

Ederington J., A. Levinson, and J. Minier, 2004, "Trade Liberalization and Pollution Havens," *Advances in Economic Analysis & Policy* 4(2), Article 6.

Engel, K., 1997, "State Environmental Standard-setting: Is There Race and Is It to the Bottom?" *Hastings Law Journal* 48(2):271–398.

Kahn, M.E., and Y. Yoshino, 2004, "Testing for Pollution Havens inside and outside of Regional Trading Blocs," *Advances in Economic Analysis & Policy* 4(2), 1288.

Lempert, R., S.W. Popper, and S.C. Bankes, 2003, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis,* Santa Monica, CA: RAND.

List, J.A., and S. Gerking, 2000, "Regulatory Federalism and US Environmental Policies," *Journal of Regional Science* 40(3):453–471.

Oates, W.E., and R.M. Schwab, 1988, "Economic Competition among Jurisdictions: Efficiency Enhancing or Distortion Inducing?" *Journal of Public Economics* 35(3):333–354.

Oates, W.E., and R.M. Schwab, 1996, "The Theory of Regulatory Federalism: The Case of Environmental Management," in W.E. Oates, ed., *The Economics of Environmental Regulation*, Cheltenham, UK: Edward Elgar, pp. 319–331.

Potoski, M., 2001, "Clean Air Federalism: Do States Race to the Bottom?" *Public Administration Review* 61(3):335–342.

Revesz, R.L., 2001, "Federalism and Regulation: Some Generalizations," in D. Esty and D. Geradin, eds., *Regulatory Competition and Economic Integration: Comparative Prospective,* Oxford: Oxford University Press, pp. 3–29.

Scheberle, D., 2004, *Federalism and Environmental Policy: Trust and the Politics of Policy Implementation, Second Edition*, Washington, DC: Georgetown University Press.

Smarzynska B.J., and S.J. Wei, 2004, "Pollution Havens and Foreign Direct Investment: Dirty Secret or Popular Myth?" *Advances in Economic Analysis & Policy* 3(2), Article 8.

Sussman, G., B. Daynes, and J. West, 2002, *American Politics and the Environment*, New York, NY: Longman.

Taylor, S., 2004, "Unbundling the Pollution Haven Hypothesis," *Advances in Economic Analysis & Policy* 4(2), Article 8.

Teitelbaum, D., 1998, "Technology Change and Pollution Control: An Adaptive Agent-based Analysis," Ph.D. dissertation, Department of Engineering and Public Policy, Carnegie Mellon University: Pittsburgh, PA.

EXPERIMENTAL INVESTIGATION IN MEDICAL MARKETS AND INSTITUTIONAL SOURCES OF PRICE INFLATION

C.A. JOHNSTON,* Interdisciplinary Center for Economic Science,
George Mason University, Fairfax, VA

ABSTRACT

We constructed a three-party market in which experts, non-experts, and insurers negotiated with each other for services, insurance coverage, and cash in such a way that we could observe prices over successive rounds of negotiations and determine whether or not they showed inflationary tendencies. We used agent-based software to execute the experiment. We found that three-party transactions among insurers, experts, and nonexperts showed inflationary tendencies, but two-party transactions between experts and non-experts did not. The findings suggested that institutional sources of price inflation can exist on the basis of the order of negotiations when there is an intermediary between consumer and supplier. The findings are consistent with a theoretical argument by Frech and Ginsburg (1975) showing that medical insurance reimbursement systems with certain price-control characteristics caused chronic price inflation

FRECH AND GINSBURG

Frech and Ginsburg (1975) analyzed a model in which health insurance induces chronic inflation. In the model, medical providers agreed with insurers to furnish unlimited services at rates set according to a survey of "usual and customary" fees within a geographic region. The policy paid a level of reimbursement equivalent to some percentile in the distribution, usually between the 75th and 90th percentile. Frech and Ginsburg then showed that by fixing the fees in the short run, the insurer set off a chain of events that caused prices paid by consumers to rise until the market was saturated. The speed of the increases depended on the percentile of the distribution chosen and the frequency of revision of the prevailing rates.[1]

"When fee schedules are set by a prevailing rate mechanism, certain plausible assumptions give rise to a chronic inflation in the fee schedule, with the rate depending on how often the schedule is adjusted and the percentile in the fee distribution chosen as the prevailing rate. This inflation only ends when consumers of medical care reach saturation," Frech and Ginsburg concluded.

* *Corresponding author address*: Carl A. Johnson, Ph.D. candidate, Interdisciplinary Center for Economic Science, George Mason University, 4400 University Drive, MSN 1B2, Fairfax, VA 22030; dcjohnst@earthlink.net.

1 It is frequently argued that medical price inflation has not historically been driven by price increases for existing procedures. "The main reason healthcare is continuing to absorb a larger share of the economy is innovation: that the range of things that medicine can do keeps increasing" (McClellan, heart study, Cutler et al. 2001). Procedures not covered by insurance include Lasik, dentistry, and cosmetic procedures.

An alternative possibility is that in an industry characterized by technological innovation, prices for existing procedures should *decline* in real terms. Therefore, the stable appearance of medical prices might mask a tendency for prices to be much higher than they would be in a more transparent market.

Frech and Ginsburg also showed that fixed-fee payments (conceptually similar to diagnostic related group payments currently used by Medicare) were less inflationary than reasonable and customary fees described above but were still inflationary. Since the fixed payments were also based on a percentile of a distribution of prevailing rates, the mechanics of fixed payment work much like reasonable and customary fees.

Frech-Ginsburg Model

The Frech-Ginsburg model assumed that health care providers charged fees for their services subject to conventional assumptions about utility of services and income to consumers and provider costs. The model assumed that some health care consumers had a health insurance policy that paid an indemnity *i* when the insured consumer suffered an illness that initially cost P_{t0} to treat. The cost of the treatment could be described as a function of *i* and some premium function of services consumed, *p*:

$$P_{t0} = i + p(x),$$

where *p* is the premium that could be charged in addition to the indemnity depending on the level of service provided *x*. The model assumed that $P_{t0} > 0, p(x) > 0$. Therefore, when $P_{tj} > P_{t0}$ and it is assumed that the level of *i* is set according to an average of previous period prices, the insurer sets the price for period P_{tj} at the previously prevailing market price P_{t0}. The price for the next market period then becomes

$$P_{tj} = P_{t0} + p(x).$$

The authors claimed the price of the service would continue to rise until the market became saturated. (Frech and Ginsburg 1975).

We can examine Frech's and Ginsburg's basic claim by simulating an environment in which negotiations proceed in a fixed order. The order of negotiation then forms the framework for computer simulations and laboratory experiments to study the emergence of inflation.

Negotiation Order

In our simulation, as in Figure 1, negotiations between agent types over commodities occur in a fixed sequence. First, insurer and non-expert exchange cash (Cash 1) for coverage (Coverage 1). Next, expert and non-expert exchange coverage (Coverage 2) for services. Finally, expert and insurer exchange coverage (Coverage 3) for cash (Cash 2).

Valuations

Each of the non-experts was assigned an initial level of cash. Experts had an initial stock of services to sell, and insurers had an inventory of coverage. Each agent type also had preferences over each commodity. Non-experts preferred to own services, which they could only obtain after first buying coverage. Experts wanted cash. Insurers wanted coverage.

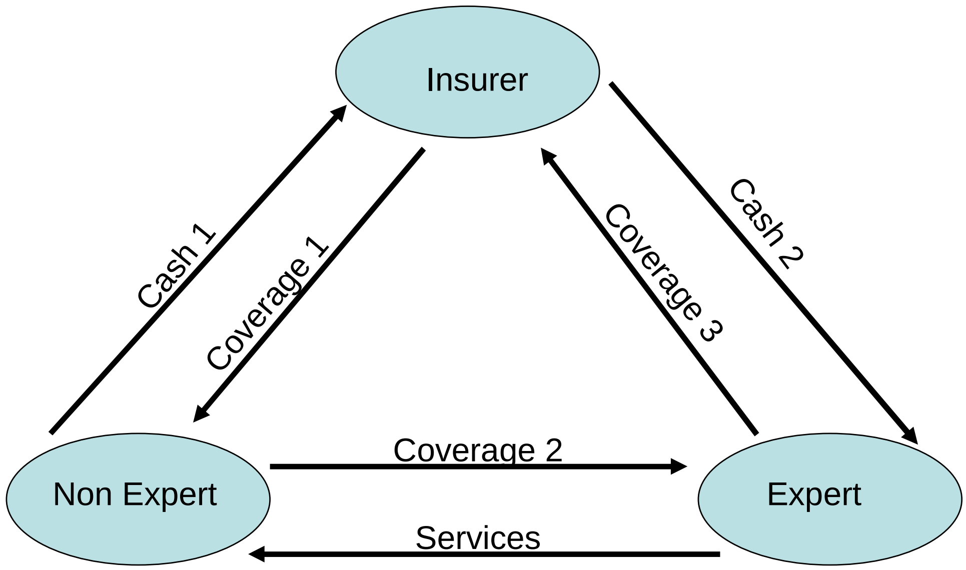

FIGURE 1 Three-party trading system with insurer

In the first set of transactions, insurers negotiated with non-experts for premiums (thereby pre-setting a constraint on any later transactions involving insurance coverage). Later transactions with insurance must be less than the value of the insurance coverage premium (Cash 1) and also less than the amount of coverage traded for the premium (Coverage 1). Each leg of the transaction was bound by the previous set of negotiations. Therefore, Cash 2 was constrained by Cash 1. Coverage 3 was constrained by Coverage 2, which, in turn, was constrained by Coverage 1. The amount of services provided to the non-expert could not exceed reimbursements to the expert (Cash 2). In each leg, both parties were assumed to trade in a fair manner, meaning that the value of coverage is at least as valuable as the cash tendered in return. Likewise, services and coverage traded between expert and non-expert have comparable value. The full set of constraints is laid out in Figure 1. The outcome of the completed negotiation round constrains the next round of negotiations.

Bargaining Rules and Price Formation

Exchange was accomplished by bilateral barter between agents at mutually-agreed-upon prices by using a posted offer mechanism. Each set of negotiations was subject to its own set of constraints. All constraints, in turn, were dependent on the outcome of previous negotiations, as outlined in Figure 2.

The constraints were not violated when the amount of services corresponded exactly to the amount of coverage purchased at an exchange rate of one unit of coverage for each unit of cash. However, because the exchange rate between coverage and services is only constrained and not fixed by the initial negotiation (between insurer and non-expert), we can see that the expert can profit by undersupplying services for a given amount of coverage. Likewise, insurers can profit by buying back coverage from the expert for less than they sold it to the non-expert in

$$\sum_{i=1}^{I}\sum_{n=1}^{N} Cash1 \geq \sum_{e=1}^{E}\sum_{i=1}^{I} Cash2 \geq \sum_{n=1}^{N}\sum_{e=1}^{E} Services$$

$$\sum_{i=1}^{I}\sum_{n=1}^{N} Coverage1 \geq \sum_{e=1}^{E}\sum_{n=1}^{N} Coverage2 \geq \sum_{i=1}^{I}\sum_{e=1}^{E} Coverage3$$

$$\sum_{i=1}^{I}\sum_{n=1}^{N} Cash1 = \sum_{i=1}^{I}\sum_{n=1}^{N} Coverage1$$

$$\sum_{i=1}^{I}\sum_{e=1}^{E} Coverage3 \geq \sum_{i=1}^{I}\sum_{e=1}^{E} Cash2$$

$$\sum_{n=1}^{N}\sum_{e=1}^{E} Coverage2 \geq \sum_{n=1}^{N}\sum_{e=1}^{E} Services$$

FIGURE 2 Constraints on three-party trading system

exchange for cash. If each agent takes advantage of the profit opportunities afforded it by the order of negotiation, then the value of coverage will steadily devalue relative to cash, resulting in inflated prices for services.

The continuous devaluation of the value of coverage throughout the order of the negotiations accounted for the price inflation that Frech and Ginsburg postulate. This also gave some extra justification for Frech's and Ginsburg's claim that the rate of inflation depended in part on how often fee schedules are adjusted (how many attempts are made in each cycle) within the insurance system.

Three-party Experiment

In this paper, we take a well-known agent-based computer simulation model, the zero-intelligence trader, and modify it to accommodate transactions between three types of agents. Gode and Sunder demonstrated that zero-intelligence trader could simulate transactions between different types of agents with heterogeneous preferences over different objects. Gode and Sunder (1993) showed that this procedure used in the context of a multi-agent program produced price and efficiency results that were comparable to those generated by a laboratory double auction, as explained in Smith (1962).

Agent-based Simulation

Although there is an established body of literature for multi-agent "zero-intelligence" bargaining, published accounts typically discuss environments with only two types or one type of agent (buyers, sellers, or buyer-sellers). It is simple to alter this existing program environment designed for two types of agents and to add a third agent to it. The three types of agents in the redesigned ZITrader are insurers, experts, and non-experts. Experts and non-experts can be thought of as generic representations of doctors and patients. The three types of agents then use the ZI Trader activation methodology to trade three quantities: cash, services, and coverage.

Agents are assigned an initial endowment of cash, coverage, and services. (Non-experts have cash; experts have services; insurers have coverage). We also give each agent preferences over each good, such that non-experts desire services; experts desire cash, and insurers want coverage. The preference for coverage is always relatively low relative to each agent's desire for the other type of commodity.[2] A unified modeling language (UML) description of the three-way trader program is shown in Figure 3.

Preferences Algorithm

We use an algorithm for representing preferences borrowed from sugar and spice trader in order to endogenously set marginal preferences of agents over commodities on the basis of their holdings. Specifically, Axtell-Epstein showed how agents can be programmed with Cobb-Douglas preferences over objects. We extend this model slightly by adding a third term to the Cobb-Douglas multipliers. The program then randomly activates agents and assigns them to buyer or seller roles depending on which agent has the higher marginal rate of substitution (MRS) for the given object with respect to the numeraire. Equation set 0.3 shows how to determine an agent's MRS for one good w_2 with respect to another good w_1. After determining respective values of MRS, the program determines a transaction price by calculating the geometric mean of the respective MRS levels, as explained by Albin and Foley (1990):

$$MRS = \frac{dw_2}{dw_1} = \frac{\dfrac{\delta W(w_1, w_2, w_3)}{\delta w_1}}{\dfrac{\delta W(w_1, w_2, w_3)}{\delta w_2}} = \frac{\dfrac{m1}{m2} w_1^{(m1-m2)/m2} w_2^{m2/m1}}{\dfrac{m2}{m1} w_1^{(m1/m2)} w_2^{(m2-m1)/m1}} = \frac{m_1 w_2}{m_2 w_1} \ .$$

Simulations were attempted by using a Cobb-Douglass function with preferences as follows:

Experts:

$$U_E = X^{0.1} C^{0.1} S^{0.8} \ ,$$

Non-experts:

$$U_N = X^{0.1} C^{0.8} S^{0.1} \ ,$$

Insurers:

$$U_I = X^{0.8} C^{0.1} S^{0.1} \ ,$$

---

2 In each case, we also assign a very small amount of nonpreferred quantities to each agent, and we also give each agent a small amount of value for nonpreferred commodities. In this way, we avoid the problem of Cobb-Douglas functions when they begin a round with zero inventory of a traded commodity that has zero preference, thus forcing a division by zero error.

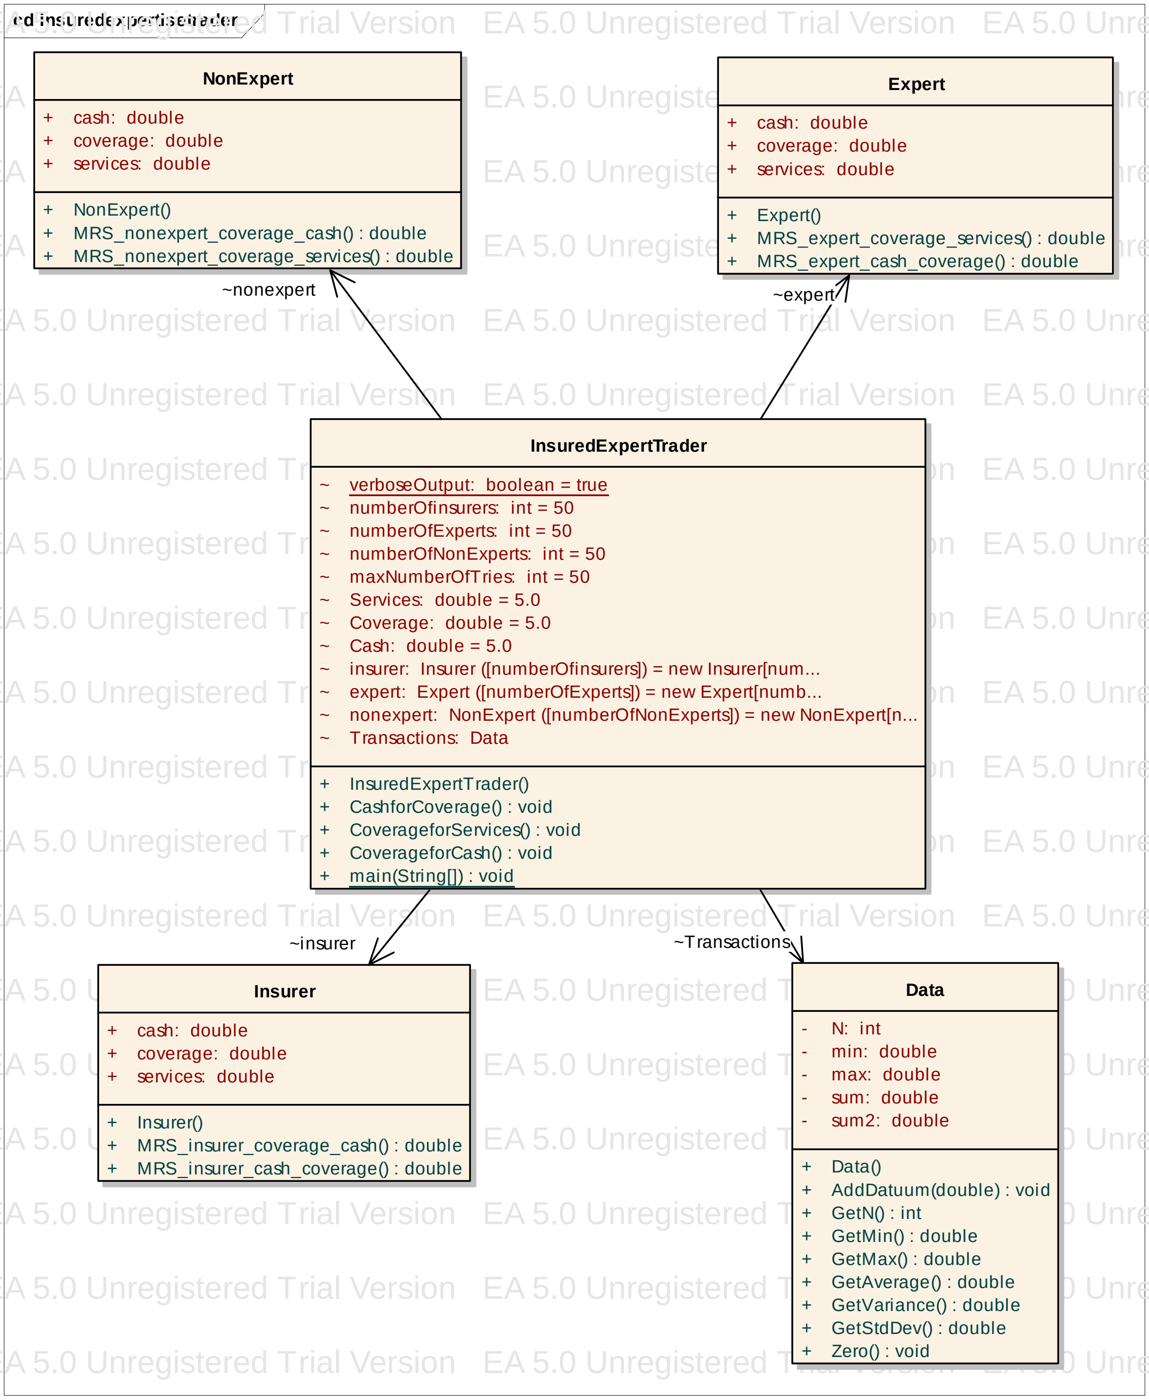

FIGURE 3 UML description of three-way trader

where *X* stands for units of insurance coverage, *C* is cash, and *S* is services. *E* stands for expert, *N* for non-expert, and *I* for insurer.

ORDER OF NEGOTIATIONS

The main program's code can be altered to contain subroutines that control how each agent would be randomly activated and made to trade with each other agent type. By sequencing the order in which the subroutines are called, one can simulate the desired order of negotiations. By repeatedly calling the routines in order, one can then simulate several rounds of negotiations.

The program calls subroutines in a specific order. First, non-experts and insurers trade cash for coverage. Then non-experts and experts trade coverage for services. Finally, insurers and experts trade coverage for cash. High-MRS buyers are matched with low-MRS sellers, and the two agents trade as described above. Trades are recorded and statistically analyzed by a standard data agent routine, also modified from the ZITrader program.

The number of attempts made to trade is a parameter in the program. In the simulations described here, the program makes 50 attempts to find qualifying buyers and sellers. The number of attempts is important to the exercise. Making a large number of attempts would satiate non-expert demand for coverage and thereby end the market process in the first round. Frech-Ginsburg, however, implies that rounds occur at a frequency such that market demand cannot be satiated in one round. Therefore, the program discussed here limits the number of potential transactions to 50, ensuring that demand is not satiated in the first round. The program executes the cycle of negotiations four times. By trial and error, we discovered this is the number of rounds required to exhaust the bulk of gains from trade.

Purpose of the Simulation

The purpose is to prove the feasibility of a test of the Frech-Ginsburg theory of inflation through simple transactions in a prescribed order of negotiations.

Simulation Results

The agents bargained with each other as explained above. Agents were randomly activated, and trade attempts were made. The software generated 50 of each type of agent, and each agent type was assigned the same number (5) of some initial stock of tradable commodity as well as tiny amounts of the other two tradable commodities. For example, non-experts had an initial stock of 5 units of cash; however, they also had initial stocks of 0.1 unit of coverage and 0.1 unit of services in order to avoid dividing by zero when conducting initial trades. The experts and non-experts were similarly furnished with initial stocks of services and coverage, respectively, as well as de minimis amounts of the other two commodities. For the reasons discussed earlier, trade attempts were limited to 50.

What We Learned from the Experiment

The simulation established that the trading scheme does produce many of the kinds of imbalances that one might expect it to show given our discussion earlier. Prices for services inflate at each stage and during each period. Reimbursements to experts for services show a decreasing trend over time, even as prices of insurance coverage go up. A progressive devaluation of the “currency” of insurance coverage seems evident. Inflation in the prior period acts as a foundation for price gains in succeeding rounds of negotiations. A graph of the average prices paid by agents in each of the three legs of the simulation is displayed in Figure 4. The different series relate to the number of transactions per leg per cycle.

Surprises

This version of the software produces extraordinarily large variances in prices and steep gains in prices from one period to the next (and steep reimbursement declines for the expert in the coverage-for-cash leg). The large variances are related to the initial trades in which there are several agents with zero holdings of the commodity being traded.

COUNTERFACTUAL EXPERIMENT

A counterfactual experiment, in which two parties trade and are found to create no inflation, would be the best way to verify the claim that the three-party institutional negotiation order causes the price inflation. If inflation occurs in the three-party structure but not in the two-party structure that is otherwise identical to the three-party experiment, then we can justify the claim that it was the procedure of negotiation and nothing else that caused the observed price increases.

The counterfactual experiment is easy to carry out. In order to turn the simulation into a two-party experiment, we simply place two slash marks in front of the lines that call for cash-for-coverage and coverage-for-cash routines, and we allow only the coverage-for-services routine to execute four times. The full findings are summarized in Table 1 and presented at length in Table 2 and in Exhibits 1 and 2 at the end of this paper.

The prices in the counterfactual experiment do not violate the notion that they all come from the same distribution and therefore no discernible inflation trend is visible. This is a sharp contrast to the previous experiments, in which inflationary gains in prices were clearly evident and significant.

WHY HUMAN EXPERIMENTS ARE NEEDED

This paper has demonstrated two different theoretical models and two sets of computer simulations that complement each other in terms of enforcing our understanding of how inflation generated by institutional sources would work and what such price inflation would look like and under what circumstances it might occur. The findings suggest that not only is such an effect possible, but it appears to be robust.

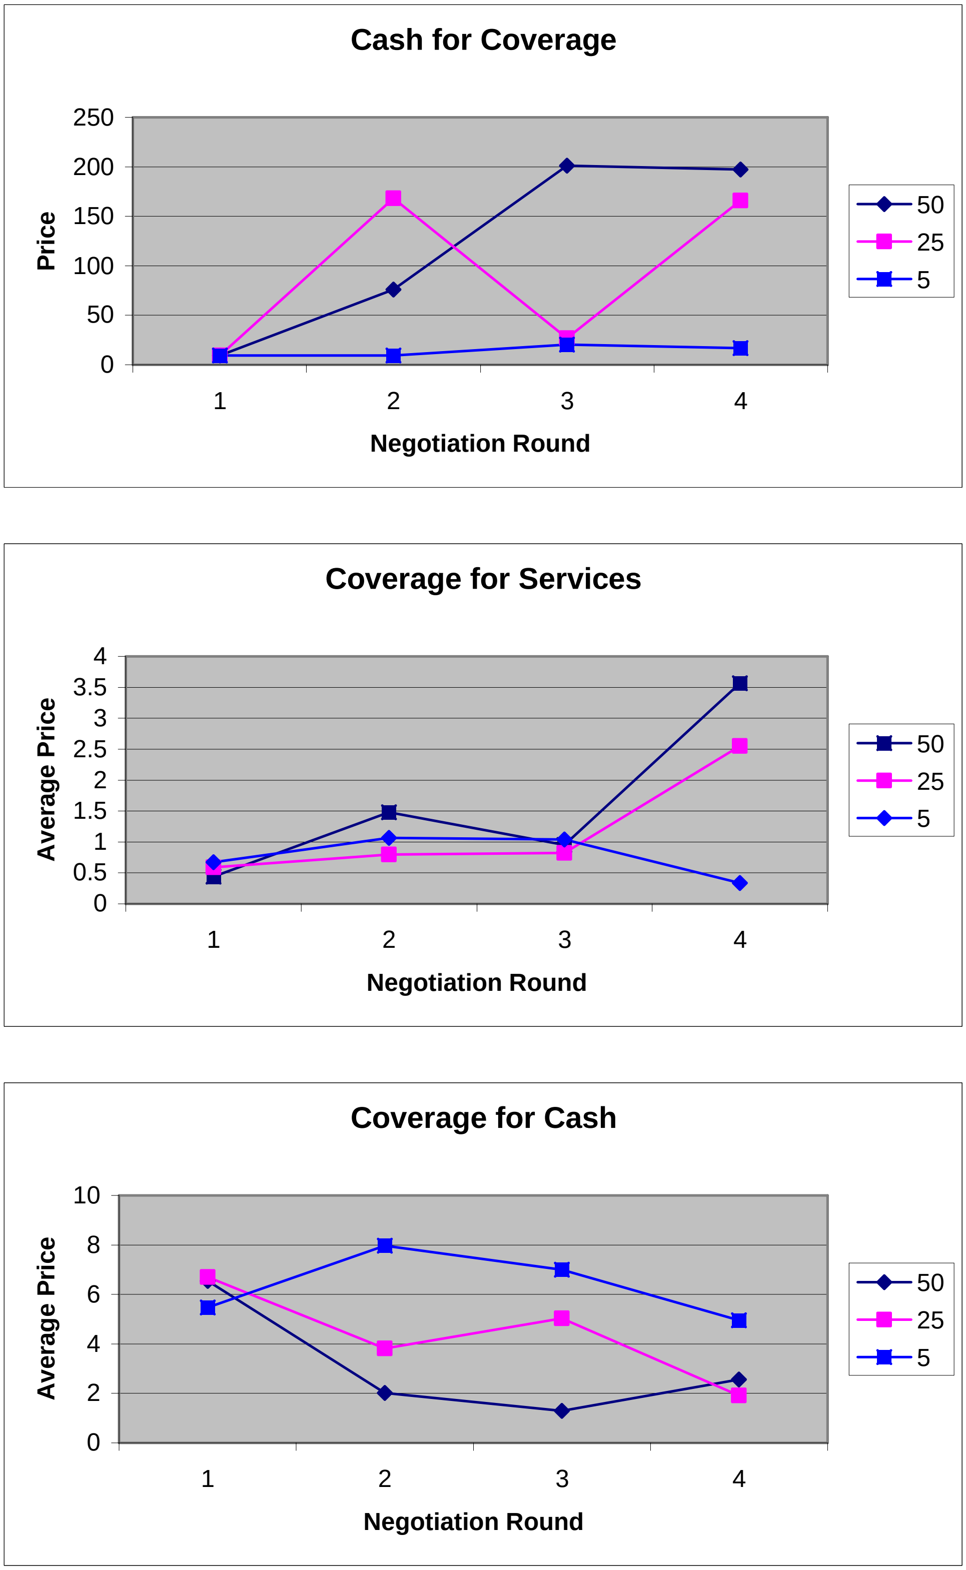

FIGURE 4 Average prices from agent-based simulation

TABLE 1 Counterfactual experiment results

	Coverage for Services			
Parameter	Round 1	Round 2	Round 3	Round 4
Average	0.63180726	0.482374001	0.399636988	0.43643578
S.D.	0.100377281	0.170497875	0.195787679	0

TABLE 2 Price evolution in three-party simulation

	Coverage for Cash			Services for Coverage			Cash for Coverage		
	1st Period	2nd Period	3rd Period	1st Period	2nd Period	3rd Period	1st Period	2nd Period	3rd Period
Avg	9.438798	118.7625	278.5622	0.442944	1.058568	2.176739	4.970369	3.630569	1.738485
Var	0	10203.68	18964.91	0.086874	0.894776	2.377667	18.5324	28.83809	1.548212

On the other hand, introduction of humans into the mechanism might change the dynamics significantly. For example, if non-experts withhold demand or if experts sell their services at less than their value, it might have a dampening effect on prices and possibly even create an (inefficient) equilibrium. Therefore, a replication of the simulation described in this paper but using humans instead of agents might yield valuable information about how human-operated, noncomputational organizations handle the dynamics of three-party negotiations.

CONCLUDING OBSERVATIONS

The health care sector draws in more and more resources each year. Medicare cost escalation is thought to be a significant threat to the financial stability of the United States. Rising health care costs are among the most frequently cited causes of job destruction. Yet despite accelerating costs, we have little to show for the expense, since health indicators have not improved much in 30+ years. The proportion of persons without health insurance, infant death rates, and male life expectancy have not improved as much as one might have hoped, given the vast resources being expended on the problem. Yet the most recent scholarship on the causes of inflation in health care dismiss the notion that inflationary pricing contributes to the escalation of health care costs. For example, McClellan et al. (2001) studied the cost of heart attack treatment. and concluded that “spending increases are mostly driven by changes in the quantity and type of services provided, not changes in the price of a given service.”

These findings have given rise to a discussion about how to curb the demand for greater intensity of treatment for diseases, such as heart attacks. One academic/policymaker has gone so far as to suggest outright that limitations on the amount of profit that developers of innovative medical products can make on new inventions constitute the best way to reign in runaway costs (presentation at George Mason University in February 2005 by Tomas J. Philipson, The University of Chicago, former Assistant Commissioner of Medicare).

It is not usually the case that technical innovations present these kinds of problems. Usually, the technical advance produces deflation. For example, computers are much more capable and much more numerous in 2005 than they were in 1980. Yet nobody thinks of computing costs as being out of control. This is because in fields characterized by innovation, it is common to see prices decline, sometimes sharply. If we were still paying the same price for a 1985 microprocessor in 2005, or a little less, we might consider it a bad trade.

Therefore, the fact that prices of old medical technologies remain stable or decline a bit does not necessarily mean that prices are behaving in a noninflationary way. The work in this report is intended to suggest that institutional sources of price inflation can exist on the basis of the order of negotiations, particularly when there is an intermediary between consumer and supplier. The simulations and theoretical work show that experimental methods could be used to explore behavioral, institutional, and economic system design that would substantially benefit our understanding of these kinds of situations.

One can also imagine many extensions of these experiments. For example, could we design a market for insurance claims that might moderate the apparent inflationary tendency inherent to U.S.-style health insurance plans?

ACKNOWLEDGMENTS

I gratefully acknowledge the advice and criticism of Robert Axtell, Dan Houser, Stephen Rassenti, and David Porter, as well as financial support from the Mercatus Center and the James M. Buchanan Center for Political Economy at George Mason University.

REFERENCES

Albin, P., and D.K. Foley, 1990, "Decentralized, Dispersed Exchange without an Auctioneer: A Simulation Study," *Journal of Economic Behavior and Organization 18*(1), pp. 27–51.

Cutler, D.M., M. McClellan, J.P. Newhouse, and D. Remler, 2001, "Pricing Heart Attack Treatments," in D.M. Cutler and E.R. Berndt, *Medical Care Output and Productivity*. Chicago, IL: University of Chicago Press.

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, DC: MIT Press.

Frech, H.E., and P.B. Ginsburg, 1975, "Imposed Health Insurance in Monopolistic Markets: A Theoretical Analysis," *Economic Inquiry 75*(1), pp. 55–70.

Gode, D.K., and S. Sunder, 1993, "Allocative Efficiency of Markets with Zero-intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy 101*(1), pp. 119–137.

Smith, V.L., 1962, "An Experimental Study of Comparative Market Behavior," *Journal of Political Economy 70*, pp. 111–137.

EXHIBIT 1 Multi-agent simulation of three-party experiments

Original Distribution of Coverage, Cash, Services

NonExpert / 0/ 250/ 0

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 250

Distribution after Trading Coverage for Cash

NonExpert / 25/ 125/ 0

Insurer / 225/ 125/ 0

Expert / 0/ 0/ 250

Update Distribution after Trading Coverage for Services

NonExpert / 9/ 125/ 37

Insurer / 225/ 125/ 0

Expert / 1/ 0/ 213

Distribution after Trading Coverage for Cash

NonExpert / 9/ 125/ 37

Insurer / 225/ 96/ 0

Expert / 1/ 29/ 213

Totals / 235/ 317/ 250

Distribution after Trading Coverage for Cash

NonExpert / 21/ 65/ 37

Insurer / 213/ 156/ 0

Expert / 1/ 29/ 213

Update Distribution after Trading Coverage for Services

NonExpert / 8/ 65/ 69

Insurer / 213/ 156/ 0

Expert / 9/ 29/ 181

Distribution after Trading Coverage for Cash

NonExpert / 8/ 65/ 69

Insurer / 219/ 135/ 0

Expert / 6/ 48/ 181

Totals / 233/ 335/ 250

Distribution after Trading Coverage for Cash

NonExpert / 16/ 25/ 69

Insurer / 211/ 175/ 0

Expert / 6/ 48/ 181

Update Distribution after Trading Coverage for Services

NonExpert / 8/ 25/ 95

Insurer / 211/ 175/ 0

Expert / 13/ 48/ 155

Distribution after Trading Coverage for Cash

NonExpert / 8/ 25/ 95

Insurer / 220/ 161/ 0

Expert / 5/ 58/ 155

Totals / 233/ 347/ 250

Distribution after Trading Coverage for Cash

NonExpert / 12/ 5/ 95

Insurer / 216/ 181/ 0

Expert / 5/ 58/ 155

Update Distribution after Trading Coverage for Services

NonExpert / 5/ 5/ 110

Insurer / 216/ 181/ 0

Expert / 6/ 58/ 140

Distribution after Trading Coverage for Cash

NonExpert / 5/ 5/ 110

Insurer / 221/ 170/ 0

Expert / 4/ 69/ 140

Totals / 230/ 345/ 250

EXHIBIT 2 Inventories of all agents after each update

Original Distribution of Coverage, Cash,/ Services

NonExpert / 0/ 250/ 0

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 250

Update Distribution after Trading Coverage for Services

NonExpert / 0/ 250/ 31

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 219

/n Stats after Trading Coverage for Services: 31 transactions took place at 0.6318072598177861 average price; 0.10037728051751019 standard deviation. /n)

Update Distribution after Trading Coverage for Services

NonExpert / 0/ 250/ 44

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 206

/n Stats after Trading Coverage for Services: 13 transactions took place at 0.4823740008931712 average price; 0.17049787522922566 standard deviation. /n)

Update Distribution after Trading Coverage for Services

NonExpert / 0/ 250/ 49

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 201

/n Stats after Trading Coverage for Services: 5 transactions took place at 0.399636988045828 average price; 0.19578767907593872 standard deviation. /n)

Update Distribution after Trading Coverage for Services

NonExpert / 0/ 250/ 50

Insurer / 250/ 0/ 0

Expert / 0/ 0/ 200

/n Stats after trading Coverage for Services: 1 transactions took place at 0.4364357804719847 average price; 0.0 standard deviation. /n)

DISCUSSION

Social Simulation Applications

(Parallel Applications Session I — Economics and Environmental Policy, Friday, October 14, 2005, 1:15–3:15 p.m.)

Chair and Discussant: *Gale Boyd, Argonne National Laboratory*

Margins and Transaction Taxes in an Intraday Continuous Double-auction Futures Market

Gale Boyd: We have been developing this resource at Argonne, The University of Chicago, and a couple other area institutions, including University of Illinois, Northwestern, and the Chicago Fed. We've all probably reached a point in our research where we say, "Oh, I could do that if only I had more data." Of course, that's usually a "be careful what you wish for" proposition, but I want to make people aware of a resource, which is relatively new.

We all know that the Census Bureau sends out individual forms to households and surveys to firms and businesses and then aggregates that data, and we think that what we're stuck with it if we ever want to use that type of economic or demographic information. We have a program with the Census Bureau that allows researchers to access the raw data that the Census Bureau collects at a secure computer lab that we operate in the Chicago area. That's certainly not going to answer all the questions that we might like to investigate with agent simulation, but it gets the information a little closer to the spirit of agent modeling. If anyone's interested in learning more about that census data center program feel free to contact me. I'm going to turn things over to Leanne.

Leanne Ussher: I'm presenting an agent-based model of a futures market, and I'm simulating *open-outcry* in a futures market. My main emphasis is looking at how regulations have institutional structure of that market that can change prices.

[Presentation]

Boyd: We have time for one or two questions.

John Sullivan: John Sullivan, Ford Motor. You might have said something in your presentation, but I'm not too familiar with this aspect of the stock market. Does this activity actually generate wealth or add value to the economy, or is this a redistribution of their own wealth?

Ussher: This is a futures market, not a stock market, and it doesn't generate wealth. All it does is incorporate information into that market, and it's meant to improve the efficiency of the price. It reduces risk if you're a hedger. You can go in there and hedge your risk, but it doesn't produce any product or increase wealth, except that it increases wealth by reducing people's risk and by allowing for an efficient price.

Sullivan: Does this actually reduce the risk for regular investors who are not margins players?

Ussher: Yes, they would represent my hedges who are in here not to make money, but to reduce the risk in their daily lives, of their daily business. The speculators are there to make money, and they want the risk because they want to try and profit from it.

Robert Reynolds: Right now you don't have learning in this model, and one way to have learning in the future would be to have a participant that could choose to be a hedger, a speculator, or a scalper. It would be interesting with a given set of parameters to let the agents come up with a particular mix of the three that would perhaps be stable or most productive for that market. That would be something that could be learned easily down the road. By looking at the actual mix of roles you could see what is emerging in a particular market. You could then look at transitions from one type of market to another in terms of the change in the roles.

Ussher: I think that's true for my speculators. My speculators could learn from each other and take on new roles. Some might be technically trading, some might be fundamentalists, and some might be imitators or contrarians. If they see one group making more money, they might imitate them and take on that role. But for the scalpers and hedgers, there are barriers for entry into getting into that.

Scalpers have to purchase a chair on the exchange, and they're located on the exchange, and not everyone is allowed to trade on that exchange. They might have to go through a broker. The hedgers are in a different business altogether, producing wheat or whatever, so you're right. There are some barriers to entry, and there are some where they could imitate each other. It's a mix there.

Unidentified Speaker: You commented in your talk about the fact that they tended to oscillate around the price. You said that you thought it might be a good thing, and it seems like it's an emergent procedure.

Assessing Ex-urban Residential Markets: An Agent-based Model

Boyd: The next paper is "Assessing Ex-urban Residential Markets in the American Mountain West: An Agent-based Model."

Li Yin: Today I'm going to talk about the application of agent-based modeling, ex-urban residential location choices.

[Presentation]

Boyd: We have time for questions and there are lots of hands.

William Rand: I'm Bill Rand from Northwestern University. I think this is an excellent model of ex-urban development. I have built a couple of suburban models in the past. There was one thing that I wasn't clear about: how much diversity among the agents was actually in the model? You briefly talked about including more types. We found that even slightly varying the amount of preference for something like a natural amenity or for nearness to roads or anything like that dramatically changed the landscape. That brings me to a second question that deals with projected versus actual development patterns. Were those projected patterns based on single runs, were they averages runs, or have you seen much change in the runs over time? Basically, is there a lot of variance in the runs?

Yin: As to the first question, at this stage, I was only able to include two types of agents, second home owners and commuters, because we find nowadays that in many places, especially in high-amenity areas, a lot of second home owners moved in, and their preferences are very much different from those of commuters. So I only looked at only two types at this stage. In the future, I may include more types.

As to the second question, I did multiple runs. Before I built this model, I worked on ex-urban growth models using this great choice of models. For this area, I had a sense of how important this variable is. We played with it, and we found that, for example, distance to route, such as within 100 meters, really matters. In other words, we played with parameters a lot. We ran a lot of times, and this is the converge of the model.

Sullivan: Ford has a joint research program with the University of Michigan in the area of sustainable mobility and accessibility, and I applaud the fact that you're addressing some of these questions. A question that comes up, especially in discussions surrounding sustainable development and what the future may hold for us, is the impact of increased transportation costs of one form or another on current settling patterns in the United States and what they might evolve to. Have you thought along those lines, or is that something that's planned for the future?

Yin: Well, for commuters, I think it's very important. If their transportation cost increased, they would need to rethink their location choices. Second home owners don't have to travel. They don't have to commute every day. That's why I chose second home owners as a separate type of agent. But, yes, that can be included as one more variable in accessibilities for commuters. I think that will affect their choices, but probably not for second home owners.

Luis Fernandez: Luis Fernandez, EPA. Actually, you said that you had only two types of agents, two classes of agents. Do you have any thoughts as to how you're going to create a typology for agents? I was working on the same project that Bill referred to. We actually had similar issues, and we used data-mining techniques, based on a large social survey. We got very interesting results when looking at the data and having the preferences emerge from the data, as opposed to just speculation or stock life-stage classes that are oftentimes assigned to these models.

Yin: Actually, we also did a survey. Before this model, we built a logistical regression that we had public meetings; we asked people what they think and then did a survey. All of my parameters are derived from the survey and from the previous model. But, yes, I think with two types of agents, it's very limited. That's what I'm going to improve in the future, hopefully.

Boyd: That reminds me of my opening comment about everything we would all do if we just had more data. I couldn't help thinking about turning this kind of model loose, if you had individual information from the Census, the American Housing Survey, the matching demographics, etc. You could actually distinguish between whether a particular home owner had commuted all the way to Denver for work or not.

Environmental Race to the Bottom: Mixed Agent-based Model

Boyd: The next paper in this session is "The Environmental Race to the Bottom."

[Presentation]

Boyd: I have one question. One of the things that is a unique characteristic of doing agent modeling is the ability to represent levels of interaction that aren't necessarily driven by traditional equilibrium assumptions. When you talk about a lot of the theoretical models, which essentially take a system that might be similar to yours, and then impose equilibrium solution concepts and try to solve them, can you talk a little on how your work might be trying to pick up different elements of that?

Alexander Alexeev: You're right. We don't need equilibrium here, although we may use things like thermodynamics or quasi-equilibrium; in this case, we go through all these quasi-equilibrium states, as well as interaction between firms and between states — two level interactions between firm and state. This is actually higher-level interaction. In this case and in the case of state and firm interaction, we have a situation like when a standard is set equal for all firms. It's the firms' decision how to use the standard because, for example, if there is a rule that states the amount of hazardous waste reported to the EPA, it's self-reporting from the firm. It may be lower and it may not be higher. In this way, we may introduce some stochastic behavior of the firm for firms.

Boyd: Thank you.

Experimental Investigation in Medical Markets and Institutional Sources of Price Inflation

Boyd: We started off this session with traders and price volatility, then moved out to the suburbs and the mountains, then to another state where the pollution taxes were lower, and now we're back to traders and prices, and in this case, inflation.

Carl Johnston: I was putting this paper together and decided that it would be about the intuition that we have at some level that middlemen are bad guys or that they're wasteful. To tell you the truth, that's how this project started. We were looking at health-care institutions in the United States, and of course we're in Virginia not far from Washington, DC, where the health-care debate is a very big topic. It's a commonplace discussion to have, at least in some circles. Many people say that if we could cut out the insurers and just allow the doctors to barter with the patients, everybody would be better off, and we'd have a more efficient system.

So with that in mind, I started looking more carefully at the subject. Now, keep in mind, that although I'm from Mason, I'm not from the Center on Social Complexity and from the experimental economics branch of the Economics Department; rather, I'm from the Interdisciplinary Center for Economic Science, which is different, but related, and we're trying to expand the overlap a little bit.

[Presentation]

Boyd: Are there any questions?

Ussher: There seem to be a lot of parallels with my model in the sense that my zero intelligence agents are also from Gode-Sunder, and market-makers or middlemen. Do you have competition between your insurers? Is there transparency and competition between them so that they wouldn't be able to keep taking more profit? Would you incorporate that?

Johnston: At the moment, this is just the Gode-Sunder random activation technique, which should get you to a competitive result, but they are zero intelligence in the sense that they are not yet forming a strategy and trying to achieve a perfect price, but that's obviously one of the things to do in the future.

Ussher: When you have your continuous double auction and your bid spread, it's the best bid and the best ask in the market. Is that right?

Johnston: Actually, it's the random draw, and then it's a geometric mean between the two. Okay, so it's not truly a double auction in this simulation. It's a comparable kind of institution, at least as was used in Gode-Sunder.

Sullivan: I believe that you said at one point in your presentation that employers like the more out-of-pocket expenses because that actually reduces the number of transactions that the insurer needs to do with the health-care provider.

Johnston: It takes more rocks out of the guy's fist, yes.

Sullivan: It also seems that an employer representing 100 or 200 or more people reduces the number of transactions between him and the provider, or *it* and the provider as well, and therefore economies occur there. Is that correct?

Johnston: Yes. So you can imagine that the system has come up with a variety of these kinds of things to keep it from exploding. That's why it has lasted for 50 years. I'm not saying that the real-world system is going to explode. You do have compensations that occur. This is just a very special case.

Boyd: I'd like to thank all of our speakers and everyone for their patience to stay strictly on the schedule.

Parallel Applications Session II —

Social Policy and Computational Knowledge

BEYOND MARKETS AND COMMUNITIES: A COMPARATIVE APPROACH TO KNOWLEDGE EXCHANGE IN ORGANIZATIONS

S.S. LEVINE,* Singapore Management University, Singapore
M.J. PRIETULA, Emory University, Atlanta, GA

ABSTRACT

While knowledge transfer has been shown to affect organizational performance, little is known about the processes of knowledge exchange between organizational agents. We propose that combinations of various modes of exchange and degrees of tie strength produce at least five different configurations: neo-classical exchange, local search, embeddedness, community exchange, and performative ties. By using an agent-based simulation of problem solving with knowledge exchange in an organizational setting, we find that embeddedness and community exchange provide results that are superior to neo-classical exchange. Performative ties, however, outperform both, even if just a minority of the organizational agents are able to extend such ties. In addition, we find that the marginal returns on performative ties are greatest when such ties are relatively rare, suggesting that the cost associated with encouraging them can be minimized with few adverse effects. We conclude by discussing managerial implications for team setup and facilitation of knowledge transfer.

Keywords: Knowledge, social network, exchange, reciprocity, performance

INTRODUCTION

Knowledge has grown to occupy a major role in the discussion on firm performance and survival (Winter 1987). In the management literature, knowledge has been recognized as a valuable resource (Kogut and Zander 1995), a source of lasting competitive advantage (Teece et al. 1997), and even the very foundation for the existence of a firm (Grant 1996). But unlike other economic resources such as capital or land, knowledge is a social entity. In the last two decades, much ink has been spilled to argue and show that organizational knowledge — the kind that is necessary to create a competitive advantage — does not reside in manuals or training books but in individuals and the regularities by which they cooperate (i.e., routines) (Cohen and Bacdayan 1994; Kogut and Zander 1992; Nelson and Winter 1982).

Consequently, the organizational literature has been devoting much attention to the transfer of knowledge between agents, either between individuals within the same organization (intraorganizational knowledge transfer) or between organizations (interorganizational), and much has been achieved. We now have established an understanding of the flows of knowledge between agents, the consequences for various phenomena of managerial interest (e.g., innovation), and the obstacles to knowledge flows.

* *Corresponding author address*: Sheen S. Levine, Singapore Management University, 50 Stamford Road, Singapore 178899; e-mail: sslevine@sslevine.com.

Valuable knowledge is knowledge that is unique and protected from easy imitation. As such, it is likely to be intangible (Teece 1998; Winter 1987) and to reside in people and the patterns of interactions among them (i.e., routines) rather than in manuals or textbooks (Winter 1995). However, it is difficult to observe or directly identify the presence of knowledge or measure its quality (Arrow 1962, 1969). For instance, it is difficult to judge whether a person does indeed carry the knowledge that she claims to possess without requiring her to reveal that knowledge. Because knowledge resides in people and routines, it cannot be easily taken from one location and transferred to another. Knowledge cannot be extracted against one's will, nor can it be force-fed to others. A recent front-page story in the *Wall Street Journal* described how experienced employees in a manufacturing facility refused to share their knowledge with newcomers or managers, because this knowledge made the veterans irreplaceable (Aeppel 2002).

If we agree that knowledge is contained in routines (Winter 1995), then it naturally follows that knowledge is a social entity and requires social interaction for transfer. Our primary interest here is modeling processes of knowledge interactions between individuals (i.e., intraorganizational agents). Specifically, we are interested in linking individual choices to organizational performance. When performing their professional tasks, agents decide who to turn to in their search for knowledge and how to negotiate the transfer. These micro choices — mode of exchange and the characteristics of the exchange partner — can eventually affect organizational performance, a macro variable. We are thus interested in linking individual choices to organizational performance. In our model, we examine the gamut of modes of exchange and the nature of relationships between the agents. Our model allows an agent to interact purely in one mode of exchange or, more realistically, in a combination of them, depending on the characteristics of the partner.

CHOICE OF PARTNERS AND MODES OF EXCHANGE

Tie Strength

From empirical research on knowledge processes, we know that individuals often search for knowledge that is necessary to perform their professional tasks, whether they are service technicians (Orr 1990, 1996), high-technology engineers (Bechky 2003), or white-shoe professionals (Haas and Hansen forthcoming; Hansen 1999).

In a study of a global professional service firm, Levine (2004) found some use of codified sources, such as textbooks and internal publications. These sources were typically used when an employee was completely unfamiliar with the industry or the problem at hand and was interested in an introductory overview. More commonly, professionals in the firm turned to their social network for knowledge. They approached *strong ties* — office mates, close friends, and family members — making a variety of requests, from asking quick questions about statistical functions in Microsoft Excel to spending several hours sourcing an insider view on an industry. Individuals also called on *weak ties* — acquaintances in their office and in other offices — when seeking references and advice (cf. Granovetter 1973). However, strong and weak ties were not the only sources for knowledge. Employees often sought knowledge from strangers — others they had neither met nor been referred to by a mutual contact. A *performative tie* involved two or more individuals that became linked following a process of wide search. While the individuals had no transaction history, nor did they expect to develop one, the transaction was carried out in a mode of generalized exchange, without expectations of reciprocity from the beneficiary to the benefactor. Table 1 summarizes the various sources of knowledge and provides examples and theoretical and empirical referents.

A knowledge transaction involved several steps: identification of the potential knowledge carrier, initiation of contact with that individual and negotiation of terms, and transfer of content. An important feature is the choice of the mode of exchange, which determines the expectations of both partners for arrangements of reciprocity.

Partner Identification: Local and Global Search

A common way to identify an exchange partner is to examine proximate individuals. Empirically, we know that people are more likely to choose as exchange partners those that belong to the same cluster (Levine and Kurzban forthcoming). The search criteria can be geographic propinquity (Marsden and Campbell 1984; Newcomb 1961; 1966), similar characteristics such as ethnicity or age (Ibarra 1992; Marsden 1988; McPherson et al. 2001), or membership in a focal group (Burt 2004; Feld 1981). In a great number of search situations, individuals (and organizations) exercise *local search*, turning to a neighbor or an acquaintance, and neglect searching distant prospective exchange partners (Levinthal and March 1993).

Global Search

Normative approaches for search call for a comprehensive search of the gamut before deciding on an exchange partner, thus achieving optimality. This approach is the hallmark of

TABLE 1 Knowledge sources, examples, and theoretical referents

Source	Contribution	Example	Theoretical and Empirical Referents
Codified sources	Provide an overview of industry, typical problems, and frameworks	Internal manuals, textbooks, proceedings of professional conferences	cf. Arrow 1969; Kogut and Zander 1995; Teece 1977
Strong ties	Vary, from help in using computer software to getting an insider view on industry	An associate consulted her knowledgeable spouse about a professional problem	Bian 1997; Coleman 1988; Nelson 1989; Straits 1991; Wellman and Wortley 1990
Weak ties	Provide an account of previous projects in the same or a similar industry	A senior analyst helped an associate sitting nearby to program a statistical routine for data analysis	Granovetter 1973; Hansen 1999, 2002; Levin and Cross 2004
Performative ties	Recall specific similar cases, suggest ways to think about issues, provide proprietary data	A manager called an unacquainted partner in Australia	Constant et al. 1996; Levine 2005; Saxenian 1996; von Hippel 1987

neo-classical economics (Wilson 1987) and the early decision-making literature (cf. Simon 1957).

While an optimal search may not be possible, organizations have created organizational indexes, which list members and their experience, to be used when searching for knowledge. Students of knowledge transfer in organizations have documented knowledge management systems (KMSs) — organizational indices that contain summaries of projects and contact information on those involved (Hansen 1999; McDermott 1999). In the case described by Levine (2004), the KMS did not attempt to capture much knowledge, but it did contain short descriptions of many of the projects performed in the firm worldwide. In addition to a sketch of the problem and the solution, it also contained contact information for all of the project team members, even if they were no longer employed at the firm, allowing direct contact. Rather than a library of codified knowledge, the KMS served more as a collection of pointers — an organizational index that identified individuals who possibly possessed relevant knowledge. Knowledge seekers used the information contained in the KMS to filter through the list of potential knowledge carriers and decide which ones to contact.

If the KMS did not provide a lead, employees sometimes sent out a mass electronic mail message to the whole office, to employees in a certain geographical region, or to all of the associates worldwide, for instance. The message detailed the knowledge needed and asked for assistance. A similar pattern of sending out mass electronic messages (with considerable success) was documented among sales personnel in a large computer hardware manufacturer (Constant et al. 1996).

Whether through the use of a central KMS or by sending a mass message to colleagues, knowledge seekers attempted to perform a global search (under constraints) for prospective exchange partners.

Typology of Exchange Modes and Tie Strengths

Once a prospective exchange partner is identified, the seeker makes contact, either directly or through a broker — a third party who can introduce both individuals to each other (for more on the role of brokers, see Burt 1992, 2005; Hargadon and Sutton 1997). Then, the sides must agree on the terms of the exchange. As detailed below, the transfer can be arranged as a spot barter (neo-classical exchange), as a favor that must be returned in the future (social exchange), or as a nonreciprocated donation (generalized exchange).[1] These modes of exchange have different meanings when used on ties of varying strength, leading to five types of ties.

Table 2 categorizes patterns of exchange according to two criteria: tie strength (strong, weak, no prior) and mode of exchange (spot, social, generalized).

[1] We do not discuss altruism as a mode of exchange because there is no substantial empirical evidence that shows it to serve as a sustainable mode of exchange within organizations.

TABLE 2 Typology of exchange

	Mode of Exchange		
Tie Strength	Spot Exchange Autonomous economic agents repeatedly search for best price and quality combination (Wilson 1987)	Social Exchange Repeated exchange with stable partners: goods, gifts, favors (Blau 1964; Bourdieu 1977; Emerson 1981; Homans 1958)	Generalized Exchange Direct reciprocity is not expected (Ekeh 1974; Lévi-Strauss 1969 [1949]; Malinowski 1920)
Strong tie (frequent, close, and intensive)	2. Local search Limited scope of search and the need for simultaneous availability of exchange items mean that transactions are executed at a suboptimal level.	3. Embeddedness Allows access to benefits that are not available in the marketplace, but limited scope of search can lead to suboptimal results. Requires pre-existing personal trust in the exchange partner.	4. Community/clan Allows access to benefits that are not available in the marketplace, and allows resource-poor agents to participate. But relationships tend to be primary and oriented for the long term, and the breadth of the search is limited.
Weak tie (not frequent, close, or intensive)	1. Neo-classical market Wide search for global optimum, but hinders customization, increases risk, and negotiation costs. Requires simultaneity in exchange.	(Weak emergence)[a]	5. Performative ties Allow wide search for potential exchange partners. Generalized exchange allows for transactions that are unbalanced in the dyadic level, and does not require simultaneity.
No prior tie (no history of dyadic interaction)			

[a] Social exchange requires repeated exchange with stable partners, and it is unlikely to take place with strangers. Closest to social exchange with strangers would be generalized exchange (or pure altruism, where reciprocity is not expected at all).

Tie Strength

As discussed earlier, tie strength measures the degree of frequency of interaction, closeness, and intensiveness of the relationship between the agents. The definition is based on the empirical work of Marsden and Campbell (1984), who concluded that a measure of closeness or intensity is the best indicator of strength. It also includes frequency, which has been commonly used to measure tie strength (e.g., Granovetter 1973).

Tie strength is a continuum here, running from strong ties, as those between close friends or family, to "no prior ties," which indicates interaction between people who are completely unacquainted directly and indirectly.[2] The label "weak tie" is added for convenience, to show the drop in tie strength, but its location along the continuum is arbitrary.

Mode of Exchange: Spot, Social, and Generalized

Mode of exchange refers to the principle underlying the transaction — what is transferred from each agent to the other. The three categories along the mode of exchange axis begin with a spot market exchange, where, in a neo-classical sense, autonomous economic agents repeatedly search for the best price-quality combination and where exchange is price-based and simultaneous, as both sides agree on a price and proceed to give and receive at the same time.[3]

Social exchange can be used for tangible goods as well as for some desirable social goods that are not easily tradable in neo-classical markets because of the difficulty in evaluating them or their unavailability for simultaneous transaction (such as exchange of prestige and friendship). The problem could be, for instance, due to the difficulty in evaluating them or their unavailability for simultaneous transaction. In the classical work of Homans (1958) and Blau (1964), individuals attain their personal or group goals by exchanging with others. Bourdieu (1977, 1997) employs this logic to analyze the practice of gift giving as a system of direct exchange, which differs from lending or swapping because of the requirements that the exchanged items be different from each other and the exchange be serial rather than simultaneous.

An important distinction of social exchange is that the transactions are repeated, or expected to be repeated; thus, the need for simultaneity is reduced, as in the case of a gift. Unlike the case of the neo-classical economic model, repeated transactions between agents are expected here. For instance, relations of power and dependency are created between two agents when one repeatedly needs a resource that the other controls but has no way of paying back. Thus the needy one "must subordinate himself to the other and comply with his wishes, thereby rewarding the other with power over himself as an inducement for furnishing the needed help" (Blau 1964, page 21). Naturally, subordination is not a behavior that lends itself to an instantaneous market transaction. The exchange is also not price-based but is governed by other rules pertaining to aspects such as value, similarity, and timing (Bourdieu 1977, 1997).

---

2 Indirect ties are those between ego and those that are tied to the people whom the focal individual already knows, such as a friend of a friend. Such ties can be important in attaining certain resources, as Boissevain (1974) showed, and therefore are quite different from the complete absence of ties.

3 If payment (or supply) does not happen immediately, the seller (buyer) expects formal assurance, such as a guarantee from the third side (e.g. credit card company).

An important commonality of neo-classical and social exchange is that both are cases of direct exchange. In the either case, exchange is a transaction between two agents, where both give and receive from each other, either identical or different goods, either immediately or sequentially. Generally, direct exchange "includes any system which effectively or functionally divides the group into a certain number of pair of exchange units so that, for any one pair X-Y there is a reciprocal relationship" (Lévi-Strauss 1969 [1949], page 146).

Generalized exchange, in contrast, occurs when a beneficiary is not obliged to reciprocate directly to her benefactor, but is to any other agent: "An individual feels obliged to reciprocate another's action, not by directly rewarding his benefactor, but by benefiting another agent implicated in a social exchange situation with his benefactor and himself" (Ekeh 1974, page 48). Generalized exchange is quite different from direct exchange: it neither requires immediate reciprocity nor creates a direct obligation to a specific benefactor.[4] Several subcategories of generalized exchange have been identified (Bearman 1997; Ekeh 1974; Lévi-Strauss 1963; Malinowski 1920; Sahlins 1965), and they differ from each other in the way the goods exchanged are transferred (for a recent review, see Takahashi 2005).

While generalized exchange often takes place in communities, where the agents eligible to participate are known and boundaries are drawn, it can also guide transactions among strangers (Befu 1977, 1980; Emerson 1981; Molm and Cook 1995). Helping a stranded driver on a remote mountain road, for example, is such an instance, because the benefactor does not expect the beneficiary to return the favor in the future.

It is important to note that generalized exchange is a term that describes a mode of exchange. It is neutral, however, as to the underlying motivation leading to the adoption of this mode. Scholars have attributed the emergence of generalized exchange to altruism (Sahlins 1972; Takagi 1996) and social norms (Ekeh 1974; Lévi-Strauss 1969 [1949]). However, it has been recently shown that generalized exchange can develop without general norms but just with individual notions of fairness (Takahashi 2000).

Cell 1: Neo-classical Market

The interaction of the axis produces several terms of exchange, some of which are more familiar than others. One very familiar case is cell 1, which combines spot exchange and the absence of social ties, which is essentially the case of neo-classical markets, where autonomous economic agents repeatedly search for the best price-quality combination and then engage in a transaction with whoever happens to offer that combination (Wilson 1987). The agents have no preference as to the identity of their exchange partners. The strength of a neo-classical market is that it vastly expands the choice of exchange partners: one goes to an (efficient) open-air vegetable market, searches for the best-priced (or highest-quality) tomatoes, bargains a deal with the seller, purchases a pound or two, and then goes home to make soup (Rombauer and Rombauer-Becker 1985). The following day, one can choose to search again for the best price or the highest quality, return to the same merchant, or go elsewhere if a better deal is known. The search is wide, and the gains are potentially greater. Drawbacks lie in the arms-length nature of

---

4 The sociological literature uses "generalized exchange" (Takahashi 2000) side by side with "generalized reciprocity" (Mauss 1954). After a close reading, it seems that the two terms refer to the same phenomenon. For simplicity, we use "generalized exchange" throughout this paper.

the transaction, which hinders customization, increases risk, and increases costs in haggling and negotiations. It also requires that the goods be available for simultaneous exchange; that is, both agents must have something desirable for the counterpart for the transaction to take place.[5]

Cell 2: Local Search

In cell 2, the search is narrower, as one searches only along her direct and immediate indirect ties; that is, among her acquaintances and her acquaintances' acquaintances.[6] Such a local search is a common deviation from the ideal, and it can occur as a result of the cost of a search, unavailability of information, or cognitive limitations (Levinthal and March 1993). Local search is inefficient, for it settles on the local maximum (best combination of quality and price), which is not likely to also be the global maximum (Levinthal 1997). In other words, one may find a better deal by extending the search. Local search suffers from the drawbacks mentioned above, and it does not benefit from the possibility of conducting a wide search, for which markets are especially conducive.

Cell 3: Embeddedness

Embedded exchange takes place when social and economic relations are intertwined (Granovetter 1992; Portes and Sensenberger 1993; Zukin and DiMaggio 1990). It is a combination of repeated exchange partner and social exchange logic, which allows for non-price-based transactions under less-specified terms, in comparison to the contract-based transactions in markets. Some of the transfers may be done as favors. Others do not carry a specific price tag but have to be repaid. Others may be market-like transactions but are performed in a more trusting environment, under less formal conditions. In embedded exchange, the need for simultaneity in exchange is reduced. These conditions allow the exchange partners to reap benefits that are not possible in a market exchange (Baker et al. 1998; Granovetter 1985; Gulati and Gargiulo 1999; Uzzi 1999; Uzzi and Gillespie 2002). In one of the first empirical accounts, Uzzi (1997) described some of the benefits: fine-grained information transfer, joint problem-solving arrangements, economies of time, integrative agreements, and greater willingness to invest and take associated risks. In a separate article (1996), Uzzi provided evidence on the financial benefits of embedded exchange vis-à-vis market-based exchange. Embeddedness, however, can result in several risks, primarily because the search for exchange partners is "deep rather than wide" (page 51). Agents repeatedly turned to the same exchange partners rather than searching widely for the best price-quality combination. Access to nonembedded agents may be limited, and an unforeseeable exit of a major network partner can be disastrous, as it may be difficult to replace. By interfering with the propagation of information from diverse sources, embeddedness can also disguise changes in the environment (Sorenson et al. 2002). For instance, a phase of "creative destruction" (Schumpeter 1942) can go unnoticed and be realized only when it is too late to adjust. Such contextual disruptions can be detrimental

---

5 In modern societies, one of these things is commonly money.

6 As a result of rapid attrition, the search is unlikely to succeed beyond direct and immediately indirect ties. Some research has shown that the completion rate is less than 12% even for a second-degree tie. Beyond that, more than 95% of attempts to extend a tie fail (Watts et al. 2002).

to the trust needed for embedded exchange (Axelrod 1984) and therefore lead to the collapse of the system rather than promote a successful adjustment.

Cell 4: Community/Clan

Many, if not most, of the documented cases of generalized exchange take place among specific exchange partners through direct and indirect ties (Bearman 1997; Lazega 2001). For generalized exchange along direct ties, Uehara (1990) discusses the relations between generalized exchange, solidarity, and social support, using data on a small network of low-income black women who support each other directly as they go through job loss. Here support flows to the needy — the ones who are unemployed — from their circle of supporting friends and relatives. The goods offered do not necessitate direct reciprocity, so the need for simultaneity is gone.

Communal exchange can also be lineal. A case of generalized exchange through indirect ties is intergenerational altruism: the transfer of assets from parents to their children without a direct return. Instead, the children are expected to make a transfer, in turn, to their children, the original parents' grandchildren, who are indirectly connected to the original giver. Intergenerational altruism has been used in sociobiology to explain the seemingly unreciprocated investment of parents in their children (Boorman and Levitt 1980). In the field of public policy, it was argued that intergenerational altruism can cause the family to behave financially as if it was a single, infinite-lived individual (Barro 1974). The potentially significant implications for governmental debt, retirement programs, and social security has been debated in detail (Abel and Bernheim 1991; Bernheim and Bagwell 1988).

A community that establishes generalized exchange as its mode of exchange benefits from the reduced need for simultaneity and direct reciprocity. It enjoys many of the benefits of embeddedness, plus the added ability to transact with agents who have little to offer in the short run, as calculations for direct reciprocity are eliminated.

However, it is important to distinguish between the environment where performative ties appear and the environment of communities and clans. In brief, "The Firm" is a collective that is quite different from a community or a clan, but this difference does not seem to interfere with the appearance of performative ties. A more detailed discussion follows.

Cell 5: Performative Ties

When resources are heterogeneous, such as knowledge in the cases presented here, a wide search becomes uniquely more efficient in facilitating transaction. Compared to embeddedness or community, performative ties allow a much wider search — wider than that achieved through direct and indirect ties. Even if indirect ties are used to extend the search beyond immediate locality, extension of such ties requires mediation of at least one other individual. As Burt (1992) showed, mediators of network ties gain from their control over transactions in the network. While this can be beneficial for an individual, it may interfere with organizational tasks. Even if only benevolent individuals are involved, the message can still become jumbled as it passes from one to another (cf. Winter and Szulanski 2001). Indirect ties are also likely to consume more time and resources because they require a seeker to contact not only the carrier but also everyone in between them. There is also a rapidly decreasing likelihood of reaching the carrier because of very strong cumulative attrition (Watts et al. 2002). Finally, indirect ties are still limited in reach — ego can reach only other individuals with whom she has indirect ties but not others who are not tied to ego's ties. Thus, performative ties greatly extend the scope of a search and decrease its cost, compared to the alternative of indirect ties.

Performative ties may feature the search benefits of a market transaction, but a wide search is insufficient to assure transaction, because once a potential carrier is found, the terms have to be negotiated. The data indicate that transfer between employees of unequal status was quite common in The Firm. Thus, the mode of generalized exchange provides benefits that are comparable to embeddedness and community, absent from markets, and especially beneficial in a knowledge-intensive organizational setting, such as the one studied here. First, knowledge in The Firm was typically amassed through work experience and in-house tenure. As noted above, most promotions came from within The Firm and were based on tenure, so those who were knowledge-rich tended to be rich in other resources, such as power and prestige. It was difficult for an analyst to return a favor to a manager, simply because analysts, as junior employees in The Firm, commanded few valuable resources. Second, even if an employee had the resources to pay back a favor, the structure of project work made it difficult to enforce direct reciprocity, even through subordination or deference (cf. Blau 1964). If the benefactor and the beneficiary worked in a team together, they both knew that the team would be disbanded at the end of the project, and they could never work together again. Furthermore, transfers were observed to come from knowledgeable individuals who were not members of the same team. To be sustainable, a system of favors requires sufficient trust in future reciprocation, but in the fluid environment of project work, such favors had to be returned quickly, as one never knew whether he would meet the beneficiary again. Not only did project-based work made it difficult to create a favor system, but The Firm, as do many professional service firms (Lorsch and Tierney 2002; Maister 1993), has a high turnover rate among employees, which makes future interaction even less certain. Third, the data also show that nonroutine projects, such as work in a new industry, tended to be concentrated in main offices, so that employees in main offices accumulated knowledge that was not available in smaller offices. This led to repeated cases of employees in smaller offices calling others in main offices, asking for help. Such a pattern exacerbated the nonreciprocity risk in helping a resource-poor or transient team member. Not only were main office employees approached by others they were not likely to meet again, but also there was little that a beneficiary in a peripheral office could offer in return. Indeed, knowledge transactions often benefited individuals who were unlikely to be able to reciprocate, such as in transactions between senior and junior employees, between employees located in faraway locations, and between individuals who were unlikely to meet again or who even had never met.

THE MODEL

We investigate the effect of choice of mode of exchange and partner on the efficacy of organizational problem solving. In addition, we investigate the robustness of a hybrid mode, combining performative ties with community relations, and examine the returns on firm-level investment in nurturing performative ties.

The model simulates agents that are embedded in local groups of direct ties, such as project teams, who work to solve a large overall problem. A problem is decomposed into assignable tasks for the agents. Each agent has a set of skills suitable for a set of tasks, which may or may not be the tasks assigned to it. For each of those skills, the agent has an attained competence level. Tasks are completed through the application of these task-relevant skills by the agent. Knowledge (as skill development) is attained either through self-learning, acquisition through exchange with another agent, or both. Acquisition from other agents is driven by the nature of the network exchange environment within which the agent resides and the ties it can exploit. By simulating the various types and parameters of network exchange environments and ties, as in the table above, we explore their impact on the dynamics of knowledge growth, distribution, and decline within an organization.

Agent Behavior

At the beginning of the simulation, each agent is randomly assigned tasks for solution. When an agent receives a task, it first checks to see whether it has any knowledge of the task. If so, it applies the knowledge. If the knowledge is insufficient to solve the task, the agent will endeavor to acquire the remaining knowledge through self-learning. However, if the agent possesses no knowledge about the task, it must acquire that knowledge through a process of search among the other agents present. The nature of that search process and the mechanisms of acquisition are largely determined by the pattern or patterns of exchange set by the simulation operator.

For the model, we represent the horizontal line of Table 2 (i.e., Mode of Exchange) in the following way:

1. *Spot exchange.* Agents will search to maximize the knowledge gained through the exchange. For knowledge acquisition to ensue, it is necessary that both agents agree and exchange knowledge under a strict requirement of simultaneity and direct reciprocity without incurring debt. As all exchanges are immediate (either agreed or declined), no social memory of agents or events is required.

2. *Social exchange.* Agents will engage in a knowledge exchange, but one agent can endure a debt of exchange to another agent if it is in good standing (i.e., without current debt to that agent). Therefore, agents must possess a social memory capable of distinguishing individual agents and their obligations. Again, direct reciprocity is expected, but social debt is permitted, and thus simultaneity in exchange is not required.

3. *Generalized exchange.* Agents will engage in a knowledge exchange, and, as in social exchange, one agent can endure a debt. However, the debt is one of indirect reciprocity, where the debt is obligated to a group and not to a specific individual. As such, there must be a mechanism to identify the extent to which an agent has or has not completed an obligation to the group.

The second axis of this typology (i.e., Interaction History) specifies the extent to which there is an existing social link, as a degree of familiarity, between the transaction partners. Both of these require social memory of specific agents.

1. *Extant social ties.* Extant social ties exist between transaction partners who have an existing (direct tie) relationship, are neighbors, or can be referred to each other by a common acquaintance (indirect tie).

2. *No social ties.* The agents in the transaction have no direct or indirect social ties, nor do they have a prior history of transactions.

Searching Existing Social Ties

The following three contexts are based on existing social ties, which are defined as direct ties as members of the immediate group, or indirect ties as agents who have direct ties with the immediate group. In local search, agents search only their local ties for knowledge. Knowledge will be acquired by any given agent only if the two agents can agree on an exchange, that is, if one agent has task-knowledge that is immediately useful for the other agent's task. Agents will attempt to maximize the exchange, but it is restricted to the local/group "market-like" environment. In embeddedness, agents search only their local ties in an attempt to maximize the acquisition, but unlike pure local search, this context sets the agent within a relatively stable social group that tolerates debt. Thus, opportunities for exchange revealed by this local group search are expanded by the acceptance of obligation as determined by its individual members. Here agents do not seek the optimal, but sacrifice by engaging in the first acceptable exchange condition. In community, although agents are again restricted in their search to local group ties, the nature of the social environment now transcends direct reciprocity requirements for individual transactions and affords opportunities for asymmetric exchanges without direct debt to specific individuals. Rather, agents are in the debt of the group, and that obligation can be managed by using a variety of social mechanisms, such as R-scores, standing, image scores, and altruism.

Searching Unacquainted Others

The remaining three contexts are based on agents who seek out others with whom they have no previous social ties when the examination of existing ties fails. As discussed earlier, the search is facilitated by KMS, which links task descriptions, solution descriptions, and contact information for all of the agents. In neo-classical markets, similar to the agents engaged in local search, these agents seek to maximize their exchange opportunities. However, these agents elect to search beyond the group and engage the KMS to spot potential opportunities on a firm-wide basis, and they attempt an exchange with the source who provides the most value. Similarly, an exchange requires simultaneous direct reciprocity without debt. An agent needs to have knowledge of value to the other agent. A performative tie is enabled by finding agents in the KMS who may have knowledge of potential value and have agreed to participate in this type of use of the KMS. As discussed earlier, direct reciprocity is not expected.

Agent, Group, Organizational, and Problem Structures

The agent structure includes the size (i.e., number of slots) in its knowledge memory, the decay rate of skill loss, learning rate parameters, and its strategy for skill replacement, where newly acquired skills must replace existing memory slots. The group structure is simple. One can vary the size of the teams, number of teams, turnover rate, bias in the knowledge for replacement agents, and attrition rate for direct, indirect, and performative ties. The organizational structure is a simple hierarchy with dispersed teams. The organizational problem structure, P, is represented as a vector of integers. Each element, Pi, is an assignable task to an agent, and the value of the element indicates the task difficulty in terms of required competence level. The problem structure can be manipulated as follows: difficulty (increasing or decreasing the competence level for each task); complexity (where there is a strict precedence order for implementing [i.e., posting] a solution such that the solution to Pi must be completed and posted before Pi+1 can be implemented); problem size (where the number of task elements in a problem can be changed); precision (which describes the precision required to achieve competency [i.e., slack in the competence level]); number of problems to be solved; and the redundancy each new problem has with the prior problems.

RESULTS AND DISCUSSION

Performance and Neo-classical Exchange, Embeddedness, and Community/Clan

We begin by comparing the efficacy of the modes of exchange against each other. To obtain variance, we let the agents handle difficult problems that contain a large number of subtasks (Figure 1 and Table 3).

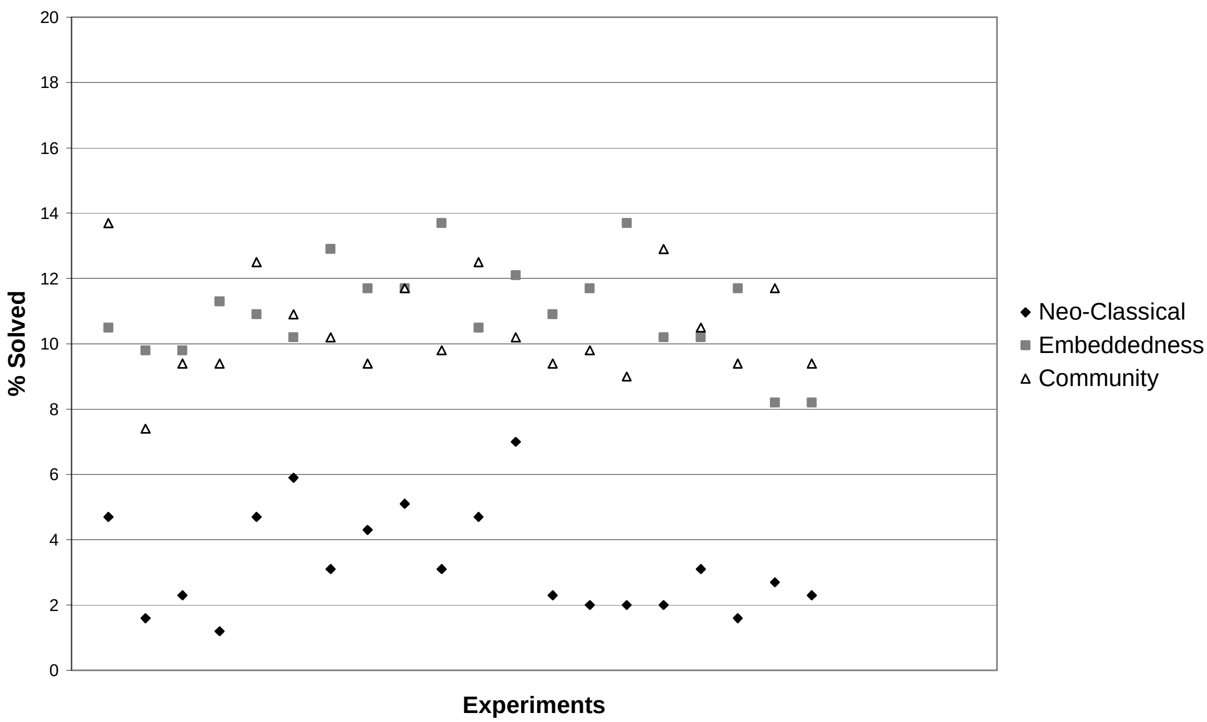

FIGURE 1 Effect of neo-classical exchange, embeddedness, and community/clan configurations on performance

TABLE 3 Effect of neo-classical exchange, embeddedness, and community/clan modes on performance

Configuration	Performance		
	Average	Median	Standard Deviation
Neo-classical exchange	3.29***	2.90	1.61
Embeddedness	11.00	10.90	1.49
Community/clan	10.46	10.00	1.58

† $p \leq 0.1$, * $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$, in a two-tailed paired t-test, n = 20.

The first experiment shows that neo-classical exchange performs poorly in comparison to either embeddedness or community/clan configurations. The difference are statistically significant between the first and the latter two, but not between embeddedness and community/clan. This finding serves to validate the model, because it confirms predictions and empirical findings. For instance, embeddedness was found to provide fine-grained information transfer, joint problem-solving arrangements, economies of time, integrative agreements, and greater willingness to invest and take associated risks (Uzzi 1996), all of which are shown to result in financial benefits of embedded exchange vis-à-vis market-based exchange (Uzzi 1997). In the context of bank loans, embedded bank/firm ties provide special governance arrangements that facilitated the firm’s access to bank-centered informational and capital resources, which uniquely enhanced the firm’s ability to manage trade credit and resulted in better financial performance (Uzzi and Gillespie 2002).

Community/clan configuration differs from embeddedness in that it employs generalized exchange as the mode of exchange rather than social exchange. Research on communities as sources of support has expanded in the last decade, examining their role in providing economic, social, and emotional benefits, which are difficult to contract in a neo-classical economic market. For instance, informal community ties are essential for survival in the impoverished barrios of Santiago, Chile. They provide food and shelter and job leads and help in dealing with bureaucracies and even short-term loans (Espinoza 1999). In Hungary, members rely on their neighborhood community for large projects, such as erecting houses (Sik and Wellman 1999). Similarly, informal ethnic associations fulfill an important mediating function between new immigrants, their societies of origin, and their new homelands. Their expected support plays a role in the decision to immigrate, and recent immigrants help new immigrants find their way in a new country once they arrive, as exemplified in work on immigrants from Hong Kong to North America (Fong et al. 1995; Salaff et al. 1999). Community institutions provide revolving credit arrangements and allow immigrants to start their own businesses, even when their commercial credit worthiness is low (Portes 1995). For the unemployed, communities can provide referrals to jobs and material and emotional support (Uehara 1990).

Similarly, clan organizations have been hailed as distinctly different from traditional hierarchical organizations, with many benefits. In Theory Z, Ouchi (1980) proposed that

Japanese firms are based on a clan logic, which is different from bureaucracies (and certainly markets), and it called for American managers to follow in changing their organizations to be more clan-like. Clan organizations are supposed to provide superior performance through stable membership of life-time employment, high interpersonal contact that facilitated nonspecialized careers and collective decision making, and organizational myths and ceremonies (Ouchi 1980; Sullivan 1983; Wilkins and Ouchi 1983).

Robustness of Performative Ties

Levine (2004) argued that performative ties are superior to other modes of exchange, because such ties combine the wide-search of neo-classical exchange with the low transaction costs of community exchange. Performative ties also allow resource-poor agents to remain productive, because dyadic reciprocity is not required. Preliminary runs of the simulations have adhered to this logic; the performance of organizations composed of performative ties agent were vastly superior to any combination.

However, it may not be realistic to assume that an organization would be entirely composed of agents that are always willing to benefit strangers. In real-life organizations, employees may be absent, busy, uncooperative, or straight-out "free riders." Hence, it is valuable to explore whether performative ties can still affect organizational performance. In this experiment, we allowed a blend of agents, the majority of which followed the community mode (i.e., within-group generalized exchange). A certain percentage of the agents were hybrid; when seeking a piece of knowledge, they began by a local search, following the logic of the community mode. However, when the knowledge was not available locally, the agents employed performative ties: they turned to the KMS, and searched for another *hybrid* agent that would be willing to exchange knowledge. It is important to note that hybrid agents could extend a performative tie only to other hybrid agents and not to the majority, which used community as the sole mode.

As evident in the results above (Figure 2 and Table 4), the hybrid configurations performed significantly better than the community configuration. In addition, an increase in the percentage of agents that were able to extend performative ties led to significantly increased performance. The results were unchanged when we included the 10% and 30% levels, as well as the levels above 50%. These results lead to a potentially important theoretical proposition: the benefits from performative ties are significant, even if the majority of the agent population adhere to another mode (in this case, community). Organizationally, the results suggests that organizations will begin to see benefits from performative ties even with low levels of individuals that are able to extend and willing to receive such ties.

Marginal Returns from Performative Ties

While performative ties can boost organizational performance above other modes of exchange, the organizational setting that is necessary to enable them may be prohibitively expensive. In this experiment, we investigate the marginal returns, in terms of organizational performance, to additional levels of agents who practice performative ties in the organizational population, under different problem structures (Figure 3).

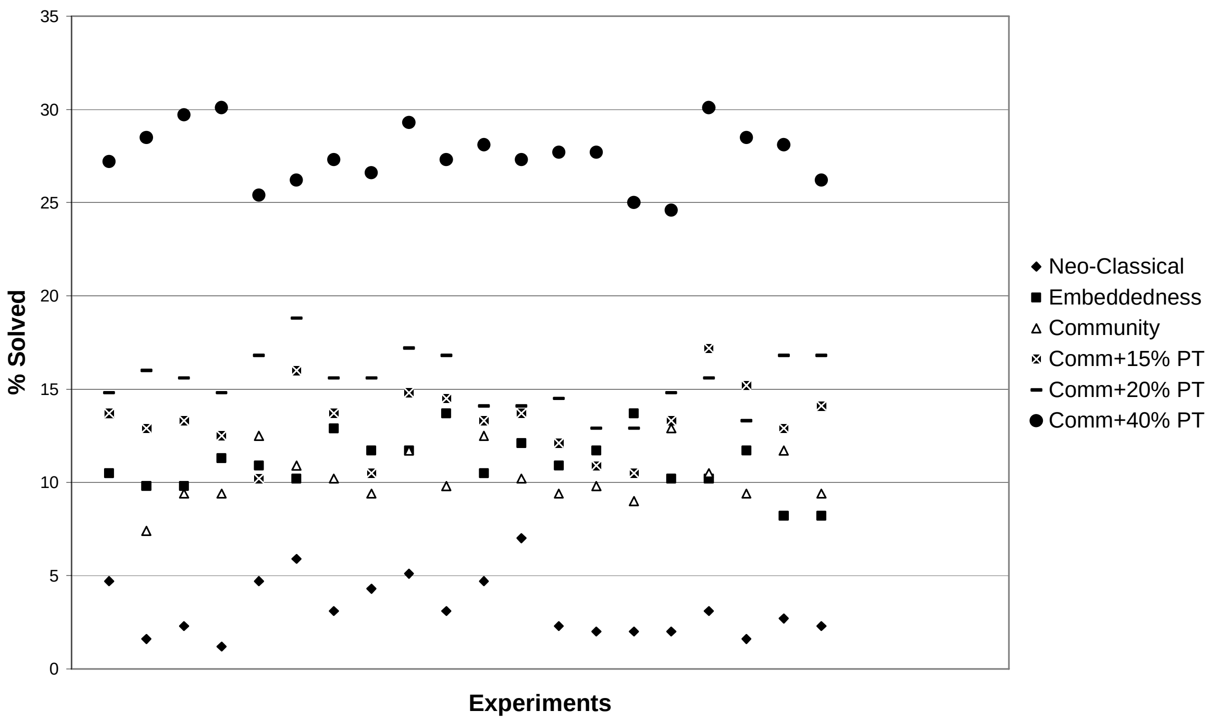

FIGURE 2 Effect of community hybrid mode on performance

TABLE 4 Effect of community hybrid mode on performance (prior tests included for comparison)

	Performance		
Configuration	Average	Median	Standard Deviation
Community + 15% hybrid	13.27**	13.30	1.84
Community + 20% hybrid	15.39***	15.60	1.54
Community + 40% hybrid	27.55***	27.50	1.59

† $p \leq 0.1$, * $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$, in a two-tailed paired t-test, n = 20.

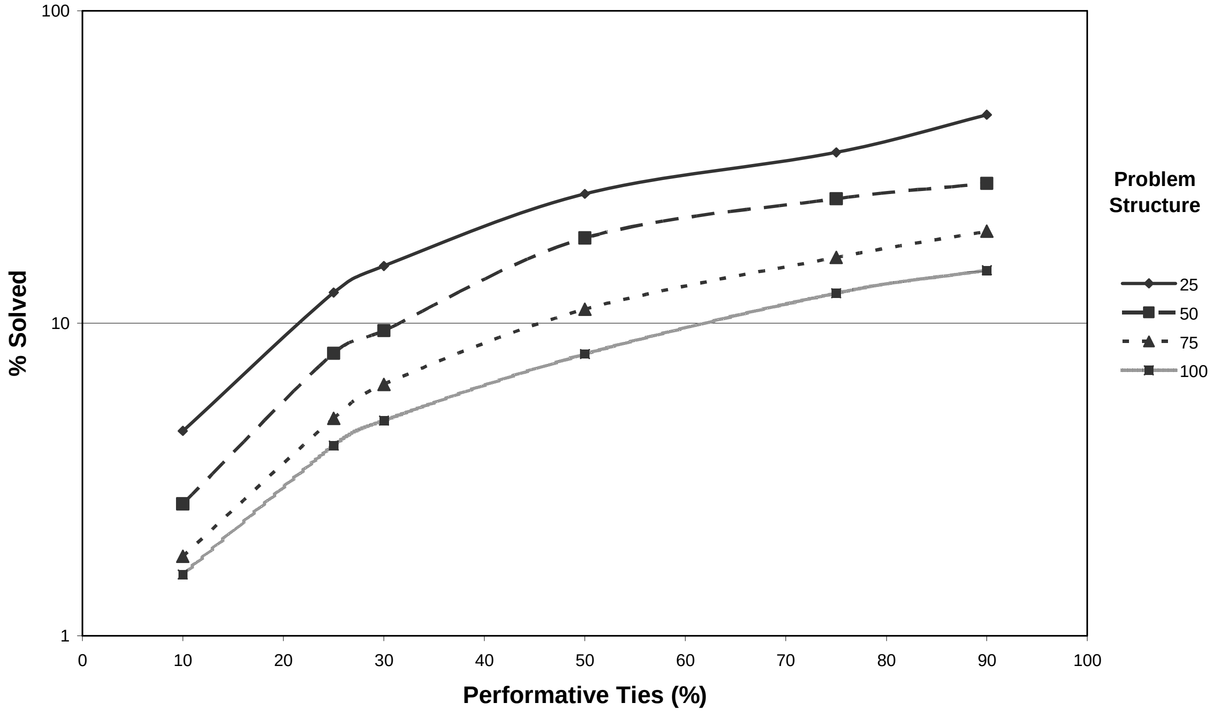

FIGURE 3 Effect of performative ties on performance under different problem structures

Performative ties are thought to be associated with a variety of organizational practices, some of which are fairly costly to implement. Most elementary, the creation and maintenance of a KMS is necessary for a wide search. A practice that may contribute to the appearance of performative ties is embeddedness of workers in multiple networks, so that the network appears dense (cf. Levine and Kurzban forthcoming). This can be achieved, for instance, by employing teams that are transient and composed of nonspecialized employees, so as to increase the perceived chances that each one of them could eventually be in the same team with any other member of the firm. Mentoring of junior employees by senior ones and cross-hierarchical teams may increase the chance of ties across levels, and multiple simultaneous team assignments immediately increase one's organizational social network.

Spatial arrangements can also increase interaction and the perception of a dense network: the rotation of employees allocated space and the generous provision of public spaces may be two ways to achieve that, but both require the expenditure of capital on extra moving and space.

Also associated with performative ties are cross-site occasions, where employees from multiple offices meet together for training or retreat, or they move to another office for a while. As Levine (2004) notes, flying even junior employees across the continent to attend a routine training session may allow them to meet more peers, but the associated cost can be borne only by wealthy firms.

Our results show that firms see the great return on their investment in such practices early on, because the marginal returns from performative ties tend to be higher at lower levels, with the steepest improvement occurring in the range of 10% to 30%. The marginal returns taper off as the percentage of hybrid agents increases, although they always remain significantly higher than zero. Similar behavior was observed along problem structures of increasing difficulty.

CONCLUSION

With the understanding that knowledge is crucial for organizational performance comes the desire to facilitate intraorganizational and interorganizational processes of knowledge exchange. We began by explicating the approaches to knowledge exchange and showing how they interact with agent characteristics. While the first study replicated prior empirical findings and theoretical propositions, the second study showed that performative ties provide superior returns, even when those who practice them are a tiny minority in the organization. We noted that this finding suggests that management will see benefits from performative ties even if only a small number adopt them.

The ability of performative ties to generate significantly higher organizational performance even at low levels may suggest that they contribute to the appearance of an organizational "small world" (Milgram 1967; Travers and Milgram 1969). Recent research on small-world networks has generated interest and interpretations of how global and local structural properties interact in dynamical systems. For example, it has been demonstrated that connecting disparate, clustered worlds by shifting a local edge in a cluster to link to a distant node has little impact on the clustering (a local property) but has a distinctly nonlinear effect on the characteristic path length (a global property) (Watts and Strogatz 1998). Because individuals are likely to communicate easily with their immediate neighbors, it is sufficient if just one of these neighbors is able to extend performative ties to enable the entire team (local cluster) to enjoy the benefit of performative ties.

Furthermore, it has interesting applications with regard to organizational diversity. It has long been argued that diversity (e.g., in gender and race) can increase organizational performance because it allows the organization to choose from a greater variety of approaches to a given problem. However, it has been recently questioned whether creating truly diverse teams is likely, or even possible (Reagans et al. 2004). At the same time, we know that individuals have social networks that are largely homogenous (i.e., composed of people that are similar to self) (McPherson et al. 2001). Our results lead us to think that creating diverse teams may not be entirely necessary. It is sufficient to have a small number of agents who can extend and receive performative ties. These may be individuals who belong to a distinct group, such as alumni of the same university. As noted earlier, their ability to extend performative ties will cascade to their neighbors. Thus, creating teams that are homogenous but include at least one agent with performative ties capability may be a sufficient substitute for the Holy Grail of full diversity in teams.

While knowledge transfer has been hailed as a means for improved organizational performance, the question of the cost of the transfer has rarely been addressed (Haas and Hansen forthcoming). This is not a moot question by any means; as is the case with any other organizational resource, the benefit from knowledge transfer should exceed the cost of facilitating the process. Our last experiment, which examines the marginal returns to performative ties, suggests that these returns are higher when performative ties are scarcer. If one bears in mind the cost of facilitating performative ties, and adds to that the fact that they immediately show a greater jump in performance early on, management may choose to keep performative ties at levels that are lower than maximum, thus economizing on costs without sacrificing a great deal of organizational performance. Future research is likely to dwell longer on the question of benefits and costs associated with knowledge.

ACKNOWLEDGMENTS

We are thankful to Esther Yansi Ye and Projjal Ghatak for able research support. Levine acknowledges a grant from the Office of Research at Singapore Management University. We benefitted from comments by participants at the Organization Science Winter Conference 2005 held in Steamboat Springs, Colorado; North American Association for Computational Social and Organizational Science (NAACSOS) 2005 held in Notre Dame, Indiana; and Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms held in Chicago, Illinois.

REFERENCES

Abel, A.B., and B.D. Bernheim (1991). Fiscal Policy with Impure Intergenerational Altruism. *Econometrica 59*(6), 1687–1711.

Arrow, K.J. (1969). Classificatory Notes on the Production and Transmission of Technological Knowledge. *The American Economic Review 59*(2) (Papers and Proceedings of the Eighty-first Annual Meeting of the American Economic Association), 29–35.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York, NY: Basic Books.

Baker, W.E., R.R. Faulkner, and G.A. Fisher (1998). Hazards of the Market: The Continuity and Dissolution of Inter-organizational Market Relations. *American Sociological Review 63*, 147–177.

Barro, R.J. (1974). Are Government Bonds Net Wealth? *The Journal of Political Economy 82*(6), 1095–1117.

Bearman, P. (1997). Generalized Exchange. *American Journal of Sociology 102*(5), 1383–1415.

Bechky, B.A. (2003). Sharing Meaning across Occupational Communities: The Transformation of Understanding on a Production Floor. *Organization Science 14*(3), 312–330, May–June.

Befu, H. (1977). Social Exchange. *Annual Review of Anthropology 6*, 255–281.

Befu, H. (1980). Structural and Motivational Approaches to Social Exchange. In K.J. Gergen, M.S. Greensberg, and R.H. Wills (Eds.), *Structural and Motivational Approaches to Social Exchange: Advances in Theory and Research*, 197–214. New York, NY: Plenum Press.

Bernheim, B.D., and K. Bagwell. (1988). Is Everything Neutral? *The Journal of Political Economy 96*(2), 308–338.

Bian, Y. (1997). "Bringing Strong Ties Back in Indirect Ties, Network Bridges, and Job Searches in China." *American Sociological Review 62*(3), 366–385.

Blau, P.M. (1964). *Exchange and Power in Social Life*. New York, NY: J. Wiley.

Boissevain, J. (1974). *Friends of Friends: Networks, Manipulations, and Coalitions*. London, England: Blackwell.

Boorman, S.A., and P.R. Levitt (1980). *The Genetics of Altruism*. New York, NY: Academic Press.

Bourdieu, P. (1977). *Outline of a Theory of Practice*. New York, NY: Cambridge University Press.

Bourdieu, P. (1997). Marginalia: Some Additional Notes on the Gift. In A.D. Schrift (Ed.), *The Logic of the Gift: Toward an Ethic of Generosity*. New York, NY: Routledge.

Burt, R.S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, R.S. (2004). Structural Holes and Good Ideas. *American Journal of Sociology 110*(2), 349–399.

Burt, R.S. (2005). Second-hand Brokerage. Working Paper. The University of Chicago.

Coleman, J.S. (1988). Social Capital in the Creation of Human Capital. *American Journal of Sociology 94* (Supplement: Organizations and Institutions: Sociological and Economic Approaches to the Analysis of Social Structure), S95–S120.

Constant, D., L. Sproull, and S. Kiesler (1996). The Kindness of Strangers: The Usefulness of Electronic Weak Ties for Technical Advice. *Organization Science 7*(2), 119–135.

Ekeh, P.P. (1974). *Social Exchange Theory: The Two Traditions*. Cambridge, MA: Harvard University Press.

Emerson, R.M. (1981). Social Exchange Theory. In M. Rosenberg and R.H. Turner (Eds.), *Social Psychology: Sociological Perspectives,* 30–65. New York, NY: Basic Books.

Feld, S.L. (1981). The Focused Organization of Social Ties. *American Journal of Sociology 86*(5), 1015–1035.

Granovetter, M. (1973). The Strength of Weak Ties. *American Journal of Sociology 78*(6), 1360–1380.

Granovetter, M. (1985). Economic Action and Social Structure: The Problem of Embeddedness. *American Journal of Sociology 91*(3), 481–510.

Granovetter, M. (1992). Problems of Explanation in Economic Sociology. In N. Nohria and R.G. Eccles (Eds.), *Networks and Organizations: Structure, Form and Action,* 25–56. Boston, MA: Harvard Business School Press.

Gulati, R., and M. Gargiulo (1999). Where Do Interorganizational Networks Come from? *American Journal of Sociology 104*(5), 1439–1493.

Haas, M.R., and M.T. Hansen (forthcoming). When Using Knowledge Can Hurt Performance: The Value of Organizational Capabilities in a Management Consulting Company. *Strategic Management Journal*.

Hansen, M.T. (1999). The Search-Transfer Problem: The Role of Weak Ties in Sharing Knowledge across Organization Subunits. *Administrative Science Quarterly 44*(1), 82–111.

Hansen, M.T. (2002). Knowledge Networks: Explaining Effective Knowledge Sharing in Multiunit Companies. *Organization Science 13*(3), 232–248.

Hargadon, A.B., and R.I. Sutton (1997). Technology Brokering and Innovation in a Product Development Firm. *Administrative Science Quarterly 42*(2), 716–749.

Homans, G.C. (1958). Social Behavior as Exchange. *American Journal of Sociology 63*, 597–606.

Ibarra, H. (1992). Homophily and Differential Returns: Sex Differences in Network Structure and Access in an Advertising Firm. *Administrative Science Quarterly 37*(3), 422–447.

Kogut, B., and U.B. Zander (1995). Knowledge and the Speed of the Transfer and Imitation of Organizational Capabilities: An Empirical Test. *Organization Science 5*(1), 76–92.

Lazega, E. (2001). *The Collegial Phenomenon: The Social Mechanisms of Cooperation among Peers in a Corporate Law Firm*. New York, NY: Oxford University Press.

Levin, D.Z., and R.L. Cross (2004). The Strength of Weak Ties You Can Trust: The Mediating Role of Trust in Effective Knowledge Transfer. *Management Science 50*(11), 1477–1490.

Levine, S.S. (2004). The Strength of Performative Ties: Dynamics of Knowledge Network Exchange, *Status and Reputation in Networks*. New Orleans, LA: Academy of Management.

Levine, S.S. (2005). A Framework for Collective Open Source Innovation. Working Paper. Singapore Management University.

Levine, S S., and R. Kurzban (forthcoming). Explaining Clustering in Social Networks: Towards an Evolutionary Theory of Cascading Benefits. *Managerial and Decision Economics*.

Levinthal, D.A. (1997). Adaptation on Rugged Landscapes. *Management Science 43*(7), 934–950.

Levinthal, D.A., and J.G. March (1993). The Myopia of Learning. *Strategic Management Journal 14* (Special Issue: Organizations, Decision Making and Strategy), 95–112.

Lévi-Strauss, C. (1963). *Structural Anthropology*. New York, NY: Basic Books.

Lévi-Strauss, C. (1969). *The Elementary Structures of Kinship* (J.H. Bell, J.R. von Sturmer, and R. Needham, Trans. Rev. ed.). Boston, MA: Beacon Press.

Lorsch, J.W., and T.J. Tierney (2002). *Aligning the Stars: How to Succeed When Professionals Drive Results*. Boston, MA: Harvard Business School Press.

Maister, D.H. (1993). *Managing the Professional Service Firm.* Free Press.

Malinowski, B. (1920). Kula: The Circulating Exchange of Valuables in the Archipelagoes of Eastern New Guinea. *Man 20*, 97–105.

Marsden, P.V. (1988). Homogeneity in Confiding Relations. *Social Networks 10*, 57–76.

Marsden, P.V., and K.E. Campbell (1984). Measuring Tie Strength. *Social Forces 63*(2), 482–501.

Mauss, M.-I. (1954). *The Gift; Forms and Functions of Exchange in Archaic Societies* (I. Cunnison, Trans.). London, England: Cohen & West.

McDermott, R. (1999). Why Information Technology Inspired but Cannot Deliver Knowledge Management. *California Management Review 41*(4), 103–117.

McPherson, M., L. Smith-Lovin, and J.M. Cook (2001). Birds of Feather: Homophily in Social Networks. *Annual Review of Sociology 27*(1), 415–444.

Milgram, S. (1967). The Small World Problem. *Psychology Today*(1), 60–67.

Molm, L.D., and K.S. Cook (1995). Social Exchange and Exchange Networks. In K.S. Cook, G.A. Fine, and J.S. House (Eds.), *Sociological Perspectives on Social Psychology*, 209–235. Boston, MA: Allyn & Bacon.

Nelson, R.E. (1989). The Strength of Strong Ties: Social Networks and Intergroup Conflict in Organizations. *Academy of Management Journal 32*(2), 377–401.

Newcomb, T.M. (1961). *The Acquaintance Process*. New York, NY: Holt.

Newcomb, T.M. (1966). The General Nature of Peer Group Influence. In T.M. Newcomb and E.K. Wilson (Eds.), *College Peer Groups*, 2–16. Chicago, IL: Aldine.

Orr, J.E. (1990). Sharing Knowledge, Celebrating Identity: Community Memory in a Service Culture. In D.S. Middleton and D. Edwards (Eds.), *Collective Remembering*, 196–189. Beverly Hills, CA: Sage Publications.

Orr, J.E. (1996). *Talking about Machines: An Ethnography of a Modern Job*. Ithaca, NY: Cornell University Press.

Portes, A., and J. Sensenberger (1993). Embeddedness and Immigration: Notes on the Social Determinants of Economic Action. *American Journal of Sociology 98*(6), 1320–1350.

Reagans, R.E., E.W. Zuckerman, and B. McEvily (2004). How to Make the Team: Social Networks vs. Demography as Criteria for Designing Effective Teams. *Administrative Science Quarterly 49*(1), 101–134.

Rombauer, I.S., and M. Rombauer-Becker (1985). *Joy of Cooking*. New York, NY: MacMillan.

Sahlins, M. (1965). *Essays in Economic Anthropology*. Seattle, WA: University of Washington Press.

Sahlins, M. (1972). *Stone Age Economics*. Chicago, IL: Aldine Atherton.

Saxenian, A. (1996). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128.* MA: Harvard University Press.

Schumpeter, J. (1942). *Capitalism, Socialism and Democracy* (2nd ed.). London, England: George, Allen & Unwin, Ltd.

Simon, H.A. (1957). *Models of Men*. New York, NY: Wiley.

Sorenson, O., J.W. Rivkin, and L. Fleming (2002). Complexity, Networks and Knowledge Flow. Working Paper. Harvard Business School, 02-09. Available at http://ssrn.com/abstract=310001.

Straits, B.C. (1991). Bringing Strong Ties Back in Interpersonal Gateways to Political Information and Influence. *Public Opinion Quarterly 55*(3), 432–448.

Takagi, E. (1996). The Generalized Exchange Perspective on the Evolution of Altruism. In W.B.G. Liebrand and D.M. Messick (Eds.), *Frontiers in Social Dilemmas Research*, 311–336. Berlin, Germany: Springer-Verlag.

Takahashi, N. (2000). The Emergence of Generalized Exchange. *American Journal of Sociology 105*(4), 1105–1134.

Takahashi, N. (2005). Generalized Exchange. In G. Ritzer (Ed.), *Encyclopedia of Social Theory*. New York, NY: Russell Sage.

Teece, D.J. (1977). Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-how. *Economic Journal 87*, 242–261.

Teece, D.J. (1998). Capturing Value from Knowledge Assets: The New Economy, Markets for Know-how, and Intangible Assets. *California Management Review 40*(3), 55–79.

Travers, J., and S. Milgram (1969). An Experimental Study of the Small World Problem. *Sociometry 32*(4), 425–443.

Uehara, E. (1990). Dual Exchange Theory, Social Networks, and Informal Social Support. *American Journal of Sociology 96*(3), 521–557.

Uzzi, B. (1996). The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Journal of Sociology 61*(4), 674–698.

Uzzi, B. (1997). Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness. *Administrative Science Quarterly 42*, 35–67.

Uzzi, B. (1999). Embeddedness in the Making of Financial Capital: How Social Relations and Networks Benefit Firms Seeking Financing. *American Sociological Review 64*(4), 481-505.

Uzzi, B., and J.J. Gillespie (2002). Knowledge Spillover in Corporate Financing Networks: Embeddedness and the Firm's Debt Performance. *Strategic Management Journal 23*(7), 595–618.

von Hippel, E. (1987). Cooperation between Rivals: Informal Know-how Trading. *Research Policy 16*, 291–302.

Watts, D.J., P.S. Dodds, and M.E.J. Newman (2002). Identity and Search in Social Networks. *Science 296*, 1302–1305.

Watts, D.J., and S.H. Strogatz (1998). Collective Dynamics of "Small World" Networks. *Nature 393,* 440–442.

Wellman, B., and S. Wortley (1990). Different Strokes from Different Folks: Community Ties and Social Support. *American Journal of Sociology 96*(3), 558–588.

Wilson, R.B. (1987). Exchange (Entry on Walrasian Models of Exchange). In J. Eatwell, M. Milgate, and P. Newman (Eds.), *The New Palgrave: A Dictionary of Economics* (Vol. 2). London, England: Macmillan Press Ltd.

Winter, S.G. (1987). Knowledge and Competence as Strategic Assets. In D.J. Teece (Ed.), *The Competitive Challenge: Strategies for Industrial Innovation and Renewal,* 159–184. Cambridge, MA: Ballinger Publishing Co.

Winter, S.G., and G. Szulanski (2001). Replication as Strategy. *Organization Science 12*(6), 730–743.

Zukin, S., and P. DiMaggio (1990). Introduction. In S. Zukin and P. DiMaggio (Eds.), *Structures of Capital: The Social Organization of the Economy*, 1–37. New York, NY: Cambridge University Press.

AGENT-BASED MODELS OF URBAN INDUSTRIAL SPECIALIZATION

M. ZEIDENBERG,* Department of Sociology, University of Wisconsin, Madison, WI

ABSTRACT

Urban industrial specialization is advantageous because a set of related industries can share suppliers, expert workers, managers, and engineers, and because there are positive economies of scale and scope associated with such specialization. In the face of the limitations of anything but the most simplistic mathematical model of such a process, an agent-based model is a good way to gain insight into the evolution of such urban specializations, if related agents are attracted to one another. After an initially random placement of establishments, the system evolves so that related agents move near one another. In the version of the agent-based model described herein, there are only two types of establishments: core establishments and supplier establishments. Each establishment belongs to a particular industrial sector. There is a directed bipartite graph connecting supplier sectors to core sectors. The result of this is the self-organized emergence of urban industrial specializations. The "cities" that emerge are not of uniform size, but vary substantially in size, as they do in the real world.

Keywords: Urban geography, supply chains, urban economics, economic simulation

INTRODUCTION

Actual markets function much differently than the idealized markets of neoclassical economic theory. In neoclassical theory, space does not exist, nor do social networks; buyers meet sellers in a perfect auction market. In real markets, there is a distribution of sizes of firms, and firms operate in social networks and in space. Relatively few large firms tend to dominate particular industries. For instance, banking is dominated by large banks such as Citicorp and Chase, aerospace by Boeing and Airbus. Even in industries such as automobiles, a dozen or so huge firms dominate the world market. And trends constantly move toward further consolidation. This is because larger firms can take advantage of economies of scale, scope in production, and, equally if not more importantly, scope in marketing. Thus a large firm like General Electric is constantly acquiring smaller, more entrepreneurial firms, so that the advantages that GE has in both production and marketing can be wedded with the innovations of these smaller firms in order to take more profitable advantage of these innovations.

Urban history is inextricably intertwined with the histories of large firms. Of course, Detroit is historically associated with automobiles. Hollywood is associated with the large motion picture studios. Silicon Valley is associated with the computer industry. New York City is associated with many industries, notably investment banking and the stock exchange (Wall Street), the clothing/design industry (Seventh Avenue), the advertising industry (Madison

* *Corresponding author address*: Matthew Zeidenberg, Department of Sociology, University of Wisconsin, Madison, WI 53706; e-mail: zeidenbe@ssc.wisc.edu.

Avenue), and the publishing industry. Other high-wage industries, such as legal, accounting, and medical services, have grown up to serve its diversified economy.

Larger cities tend to have more diversified economies and therefore are less vulnerable to downturns in particular industries. In order to understand why larger cities have more diversified economies, we need to understand the relationship between industrial location and urban growth.

Despite the hype about the importance of small business, large firms still dominate the economies of the economically advanced countries. In fact, many smaller firms exist mainly to serve the needs of the larger firms or of their workers. Thus, a cluster of large hospitals and universities is serviced by many small service firms nearby, such as restaurants, bookstores, and photocopy shops. Or, on an industrial model, a large automobile plant is serviced by a large number of suppliers, making all sorts of components that go into the car, such as the seats, the dashboard plastic molding, or precision parts that are used to assemble the engine.

THE ECONOMIES OF PROXIMITY

The basic idea of economic geography, even in the information age, is that there are economies associated with proximity. (Ironically, one of the best-known industrial districts, Silicon Valley, is itself the quintessential information-based economy, giving the lie to the idea that geography doesn't matter for the most "advanced" parts of the economy.)

Of course, firms do not respond to proximity alone. In fact, they primarily respond to the availability of markets for their goods. But proximity minimizes transportation costs and often transaction costs as well. Even in the age of low communication and computation costs, there are savings associated with face-to-face communications, especially in coordinating complex activities. If economic activities are highly standardized, are not heavily dependent on human knowledge and on the interactions between skilled and/or educated workers, and do not change rapidly over time, then they typically can be done at a distance. However, if one or more of these situations do not hold, there is still typically an advantage to proximity. Firms in a particular industry or set of related industries still tend to locate close to one another. Locating near one another allows suppliers and large firms (OEMs, or original equipment manufacturers) to tap into a shared labor market with specific knowledge of that particular industry.

Workers with skills relevant to this labor market also tend to move to the industrial district as well to take advantage of the "ideas in the air." Often industrial districts are so rich with knowledge of an industry that it is difficult to determine who originated particular ideas, and new firms within the district tend to be born and die frequently. In addition, particular large firms grow up, often rapidly (e.g., Google, Ford, Hewlett-Packard), and come to play dominant roles. Industrial districts have been recognized by economists for a long time; the term "industrial districts" appears to have been first used by the British economist Alfred Marshall (1890, 1920). There is a substantial contemporary literature on industrial districts and the similar concept of industrial clusters, most of which is qualitative, descriptive, prescriptive, and analytical; see, for example, Piore and Sabel (1992), Porter (1998), and Harrison (1992), among many others.

Supply Chains and the Interaction of Location Decisions

Location decisions of firms are based on location decisions of other firms. If a large firm locates in a particular place, firms that supply it are likely to locate nearby. If one OEM moves into a particular location and attracts workers and suppliers, this may attract another OEMs in the same or a similar industry. This is also true of other large firms in the service industries, such as health care, financial services, and education.

Empirical studies of city sizes have shown that city populations tend to obey a power law. A simple model of population growth can account for such a power law (Krugman 1996). However, a more complex process also underlies this, as people move where the jobs are, and firms move where the people and the other firms are.

Real economies function on a supply chain; that is, a series of inputs creates the final product. Multiple supply chains converge on a single point: the final product. Labor and capital are the inputs to these supply chains. Many workers are skilled and specialized and associated with only one or a few supply chains. This is true of the service industries as well as manufacturing.

Since supply chains have more than one level of supplier, it is an oversimplification to say that the economy can be modeled in terms of OEM-supplier relationships. More realistically, there are relationships between suppliers at various tiers of the supply chain, and then the final supplier-OEM relationships.

The economy could be modeled with agents representing individuals, some of whom are workers, who can hold multiple jobs. There are also nonworking individuals, such as children, the retired, the unemployed, the disabled, and stay-at-home parents, who do not contribute to production that is captured in the market but nevertheless contribute to consumer demand. Each of these could be modeled with an agent, and person-level agents could be grouped into households.

On the firm side, one could model the entire supply chain, but one would need to know the topology of the chain and the relative numbers of establishments at each level. One would also need to know the demand flows that run between each pair of establishments in the chain. Much of this information can in fact be determined from input-output models of regional economies and establishment data, such as those from the U.S. Census Bureau (2002) and U.S. Bureau of Economic Analysis (1997), but doing so would be a rather Herculean task. Finally, one would need to know final demand which is also available from the input-output tables.

Simple Model of the Geography of Supplier Relations

In an initial, simple model, there would be two main types of agents — establishments and workers — and each agent would make decisions about where to locate based on access to markets and proximity, which are related to one another. This is based on the theory that relations between firms, like other social relations (e.g., between friends or establishments), are "sticky" and that firms like to do business with other firms with which they have longer-term relationships and with which they have established relationships of trust. They do not want to constantly be switching suppliers in order to get a rock-bottom price, unless they are buying a commodity that is readily available with a reliable price and quality from a large number of vendors. In a fully fleshed-out model, it also would be necessary to model commodities and the price mechanism, perhaps by using a model of trade similar to that used in Sugarscape (Epstein and Axtell 1996).

However, for an initial model of supplier-OEM relations in an urban landscape, it is not necessary to have this much complexity. In fact, it makes sense to start with a simpler model that only has a few features of the more complex one just described. A simple agent-based model can capture the utility associated with proximity by creating agents that respond to proximity alone.

My initial model has just two levels in the supply chain — suppliers and OEMs — and does not model workers at all. It also does not model final demand or the price mechanism. All it does is model the responsiveness of the agents to proximity. However, just because the model is simple doesn't mean that it cannot generate insight. To take this position is similar to criticizing Schelling's (1978) famous model of segregation because it did not take account of housing prices or social class.

Generally, there are more suppliers than OEMs, although the relative number varies in a complex way through the supply chain and is dependent on the industry. Thus, if we are, as a first cut, going to model the economy in terms of relations between suppliers and OEMs, we need to have more suppliers than OEMs.

Analytically, we have a distinction between a supplier *type* and an individual supplier *establishment*. A bipartite graph represents the relations between supplier types and establishment types. Each node in that graph represents a set of individual establishments of that type. When a supplier type node is connected to an OEM type node, this indicates that there are one or more instances of that type of supplier that supply that type of OEM. Thus a single link between nodes in the type graph represents one or more links between instance nodes, which may (and usually are) multiple on each side of the bipartite graph.

Given a fixed number of supplier types and a fixed (different) number of OEM types, there are many possible topologies of such a bipartite graph. For instance, one of the simplest possibilities is the following: we have two supplier types 1 and 2 and two OEM types C and D, where suppliers of type 1 only supply to OEMs of type C, and suppliers of type 2 only supply to OEMs of type D. Thus the bipartite graph consists of two disjoint pieces. Alternately, suppliers of types 1 and 2 both supply to suppliers of types C and D. In this case, the bipartite graph is as completely connected as such a four-node bipartite graph can be.

In my simplified model, with just one layer in the supply chain represented by this bipartite graph, the inputs to the model include the following: the number of distinct supplier types, the number of establishments for each supplier type, the number of distinct OEMs, and the number of distinct OEM establishments for each OEM type. In addition, the graph itself is input to the model.

The urban landscape is a simple square grid of cells. All of the suppliers and OEMs are initially placed in random locations on the grid. Each supplier and OEM agent, taken in turn, is given the choice of moving from its current location to another location in the grid, where a fixed number of random, unoccupied locations is considered. The move is made from the current location to the new location with the maximal utility for that agent, if that maximal utility exceeds the current utility. Otherwise, the agent stays put.

There are obviously many options with regard to the utility function. One of the simplest functions that rewards supplier-OEM proximity is to count the number of agents that would be neighbors (in the Moore neighborhood) and are also adjacent in the bipartite graph — that is, a supply relationship could exist between the agent in question and the neighbor agent. The more neighbors and potential supply relationships, the better. A possible enhancement is to attach some disutility to the presence of neighbors of one's own particular type, perhaps after some threshold is reached. This would amount to creating a disincentive for firms to move to a locale if there is too much competition and congestion. If supply relations are in fact "sticky," it is realistic to think that a market would be hard to break into once it is saturated. Such a disincentive is probably a major factor that prevents all the firms and population in the country from "lumping up" in one place (others are the availability of natural resources and the climatic preferences of the population).

Model Results

I have not yet implemented a more complex utility function. Instead, I have limited my experiments so far to experimentation with the topology of the bipartite graph. In the first experiment, there are two supplier types and two OEM types. Each supplier of type 1 supplies to OEMs of type A; likewise for 2 and B. There are 50 suppliers of each type, and 10 OEMs of each type (reflecting the fact that suppliers tend to exceed OEMs in number). Thus there are a total of 120 agents in the system. On one particular run (the runs can differ because of differences in the random initialization of the grid), the 120 agents are found in 99 clusters of adjacent agents across the grid; thus, most agents are initially in singleton clusters.

On this run, the agents are all initially scattered throughout the grid. After 5,000 updates of the grid, in which the system attempts to move each agent in turn to another location, moving it if the new location has more supply chain graph neighbors than the old location, the grid has been updated into 15 "cities," which are either 1A cities or 2B cities. The mean city size is therefore eight agents, and there are a variety of cities of different sizes. The results at the end are shown in Figure 1.

In the second version of the model, there are four types of suppliers, labeled with the numbers 1–4, and four types of OEMs, labeled with the letters A–D. Each of 1 and 2 supply to both A and B, and each of 3 and 4 supply to both C and D. As before, there are 50 suppliers of each type, and 10 OEMs of each type. There are a total of 240 agents, and there are initially 165 "cities," so singletons are somewhat less common than before, because the grid is initially more densely populated.

After 5,000 grid updates, we have 21 cities, as shown in Figure 2. Thus the mean city size is of the same magnitude as before; here it is around 11.4 as opposed to the prior 8. Unlike before, we have some cities that cross over the supplier relations; that is, they contain suppliers and OEMs that are not connected by a relation. This is simply due to the increased overall congestion. However, within these cities, there are neighborhoods that are governed by the supplier relation.

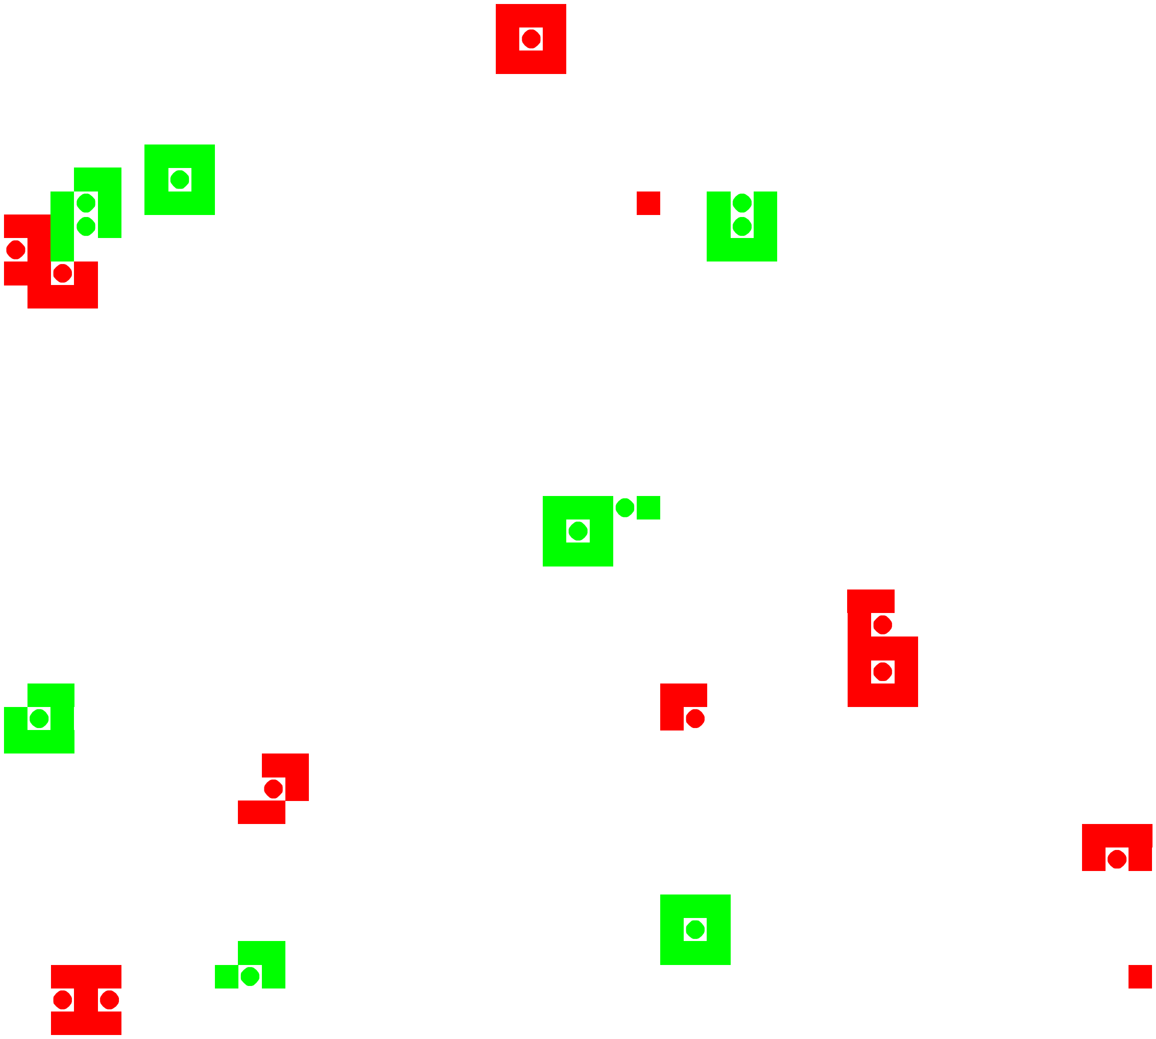

FIGURE 1 State of a grid of suppliers and OEMs after 5,000 iterations; graph topology consisting of two supplier-OEM pairs

Legend:
Red square: Supplier 1
Green square: Supplier 2
Red circle: OEM A
Green circle: OEM B

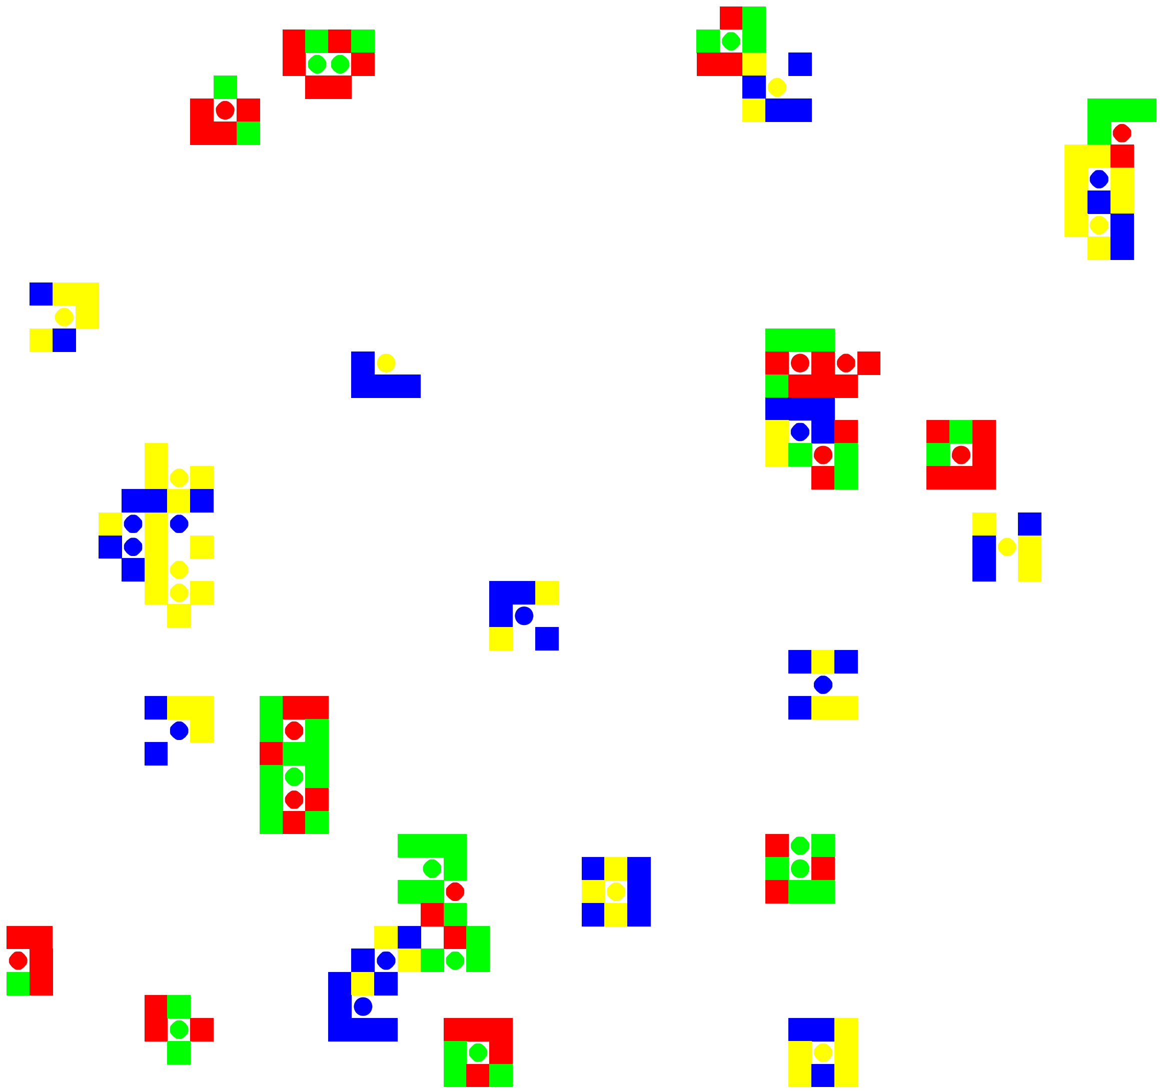

FIGURE 2 State of a grid of suppliers and OEMs after 5,000 iterations; relations between OEMs and suppliers are characterized by two disjoint fully connected bipartite graphs, each consisting of two supplier types and two OEM types

Legend:
Red square: Supplier 1
Green square: Supplier 2
Blue square: Supplier 3
Yellow square: Supplier 4
Red circle: OEM A
Green circle: OEM B
Blue circle: OEM C
Yellow circle: OEM D

The third version of the model is the same as the second, except that an additional supplier/OEM relation is given, between supplier 3 and OEM B. The result after 5,000 grid updates is shown in Figure 3. This slightly increases the probability of urban clumping, and the number of cities falls to 19, raising the mean city size to about 12.6.

Thus we see that a relatively simple model can account for the emergence of "urban industrial districts." Further refinements, as I have described, should account for more details of urban economics and geography.

REFERENCES

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom Up*, Washington, DC: Brookings Institution Press.

Harrison, B., 1992, "Industrial Districts: Old Wine in New Bottles?" *Regional Studies*, Vol. 26, pp. 469-483.

Krugman, P., 1996, *The Self-Organizing Economy*, Cambridge, MA: Blackwell Publishers.

Marshall, A., 1890, 1920, *Principles of Economics,* Amherst, NY: Prometheus Books.

Piore, M., and C. Sabel, 1984, *The Second Industrial Divide: Possibilities for Prosperity*, New York: Basic Books.

Porter, M.E., 1998, "Clusters and the New Economics of Competition," *Harvard Business Review*, November-December.

Schelling, T., 1978, *Micromotives and Macrobehavior*, New York: Norton.

U.S. Bureau of Economic Analysis, 1997, *Benchmark Input-Output Account*s; available at http://www.bea.gov/bea/dn2/home/benchmark.htm.

U.S. Census Bureau, 2003, *County Business Patterns;* available at http://www.census.gov/epcd/cbp/view/cbpview.html.

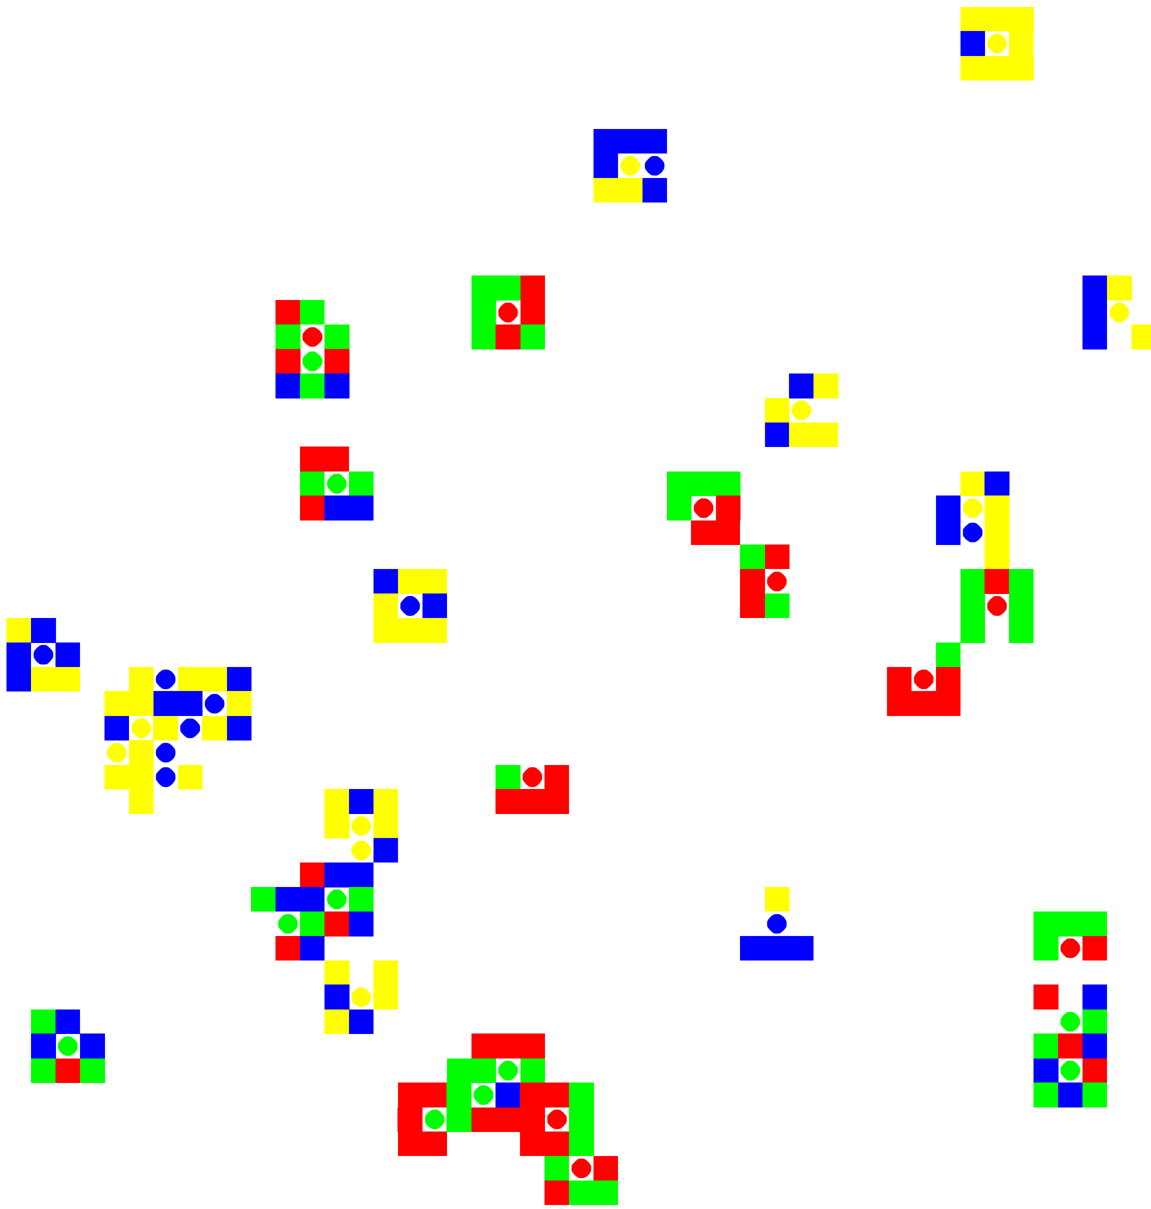

FIGURE 3 State of a grid of suppliers and OEMs after 5,000 iterations; relations between OEMs and suppliers are characterized by two disjoint fully connected bipartite graphs, each consisting of two supplier types and two OEM types, and one additional connection between the two bipartite graphs that would otherwise be disjoint

Legend:
Red square: Supplier 1
Green square: Supplier 2
Blue square: Supplier 3
Yellow square: Supplier 4
Red circle: OEM A
Green circle: OEM B
Blue circle: OEM C
Yellow circle: OEM D

MODULAR BAYESIAN INFERENCE AND LEARNING OF DECISION NETWORKS AS STAND-ALONE MECHANISMS OF THE MABEL MODEL: IMPLICATIONS FOR VISUALIZATION, COMPREHENSION, AND POLICY MAKING

K. ALEXANDRIDIS* and B. PIJANOWSKI,
Department of Forestry and Natural Resources, Purdue University, West Lafayette, IN

ABSTRACT

This paper describes a modular component of the MABEL model agents' cognitive inference mechanism. The probabilistic and probabilogic representation of the agents' environment and state space is coupled with a Bayesian belief and decision network functionality, which in fact holds Markovian semiparametric properties. Different approaches to modeling multi-agent systems are described and analyzed; problem-, model-, and knowledge-driven approaches to agent inference and learning are emphasized. The notion of modularity in agent-based modeling components is conceptualized. The modular architecture of the decision inference mechanism allows for a flexible architectural design that can be either endogenous or exogenous to the agent-based simulation model. A suite of decision support tools for modular network inference in the MABEL model is showcased; the emphasis is on the component object model versus interoperability development interfaces. These tools provide the complex functionality of developing "models within models," thus simplifying the need for extensive research support and for a high-end level of knowledge acquisition from the end-users' perspective. Finally, the paper assesses the validity of visual modeling interfaces for data- and knowledge-acquisition mechanisms that can provide an essential link between an in vitro research model, and the complex realities that are observed and processed by decision-makers, policy-makers, communities, and stakeholders.

Keywords: Agent-based model, MABEL, Bayesian belief networks, Bayesian decision networks, visualization, decision-theoretic inference, policy making

INTRODUCTION

Agent-based systems and models have passed through many stages in their historical evolution: from experimentation leading to discovery; to architectural modeling and the development of models and mathematical representations; to game-theoretic or mental simulation modes; to more realistic and robust applications; and to theory construction and the study of complex, robust, and resilient structures and patterns. Recent advances in the multi-disciplinary research and modeling of complex systems (e.g., spatial complexity, complex network dynamics) laid out the roadmap for advancing the comprehensibility, usability, and applicability of agent-based models and mechanisms for a wide variety of applications, decision makers, and policy makers (Ma and Nakamori 2005; McIntosh et al. 2005).

* *Corresponding author address*: Konstantinos Alexandridis, Human-Environment Modeling and Analysis Laboratory (HEMA), Department of Forestry and Natural Resources, Purdue University, 204 Forestry Building, 195 Marsteller Street, West Lafayette, IN 47907; email: ktalexan@purdue.edu.

In the special case of spatially explicit agent-based models, advances in computational methods and resources and in complex multi-disciplinary ecological and natural resource research methodologies — in conjunction with advances in more specialized statistical and probabilistic approaches to modeling, estimation, and assessment at the spatial level — allow researchers to explore an additional dimension of research related to model utilization by decision and policy makers (Foley et al. 2005; Heemskerk et al. 2003; van Vuuren and Bouwman 2005; Verburg et al. 2004). Agent-based models are becoming more and more commonly encountered not simply as valuable research tools for the discovery and analysis of complex systems but also as end-user mechanisms for enhancing decision-making capabilities, testing policy-specific dimensions of spatial problems, and exploring additional scenarios and alternative futures never realized in the past. These new dimensions of research are part of a wider goal and vision of integrating science into society in ways that enhance the well-being and welfare of citizens and society as a whole (Brown et al. 2002; Drennan 2005; Ehrlich and Kennedy 2005; Hammersley, 2005).

Compared to traditional models (i.e., economic, statistical, and equation-based models) or to single-disciplinary methods, spatially explicit agent-based systems present greater challenges to the integration scheme described above. A varying degree of embedded systemic complexity, coupled with a degree of uncertainty embedded in the natural and socioeconomic systems of our real world, is often viewed and misunderstood as a form of “black box” science. Methods and techniques that may be widely accepted among scientists and researchers are not easily adopted by stakeholders, decision makers, or policy experts because their lack of simplicity and transparency renders their comprehension and diffusion a challenge.

This paper proposes and demonstrates the usability and value of modular components in an agent-based framework, specifically, the Multi-Agent-Based Economic Landscape (MABEL) model. One important modular component of MABEL — namely, the cognitive inference mechanism — is described. The functionality of the inference mechanism as it relates to the mechanics of the modeling architecture is defined and encapsulated. This paper also demonstrates the value of utilizing this core mechanism for building a user-oriented interactive interface that enhances user experience and, at the same time, integrates user inputs back into the modeling process. The emphasis is on the ability of the interface to interact with the simulation framework to provide useful analysis results and graphical operations that can be used directly in a policy-making exercise.

BAYESIAN AGENT INFERENCE AND LEARNING

Agent inference is the ability of agents to make complex decisions, adapt to their environment, and learn from their decisions or decisions made by other agents. Although contextually, agent inference is not difficult to comprehend, capturing its symbolic and semantic formulation is quite a challenge for the researcher or analyst. Inferential modeling in agent-based systems is the “heart” of developing artificial intelligence and complex computational methodologies (Calmet et al. 1996; Edmonds et al. 2000; Gupta and Sinha 2000). Agent inference must accomplish a series of tasks, such as:

1. Provide the model with a mathematically sound representation of agent decisions, and learning;

2. Establish a sensible network of relationships or relational links not only among agents and their classes but also between causes and actions;

3. Provide the simulation environment with an adequate level of stochasticity and dynamic character so it is able to capture the magnitude and patterns of change that it is designed to replicate;

4. Bound the agents and their computational environment within the level of rationality and rules that natural, historical, and scientific observation and analysis dictate; and

5. Allow complex system properties, such as emergence, adaptivity, resilience and robustness, to be explored as integral parts of the dynamic simulation framework.

Bayesian inference is a special method of nonparametric, probabilistic, and stochastic evaluation of noisy data (Ahmed and Reid 2001; Pearl 1988). Probabilistic inference methods are extremely useful for cases or situations in which a high or deep level of uncertainty is embedded in the data or in which the decision maker is faced with incomplete observations to access the future. The power of the Bayesian nonparametric methods lies in the ability of the researcher to assess, quantify, and analyze qualitative and evaluative statements related to the way that decisions are made and the relationship between state-space and actions taken. Furthermore, Bayesian inference methods allow for complex hierarchical network scale development (Conte and Castelfranchi 1995; Eagly and Chaiken 1993; Stocker et al. 2002) and the elicitation of likelihood measures of indented decisions. Bayesian nonparametric assessment involves three fundamental concepts:

1. Identifying parameters for eliciting decisions (Muller et al. 2005; Zhu and Morgan 2004),

2. Evaluating the prior degree or probability of occurrence and developing empirical probability density distributions (Bohning and Schon, 2005; McIver and Friedl 2002; Sen 1981), and

3. Estimating and learning conditional posterior probabilities for actions performed given a constructed and estimated network structure (Hall and Yatchew 2005; Stewart 2005; Tiku et al. 1986).

With regard to *Bayesian agent inference*, Bayesian inference provides an alternative, probabilogic, nonparametric estimation of an agent's beliefs, desires, and intentions (BDI), especially when the modeling design and architecture favor the character of BDI agents. The BDI approach to agent-based modeling presents a highly robust and theory-grounded methodology to address agent intelligence and elaborative human-like agent character (Feng et al. 2003; Norling and Sonenberg 2004; Rao and Georgeff 1995). Within this context of agent inference, Bayesian methods or nonparametric decision estimation include BBNs (complex Bayesian belief networks of agent beliefs and intentions) or BDNs (complex Bayesian decision networks of agent decisions for actions), or both. BBNs and BDNs combine the mathematical parameterization methodology of Bayesian inference with the intelligent and learning character of multi-agent systems (Alexandridis et al. 2004; Alexandridis et al. 2006; Alexandridis and Pijanowski 2006; De Cooman and Zaffalon 2004; Korb and Nicholson 2004; Neapolitan 2004).

While the volume of literature on the Bayesian methods of inference is quite extensive, the utilization of these nonparametric methods in systems of spatial complexity and environmental modeling applications is quite limited. Some of the main reasons for this inconsistency are that (a) Bayesian artificial intelligence is a relatively new field of research, and the transition from theory to application and problem-oriented research has not been realized fully; (b) analysis of spatially complex structures requires multi-disciplinary applications and research skills, a fact that slows up the development and progression of such modeling research; (c) spatially complex agent interactions emerge at a magnitude of scales, both spatial and temporal; thus, estimating modeling parameters involves arrays or matrices of interactions instead of single parameter estimation (the latter point renders estimation properties a mathematical and statistical challenge); (d) coping with uncertainty and incomplete information, while commonly encountered in the real world, requires a departure from traditional statistical theory and comprehension of the fact that systems might display unpredictability and instability of patterns under such conditions.

In the MABEL model (Alexandridis et al. 2004; Alexandridis et al. 2006; Alexandridis and Pijanowski 2006; Lei et al. 2005), such an architecture is employed to simulate agent intentions for decisions on land use change. Four main components are essential for such decisions: (1) the state-space (agents' environment), (2) a transition modeling mechanism (mapping state-space to actions), (3) the agents' expectations for utility of his actions (expected utility elicitation), and (4) the expected rewards that the agents anticipate for their intended actions. In addition, agents face evidence entering their perceptual environment (in the form of prior decisions, or decisions made by neighboring agents), and a learning mechanism combines their prior beliefs with the new evidence as they enter their systemic sensoring mechanism. Combining the agent's BBN intentional learning mechanism with the agent's BDN action learning mechanism enhances agent and simulation behavior over space and time. A schematic representation of such a coupling is shown in Figure 1.

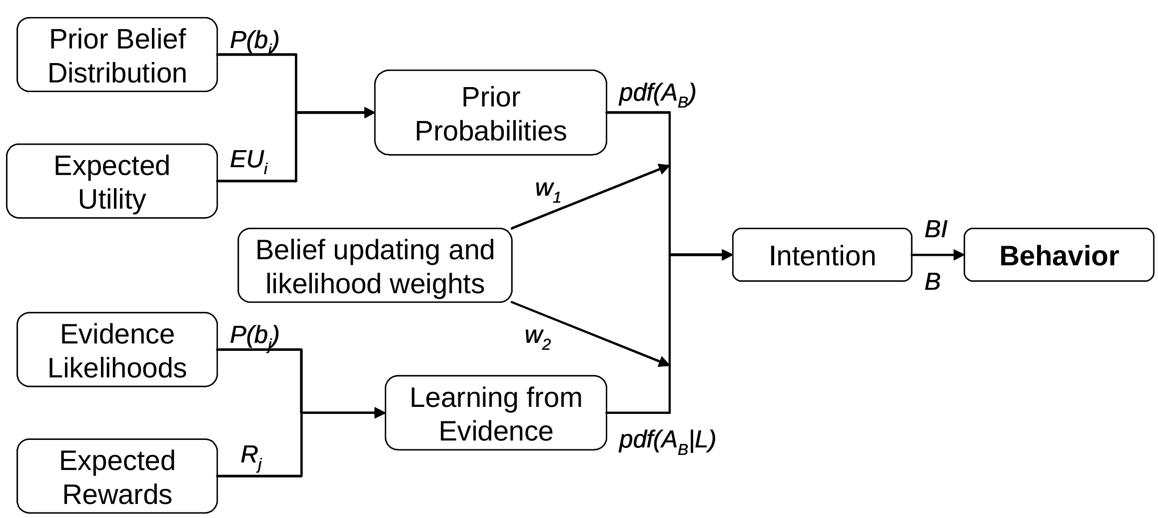

FIGURE 1 Elaborative mechanism for agent learning in the MABEL model

Figure 1 illustrates the process of decision making in the MABEL model in terms of the Bayesian underlying structure. Each agent *i* is faced with a prior belief distribution, denoted as *P(bi)*. This belief distribution can be conceptualized as a three-dimensional array with dimensions $n \times m \times k$, where

$$\{n,m,k\} = \{BeliefStates, BeliefNodes, Actions\}. \tag{1}$$

In a backward propagation, the previous statement implies that for each potential action, multiple nodes (variables) exist, and for each node, multiple states (probabilistic) exist. The multi-dimensional array of the prior belief distribution is actually a complex BBN representing this prior distribution structure.

Each element of the multi-dimensional array of the prior belief distribution has an expected utility value (EU — the expectation that an agent has if a given combination of state, node, and action was to be undertaken). Combining the prior belief structure with the agent's utility expectations provides us with a probabilistic distribution measure of the agent's expected next state. This probability distribution of expected utility is what an agent faces without any new information entering the inference system. In complex reality, agents, as well as decision makers, do obtain dynamically new information, learn from decisions made in previous time steps, and face potential rewards for their actions. This process is often called reinforced learning. The probabilistic structure for the evidential mechanism is a two-dimensional array with dimensions $n' \times m'$, where

$$\{n' \times m'\} = \{LikelihoodStates, LikelihoodNodes\}. \tag{2}$$

Similarly, for each of the nodes of the network and their associated states, there is a probability (likelihood) that evidence or experiences would indicate that they would change in the near future. Mapping the likelihood probability distribution to the expected rewards (gains or losses) that these changes entail for the agents provides us with a conditional probability distribution for intended actions, given the evidence likelihoods.

The Bayesian learning algorithms are designed to estimate the optimal weights with which the intentions for the next time step of each agent are calculated. In other words, they are designed to estimate the strength and degree to which new evidence entering the inferential system of the simulation alters the intentions of agents for action. In the MABEL model, this process is performed by using the expectation maximization (EM) algorithm (Beal et al. 2003; Bohning and Schon 2005; Dellaert 2002; Friedman 1998; Hutter and Zaffalon 2005). The EM algorithm utilizes an iterative and dynamic maximum likelihood estimation technique in order to approximate the posterior learning distribution for agents' actions.

TYPES OF AGENT INFERENCE

In the context of the simulation design and modeling procedure, we can identify three variations of agent inference:

1. *Problem-driven inference.* Agent properties and decisions vary across problem and modeling domains. The emergent ability of the agents to achieve complex problem solving and optimization is essential.

2. *Model-driven inference.* Agent behavior drives the evolution of simulation ensembles and the problem domains for applications. The emergent ability of the agents for capturing complex patterns and processes is essential.

3. *Knowledge-driven inference.* Agent knowledge-base and learning capabilities help in identifying problem and application domains. The emergent ability of the agents for complex learning and adaptation is essential.

An empirical assessment of the essential inference mechanisms is shown in Figure 2. We can consider three group categories that isolate particular characteristics within a common pool of identifiers or modes. A mixture of these modes can help us characterize the type and qualitative characteristics of the agent type inference. The first group is the driving force for the agent inference and can be driven by real-world processes, research questions, or policy-related drivers. The second group emphasizes the characteristics of the modeling process, such as the observational or state-space characteristics, the hypotheses or assumption-bases that the model utilizes, and the scenarios that the model assesses. The third group is composed of the elements of the knowledge-base or learning components of the model. Knowledge bases can be characterized by the purpose they are designed to test (i.e., databases designed for uncertainty or error-testing, testing specific hypotheses, or testing developed scenarios for simulations).

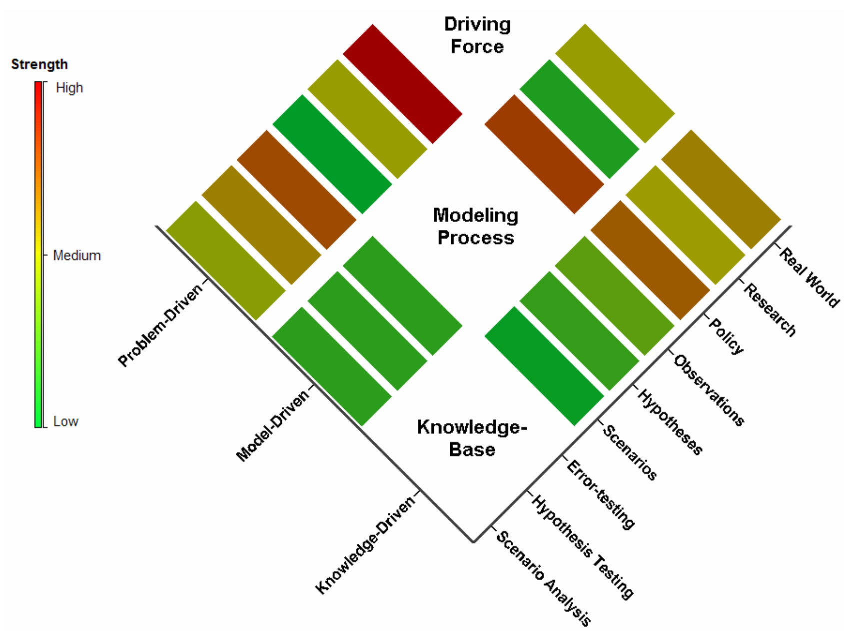

FIGURE 2 Process versus mode in agent type inference (The strength of the relationship indicates an empirical assessment.)

Within the above described framework, problem-driven agent inference mechanisms depend mainly on the problem-specific driving forces, emphasize the modeling of observations and hypotheses, and utilize knowledge bases designed mainly for error and hypothesis testing. An example of such an inference mechanism is finding optimal environmental policies from a pool of available solutions within a finite pool of available resources.

Model-driven agent inference mechanisms depend mainly on the focus of the modeling processes used and emphasize the modeling of policy and real-world applications. The value of these types of inference mechanisms is not particularly high for theoretical or research-driven applications, as they aim to explore the emergence or generative dynamics of the modeling and simulation mechanisms. An example of such an inference mechanism is simulating the emergence of social and economic phenomena within a finite population or with the use of simple rules.

Knowledge-based inference mechanisms depend mainly on the availability of knowledge and observational learning techniques and emphasize mainly policy and real-world types of applications. An example of such an inference mechanism is predicting changes in land use given a set of policy directives or scenarios embedded in the real landscapes.

While these types of agent inference are not mutually exclusive and can be present simultaneously in some combination in our modeling enterprises, the above classification can help researchers understand the strengths and weaknesses of the models, tools, and methods employed. It can also help in assessing the validity and appropriateness of databases and knowledge bases, scenarios and policies, applications to be simulated, or research questions that can be answered within a given study.

THE NEED FOR MODULARITY

Designing and implementing spatially explicit agent-based modeling enterprises require a relatively high level of expertise, such as an understanding of the complexity of agent behavior and agent interactions (Baroni et al. 2005; Parker et al. 2003). Often the ability to communicate and diffuse the assumptions, mechanisms, and results of the simulations to the various stakeholders, policy makers, and end-user communities encounters a number of difficulties (Carley 2002). Additional difficulties emerge when the research need for abstract symbolic representation (e.g., mathematical, statistical, computer-language-dependent) conflicts with the need for simplicity and transparency for comprehension and cognition (Burley 2004; Cox 2005). Also, as seen in the previous section, such modeling enterprises are often domain specific and thus not always “one size fits all.”

The degree of robustness of the agent-based modeling enterprises that is achieved thus often depends significantly on the subjective researcher’s skills. The researcher needs to anticipate outcomes or problems in which stakeholders and policy makers have a potential interest. Specific databases need to be constructed and data collected in advance. Models and simulation experiments need to be calibrated and assessed, and full-models can be very computationally intensive (a fact that affects the ability to replicate or assess the validity of the simulation results by the end-user community).

For these reasons, modularity in agent-based modeling mechanisms is a desired system and modeling property. Modularity in agent-based simulations:

- Provides enhanced visualization capabilities beyond the input-output (I/O) process of the modeling mechanism;

- Enables comprehensibility of patterns and processes emerging at the problem-, model-, and knowledge-driven levels;

- Provides decision and policy-makers with the ability to control assumptions and inference mechanisms and minimizes the researcher's subjective bias in the simulation and analysis process; and

- Allows for the capability to attract and collect expert judgments, scenarios, and hypotheses that emerge from the "ground up."

Modular components and agent-based mechanisms can be stand-alone. End-users do not need to run full-model simulations (with their limitations in comprehensibility and transparency). The expertise required for full comprehension of the agent-based model can be modularized as well, by focusing exclusively on inference mechanisms, learning, and problem solving as separate modular components of the modeling enterprise. Finally, the modularization of agent-based model components provides easy access to calibration, assessment and scenario development techniques, thus reducing the perceived complexity of the mechanism from the end-user's perspective.

MODULAR DECISION SUPPORT TOOLS FOR THE MABEL MODEL

This section provides a descriptive demonstration of three examples of inference tools for the MABEL model: decision net inference tool (problem-driven), agent spillover effects tool (model-driven), and MABEL scenario generator tool (knowledge-driven). All three components operate as stand-alone user interfaces and provide added modularity and comprehensibility for the full MABEL model simulations.

MABEL Full-model versus Stand-alone Inference Tools

Modular inference capabilities exist within the architectural framework of the MABEL model (Lei et al. 2005). Specifically, the MABEL model embeds the Netica C++ API (Norsys 2005) within its architecture. The decision-theoretic inference and Markov decision-making mechanisms in MABEL utilize this framework to perform diagnostic and learning inference for utility acquisition and optimization tasks (Alexandridis et al. 2004; Alexandridis and Pijanowski 2006). Nevertheless, visualizing these mechanisms and their dynamics is not possible without the use of new tools.

These added visualization capabilities are achieved via the development of new inference tools within the MABEL framework. These tools are developed in the Visual Basic.NET developing framework (Microsoft 2003) and utilize the Netica VB API (Norsys 2005) and the Microsoft *Office.Interop* components of the .NET framework. Each of these tools compiles and utilizes existing or modified MABEL decision networks and decision mechanisms and can re-compile revised networks back to the MABEL model for simulation runs.

The *Interop.Netica* component within the .NET framework provides high usability of objects, classes, functions, and class members for use within the development environment (Figure 3). Each of these methods, when called within the Netica application framework, allows for a comprehensive visualization of the inference mechanism and the BBN or BDN.

Decision Net Inference Tool Example

The *decision net inference tool* is an example of a problem-driven inference mechanism that provides a comprehensive and veridical end-user interface for BDN inference. It provides a robust query of agent BDN enhanced with the full network visualization capabilities of the Netica.interop application (Norsys 2005) through the VB API .NET framework. Any BBN or BDN can be loaded and queried through the main control panel of the user's interface (Figure 4A).

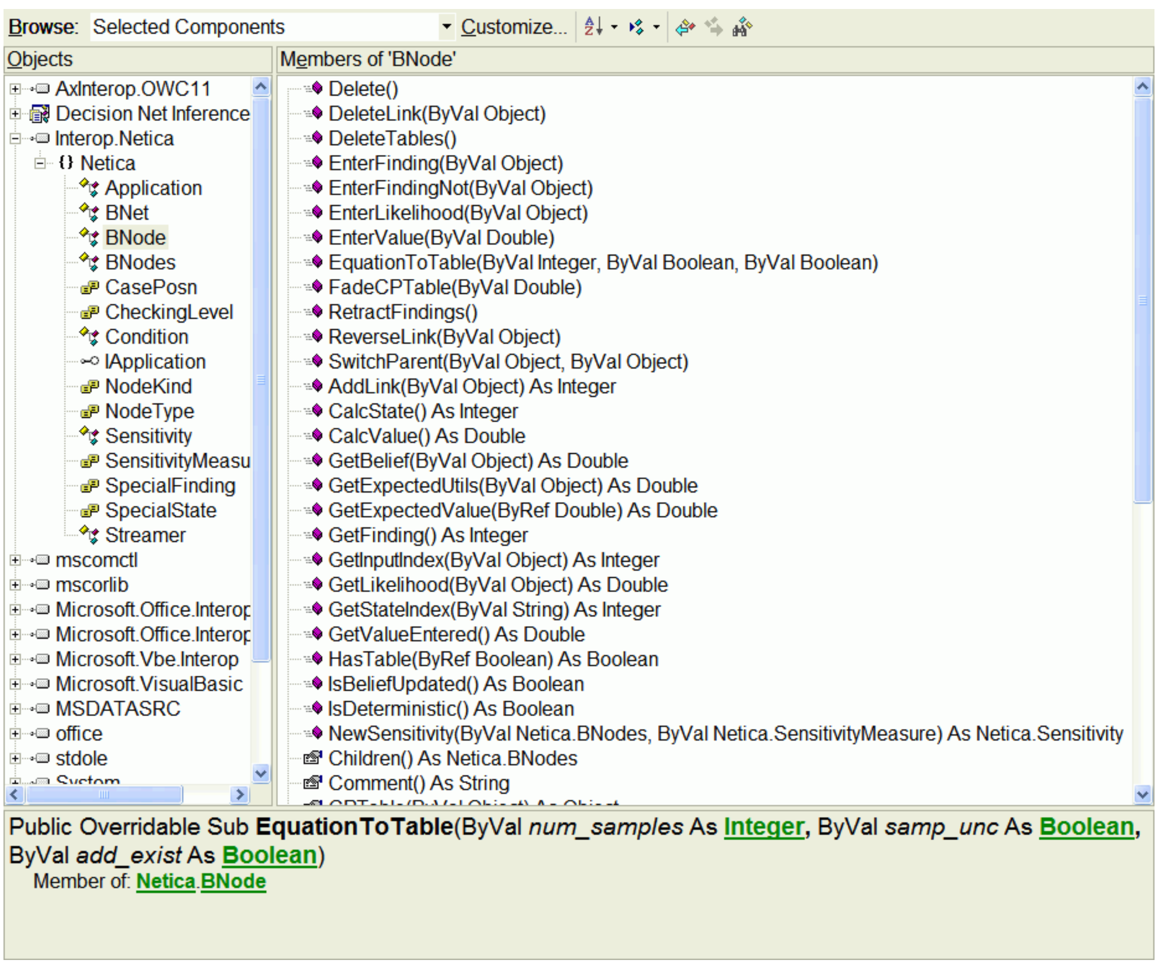

FIGURE 3 Functionality of the Interop.Netica component in the .NET framework: Objects, classes, members

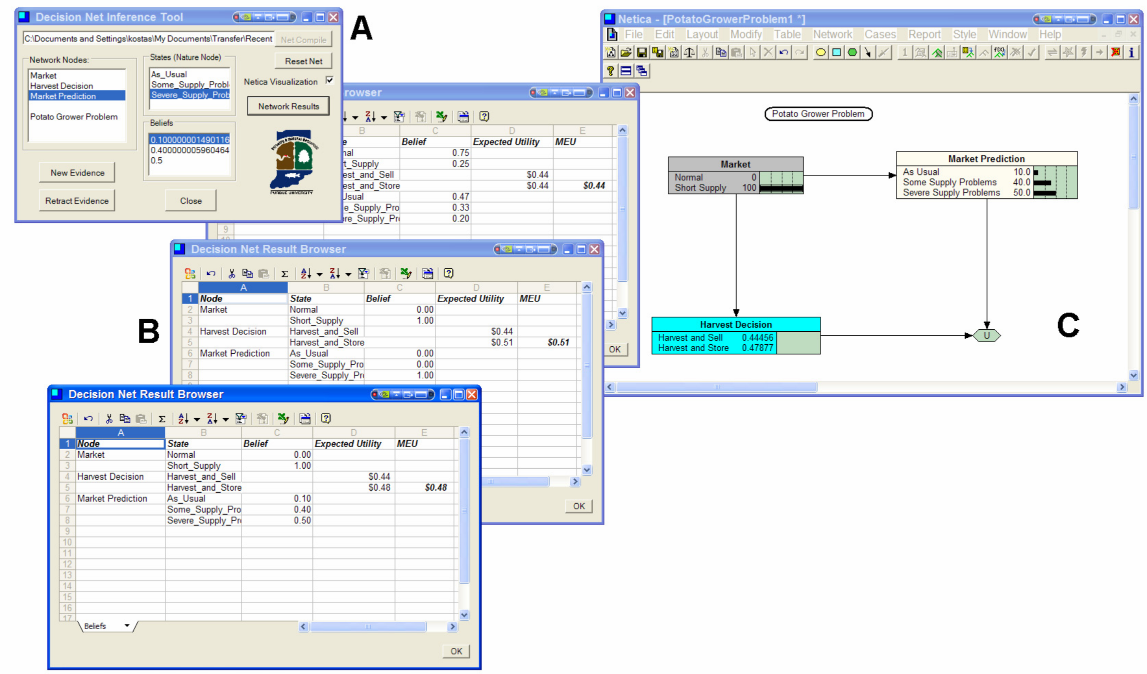

FIGURE 4 User interface of the decision net inference tool for MABEL inference (Figure 4A shows the main control panel, 4B shows different output windows using the office.interop interface, and 4C shows the Netica decision network visualization called from the main control panel.)

For this example, a real-world application of a decision network is provided: the *potato grower problem* described by Hardaker et al. (1998). In short, the decision problem copes with a farmer's uncertainty about a decision whether to harvest and sell or harvest and store his potato yield. The farmer faces two sources of uncertainty: market conditions (i.e., normal versus short supply) and the acquisition of market prediction estimates (i.e., estimates of the supply problems in the market). In the latter, a cost is involved for acquisition of a market prediction (e.g., a private company providing estimates), and each of the harvest decisions results in a monetary utility acquisition. While the decision network layout looks simple (Figure 4C), elicitation of the prior and posterior probabilities of the network is somewhat complex (Hardaker et al. 1998).

By using the main control panel (Figure 4A), a user can enter inference evidence into the decision network or test inferential assumptions about his knowledge about the conditions of the market. Then estimates of the posterior probability distribution and the potential monetary value of the decision can be assessed via a streamer process that outputs the network results into the office.interop spreadsheet component (Figure 4B). Multiple evidence-assessment pairs can be performed, and multiple results can be obtained and analyzed simultaneously.

The decision net inference tool provides a useful interface that gives end-users the ability to test and analyze the implications of their assumption-based and inferential decision-making capabilities. The tool can also be used for training/learning inference at the decision-making level. Finally, network evidence and experiences by users can be collected and saved as case studies for agent training and learning by using the empirical user's assessment and utilized further in realistic MABEL simulations.

Agent Spillover Effects Tool Example

The *agent spillover effects tool* is an example of a model-driven inference mechanism that allows simple rules of spatial inference to be incorporated into the agent-based simulation framework for MABEL. It focuses on aiding understanding of spatial spillover effects of client simulations in MABEL and of the immigration-emigration effects of land use change at local and regional spatial scales. Specifically, when an ensemble of local or regional spatial simulations is performed, traditional agent inference does not account for in- and out-migration flows of agents into the spatial area of the simulation. Often an endogenous assumption is being made at the simulation design level either to ignore the spatial and socioeconomic effects of connectivity across spatial scales or to assume a nonspatial, deterministic flow exchange. The reality of landscape and spatially explicit socioeconomic histories renders such assumptions naïve at best. Thus, there is a need to develop a more knowledge-driven and spatially explicit mechanism for determining the effects and the degree and strength of these effects within a simulation ensemble framework in MABEL.

The agent spillover effects tool utilizes an eight-nearest-network distributional approach to spatial configuration of a landscape (Figure 5). Assuming that the simulation area represents the central community of the framework, an ensemble of eight neighboring simulations can be identified (Figure 5A). A stochastic decision network based on hypothetical normal distributions can then been assessed via the Netica.interop application (Norsys 2005). A series of evidence- or assumption-based inferences can be performed to determine the mean value, standard deviation, and skew of the joint distributions of the eight neighbors. The elicitation of the spatial distributional effects depends on the categorical (Likert) density estimation of the central simulation area, the categorical (Likert) joint density estimation of the neighbor spatial area, and the spatial directional effects of the joined spillover distributions.

By providing knowledge or evidential information about these three input rules and performing a diagnostic inference of the network, a percentage estimation of the degree of the spillover effects that can be transitioned across neighboring simulations can be done (upper window of part A in Figure 5). In many simulation cases, we might know in advance the number of agents exceeding capacity in the current simulation area (i.e., from new agent emergence from a parcelization algorithm in MABEL), in which case a quantitative assessment of the number of agents to be transitioned to neighboring townships can be assessed (lower window of part A in Figure 5).

By performing sequential assessments for a finite set of a simulation area ensembles, we can identify a stochastic assessment of the in- and out-migration dynamics of the MABEL framework and establish rules for cross- and within-scales modeling assessments. Model calibration and dynamics related to different or alternative spatially explicit hypotheses (e.g., number of agents, theoretical distributions) can be performed. In addition, end-users and policy makers can use their experience and evidential knowledge to train and quantitatively elicit better network distributions (i.e., training for distribution statistical moments) or determine other non-normal distributions that can be present in our landscapes. Finally, the tool can be used to understand residual dynamics of land use change that are not directly related to land use changes within the simulation area but are transfer effects from spatial changes within wider scales.

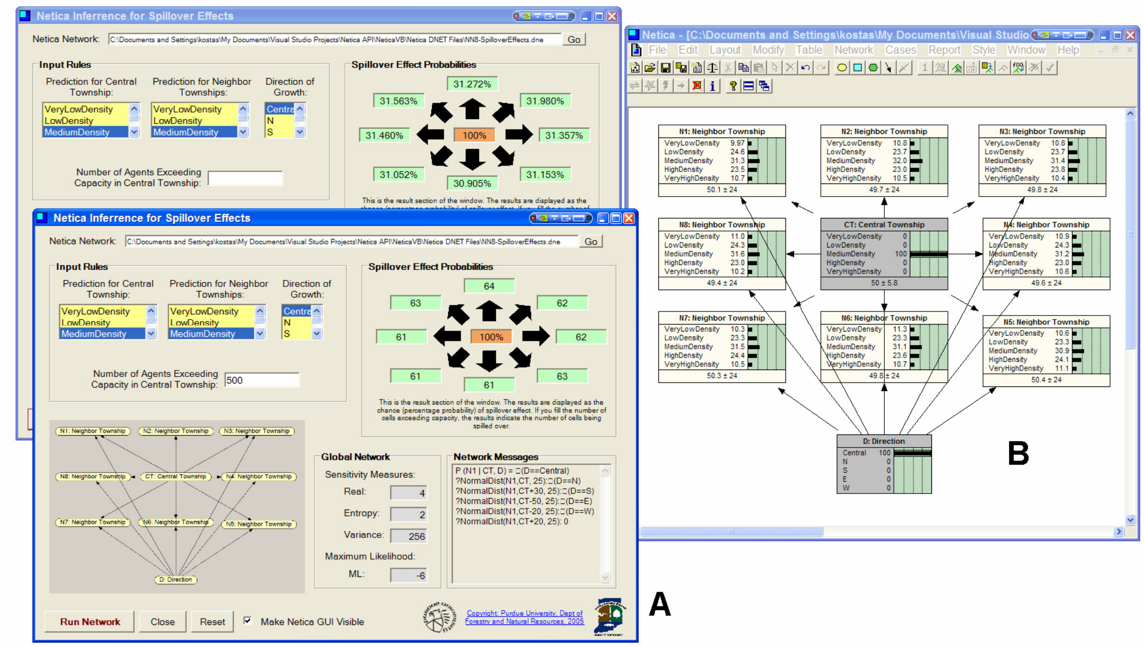

FIGURE 5 User interface of the agent spillover effects tool for MABEL inference (Figure 5A shows instantiations of the main control panel, and 5B shows the Netica decision network visualization called from the main control panel.)

MABEL Scenario Generator Tool Example

The *MABEL scenario generator tool* provides an example of a knowledge-driven inference mechanism. It uses a mechanism for generating knowledge-driven scenarios for modeling assessments and uses the e*xpectation maximization* (EM) algorithm (Beal 2003; Friedman 1998; Moon 1996) for agent training and learning. In the specific example provided in Figure 6, the tool provides visualization and inference for alternative scenarios related to farmer-class agents and the capability of performing EM learning for agent classes on the basis of existing empirical evidence from real farm decisions observed in the landscape.

The main control panel of the user's interface provides the ability for performing simple evidence elicitation of the scenario decision network (Figure 6A) or performing a learning simulation over several time-steps of agent learning via the EM algorithm, given evidence (Figure 6B). The elicited results, including the simulation learning process visualization, can be called and viewed through the Netica.interop application window (Figure 6C). The lower part of the main control panel provides a real-time office.interop spreadsheet streamer that monitors beliefs, likelihoods, and simulation step network results.

Results obtained from EM training and learning via the MABEL scenario generator tool can provide useful insight on the agent-learning mechanism in MABEL. Agent- learning results and simulation patterns like the one provided in Figure 7 can be easily assessed and evaluated by using the tool's user interface. Additional scenarios and model elicitation can be done with the use of user's knowledge and evidential information. Finally, the tool can be used for training and comprehension of the dynamics present in our real-world decision dynamics.

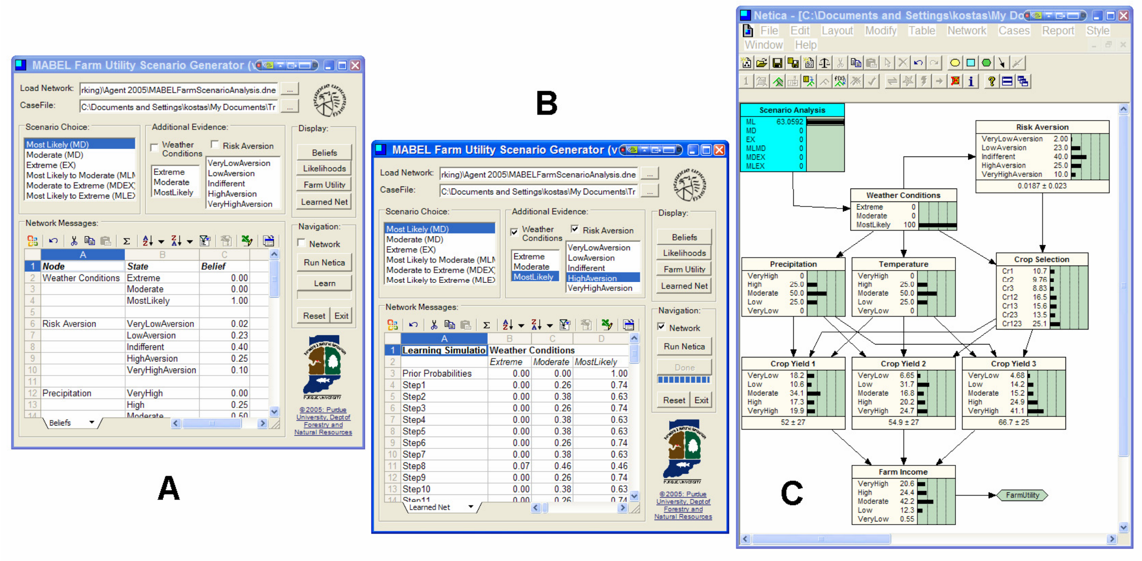

FIGURE 6 User interface of the farm utility scenario generator for MABEL inference (Figure 6A shows the main control panel for simple evidence elicitation, 6B shows the main control panel for network learning simulations, and 6C shows the Netica decision network visualization called from the main control panel.)

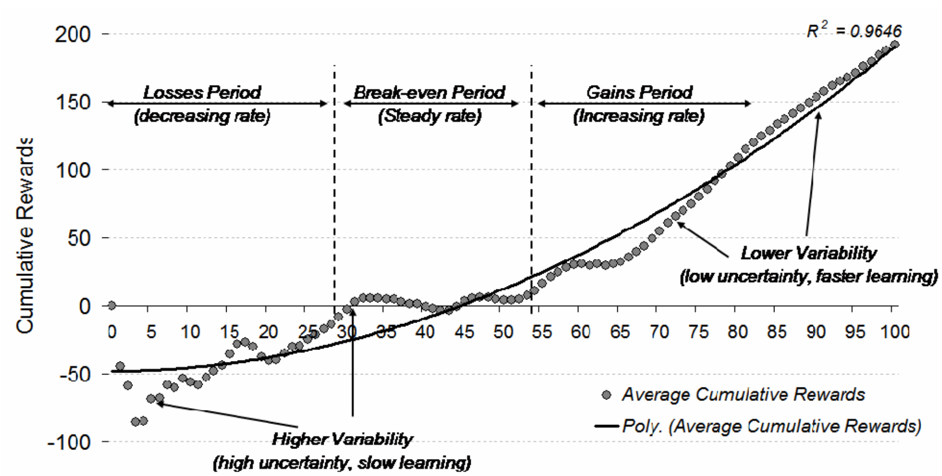

FIGURE 7 Example of the inference learning results obtained for MABEL learning simulations by using the EM algorithm training for 100 agent time-steps

VISUALIZATION, COMPREHENSION, AND POLICY MAKING REVISITED

The examples provided in the above sections, combined with the theoretical arguments and discussion preceding them, allow us to assess a set of wider implications for a *visualization, comprehension and policy-making (VCPM)* framework. This set can be expressed as a sequence of propositions or lessons learned, as follows:

1. Visualization and comprehension cannot be achieved through black-box-type complex agent-based models. They require a clear and informational-rich understanding of the mechanisms, dynamics, and patterns of complexity present in our models.

2. Agent-based models and simulations are the most appropriate research and methodological approaches to building communities and community understanding from the ground up. They have the power and ability to enhance decision-maker and stakeholder understanding of complex spatial mechanisms, drivers, and processes of change.

3. Enhancing collaboration and coordination at the stakeholder and decision-maker level can lead to more accurate predictions and more influential decisions and reduce uncertainty and risk in decision making. It can also boost dialogue and communication across and within communities for achieving sustainable and attainable alternative futures.

4. Combining quantitative and nonparametric assessment methods with qualitative and scenario assessments, often provided interactively, can enhance the value of information in our models and modeling enterprises. It can also shed some light on the complex cross-scale and cross-component mechanisms present in our simulation domains.

5. Reconciling the needs for transparency and veridicality in modeling does not have to result in abstract and often inaccurate modeling endeavors but can instead result in informational content-rich enterprises.

6. Developing simulation and inference tools that can serve also as training, learning, and educational exercises builds future community capabilities for decision and policy making that can have significant impact.

Beyond the specific tools and modeling enterprises examined in this paper, a whole array of quantitative and qualitative assessments for network inference mechanisms can be employed to enhance our multi-layered representation of reality (Figure 8).

Additional qualitative methods for probabilistic and stochastic inference — such as role-playing games, participatory rural assessment techniques, interviews and surveys, and qualitative scenario assessment methods — are examples of techniques that can be used for eliciting BBNs and BDNs in the MABEL framework. Combined with traditional quantitative assessment methods and knowledge bases, they present a powerful mechanism of inference in multi-agent simulation systems.

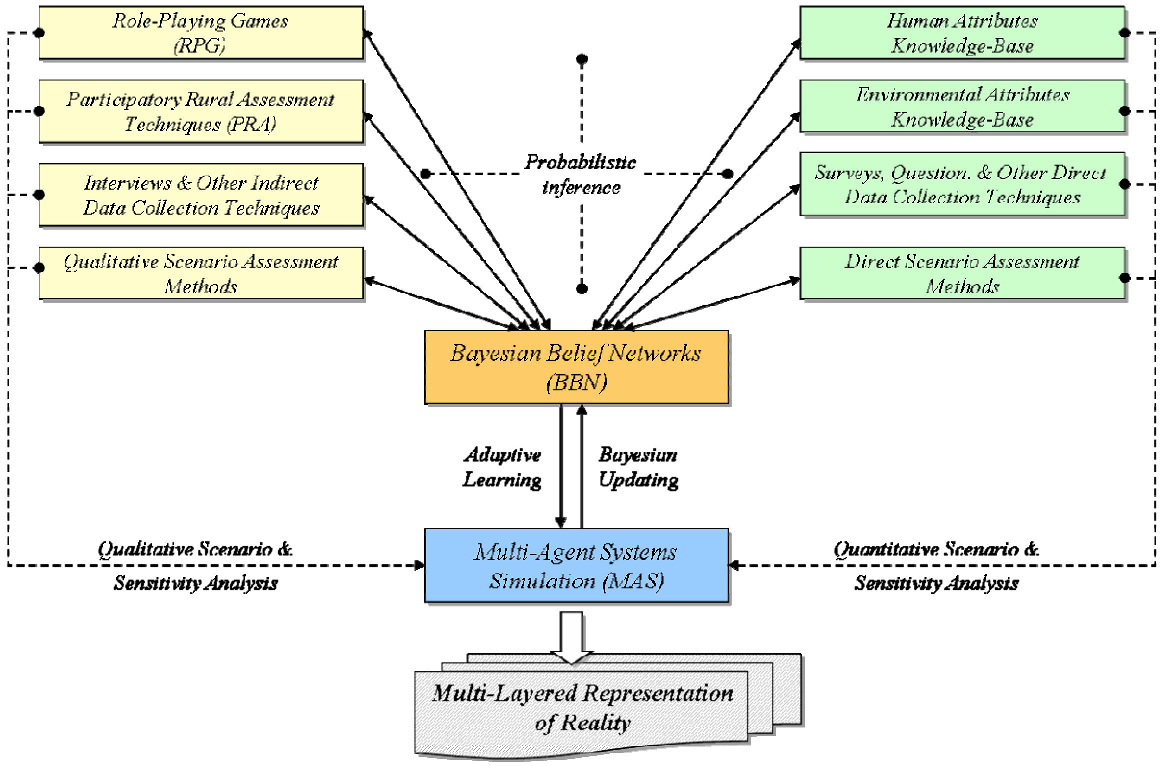

FIGURE 8 Schematic model for Bayesian artificial intelligence modeling: Quantitative and qualitative elicitation and probabilistic inference in Bayesian belief networks

BIBLIOGRAPHY

Ahmed, S.E., and N. Reid. 2001. *Empirical Bayes and Likelihood Inference*. Lecture notes in statistics (Springer-Verlag). New York: Springer.

Alexandridis, K., B.C. Pijanowski, and Z. Lei. 2004. “The use of robust and efficient methodologies in agent-based modeling: Case studies using repeated measures and behavioral components in the MABEL simulation model.” In C.M. Macal, D. Sallach, and M.J. North (editors), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity, and Emergence*, ANL-DIS/05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Alexandridis, K., B.C. Pijanowski, and Z. Lei. 2006. “Assessing multi-agent parcelization performance in the MABEL simulation model using Monte-Carlo replication experiments.” *Environment & Planning B: Planning and Design*, Special Issue 33 on Agent-based Modeling (in press).

Alexandridis, K.T., and B.C. Pijanowski. 2006. “A Multi Agent-based Behavioral Economic Landscape (MABEL) model of land use change” (in revision).

Baroni, P., M. Giacomin, and G. Guida. 2005. "Self-stabilizing defeat status computation: Dealing with conflict management in multi-agent systems." *Artificial Intelligence 165*(2):187–259.

Beal, M.J. 2003. *Variational Algorithms for Approximate Bayesian Inference*. Ph.D. Dissertation. London, UK: The Gatsby Computational Neuroscience Unit, University of London.

Beal, M.J., Z. Ghahramani, J.M. Bernardo, M.J. Bayarri, A.P. Dawid, J.O. Berger, D. Heckerman, A.F.M. Smith, and M. West. 2003. "The variational Bayesian EM algorithm for incomplete data: With application to scoring graphical model structures." In *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, pp. 453–464. Oxford, UK: Oxford University Press.

Bohning, D., and D. Schon. 2005. "Nonparametric maximum likelihood estimation of population size based on the counting distribution." *Journal of the Royal Statistical Society: Series C (Applied Statistics) 54*(4):721–737.

Brown, D.G., S.E. Page, R. Riolo, and W. Rand. 2002. *Modeling the Effects of Greenbelts at the Urban-Rural Fringe*, June 24–27, Lugano, Switzerland.

Burley, J. 2004. "The restoration of research." *Forest Ecology and Management 201*(1):83–88.

Calmet, J., J.A. Campbell, and J. Pfalzgraf. 1996. *Artificial Intelligence and Symbolic Mathematical Computation: International Conference, AISMC-3, Steyr, Austria, September 23–25, 1996, Proceedings*, Lecture Notes in Computer Science, 1138. Berlin; New York: Springer.

Carley, K.M. 2002. "Simulating society: The tension between transparency and veridicality." In C. Macal and D. Sallach (editors), *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, ANL/DIS-03-1, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 11–12.

Conte, R., and C. Castelfranchi. 1995. *Cognitive and Social Action*. London: UCL Press.

Cox, M.T. 2005. "Metacognition in computation: A selected research review." *Artificial Intelligence 169*(2):104–141.

De Cooman, G., and M. Zaffalon. 2004. "Updating beliefs with incomplete observations." *Artificial Intelligence 159*(1–2):75–125.

Dellaert, F. 2002. *The Expectation Maximization Algorithm*, Technical Report GIT-GVU-02-20. Atlanta, GA: College of Computing, Georgia Institute of Technology.

Drennan, M. 2005. "The human science of simulation: A robust hermeneutics for artificial societies." *Journal of Artificial Societies and Social Simulation* 8(1); available at http://jasss.soc.surrey.ac.uk/8/1/3.html.

Eagly, A.H., and S. Chaiken. 1993. *The Psychology of Attitudes*. Fort Worth, TX: Harcourt Brace College Publishers (HBJ).

Edmonds, B., S. Moss, and P. Davidsson. 2000. "The use of models: Making MABS actually work." In *Multi Agent Based Simulation,* pp. 15–32. Springer-Verlag.

Ehrlich, P.R., and D. Kennedy. 2005. "Sustainability: Millennium assessment of human behavior." *Science 309*(5734):562–563.

Feng, S., L. Da Xu, C. Tang, and S. Yang. 2003. "An intelligent agent with layered architecture for operating systems resource management." *Expert Systems 20*(4):171–178.

Foley, J.A., R. DeFries, G.P. Asner, C. Barford, G. Bonan, S.R. Carpenter, F.S. Chapin, M.T. Coe, G.C. Daily, H.K. Gibbs, J.H. Helkowski, T. Holloway, E.A. Howard, C.J. Kucharik, C. Monfreda, J.A. Patz, I.C. Prentice, N. Ramankutty, and P.K. Snyder. 2005. "Global consequences of land use." *Science 309*(5734):570–574.

Friedman, N. 1998. "The Bayesian structural EM algorithm." In *UAI'98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence,* July 24–26, University of Wisconsin Business School, Madison, WI.

Gupta, M.M., and N.K. Sinha. 2000. *Soft Computing and Intelligent Systems: Theory and Applications*. San Diego, CA: Academic.

Hall, P., and A. Yatchew. 2005. "Unified approach to testing functional hypotheses in semiparametric contexts." *Journal of Econometrics 127*(2):225–252.

Hammersley, M. 2005. "Should social science be critical?" *Philosophy of the Social Sciences 35*(2):175–195.

Hardaker, J.B., R.B.M. Huirne, and J.R. Anderson. 1998. *Coping with Risk in Agriculture*. Wallingford, UK: CAB International.

Heemskerk, M., K. Wilson, and M. Pavao-Zuckerman. 2003. "Conceptual models as tools for communication across disciplines." *Ecology and Society 7*(3); available at http://www.consecol.org/vol7/iss3/art8/.

Hutter, M., and M. Zaffalon. 2005. "Distribution of mutual information from complete and incomplete data." *Computational Statistics & Data Analysis 48*(3):633–657.

Korb, K.B., and A.E. Nicholson. 2004. *Bayesian Artificial Intelligence*. Series in *Computer Science and Data Analysis*. London, UK: Chapman & Hall/CRC.

Lei, Z., B.C. Pijanowski, and K.T. Alexandridis. 2005. "Distributed modeling architecture of a Multi Agent-based Behavioral Economic Landscape (MABEL) model." *Simulation: Transactions of the Society for Modeling & Simulation International 81*(7):503–515.

Ma, T., and Y. Nakamori. 2005. "Agent-based modeling on technological innovation as an evolutionary process." *European Journal of Operational Research 166*(3):741–755.

McIntosh, B.S., P. Jeffrey, M. Lemon, and N. Winder. 2005. "On the design of computer-based models for integrated environmental science." *Environmental Management 35*(6):741–752.

McIver, D.K., and M.A. Friedl. 2002. "Using prior probabilities in decision-tree classification of remotely sensed data." *Remote Sensing of Environment 81*(2–3):253–261.

Microsoft. 2003. *Microsoft Visual Studio.NET.*

Moon, T.K. 1996. "The expectation-maximization algorithm." *IEEE Signal Processing Magazine 13*(6):47–60.

Muller, P., G.L. Rosner, M.D. Iorio, and S. MacEachern. 2005. "A nonparametric Bayesian model for inference in related longitudinal studies." *Journal of the Royal Statistical Society: Series C (Applied Statistics) 54*(3):611–626.

Neapolitan, R.E. 2004. *Learning Bayesian Networks*. Upper Saddle River, NJ: Pearson Prentice Hall.

Norling, E., and L. Sonenberg. 2004. "Creating interactive characters with BDI agents." Paper presented at the Australian Workshop on Interactive Entertainment, Sidney, Australia.

Norsys. 2005. *Netica: Advanced Bayesian Belief Network and Influence Diagram Software (Version 2.17).* Vancouver, Canada: Norsys Software Corp.; available at www.norsys.com.

Parker, D.C., S.M. Manson, M.A. Janssen, M.J. Hoffmann, and P. Deadman. 2003. "Multi-agent systems for the simulation of land-use and land-cover change: A review." *Annals of the Association of American Geographers 93*(2):314–337.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann Publishers.

Rao, A.S., and M.P. Georgeff. 1995. "BDI agents: From theory to practice." In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, San Francisco, CA.

Sen, P.K. 1981. *Sequential Nonparametrics: Invariance Principles and Statistical Inference*. New York: Wiley.

Stewart, M.B. 2005. "A comparison of semiparametric estimators for the ordered response model." *Computational Statistics & Data Analysis 49*(2):555–573.

Stocker, R., D. Cornforth, and T.R.J. Bossomaier. 2002. "Network structures and agreement in social network simulations." *Journal of Artificial Societies and Social Simulation* 5(4); available at http://jasss.soc.surrey.ac.uk/5/4/3.html.

Tiku, M.L., W.Y. Tan, and N. Balakrishnan. 1986. *Robust Inference*. New York: M. Dekker.

Van Vuuren, D.P., and L.F. Bouwman. 2005. "Exploring past and future changes in the ecological footprint for world regions." *Ecological Economics 52*(1):43–62.

Verburg, P.H., P.P. Schot, M.J. Dijst, and A. Veldkamp. 2004. "Land use change modelling: Current practice and research priorities." *GeoJournal 61*(4):309–324.

Zhu, J., and G.D. Morgan. 2004. "A nonparametric procedure for analyzing repeated measures of spatially correlated data." *Environmental and Ecological Statistics 11*(4):431–443.

SOCIAL COGNITION IN COMPLEX TEAM NETWORKS

S.M.D. SEAVER, R.D. MALMGREN, M. SALES-PARDO, and
L.A. NUNES AMARAL,* Department of Chemical and Biological Engineering,
Northwestern University, Evanston, IL
A.A. MOREIRA, Department de Física, Universidade Federal do Ceará, Ceará, Brazil
D. DIERMEIER, Managerial Economics and Decision Sciences,
Kellogg School of Management, Northwestern University, Evanston, IL

ABSTRACT

We use the density classification task to model social cognition in complex networks. Agents are embedded in a team network (Guimerá et al. 2005) and use simple social learning heuristics to adjust their current state. We investigate various agent-specific biases and system performance measures. We show that team networks perform less well than small-world networks.

Keywords: Social cognition, complex networks, agent-based modeling

KNOWLEDGE NETWORKS

Innovation and problem solving are almost always the consequence of collaborations in teams or large groups. Here we model creative interaction in organizations as a complex network (e.g., Guimerá et al. 2005, Moreira et al. 2004, Uzzi and Spiro 2005) and systematically investigate whether and how the structure of networks affects an organization's ability to solve problems and find innovative solutions.

Recent research has begun to uncover the factors that lead to successful team networks. An important example is an empirical study of the Broadway industry (Uzzi and Spiro 2005). To conduct their analysis, Uzzi and Spiro used a newly created data set on creative networks in Broadway musicals (all 2,258 Broadway productions from 1877 to 1990). This industry is a good choice to investigate this question for two reasons. First, there are no firms. Teams consist of producers, directors, lyricists, composers, etc. and form on a project-by-project basis. Dense social clusters form if the same people repeatedly collaborate. New links form as newcomers join the system or as incumbents collaborate on a project for the first time. The second reason is that success in the Broadway industry can easily be measured by box office revenues, critical acclaim, or the length of runs.

Uzzi and Spiro show that the connectedness of the network (measured by clustering coefficients) has a strong effect on the average yearly success of the industry, even if controlled for various other independent variables, such as market characteristics. Surprisingly, the relationship is curvi-linear. Uzzi and Spiro suggest that if networks are too dense, too much imitation may lead to the dominance of conventional ideas. On the other hand, if networks are

* *Corresponding author address*: Luís A. Nunes Amaral, Department of Chemical and Biological Engineering, Northwestern University, Evanston, IL 60208; e-mail: amaral@northwestern.edu.

not sufficiently connected, fresh ideas may emerge, but they will remain isolated. From a managerial point of view, this means that there is a “sweet spot” in the structure of the networks. An organization could potentially measure the connectivity of its internal network and then attempt to increase or decrease its connectivity.

Influenced by this line of work, Guimerá et al. (2005) studied a model of team formation characterized by the propensities that incumbents continue to collaborate on a new project or are matched with other incumbents or with newcomers. They show that different matching probabilities lead to different network topologies and therefore different expected performances. Surprisingly, the model is able to replicate the structure of various collaboration networks from different scientific communities.

What is missing from these papers is a behavioral model of collective problem solving that then can be used to examine the effects of network topologies on system performance in more detail.

MODEL OF COLLECTIVE PROBLEM SOLVING

In a recent series of papers (Moreira et al. 2004, Seaver et al. 2006), we demonstrated how a complex network topology combined with noise can alter the efficiency with which a system of decentralized units performs a global cognition task. The idea is to consider networks of agents who can be in one of two states denoted here $\sigma_j = \pm 1$. Moreover, to capture the effects of misunderstandings and other forms of miscommunication, there is a probability η of miscommunication; that is, with probability η, agents perceive the state of any connected agent to be randomly ± 1. Collective problem solving is modeled as a density classification task — a widely used measure of coordination and global information processing (Crutchfield and Mitchell 1995). For a system made up of units whose state is a binary variable, the density classification task is completed successfully if all units converge to the same state and the coordinated state is identical to the majority state in the initial configuration. The initial state of the system can be interpreted as receiving a partial signal about the correct solution.

Agents may use various social learning heuristics. Specifically, we consider a population composed of N independent agents and assign to each agent a bias $b_j \in \{-1, \sigma_j(t), 1\}$ with $j = 1, \ldots, N$. If $b_j = \sigma_j(t)$, the agent is “conservative.” A conservative agent requires a “qualified” majority of his neighbors to convince him to change his state. If $b_j = -1$, the agent is “partisan” toward -1. A partisan agent will “stick” to his preferred state until convinced otherwise by a qualified majority of his neighbors. The update rule for an arbitrary agent implementing a qualified majority heuristic is

$$\sigma_j = (t+1) = \begin{cases} +1 & \Delta_j(t) > -b_j s_j \\ \sigma_j(t) & \Delta_j(t) = -b_j s_j \\ -1 & \Delta_j(t) < -b_j s_j \end{cases} , \tag{2}$$

where $\Delta_j(t)$ is the average noisy “signal” an agent receives from his k_j neighbors.

We further assume that agents can have different levels of bias strength $s_j \in [0, 1]$, which measures the fraction of an agent’s neighbors that must be in the state opposing the state preferred by the agent in order to "force" him to change state. If $s_j = 0$, the agent uses a majority rule, whereas if $s_j = 1$, the agent is frozen and will never change his state.

The efficiency $E_\phi(p,\eta,N)$ of an updating rule ϕ is a function of noise intensity η, rewiring probability p, and system size N. Specifically, we have

$$\varepsilon(t) \equiv \frac{N_+(t) - N_-(t)}{N}, \tag{2}$$

where N_+ is the number of agents that are in state "+1" and N_- is the number of agents that are in state "–1." For each realization, we let the system evolve for $2N$ time steps. We define the efficiency $\bar{\varepsilon}$ of a single realization as

$$\frac{1}{N/4}\int_{2N-N/4}^{2N} \varepsilon(t)dt \tag{3}$$

The efficiency E_ϕ for each set of parameter values is the average of $\bar{\varepsilon}$ over 1,000 realizations.

Moreira et al. (2004) consider small-world networks (Watts and Strogatz 1998) with $k = 6$ neighbors, rewired with a probability p (where a small-world topology occurs at p of ≥ 1) and a probability η of miscommunication. They show that if agents use the majority heuristic (i.e., set $s_j = 0$), the system exhibits rapid and robust convergence to the correct state, provided the interaction structure is characterized by moderate noise and constitutes a small-world-network. Importantly, both conditions are necessary for effective problem solving. Moreover, more complicated decision rules that work well in the case of $\eta = 0$ and $p = 0$ (e.g., Gacs-Kurdyumov-Levin [GKL] rule) (Crutchfield and Mitchell) fail to function when communication is noisy or interaction occurs in an asynchronous fashion. Intuitively, this implies that not only are random connections to other members of the network important for problem solving, but that, if they exist, even extremely simple decision heuristics can be successful. Surprisingly, for the case of $k = 6$ (the case for which the GKL rule was optimized), the efficiency of the majority rule reaches a value of 0.85, which is *greater* than what is obtained with the GKL rule under the idealized conditions $\eta = 0$ and $p = 0$. Strikingly, unlike the GKL rule, the majority rule yields efficient coordination even for asynchronous updating. One can thus understand the role of each of the components in this condition. The noise enables the system to escape conformations with multiple domains. The rewiring of the connections allows fast access to information from across the system; that is, the long-range connections make the system a small world. Importantly, the small-world phenomenon is *not enough* to ensure the convergence to the correct classification. Only with the combined effect of the noise and the small-world topology can the system reach a consensus within the permitted evolution time.

Moreira et al. originally used a value of $c = 0.5$. However, with the use of agents that have a preferred state, we need to ensure that the majority state — chosen to be +1 — is always the same. To find the "best" value for c, we explore, for a variety of system sizes, the resulting efficiency with which a system reaches consensus within $2N$ time steps, as a function of c (Figure 1A). We choose a combination of a system size that is computationally inexpensive ($N = 401$) and the lowest value of c that will lead to an average efficiency for $\geq 0.95\%$ of the time

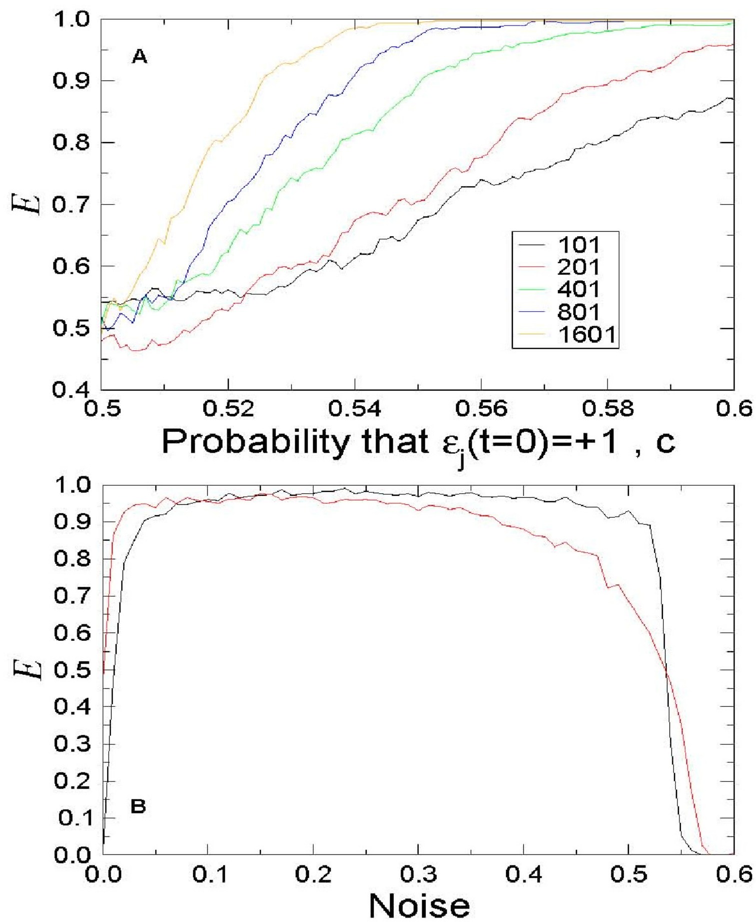

FIGURE 1 Exploration of initial distribution of state +1. The probability that an agent's state would initially be +1 (*c*) in Moreira et al. (2004) is 0.5. The initial majority was always tiny, and the final consensus was either 1 *or* –1. In order to force a consistent consensus in the context of which we could explore biased agents, *c* was chosen to be 0.57. Figure 1A highlights why $c = 0.57$ *and* $N = 401$ were chosen. It shows the response of systems to the increase in the starting fraction of agents with a positive state ($p = 0.15$ and $\eta = 0.2$). A system size of 401 was chosen as a compromise because of the performance time; thus, 0.57 was used as the starting fraction where each realization would, on average, achieve an efficiency of $\geq$0.95. By using the parameter in this manner, the majority will always be +1 and almost always become the final consensus of the system, and the dynamics that resist the consensus will be more obvious and easier to understand. Figure 1B shows a comparison between the original (old) method of determining efficiency (Moreira et al. 2004) and the new method where $N = 401$. In addition, $c = 0.5$ for the old method and $c = 0.57$ for the new method.

(c = 0.57). The resulting efficiencies calculated by the original and new methods are compared as a function of noise in a homogenous population and found to be similar (Figure 1B).

Seaver et al. (2006) considered more general interaction structures with s_j of >0. In this case, agents may exhibit various forms of decision biases, as discussed above. Adding partisans to the model (even if they are distributed evenly between partisans for +1 or −1) dramatically decreases the performance of the system. In the case of conservative agents, the model yields a surprising conclusion. For moderate levels of bias (s_j of <5/7), not only does the system show remarkable levels of efficiency, but system performance actually increases as the fraction of conservatives increases, provided the noise level is sufficiently high. However, there is an important tradeoff with the speed which a solution is reached; that is, the time to reach consensus grows in the fraction of conservatives. If the fraction is larger than 30%, consensus cannot be reached within $2N$ time steps.

To estimate the "typical" number of time steps it takes for a system to reach the steady state, we use the following method. We generate 5,000 time series of T time steps with different initial conditions using the same parameters (Figure 2). We compute the average time series $\overline{\varepsilon(t)} = \frac{1}{5000}\sum_{i=1}^{5000} \varepsilon^i(t)$. Then for each time step t, we compute the average efficiency for the remaining time steps $\overline{\varepsilon_{av}(t)} = \sum_{t'>t} \overline{\varepsilon(t')}\frac{1}{T-t}$. We fit the resulting curve (Figure 2B) to a stretched exponential:

$$f(x) = A(1 - Be^{-(\frac{t}{\tau})^{\beta}}). \tag{4}$$

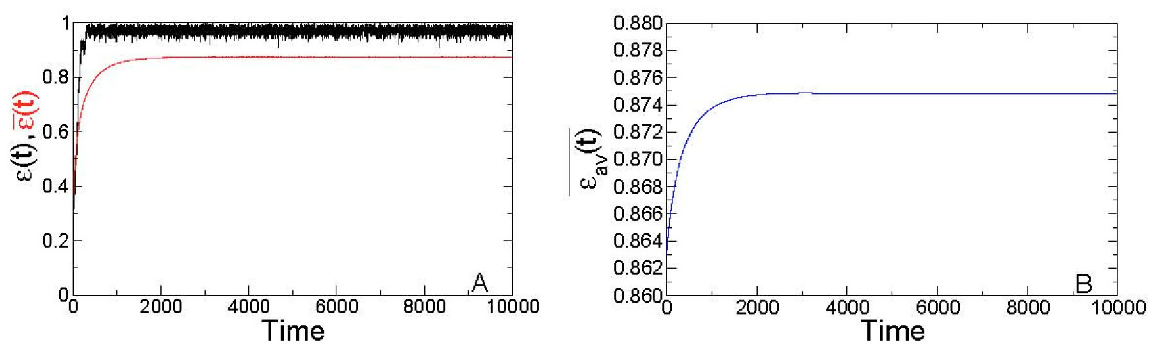

FIGURE 2 Method used to estimate t^*. Figure 2A shows examples of one of 5,000 time series generated from a system where N = 401, η = 0.4, p = 0.15, f_C = 0.4, and s = 4/7 and the average time series $\overline{\varepsilon}(t)$. Figure 2B shows the resulting curve when, for each time step t in $\overline{\varepsilon}(t)$, we compute the average efficiency for the remaining time steps $\overline{\varepsilon_{av}(t)} = \sum_{t'>t} \overline{\varepsilon(t')}\frac{1}{T-t}$. This curve is fitted to a stretched exponential $f(x) = A(1 - Be^{-(\frac{t}{\tau})^{\beta}})$. $t^* = \tau(-\ln\frac{0.05}{B})^{\beta}$ is an estimation of the time it takes for the efficiency to vary by 95% of the total variation at $t = \infty$.

We estimate the convergence time t^* as the time step at which the efficiency variation equals 95% of the total variation of $\overline{\varepsilon_\infty(t)}$:

$$t^* = \tau(-\ln\frac{0.05}{B})^{\beta}. \tag{5}$$

We explore a system with $N = 401$, $\eta = 0.4$, $p = 0.15$, $f_C = 0.4$, and $s = 4/7$. We obtain the convergence time t^* for the system at $\eta = 0.4$ and an increasing f_C. Figure 3A shows that the time needed for a system to reach a steady state increases with f_C. The steady state is dominated by agents in the +1 state, meaning that a system is highly efficient at reaching a consensus state of +1 provided that it is given enough time to evolve.

Finally, from Moreira et al. (2004), we know that a larger k will make a system more robust to noise. Our results have shown that the presence of conservatives will also make a system more robust to noise. Therefore, to observe whether conservatives would make a system more robust to noise independently of k, we explore the efficiency for $k = 4$, $N = 401$, $p = 0.15$, $s = 2/5$, and $s = 3/5$ (Figure 4). Although these results are not directly comparable to the results for $k = 6$, the same trend can be seen with $s = 2/5$, where the presence of conservatives still makes the system more robust to noise.

TEAM NETWORKS

In addition to considering a small-world topology, we also consider more realistic social networks generated according to the model of Guimerá et al. (2005). The model is built on the concept that collaboration usually occurs in teams. A network is built by using teams of a fixed number of agents (m) as building blocks. For each agent in a new team, there is a probability ρ that it will be an incumbent (i.e., an agent already present in the network) and a probability $1 - \rho$ that the agent will be a newcomer. If there are more than one incumbents in a team, there is a probability q that an extra incumbent will be a previous collaborator of the incumbents already

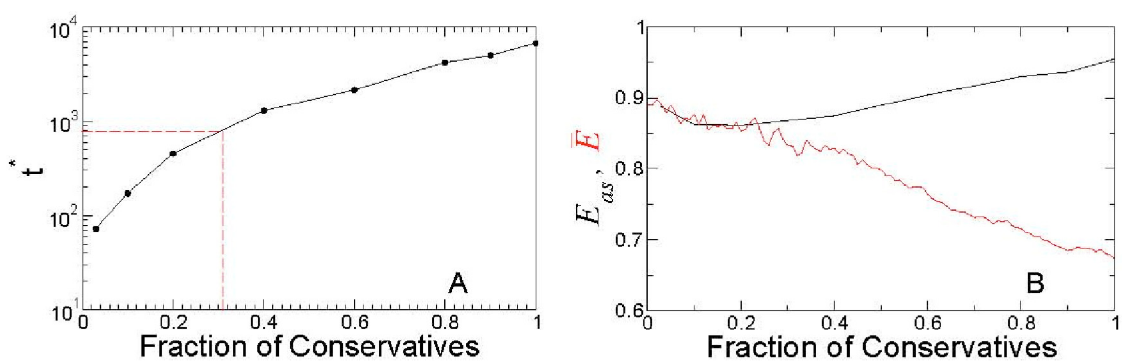

FIGURE 3 Figure 3A shows time t^* for a population of agents to reach the stationary state. Figure 3B compares the asymptotic efficiency (E_{as}) and the efficiency E at $t = 2N$ as a function of the fraction of conservatives. Note that for an f_C of >0.3, the system cannot reach the stationary state in the $2N$ time steps used in the simulations.

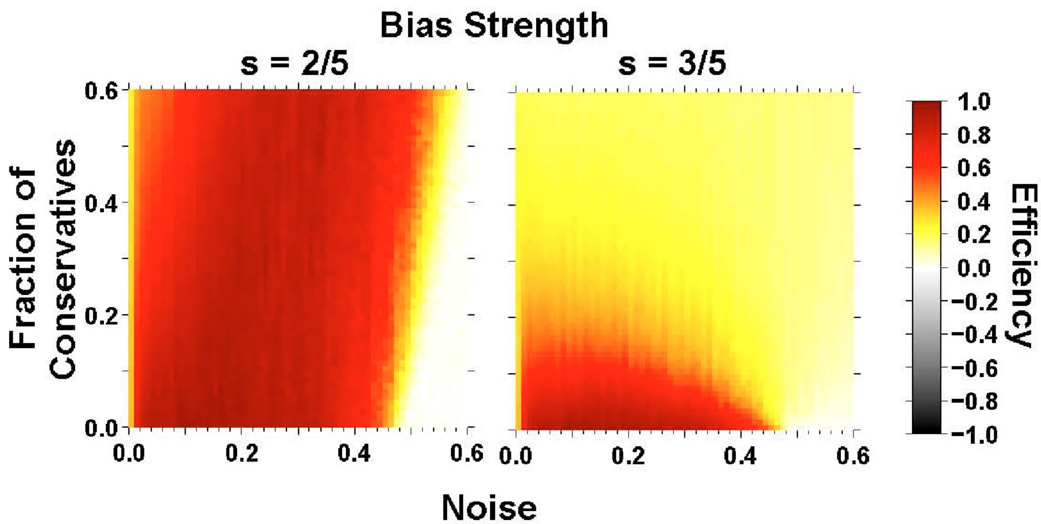

FIGURE 4 Efficiency of the system at reaching consensus as a function of the fraction of conservatives and of the noise intensity. The parameters are N = 401, c = 0.57, p = 0.15, and k = 4. Note that the change in k means a change in the range of s. It is clear that the presence of conservatives will increase the robustness of the system.

in the team. Otherwise, with a probability $1 - q$, the extra incumbent is *any* incumbent in the network. Finally, agents that do not participate in any team for a long time are considered to have retired and are removed from the network, allowing the size of the network to reach a steady state.

Almost every network generated by using this model is fragmented; thus, for a more reliable comparison to idealized small-world networks, we use the giant components of the networks. For one set of giant components, we use ρ = 0.364 and a retirement rate τ = 250 to reach both a mean network size of ~ 800 and an average giant component size that is ~50% of the network. For the second set of giant components, we use ρ = 0.5 and τ = 140 to reach a mean network size of ~440 and an average giant component size that is ~90% of the network. For both set of giant components, the average size is ~400 agents. Figure 5 contains examples of the giant components we use, displaying the differences in the network organization. In order to quantify the differences in sparsity, we measure the average degree of every node in these two sets. In the giant components that make up only 50% of their networks, the average degree is $\simeq$4.06. In the giant components that make up 90% of their networks, the average degree is $\simeq$4.45. We explore the efficiency of these networks at reaching consensus as a function of the fraction of conservatives and of the noise intensity. The results are present in Figure 6. The average degree of each agent in the small-world networks is 6 simply because of the number of neighbors k = 6, but this measure does not seem to be the dominant reason that the overall efficiency in the team networks is much smaller. Compare these results with the results for small-world networks where k = 4 (Figure 4). The dominant factor seems to be that these networks are not small-world networks.

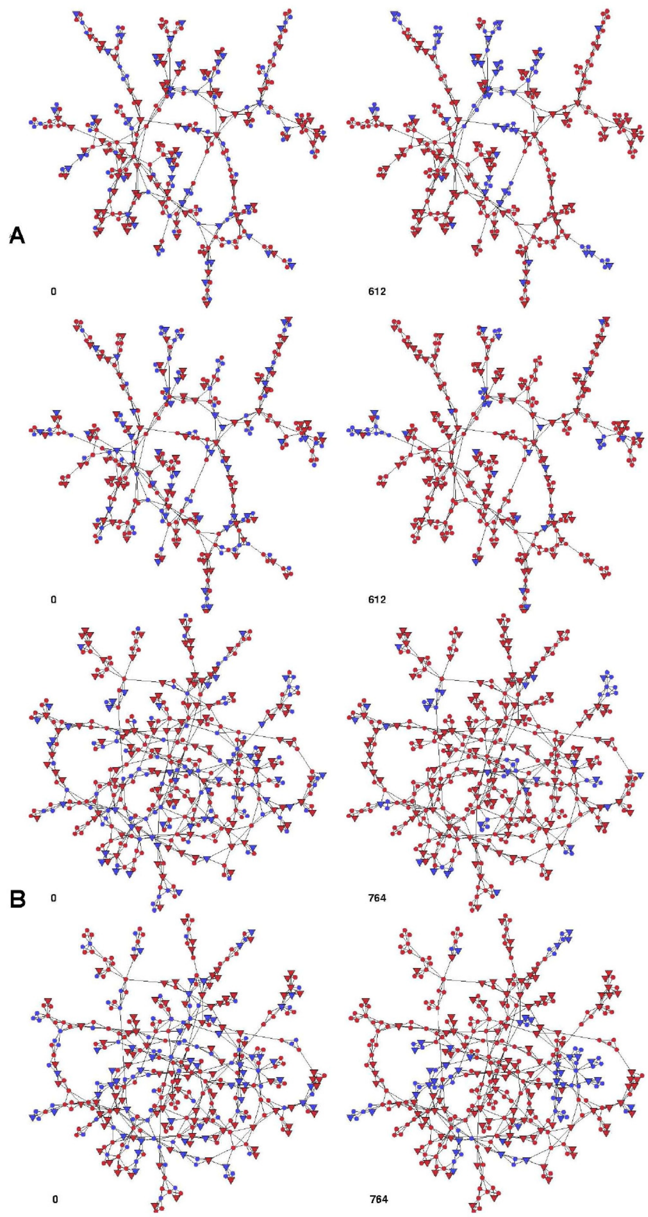

FIGURE 5 Examples of initial configurations and final states of teams networks. The networks here are built according to the model of Guimerá et al., being giant components that consist of (A) 50% or (B) 90% of the number of nodes. For each set, two different seeds were used to create the different distributions of agent types and states. Note that the final localized regions of blue (–1) nodes are clustered around conservatives of the same state, which is indicative of the short-range influence that conservatives have.

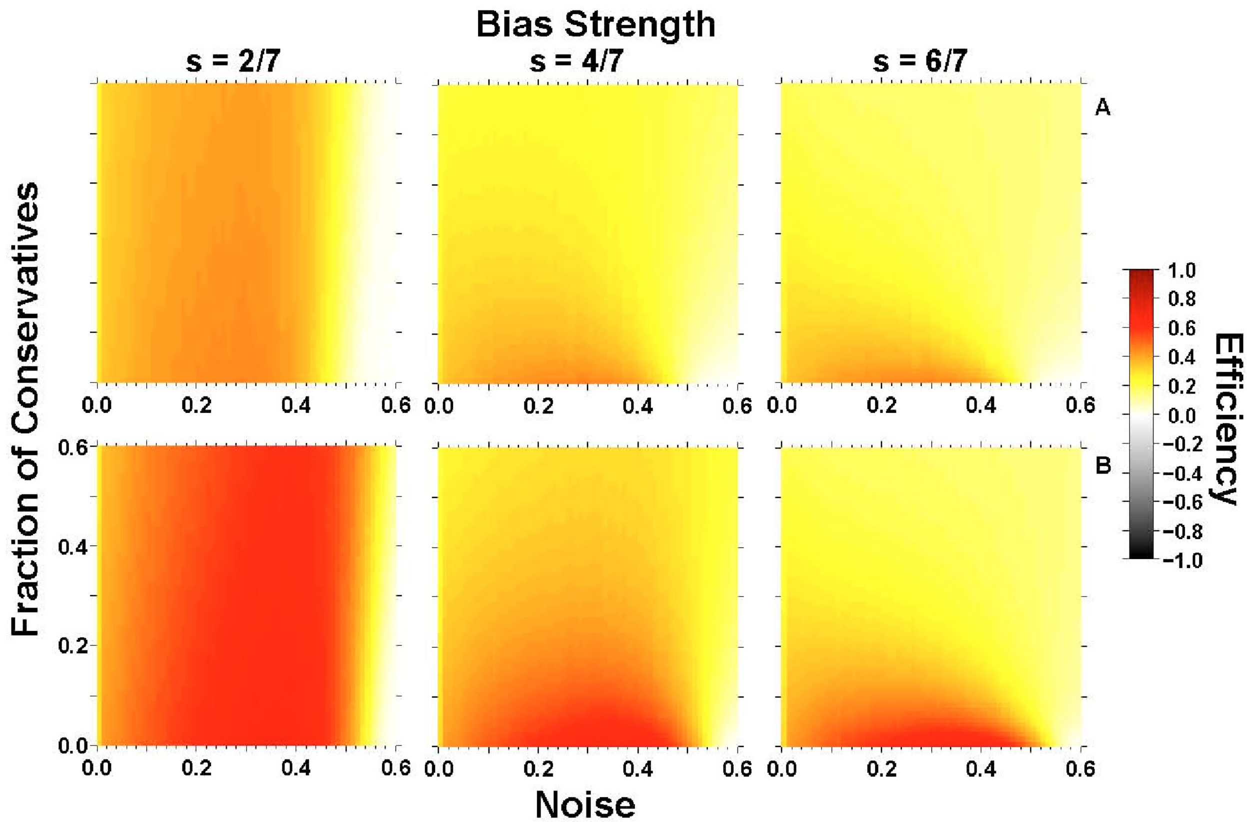

FIGURE 6 Efficiency of the teams model. This shows the efficiency of the system (built according to the teams model of Guimerá et al.) at reaching consensus as a function of the fraction of conservatives and of the noise intensity. These results are for systems with 57% of the agents being initially in state +1 and (A) a mean number of 398 agents, $\rho = 0.364$, $q = 0.5$, $\tau = 250$ or (B) a mean number of 363 agents, $\rho = 0.5$, $q = 0.5$, $\tau = 140$. The value used for ρ in Figure 6A means that the average fraction of the total number of nodes that are in the giant component is around 49% and the average degree is 4.06. In Figure 6B, they are 89% and 4.45, respectively. The other parameter value used is $c = 0.57$.

CONCLUSION

We provide a model of social cognition in complex networks. We show that very simple rules can perform very well in problem-solving tasks, such as the density classification system, provided that the networks satisfy the small-world property and that the interaction between agents is subject to moderate noise. In the case where interaction between agents is given by a self-organizing team network estimated from the data, the system performs moderately well, but not as well as in the case of small-world networks. This suggests that there is room for managerial intervention to improve system performance to increase the connectedness of the network.

REFERENCES

Crutchfield, J.P., and M. Mitchell, 1995, *Proceedings of the National Academy of Sciences U.S.A.* 92, 10742.

Guimerá, R., B. Uzzi, J. Spiro, and L.A.N. Amaral, 2005, *Science* 308, 697.

Moreira, A.A, A. Mathur, D. Diermeier, and L.A.N. Amaral, 2004, *Proceedings of the National Academy of Sciences* 101, 12083.

Seaver, S.M.D., R.D. Malmgren, A.A. Moreira, M. Sales-Pardo, D. Diermeier, and L.A.N. Amaral, 2006, *Science*, under review.

Uzzi, B., and J. Spiro, 2005, *American Journal of Sociology* 111, 447.

Watts, D.J., and S.H. Strogatz, 1998, *Nature* 393, 440.

DISCUSSION

Social Simulation Applications

(Parallel Applications Session II — Social Policy and Computational Knowledge Friday, October 14, 2005, 1:15–3:15 p.m.)

Chair and Discussant: *Ed Tanzman, Argonne National Laboratory*

Beyond Markets and Communities: A Comparative Approach to Knowledge Exchange in Organizations

Ed Tanzman: My name is Ed Tanzman, and I'm on the staff at Argonne National Laboratory. I want to welcome all of you to our split session on social policy and computation. We have four excellent papers this afternoon, which, I think, cover a range from very methodological to very applicational. I think they'll provide us with many interesting insights on how agent-based modeling can be used in the context of addressing questions of social development and social policy and the methods by which the science itself will advance.

Our first speaker, Sheen Levine, has a plane to catch, and so I'll know he'll be on time, and our last speaker is coming off of a plane. We're pretty confident he'll be on time because we haven't heard that there are any problems. In recognition of the importance of keeping our schedule, I want to turn the floor over to Sheen Levine, from Singapore Management University, with his paper entitled "Beyond Markets and Communities: A Comparative Approach to Knowledge Exchange in Organizations."

Sheen Levine: Thanks for the introduction. I'm Sheen Levin. I would love to stay here because there have been some very interesting presentations this morning, and I bet it will continue tomorrow.

[Presentation]

Tanzman: We have time for a few questions.

Doug Lauen: I'm not sure about the implications of this question, but I worked in a consulting firm at one time. When somebody wanted advice, if it was just a quick question, people would talk in the hallway, but if it involved a long-term conversation — a meeting or something like that — you had to get a job code from the person because there was an internal labor market, and if you didn't get a job code, the company ate the costs in overhead. I was wondering what your experience has been in this area. There are opportunity costs of people's time; if they're not doing something productive, the firm is not going to be very competitive.

Levine: That's right. This is a very good question because it is exactly what neoclassical exchange is about. Neoclassical is strict economic exchange. "If you have the money, I'll talk to you. If you don't have the money, don't bother me." As we saw in the simulation, these companies do not perform very well because, for example, you're starting a new project. You don't have the money. You potentially will have the money if the client buys your proposal, but you can't pay for people's time until the client buys your proposal. So at this point in time, these companies are not going to perform as well as companies that are based on either community mode or performance mode.

Mengxiao Zhu: My name is Mengxiao Zhu, and I'm a graduate student from the University of Illinois at Urbana-Champaign. It seems you could build a structure between those agents. How do you choose which mechanism you use and which mechanism is not used? For example, you use social exchanges of interests and critical mass. Why don't you include some models? I mean, we do not know each other, but we might read the same book, and then we would more likely build a connection.

So my question is why do you choose certain mechanisms and not the others? What are your criteria for choosing? Also, in talking about the map of the knowledge, who knows whom and who knows what, how many steps can continue to knowledge map with work? Maybe I know somebody, and somebody knows somebody knows some knowledge. In this simulation model, this level of knowledge would function well. In the simulation, I was also wondering if broken ties sometimes were re-established. Would you please explain the broken mechanism you used in your model and also some other attributes, like the efficiency, the willingness to be efficient, even when someone has enough knowledge, but they are not willing to be efficient, to solve this problem? So how you represent this kind of model? Thank you.

Levine: Okay. I think we should recruit you as a co-author. Let me give your questions one by one.

First, I think that you alluded to using tags, which is something that we're looking into. Right now "Homophile" is represented by the fact that in the community mode, I talk only to my friends and my coworkers. People on the same team are going to help each other out. People that are not on the same team and are not neighbors will not help each other out. We think about adding some tag, and a tag could be anything from graduating from the same university to having the same ethnicity, and although we don't know each other, we'll help each other out based on that. This is one way to tackle that.

I think your second question was about attrition. This is an interesting point. If people want to know what other people know, they can use the knowledge management system (KMS), which is an organizational index. This is a game based on empirical work. It tells people who works and what and who has certain types of knowledge. We can make this system very precise or imprecise based on what we determine in the simulation. If they want to contact a stranger, they go to the KMS to see who has experience with the task.

Steve Younger: I'm Steve Younger, from Los Alamos. A couple of things struck me when you were talking. One is that there is an attribute of altruism that contributes to someone's willingness to share. How do you fit that in? Also, the concept of demand sharing, that is, an unwillingness to say no to someone because it may influence your reputation. Finally, the concept of community can be complicated in that it can be the community of the company, but it can also be the community of scholars, so there can be an element of delayed reciprocity and an element of reputations. How do you fit those into your scheme?

Levine: Well, these are very broad topics. We include altruism, but I'm hesitant to talk about it because it's difficult to envision a company that is completely based on altruism. Just think about the recruiting problem. How do you identify altruists and free-riders?

Yes, the business model of these professional service firms is that you solve a problem once in one location for one client and then you enable the flow of knowledge within the company so that other teams at work on similar problems can use this knowledge to solve problems of other clients in other locations.

The fundamental question is how did this knowledge get around, and there are several ways it could go around. One of them is by paying people to tell you what they know. Another one is by recruiting altruists that would help you out just because they get pleasure in helping other people. Other forms are community or generalized exchange, in which people help each other, but it's not pure altruism because they expect something in return, but not necessarily from the person they helped. Based on ethnographic research, we know that there are communities that are based on generalized exchange as a mode of exchange. We think that organizations can also be based on generalized exchange. This is the reason we included these modes of exchange and not pure altruism as a mode of exchange.

Tanzman: All right. Last question, please.

Lee Hoffer: My question is similar to that in the exchange: a lot of things are transferred. Many of these things have economic value, but many don't have economic value, so it's things like respect and other things of this nature. Have you thought of ways to try to integrate that? I'm curious because I have a very similar situation with dealers and users where, with respect to the situation, it often is getting transferred along with economic utility in exchange. Also, I wonder if you have had any findings from your work on what additional outcomes you might look at other than just completing tasks.

Levine: Well, I'm a business school professor, so that's what I'm interested in. I think that there are other things going on there, like emotional support, for example, or friendship, which is a way of keeping these people in the company. I haven't been looking at it specifically, but, as we know, qualitative data are so rich that I could probably go back to the transcripts of the interviews of the operations and look into these things.

All right, I must leave, and I apologize, but my e-mail is there on the board. If you have questions or ideas or even better critique of what we did, please do write to me. We appreciate any input at this early stage of the project. Thank you.

Tanzman: Thank you, Dr. Levine. I don't know if anybody has any other questions for Dr. Levine while we're waiting for Dr. Zeidenberg to set up, but I have one. Did you attempt to model, or could you model, time as a variable in the effectiveness of different methods of exchange of knowledge? You assumed, if I heard you correctly, that there were 30-day cycles roughly.

Levine: This is the limit that we gave them. We gave the organizations a time horizon, and said that we would give them a set of problems and come back 1 month later to see how much they'd managed to solve. This is how we compare the performance of the organization.

Tanzman: Do you think that if you modeled a range of times, you might find different levels of effectiveness of the methods that you were comparing, of the practices that you were comparing?

Levine: Yes, and this is what we are looking at next. I'm actually going to Emory to meet Mike, who is my coauthor, and we're talking about the next development of this simulation, and this is one of the things that we are going to incorporate.

Tanzman: Another question in the front.

Steven Wilcox: You gave us a graph showing diminishing marginal return for the amount of precipitive ties. You had probability or percentage on a long scale. I was wondering how that result would change if you did it on a linear scale or a logistic scale.

Levine: [Unintelligible]

Neil Silbert: In the professional services industry, it's very often that knowledge is discussed, not very evenly distributed, but also times very hierarchically distributed where there may be levels of directors, partners in subject-matter areas of expertise. How does a hierarchical distribution of knowledge influence your model?

Levine: [Unintelligible]

Silbert: Yes, I had found that there were very strong forces for subject-matter diffusion in an organization and a certain level of qualitative responsibility for the dispersion of that knowledge, so that raised the issue.

Ana Carrie: Hi. Ana Carrie, Trinity College, Dublin. Can you talk briefly how you encoded the skills, the knowledge, and then go on to how someone would actually ascertain whether they had the skill to do a job or whether somebody else had the skill? For example, was it just a string of characters, or what did you do?

Levine: [Unintelligible]

Carrie: How did you represent skill?

Levine: [Unintelligible]

Carrie: So you would have had just a finite number of 16 types of knowledge, and you have Type 1, Type 2, Type 3, and then the job comes along. Do jobs require multiple types of knowledge, or one type at a time?

Levine: The jobs require multiple types of knowledge.

Tanzman: Do we have some other questions? We might as well take advantage of this extra time. Go ahead, Chick.

Charles Macal: I was intrigued by the pattern of reciprocation. The matrix you had was similar to what I've seen for nonhuman primate reciprocal relations, like from the Cornell Primate Lab and places like that. Is there any relationship between primate behavior and organizational behavior, at least in terms of what the theory is feeding?

Levine: [Unintelligible]

Macal: Well, I mean, but formally. Is it the case, formally?

Levine: [Unintelligible]

Macal: Oh, the kinds of reciprocation, there are primary relations for which there is no immediate reciprocation in terms of time?

Levine: Yes.

Macal: Yes, but it still occurs, the concept being that these primates live together in a little clan and that there's an expectation even on their part that they would be reciprocated later potentially for those basic kinds of things.

Tanzman: I thank you all for your patience while we played dueling computers here.

Agent-based Models of Urban Industrial Specialization

Tanzman: Dr. Zeidenberg comes to us from the University of Wisconsin and is going to present his talk on agent-based models of urban industrial specialization.

Matthew Zeidenberg: Actually, my talk segues pretty well from the last one because I am studying many of the same things. The presentation is quite preliminary, but there's been a large block of literature, going back to perhaps the 1980s or so, about specialization industrial districts. The seminal book that was written about this subject was by Piore and Sabel, *The Second Industrial Divide*. It's mainly been this qualitative work where people are talking about observations of well-known industrial districts like Detroit, Hollywood, Silicon Valley, and Seventh Avenue in New York. The underlying idea is that there are all these externalities that are associated with these districts.

[Presentation]

Tanzman: Let's open the floor for some questions. First, though, I have a question that relates to the idea of sticky relationships. I see that the process over time of the aggregation that occurs perhaps captures some of the …. After the supplier and the OEMs [original equipment manufacturers] are stuck together, it would take a certain amount of time for them to re-assort. Do you think it would be possible with this type of method to model urban decay, in other words, the backward process? What would you see if you did?

Zeidenberg: Are you saying that there would be firms that are stuck where they were because they have no more market? I actually have thought about that with respect to the idea of putting some historesis in the model. Essentially, if you're already some place and you find another place that looks more attractive, you don't necessarily move there because there are large fixed costs associated with the move. In terms of urban decay, though, if you're losing your market, eventually you're going to either go out of business or move.

Tanzman: Thank you very much.

Modular Bayesian Inference and Learning of Decision Networks as Stand-alone Mechanisms of the MABEL Model: Implications for Visualization, Comprehension, and Policy Making

Tanzman: Our next paper is by Dr. Kostas Alexandridis, from Purdue. The title of his talk is "Modular Bayesian Inference and Learning of Decision Networks" involving the MABEL model, to shorten it a bit.

[Presentation]

Kostas Alexandridis: I'm glad to answer any questions.

Reginald Tucker-Seeley: Reginald Tucker from the Harvard School of Public Health. My question is very simple because I'm not that familiar with agent-based modeling. When you're putting up the flow for the agents and there were intention and behavior, how do you determine how the agents go from their intention to actually implementing that behavior?

Alexandridis: The MABEL model is very complex, but the basic idea is that we have an optimization model that maximizes the utility of the agent by using the Bayesian belief networks. The belief networks actually maximize the utility for an agent, given the information or the knowledge they have in advance, and that's a utility maximization problem. It's commonly encountered in economics, but we solve that in a more dynamic way, rather than equilibrium approaches.

Macal: Out of curiosity, do you have a website for MABEL that explains more about the model?

Alexandridis: Yes, and we have a three current papers right now. They will all be available on the website, and they are stand-alone. You just need Netica to run those models. [Editors' note: see http://web.ics.purdue.edu/~ktalexan/research.htm.]

Social Cognition in Complex Team Networks

Tanzman: Dr. Diermeier comes to us from Northwestern University, and the title of his talk is "Social Cognition in Complex Team Networks."

Daniel Diermeier: Thank you for coming. I'm going to speak about some work that my colleagues and I at Northwestern have been involved in for about a year-and-a-half. We're interested in social cognition in complex networks. I want to motivate the problem a bit so that you understand why we're interested in it and what the idea is. Before I do that, however, let me tell you who's working with me on this. There are a whole bunch of people: Luis Amaral, Roger Malmgren, Julio Ottino, and Sam Seaver, and without going into details, there are papers on which this talk is based. My purpose is to give you some ideas from that.

[Presentation]

Tanzman: We have time for a couple of questions for Dr. Diermeier.

Gabriel Istrate: Gabriel Istrate, Los Alamos. Have you looked into H. Peyton Young's results on the dynamics of norms? He has a theoretical result that basically says that if the network has a small neighborhood with a high in-clustering, then dynamics of norms converges very fast. This seems quite similar although the model is slightly different.

Diermeier: We've not looked at convergence speed. You're right. There is a connection to the work in economics on herding and denying information aggregation. We're really interested in the problem. This is a different scientific culture. You look at what Payton and other people have done, and you see that they're not usually interested in the social structure, at least most of the time. They want to prove results that hold for all types of network structures. We're interested in the topology of the network, and that's why we need to model it directly, but you're absolutely right, there's a close connection with these results.

Unidentified Speaker: But there is one result …

Diermeier: There is one result, that's true. But in most cases, the focus of the results is really about proving something for the stuff on potential games, something that holds for all types of networks. We're really interested in dissecting that and trying to get a handle on it. So it's a different perspective, but the research programs are clearly related.

Zhian Li: My name is Zhian Li. I need to clarify something about the network. Regarding the network connections, does the diversity of the knowledge network really influence the final result? And regarding problem complexity, does that mean you need to go through more levels because some knowledge, in tacit knowledge, is easier to learn?

Diermeier: I would argue that you're looking at different dimensions, so we're not talking about adversity. In most of our analyses, our agents are homogeneous. In this case, they're distributed — they're pulled from one distribution — but we're not really interested in measures of diversity. The second thing you're talking about is explicit versus implicit knowledge, and our agents from a cognitive point of view are very simple.

Li: … diversity is knowledge diversity. For example, both my friend and I connect with three persons three times a day. Of the persons I'm connected with, one is in a hard science and the other is in social science. But my friend might be in contact with all the computer sciences, and this is what I mean by diversity.

Diermeier: We capture this in the structure of the network, so there are no types there. That's what I mean. There are no types, right? We have done some work; we have looked at that, but not in this type of application. In our context, there's a homogeneous context. We're interested in the design variable — the degree and type of interaction, not with what type of people interact.

Alexandridis: Do you consider leadership between teams as a factor for connectivity?

Diermeier: Yes. That's a great point. We haven't looked at leadership. There is other stuff on teams where we have looked at team structures and newcomers and so forth, but we haven't looked at leadership per se. I think that's an excellent question. There is also the question about how to model that in this context, but I think it's a good idea. You could have heterogeneity in this sense and leadership. You could say, "This is a person people are more likely to follow. If they're connected with this person, it's more likely." It's a person-specific bias, not a state-specific bias. That will work.

Joanna Bryson: This isn't really a question; it's a comment on what was just said. Some people are already starting to model this through social acquisition, looking at the use of prestige to get good knowledge propagating. So if you assume that there's good and bad knowledge out there, how do you explain the good knowledge? All you have to do is look at very simple indicators, say, for example, the longevity or the amount of energy or the amount of connections. As long as you have some way to value the individual, you can increase the probability of imitating their knowledge that way. It's a way to do leadership with very simple agents.

Diermeier: Yes, that's interesting. The only way we capture them is in the connectivity, so there's a social capital story here. That's there, but there's nothing in there; there's low intrinsic credibility in that sense. It's that they're more likely to be persuasive than others, for example. That would be interesting to look at that. Absolutely.

Tanzman: I want to thank you all for your attention, to thank our presenters for their excellent presentations, and to encourage you to continue this dialogue during the breaks and lunches and so forth tomorrow. Thank you very much.

National Security and Emergency Management

AGENT-BASED MODEL FOR SIMULATION OF WEST NILE VIRUS TRANSMISSION

Z. LI,* Decision and Information Sciences Division, and J. HAYSE,
I. HLOHOWSKYJ, and K. SMITH, Environmental Science Division,
Argonne National Laboratory, Argonne, IL
R. SMITH, University of Illinois at Urbana-Champaign, Urbana, IL

ABSTRACT

The rapidity at which the West Nile virus (WNV) spreads and its potential for serious medical consequences underscore the necessity for better understanding the virus's transmission pathways and the conditions that affect its transmission. This paper describes an attempt to build a virtual laboratory by using agent-based modeling (ABM) techniques; the laboratory is to be used by WNV epidemiology researchers to study the characteristics of WNV transmission, including transmission pathways and the conditions under which a WNV outbreak might occur. The WNV transmission model uses the Repast ABM toolkit to simulate the dynamic interactions of the entities involved in WNV transmission. The results show that ABM is an effective technique for developing simulations of the transmission of infectious diseases. The modeling approach developed in this study is also applicable to simulations of the transmission of other infectious diseases, such as severe acute respiratory syndrome (SARS), avian influenza, and malaria.

Keywords: Agent-based modeling, West Nile virus, infectious diseases, transmission, avian influenza, SARS

INTRODUCTION

West Nile virus (WNV) is a deadly disease for which there currently is no effective treatment or vaccine. In recent years, WNV has spread across the mainland of the United States, causing a great deal of concern among the public as well as within federal and state public health agencies and natural resource agencies (CDC undated). WNV has been implicated in human fatalities in most U.S. states and identified as the cause of major reductions in native bird populations in many areas. The latest reports also show that WNV, like meningitis, can cause paralysis in humans (Neergaard 2005). The rapid spread of the disease and its potential for serious medical consequences underscore the necessity for better understanding the virus's transmission pathways and the conditions that affect its transmission.

As do other natural phenomena, it is believed that transmission of WNV has its own intrinsic characteristics. A fundamental understanding of these characteristics would assist the research community, government agencies, and local communities in developing more effective monitoring approaches and preventive control measures for the virus. There is a need for a tool

* *Corresponding author address*: Zhian Li, Decision and Information Sciences Division, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439; email: li@anl.gov.

that can facilitate further understanding of the intrinsic characteristics of various transmission pathways.

The primary pathway for transmitting WNV to humans, wildlife, and domestic animals is through the bites of infected mosquitoes. Mosquitoes become infected when they feed on infected animals (especially birds and mammals) and then re-transmit the virus to other organisms during subsequent blood meals. Once the virus is transmitted to an animal host, it multiplies within that host animal, where it creates a reservoir for further infection. The interactions among mosquitoes and host animals form a cycle of WNV transmission.

Figure 1 illustrates the dynamics of WNV transmission. Agent-based modeling (ABM) techniques are especially well-suited for evaluating such processes. This paper describes an agent-based model (also ABM) that is being developed to simulate the spread of WNV. The objective of this study is to develop a virtual laboratory to be used by WNV epidemiologists to study the intrinsic characteristics of WNV transmission via computer simulations.

WEST NILE VIRUS

WNV is a vector-borne disease. This means that WNV infection is spread via intermediate hosts, such as mosquitoes. The virus is transmitted from infected mosquitoes to hosts and multiplies in the blood of the infected hosts. The most common hosts are avian and mammalian species, such as American crows, blue jays, raccoons, or chipmunks. WNV epidemiology research indicates that certain reptiles, such as crocodiles and lizards, can also serve as hosts for the virus. The virus can transmit back to mosquitoes when noninfected mosquitoes bite the infected hosts. In this way, the virus completes its transmission circle.

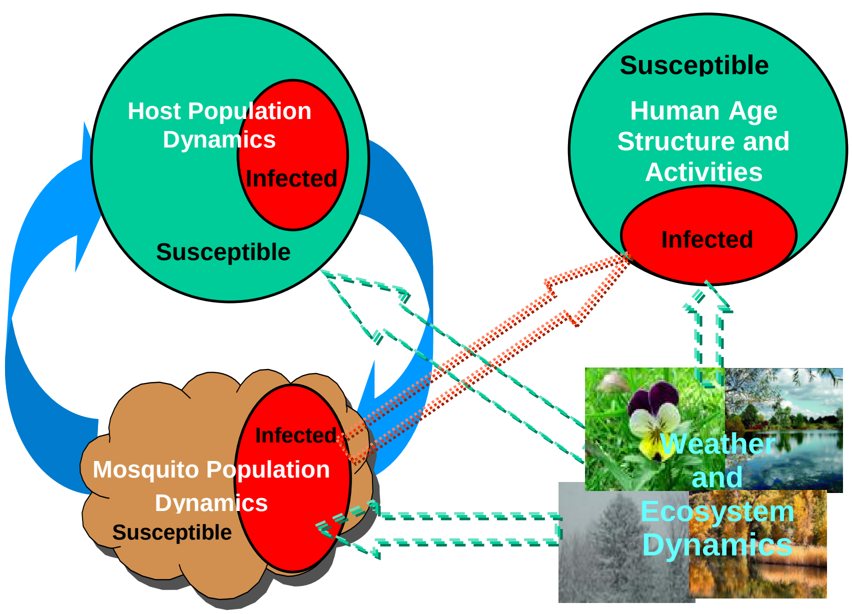

FIGURE 1 Dynamics of WNV transmission

Large-scale geographic spread of WNV has been attributed primarily to the foraging and migration of the avian hosts. Mammalian or reptile hosts do not, in general, transverse great distances and thus are not likely to be responsible for the rapid spread of the virus across the United States. In contrast, the local spread of the virus may be affected by the behaviors and movements of all three types of hosts.

People may become infected when they are bitten by infected mosquitoes. Humans are not in the transmission circle of the virus because a human's blood system cannot serve as a reservoir for the virus. Thus, the virus reaches a dead end in the human body because the virus does not attain a concentration in human blood that is sufficient to infect a mosquito that is taking a blood meal from the infected person.

The dynamics of the habitats of the hosts and mosquitoes are important factors in the WNV transmission process. The qualities of the habitats affect the mosquito and host populations by affecting their reproduction and death rates. In addition, the quality of the hosts' habitats will also affect their foraging behaviors and therefore the transmission speed of the virus, since the home ranges of birds tend to expand as abundant food becomes less available.

The last important component of WNV transmission is weather. Weather conditions are related to the dynamics of all component processes in the WNV transmission circle. First of all, temperature, humidity, and surface moisture are key conditions for mosquito reproduction. Mosquitoes are not be able to produce eggs if the temperature is low or the soil moisture is zero. The larvae and pupae must grow in water as well. Conversely, mosquitoes may have an advantage when soil moisture is relatively low. Although the mosquitoes may face some degree of adversity as moisture decreases, it has been established that paradoxically, relative to their predators and hosts, they suffer less when soil moisture decreases (Marra et al. 2004). The humidity and surface moisture are, in turn, directly related to temperature and precipitation. The soil type, slope, and aspect of a specific geographic location are all factors affecting the surface moisture and availability of water for larvae and pupae to develop. In addition, the mosquito biting rate is also a function of temperature and weather conditions, because these conditions affect the outdoor activities of human beings. It is unlikely that people will stay outside when it is raining or very hot, thus their likelihood of being bitten is reduced. Figure 2 illustrates the interactions among the different agents and dynamics processes in the WNV transmission circle.

AGENT-BASED MODEL FOR EVALUATING WEST NILE VIRUS DYNAMICS

Argonne National Laboratory is using an ABM approach to simulate the spread of WNV. ABMs are tools that can simulate the behaviors of individual entities within a complex adaptive system (Kohler et al. 2000; Woolridge 2002; Ferber 1999). In an ABM, an agent, representing an individual entity, behaves in a specific location by following a set of simple rules and with a limited knowledge of neighboring areas. The key difference between the ABM method and traditional statistical simulation methods is that the ABM method does not attempt to predict what will happen during the evolution of natural phenomena. Rather, it mimics the behaviors of the individual participants in the system to simulate the evolution of a natural phenomenon. The intrinsic characteristics and emergent behavior of the system can then be observed through simulation. Hence, this technique is well-suited for simulating the spread of diseases, in which there is no central control of the process. In the case of epidemic diseases such as WNV, all of the participating entities act as autonomous agents.

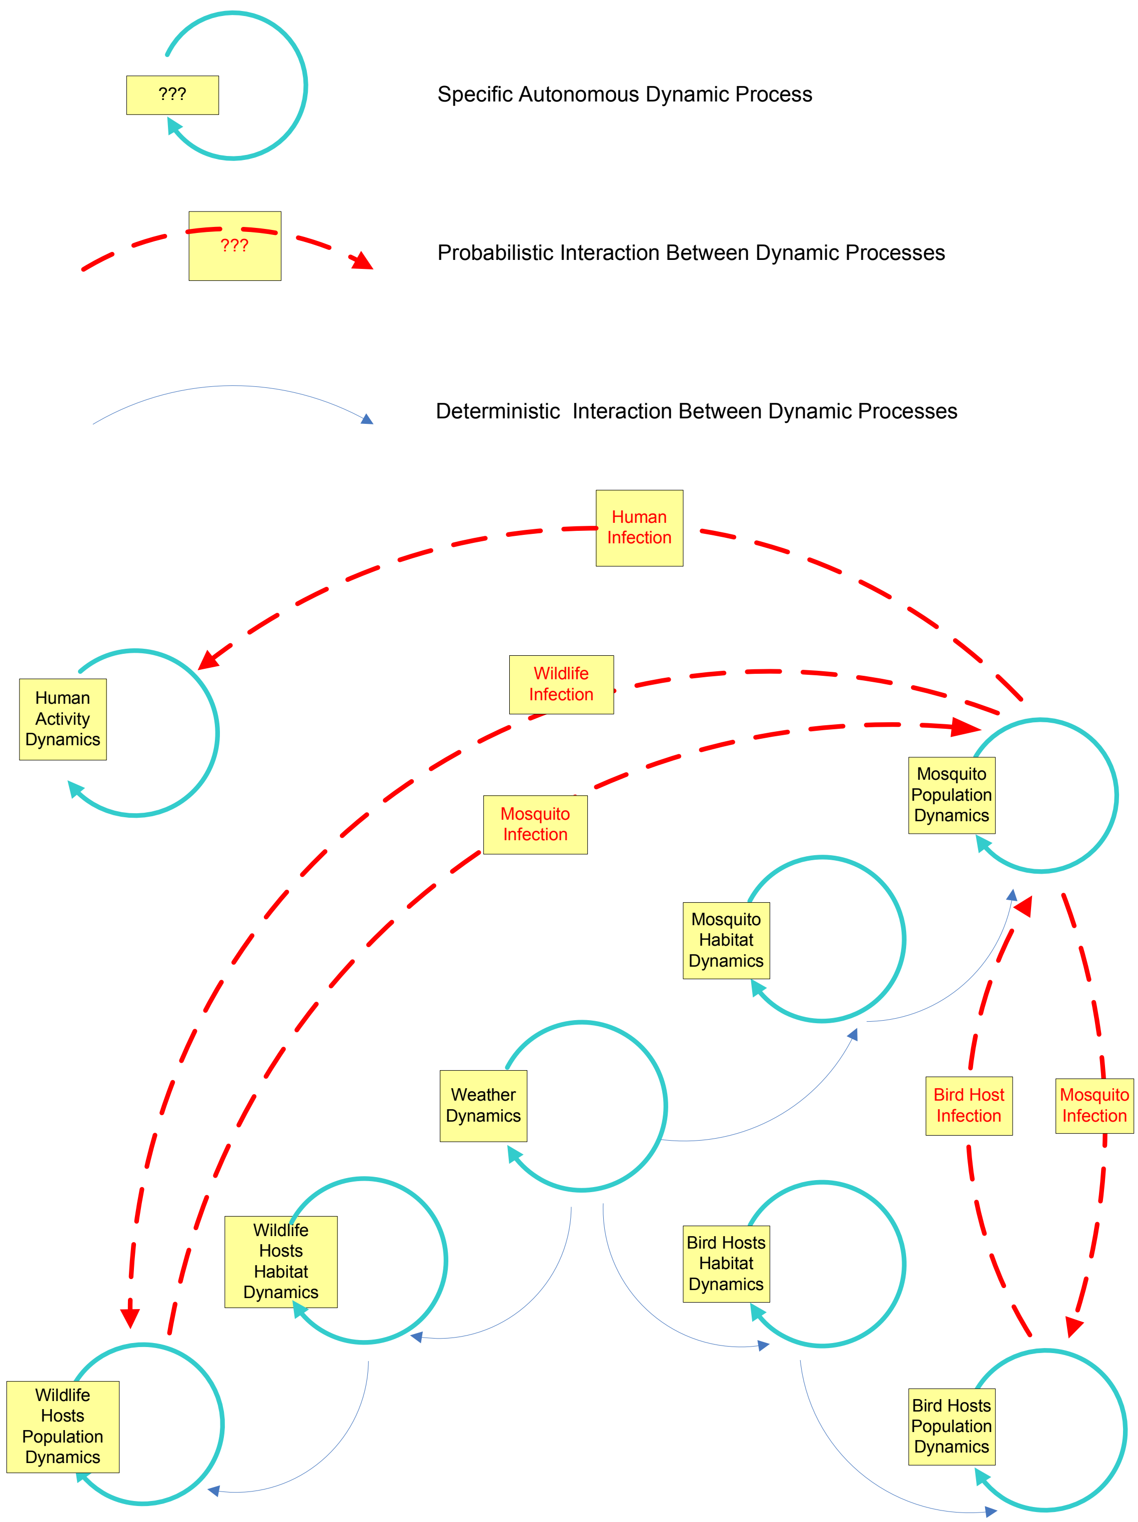

FIGURE 2 Interactions among different agents and dynamics processes in the WNV transmission circle

The ABM being developed simulates the transmission of WNV by using agents to represent mosquitoes, avian hosts, mammalian hosts, and humans. The activities and interactions of these various individual agents are simulated within a specified geographic area. A raster map is used to represent the area as a regular array of cells, which are linked to agent-specific and environmental data, such as habitat suitability values, weather conditions, vegetation cover, and other parameters. Habitat suitability values represent the suitability of specific locations (individual raster cells) for foraging and nesting by mosquitoes, birds, and mammals. A land use map is used to identify the areas where humans are likely to be at risk of being bitten by infected mosquitoes. The simulation also incorporates a weather model, which provides data on the temperature, precipitation, humidity, and surface moisture for each raster cell at each time-step. These meteorological parameters influence habitat quality for the mosquito and various host agents, as well the distribution and activities of human agents within each raster cell. The use of ecological conditions and climate parameters makes it possible to study the combined impact of weather and ecological conditions on mosquito reproduction and host population dynamics.

The ABM for simulation of WNV transmission is developed by using the Repast ABM development platform (Repast 3 undated). The geographic area of the model is an area of about 64 square miles centered on Oak Lawn Township in Cook County, Illinois. This area was chosen for the model because there is an ongoing field survey program for WNV transmission in this area. This field study is being carried out by the Spatial Epidemiology Laboratory at the College of Veterinary Medicine, University of Illinois at Urbana-Champaign, with the support of the Illinois Department of Public Health. Use of this same area for development of the WNV ABM will greatly aid in establishing quantitative relationships between different agents and the various ecological, environmental, and human behavioral parameters.

The WNV model will incorporate a geographic information system (GIS)-based visualization module for displaying the simulation results, including the distribution of infected and uninfected human populations, host populations, and mosquito populations, and the direction and rate that WNV is spreading within the 64-square-mile area at a 1-acre resolution. At this resolution, the model considers 40,600 1-acre raster cells. The selection of this resolution is a trade-off between model accuracy and computer resource availability.

Habitat conditions within each raster cell are used to estimate distributions for selected host species, while home range information for each species will be used to develop movement rules that will determine the movement of individual host agents within the modeled geographic area. Mosquitoes, host agents, and human agents in a given cell will be selected randomly for interactions. The WNV ABM will not predict what is going to happen; rather, it will simulate the self-centered processes that occur in the natural environment (under various environmental and ecological conditions) so that the characteristics that affect WNV transmission processes can be evaluated.

Note that the framework of the WNV model may also be applicable for modeling the transmission of other epidemic diseases, such as SARS, avian influenza, and malaria. Such ABMs for infectious diseases may allow researchers to better explore the conditions under which an epidemic might occur.

MODEL COMPONENTS

The WNV transmission cycle includes several interrelated complex dynamic processes, including weather dynamics, mosquito population dynamics, host agent population dynamics, mosquito habitat quality dynamics, host agent habitat dynamics, and human activities. Among these processes, weather plays a critical role because it affects all other processes, both directly and indirectly. The ABM of WNV transmission must capture these dynamic processes and put them in context. On the basis of the literature on emerging infectious diseases, mosquito entomology, and veterinary epidemiology (Hayse et al. 2005; Bernard et al. 2001; Lanciotti et al. 1999; Ruiz et al. 2004), the quantitative relationships among the different agents and processes can be defined with a flowchart, as illustrated in Figure 3.

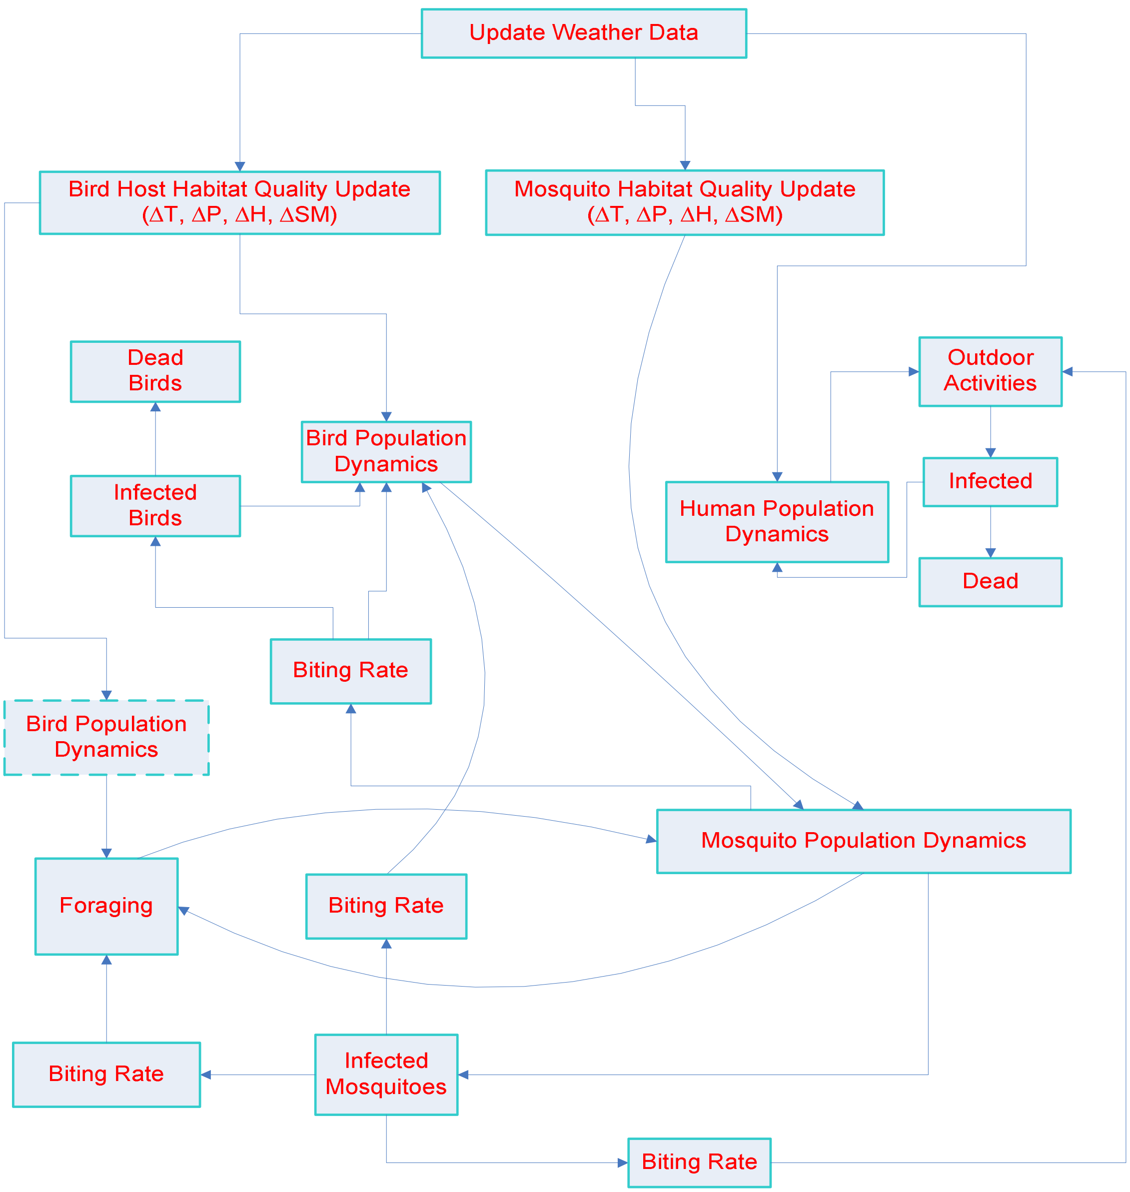

FIGURE 3 Agent-based WNV model

Mosquito Agent and Population Dynamics Model

The mosquito agent is the vector of WNV transmission. Female mosquitoes must take blood meals to develop eggs. Therefore, female mosquitoes will start to hunt for blood when they are ready to produce eggs. Female mosquitoes will take blood from a wide range of animals, including wildlife, domestic livestock, and humans. Mosquitoes may become infected with WNV when they bite infected animals, and animals and humans may become infected after they are bitten by infected mosquitoes. Infected animals may also act as reservoirs and amplifiers of the virus, because WNV can reproduce in many species. The probability of an animal being bitten by an infected mosquito depends largely on the mosquito population where the animal is located and the percentage of that mosquito population that is infected with the virus. While other factors, such as time of day, season, proximity, and the host's blood type, also play a role in the susceptibility of an animal to becoming infected, the model will use the simplified assumption that the likelihood of an animal becoming infected is based primarily on the mosquito population.

Mosquito population growth is a dynamic quantity that depends on the reproduction and death of the mosquitoes. The change in mosquito population can be defined by using the following differential equation:

$$dP(t)/dt = P(t)*(R-D) \ , \tag{1}$$

where P, R, and D are mosquito population density, mosquito reproduction rate, and mosquito death rate, respectively.

Mosquito reproduction and death rates are functions of weather conditions, especially temperature *(T)*, humidity *(H)*, and soil surface moisture *(SM)*. The soil moisture, in turn, is related to temperature, precipitation *(Pr)*, soil type, and the topographic characteristics of the study area, such as slope, aspect, and elevation. Mosquito reproduction depends also on the mosquito population density and habitat quality. Thus, Equation 1 can be expressed more explicitly as:

$$dP(t)/dt = P(t)*\{R[T(t),H(t),Pr(t),SM(t)]-D[T(t)]\} \ . \tag{2}$$

The solution to this differential equation can be obtained as:

$$P(t) = P_0 * \exp\{\int_0^t R[T(t),H(t),Pr(t),SM(t)]-D[T(t)]dt\} \ , \tag{3}$$

where P_0 denotes the initial mosquito population at a specific location. In the spring and summer, as the weather conditions and habitat quality favor mosquitoes, the mosquito reproduction rate increases and its death rate decreases. In the late fall and early winter, the reproduction rate can decrease and the death rate can increase sharply as a result of the sudden reduction of mosquitoes caused by frost. Some mosquitoes may survive over winter, but they are inactive and do not feed and are thus not capable of transmitting the virus during this time.

Host Agent Model

The hosts in the WNV transmission model are agents that maintain, amplify, and spread the virus across geographic locations. More than 100 species have been reported to be capable of serving as host species, including reptiles, birds, and mammals (CDC undated). The initial WNV ABM includes only three host species: black-capped chickadee, blue jay, and American crow. Additional host species will be included in future versions of the model.

The bird host model simulates bird reproduction, foraging, and interaction with mosquito agents. The user specifies a growth rate that does not consider impacts from WNV on the population. In general, bird reproduction is a function of the quality of the habitat in the surrounding area. The ABM assumes that bird reproduction is constant over the bird's reproduction period (i.e., from spring to early summer). Thus, the bird population can be defined as:

$$P(t+\Delta t) = P(t) * (1+r) * \Delta t \ , \tag{4}$$

where *P*, *r*, and *t* are population, reproduction rate, and time, respectively.

Within the bird host model, each bird is given a home location within the area from which all subsequent movements will occur. During daytime, birds leave their home locations to forage for food in the morning and return home at dusk. The model assigns a probability for the bird to travel from cell to cell while searching for food. Following the initial movement of the bird to another cell, the likelihood that the bird will move to another cell at the next time-step is a function of the habitat quality of that cell. The bird has a higher probability of continuing to move if the occupied cell has a low habitat quality. In other words, a bird is less likely to move from its present location to another location if it would experience a decrease in habitat quality. The distance the bird travels in each time-step is determined by the home range of the species multiplied by a movement factor that ranges from 0 to 1. The value of the movement factor is randomly selected for each time-step. For a simulation, the model centers a circular home range (the area to which an animal confines its normal activities) on a specific raster cell and then multiplies the radius of that home range by the movement factor. Thus, a movement factor of one results in the greatest distance traveled during the time-step. No bird is allowed to move outside its home range at any time. If a bird goes across the boundary of the area, a similar bird will enter the area from a randomly selected cell from among the boundary cells. The foraging movement of a bird is determined by the following equations:

$$X_{new} = X_{current} + R * \cos(\theta) , \tag{5}$$

$$Y_{new} = Y_{current} + R * \sin(\theta) \text{ , and} \tag{6}$$

$$\theta = \pi * \alpha / 180 \ , \tag{7}$$

where *R* is a random distance between zero and the radius of the maximum home range of the bird, α is a randomly selected angle between 0° and 180°, $X_{current}$ and $Y_{current}$ are the *X* and *Y* coordinates of the cell in which the bird is currently located, and X_{new} and Y_{new} are the *X* and *Y* coordinates of the cell to which the bird is moving. Figure 4 illustrates the movement of three hypothetical birds in a day in the study area.

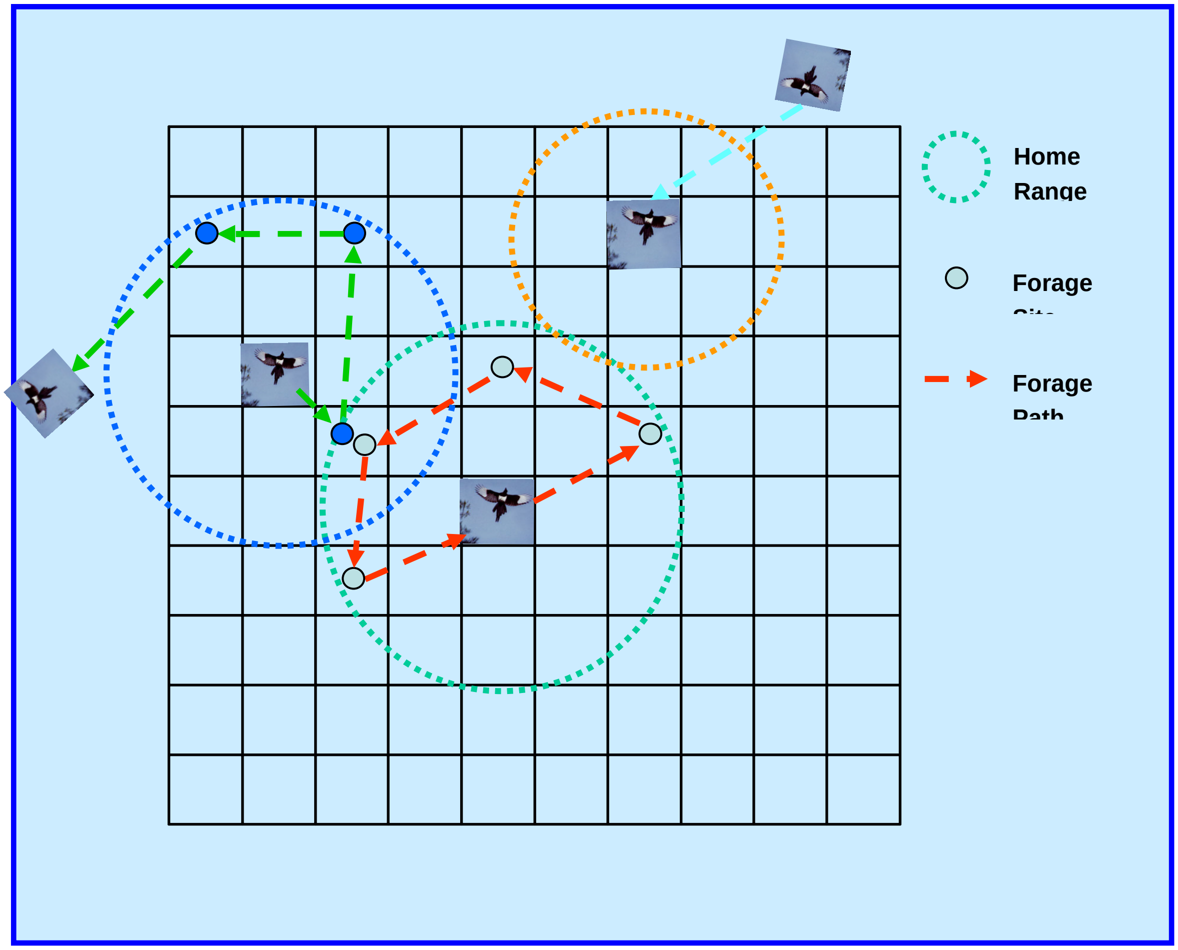

FIGURE 4 Schematic illustration of bird foraging movement

Human Agent

Humans may become infected with WNV when they are bitten by infected mosquitoes. About 6% of infected people will develop symptoms, and a small fraction of these individuals will die as a result of the infection (Hayse et al. 2005). The number of people who will become infected in cell NP_{inf} is estimated as:

$$NP_{inf} = P \times IR_p \times \frac{P_{m,inf}}{P_m} \times BR \ , \tag{8}$$

where P is the human population (density) in a cell, IR_p is the human infection rate, $P_{m,inf}$ is the population of infected mosquitoes in the cell, P_m is the total population of mosquitoes in the cell, and BR is the biting rate of mosquitoes.

The number of human deaths, NP_d, caused by WNV infection can be calculated as:

$$NP_d = NP_{inf} \times DR \, , \tag{9}$$

where *DR* is the death rate due to WNV infection.

The differences in the likelihood of becoming infected reflect differences in the time spent outdoors among the age groups and in the likelihood that individuals in each group will be outside during periods of greatest mosquito activity. Table 1 shows the age ranges of the different age groups used in the model.

In addition, the potential for a human to become infected and, if infected, to die is also a function of age (Hayse et al. 2005). In consideration of these facts, the model uses a different likelihood of infection and death. The infection and death rates for these age groups are calculated from the data published by Hayse et al. (2005).

It is also important to note that the probability for a person to be bitten by infected mosquitoes is extremely small while the person stays in the house. Therefore, we assume that people get mosquito bites only when they are participating in outdoor activities. For this reason, the model tracks the number of people and the time these people participate in outdoor activities.

Land Use

The distribution of agents, hosts, and habitats is a function of the land use in the area of interest. The 64-square-mile area encompassed by the WNV ABM consists of a heterogeneous mixture of 63 land-use types, including residential, commercial, industrial, agricultural, open space, wetlands, and water. Each raster cell may have a single land use or multiple land uses, depending on the location of the cell. The types of land uses within each raster cell are used to determine habitat quality for mosquitoes and birds, human population density, and the likelihood of human outdoor activity.

TABLE 1 Age groups and likelihoods of participating in outdoor activities

Group	Age	Likelihood of Participating in Outdoor Activities
1	0–1	These infants are very unlikely to participate in outdoor activities at the high mosquito blood-meal times, during dawn and dusk.
2	2–15	These children are often outdoors, especially in the afternoons/evenings when transmission is most likely.
3	16–54	People in this group in general spend most of their time in schools or offices and have stronger immune capability. They will hence have a lower probability of getting infected, assuming all people are at the same healthy condition.
4	55 and older	Seniors who do not need to go to work are more likely to participate in outdoor activities.

Mosquito and Bird Habitat Quality

Habitat quality plays an important role in the abundance and distribution of mosquitoes and birds; thus, it is an important parameter that governs the population dynamics of both of these biota. Each raster cell in the model is assigned a habitat quality value that may range from 1 (poor) to 3 (excellent) on the basis of the land use categories present within each raster cell.

The mosquito and the habitat quality values represent the perceived ability of the habitat to support high mosquito densities and are also used to set initial mosquito densities. The values are determined on the basis of land use. Each of the land-use categories identified for the study area is assigned a mosquito habitat quality value of 1 (poor), 2 (good), or 3 (excellent). Land-use categories (such as industrial or commercial) with little or no vegetative cover or surface water are assigned a mosquito habitat quality value of 1 (poor). Alternately, an open space land use (such as a forest preserve or wetland) is assigned a habitat quality value of 3 (excellent). In the model, the land use within each raster cell is identified, and the appropriate habitat quality value is assigned to that cell. For cells encompassing multiple land-use categories, the habitat quality value of a cell is calculated as the weighted mean of the land-use categories within the cell. A higher habitat quality rating indicates that the cell can support a higher mosquito density and reproductive rate and that it also exhibit a higher mosquito biting rate than can/does a cell with a lower habitat quality value.

Bird habitat quality values are determined in a manner similar to that used to characterize mosquito habitat quality. The habitat quality in each raster cell is then used to set the initial distribution and abundance of the bird hosts and to influence the movements of birds. At the start of a simulation, cells with higher habitat quality are assigned higher starting bird densities than cells with lower habitat quality. Bird habitat quality is then used to characterize the likelihood for a bird to move from one cell to another in the next time-step. If the habitat quality of the occupied cell is low, the likelihood that the bird in that cell will move to another cell is high, while if the habitat quality is high, the likelihood that the bird will move is low.

Human Population Density and Probability of Outdoor Activity

The likelihood of a human becoming infected with WNV is a function of the likelihood of a human being present in an area of mosquito abundance and the amount of time an individual would spend outdoors at that location, and both of these conditions are a direct function of land use. Each land-use category present in the study area is assigned a human population density value of 1 (low), 2 (intermediate), or 3 (high). For example, open spaces have a low population density (1), while a commercial shopping mall or office building may have a high density (3). Each land-use category is also assigned a value for its likelihood of outdoor human activity of 1 (low), 2 (intermediate), or 3 (high). For example, on land used for commercial or industrial purposes, most human activity would occur indoors, so the likelihood of outdoor activity is low (1). In contrast, most if not all human activity on a forest preserve or a golf course would occur outdoors; thus, the likelihood of outdoor activity in these areas is high (3). For cells encompassing multiple land-use categories, population density and the likelihood of outdoor activity are estimated as a weighted mean of the land-use-specific population density or outdoor activity likelihood values present within the cell.

Weather Dynamics

Weather is a critical component of the transmission of WNV. It affects all of the agents involved in the model, either directly or indirectly, and is especially important with regard to its effect on mosquito population dynamics and activity. The most important attributes of the weather data include:

1. Temperature,
2. Precipitation, and
3. Humidity.

These three attributes are spatially distributed and temporally variable, and they directly influence surface soil moisture, which is especially important in mosquito population dynamics. The WNV ABM uses real weather data for the study area for modeling mosquito population growth (see Equation 3).

When historic weather data are used, simulations may be conducted under three different precipitation regimes:

1. Dry (lower precipitation in 30% of years),
2. Wet (higher precipitation in 30% of years), and
3. Normal (middle-level precipitation in 40% of years).

The model includes several years of weather data for each precipitation regime. When a precipitation regime is selected for use in a simulation, the model randomly selects a year from the selected regime and uses the observed hourly weather data from that year for modeling mosquito population dynamics. In this way, the model incorporates actual yearly weather patterns to avoid using another model for generating weather data.

SUMMARY AND FUTURE RESEARCH

An ABM for simulating the dynamics of WNV transmission has been developed by using the Repast ABM development environment. The model simulates the distributions, behaviors, and population dynamics of mosquitoes and birds, and it simulates the interactions among these organisms and humans in a spatial and temporal manner.

It is important to point out that the current simulation is not a predictive model. It is a tool for researchers to use as a virtual laboratory to uncover the dynamic behaviors of the system as a result of the interactions among individual agents in the system. The accuracy and usefulness of the model depend heavily on the quantitative definitions of the behaviors of, and interactions among, the individual agents in the model. For an ABM of this type, extensive field research data are needed in order to calibrate it. Using these kinds of data for calibration is an important next step for this research.

The current implementation of the model includes only three types of representative birds as the hosts. In the next phase of our research, we will expand the model to allow for multiple host species, such as mammals and reptiles, as well as multiple species for both mosquitoes and hosts.

ACKNOWLEDGMENT

This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Bernard, K.A., et al., 2001, “West Nile Virus Infection in Birds and Mosquitoes, New York State, 2000,” *Emerging Infectious Diseases* 7(4), Aug.

CDC (Centers for Disease Control and Prevention), undated, *West Nile Virus*, Division of Vector-borne Infectious Diseases; available at http://www.cdc.gov/ncidod/dvbid/westnile.

Ferber, J., 1999, *Multi-agent Systems: An Introduction to Distributed Artificial Intelligence*, Addison-Wesley, Feb. 25.

Hayse, E.B., et al., 2005, “Epidemiology and Transmission Dynamics of West Nile Virus,” *Emerging Infectious Diseases* 11(8), Aug.

Kohler, T.A., et al., 2000, *Dynamics in Human and Primate Societies: Agent-based Modeling of Social and Spatial Processes,* Santa Fe Institute Studies in the Sciences of Complexity Proceedings, Oxford University Press, Feb. 1.

Lanciotti, R.S., et al., 1999, “Origin of the West Nile Virus Responsible for an Outbreak of Encephalitis in the Northeastern United States,” *Science* 286, Dec. 17.

Marra, P.R., et al., 2004, “West Nile Virus and Wildlife,” *BioScience* 54(5):393–401.

Neergaard, L., 2005, *U.S. Braces for West Nile Virus Season*, Associated Press, May 31.

Repast 3, undated; available at http://repast.sourceforge.net/.

Ruiz, M.O., et al., 2004, “Environmental and Social Determinants of Human Risk during a West Nile Virus Outbreak in the Greater Chicago Area,” *International Journal of Health Geographics* 3(8).

Woolridge, M., 2002, *Introduction to Multiagent Systems,* John Wiley & Sons, June 12.

MULTI-AGENT MODELING AND ANALYSIS OF THE BRAZILIAN FOOD POISONING SCENARIO

V. MYSORE,* O. GILL, R.S. DARUWALA, M. ANTONIOTTI, and
B. MISHRA, Courant Institute, New York, NY
V. SARASWAT, IBM T.J. Watson Research Center, Yorktown Heights, NY

ABSTRACT

The multi-agent modeling and analysis of catastrophic events raise many challenging problems since they involve a large, interacting, mobile population with complex behaviors. This research aims to address these problems through the analysis of simulations and to aid planning efforts for future catastrophic events through parameterized stochastic models covering the health care providers, emergency responders, and affected population. As a test case, we examine the massive outbreak of *Staphylococcus aureus* food poisoning that occurred in Minas Gerais, Brazil, in 1998 to demonstrate and evaluate our tools and techniques. In this incident, 8,000 people consumed contaminated food at a priest's ordination. Of these, 81 were admitted to intensive care units of 26 local hospitals after a triage, and 16 of them eventually expired. We capture the dynamics of such an outbreak by using two kinds of abstract agents — *hospital* and *person*, further augmented with *information* and *communication channels*. Hospital locations and current capacities are broadcast by the hospital to its patients and to persons with a radio and subsequently exchanged between neighboring persons. This "outbreak" model has been implemented in the Java version of Repast 3.0. Most attributes are scaled to be in the range of 0 to 1, with most behavior being probabilistic. We document the relative performance of the different simulations by using a range of parameter values for communication channels, personalities, and triage policies, to understand their combined effect on the overall survival rates. We also introduce the XSSYS trace analysis and model checking tool for answering complex temporal logic queries over Repast traces. We discuss how such simulation-based analysis can become a rigorous tool in aiding public health policy planning.

Keywords: Social simulation, catastrophe preparedness, emergency response, Repast

INTRODUCTION

The computer modeling and simulation of catastrophic scenarios, when enhanced with sophisticated automated reasoning, promise to be a very valuable tool for developing public health policies and disaster management strategies. In the horrific wake of Hurricane Katrina that ravaged the State of Louisiana, it became doubly shocking as word spread very rapidly about the computer models that had accurately predicted many of the ramifications of such a disaster. Indeed, the Center for the Study of Public Health Impacts of Hurricanes of Louisiana State University had conducted extensive research on this topic and constructed elaborate models of such a scenario (see Heerden and Binselam 2004). While it is much less likely that other

* *Corresponding author address*: Venkatesh P. Mysore, Courant Institute, 715 Broadway #1006, New York, NY 10012; e-mail: mysore@cs.nyu.edu.

simulation efforts can achieve such predictive fidelity, most catastrophe simulation projects (e.g., SEAS project of Chaturvedi et al. 2003, Project RESCUE of Mehrotra et al. 2004, and VISTA tool of Louie and Carley 2004) still focus on one of two nonoverlapping goals: *disaster prediction* and *disaster management*. In this paper, we do not even broach disaster prediction; instead, we focus on the analysis of simulations to aid planning efforts for future catastrophic events. We are part of the Large Scale Emergency Response (LaSER) research group of the New York University (NYU) Center for Catastrophe Preparedness and Response (CCPR), which is a partnership with the U.S. Department of Homeland Security and its Office for Domestic Preparedness. Catastrophe preparedness involves stocking and distributing resources to minimize fatalities, planning an emergency response strategy, and educating the general population. These desiderata will dictate, among other things, the distribution and use of available resources, and the means and nature of the information and instructions provided to the health care providers, emergency responders, and affected population (see Lasker 2004). This paper deals with these issues through the multi-agent modeling of catastrophic events that involve a large, interacting, mobile population with complex behaviors and goals.

We use the massive outbreak of *Staphylococcus aureus* food poisoning that occurred in Minas Gerais, Brazil, in 1998 (Do Carmo et al. 2004) to demonstrate and evaluate our tools and techniques. Although the fraction of fatalities (16/8,000) may not be regarded to be of catastrophic magnitude, the scenario is ideal for observing the effects of different instructions and policies on the behavior of the large affected population and the medical facilities. We capture the dynamics of such an outbreak by using two kinds of abstract agents — *hospital* and *person* — enhanced with *information* and *communication channels*. After exploring a number of simulation systems, this "outbreak" model has been implemented in the Java version of Repast (Collier et al. 2005). Most attributes are in the range of 0 to 1, with most of the behavior governed by random-number-based probabilities. We document the relative performance of the different simulations by using a range of parameter values for communication channels, personalities, and triage policies, to understand their combined effect on the overall survival rates. We also introduce the XSSYS trace analysis and model checking tool (Antoniotti et al. 2003) developed in our laboratory and show how it can answer complex temporal logic queries over Repast traces. We conclude by suggesting how such a schema provides a reasonable way of modeling, simulating, and analyzing other catastrophic scenarios as well.

BRAZILIAN OUTBREAK

In 1998, a massive outbreak of *Staphylococcus aureus* food poisoning occurred in the rural town of Minas Gerais, Brazil, where around 8,000 individuals attended a Catholic priest's ordination. The trace-back investigation implicated food preparers, who were culture positive for enterotoxigenic *Staphylococcus aureus*, as the source of contamination. However, it was the improper storage temperature of the food, which was prepared 2 days in advance, in the summer weather that allowed the optimal growth of bacteria and production of *Staphylococcus enterotoxin* (SE). Symptoms like intense nausea, emesis, diarrhea, abdominal pain, prostration, and dizziness were pronounced in less than 4 hours after consumption of the contaminated food in about half the population (~4,000). Almost half of them (~2,000) decided to proceed to one of the 26 nearby hospitals without letting the situation exacerbate further. However, this overwhelmed their emergency departments, forcing a triage. A triage, in medical parlance, refers to a set of policies to partition the vast number of patients into different groups (e.g., those requiring immediate intensive care, those requiring general hospitalization, and those requiring only medication or saline). This process helps the hospital distribute the available resources optimally under the time constraints imposed by the prognosis of the disease. In Minas Gerais, 396 (~20%) people required admission after triage, and of these, 81 (~20%) required admission to the intensive care unit (ICU). Patients with improving health were discharged from the ICU within 7–10 days. A total of 16 (~20%) patients subsequently developed irreversible multi-system shock and expired while hospitalized. While people of all ages (1–86) attended the ordination, the 16 fatalities occurred only in the oldest (65 and above) and the youngest (5 and under) groups. The sex of the individual was found to have no influence on the clinical outcome among those treated in the ICU.

MULTI-AGENT OUTBREAK MODEL

We capture the dynamics of such an outbreak by using two kinds of abstract agents: *hospital* and *person*. A hospital is an abstraction of any medical facility accessible in the area (26 in the Brazilian case), while a person is an abstraction of any individual who consumed the contaminated food (8,000 in the Brazilian case). The effect of the general population who did not attend the ordination is not modeled in our simulation. The model is then enhanced with *information* and *communication channels*, with the two vital pieces of information being the locations of the hospitals and their current capacities.

Food Poisoning

The food poisoning is modeled by functions that describe the time variation of the person's health, with and without treatment. Effectively, any "disease" can be modeled in terms of the (possibly time-varying) amount by which the affected agent's "health" can deteriorate or recover with and without treatment, at each time-step of the simulation. The individual's resistance or susceptibility to the specific disease is captured by a personalized variable, which modifies the disease-health-treatment functions. This can be used to abstract factors such as age, sex, health condition before food consumption, and genetic makeup. Probabilities are introduced to capture unpredictability and variability in real situations. We can use this simple but effective abstraction to model other conditions, such as Sarin gas attacks, radiation exposure, etc. Since the initial amount consumed and the dose/response relationship in human oral exposure to SE are unknown, the initial health of each person is assumed to be a random value in a meaningful range.

People's Behavior

The *persons* move toward their place of work from the site of food poisoning. Depending on their deteriorating health level and personality parameters, they choose to go to the one *hospital* they are initially aware of. Additional information is acquired by talking to neighboring agents. A time stamp of the information is maintained, so the persons update their knowledge only if more current information is available. Further, some persons are equipped with radios, which give them access to the current information about all the hospitals. People recompute the destination hospital toward which they should be moving on the basis on the distance to and the believed current capacity of each medical facility they are aware of. In addition, they always move toward the nearest free hospital, unless they are very sick and opt to go to the nearest hospital, even if it is full. The complexity of the model is increased further with *personality* parameters, which capture whether an agent chooses to go to a hospital, talk to neighbors, accept the new information, or recompute the best hospital. *Group behavior* is captured by letting adjacent people moving toward the same destination wander less.

Hospital Behavior

The *hospital* aims to admit every *person* who reaches its premises and invests its resources in the order of their admittance and proportional to their ill health. Hospital resources, consisting of infrastructure, beds, nurses, and doctors, are recovered when a patient is discharged or deceased; medical supplies, like drugs and saline, are irrecoverable. The hospitals also perform a local broadcast of complete current information to all persons who are admitted or waiting at their facility. The hospital model is enriched by identifying three different modes of operation — full, critical, and available — corresponding to the current amount of resources. With the *triage* policy in place, the hospital agent handles admitted persons as before. However, it admits new persons only if it has resources to spare (available mode). If it is operating in the critical mode, it admits only critically ill persons. No new persons are admitted in the full mode. With the *transfer* policy in place, admitted patients who have recovered reasonably are discharged earlier than usual and instructed to go to a different hospital if symptoms recur. In their place, critically ill persons who are waiting are admitted. Probabilistic parameters are used to capture the policies that govern the hospital's decisions on when to admit a new patient, in which order to treat the admitted patients, when to transfer a recovering patient to a nearby hospital, and which critically ill patient to admit in the vacancy created.

ANALYZING THE OUTBREAK

Since the modeled system involves a large number of agents, uses a vast number of parameters, and attempts to capture the stochastic nature of the infection and behavior, traditional symbolic or algebraic analyses are not immediately possible. Instead, the analyst must resort to simulation-based analysis to obtain average performance statistics over a large number of trials. Combined with individual inspection of a small number of characteristic traces, evaluation of the relative merits of different emergency response strategies becomes possible. We use the statistics-based analysis tools provided by Repast and introduce the temporal logic trace analysis tool XSSYS.

Numerical Results

Since the most significant aspect of the model is its extreme sensitivity and unpredictability, general average/comparative trends (as opposed to absolute values) in the death rate can be used to observe the effect of variations in parameters of interest (with the other dimensions fixed at justifiable values). We obtain trends (typically averaged over three runs) around the Brazilian scenario with 8,000 people and 26 hospitals leading to a death rate of 0.2%.

Effect of Hospital Resources, Communication, and Grid-Size on Death Rate

We first observe how the number of deaths varies with hospital resources (Figure 1). Shown there are the plots for a 250 × 250 grid and an 800 × 800 grid, with communication enabled and disabled, and with no triage policy implemented. From the plot, we observe that the number of deaths clearly declines when hospitals have more resources, since each hospital is able to allocate more resources (treatment) per person. Also note that in a small grid (250 × 250), where hospitals have few resources, communication works against our model. This is because people converge to the nearest hospitals, exhausting their resources quickly. By the time the hospital runs out of resources and turns people away, they are too sick to survive a trip to the next hospital. However, when the hospitals have plenty of resources, the difference in survival rates is negligible when communication is used versus when it is not. In the 800 × 800 cases, the difference in distances between the closer and farther hospitals is much greater. Hence, it works to a person's benefit to communicate and obtain information about nearby hospitals.

Effect of Number of Hospitals, Triage, and Grid Size on Death Rate

Next we analyze how the number of hospitals affects the number of deaths by using plots (Figure 2) for a 300 × 300 grid and an 800 × 800 grid, with the triage policy enabled and disabled. We first note the expected phenomenon: increasing the number of hospitals decreases the death rate, since there are fewer patients per hospital. We also note a slightly higher death rate when the grid size is larger because the average distance to a hospital is longer: people reach the hospitals when they are sicker, and more persons are not able to survive the journey. More important, this figure leads to a dramatic conclusion: the triage policy, as interpreted in the model, always works against the people. The failure of the triage policy can be attributed to a key aspect of the food-poisoning health function: a healthy person is just as likely to worsen as an already unhealthy person. Thus the patients who were discharged slightly early because of the critically ill people who were waiting end up falling sick again, and the critical ill persons themselves seldom recover. Second, the health of people who are refused admission (because they are not critically ill or because the hospital is full) worsens during their trip to a different hospital. The net effect is that the hospitals have to treat sicker people. This suggests that it is wiser for people to reach the nearest hospital, and then for the hospitals to have a system of redistributing their resources (i.e., moving equipment and doctors, as opposed to moving patients).

Effect of Number of People, Grid Size, and Initial Pattern on Survival Rate

Next, we observe how the number of people affects the fraction of people who survive by using plots (Figure 3) for a 300 × 300 grid and an 800 × 800 grid, with communication enabled and the triage policy disabled. We also inspect the effect of people starting at random positions in the grid as opposed to being concentrated at a location. From these plots, we again observe the expected trend: as the number of people increases, the fraction of people who survive declines. Similarly, the 800 × 800 grid results in a slightly larger percentage of the people dying because the average distance to the nearest hospital is longer. The difference in survival percentages for the concentrated and the random initial positions is not statistically significant. This can be understood as the average person's starting point's distance to the nearest hospital being roughly

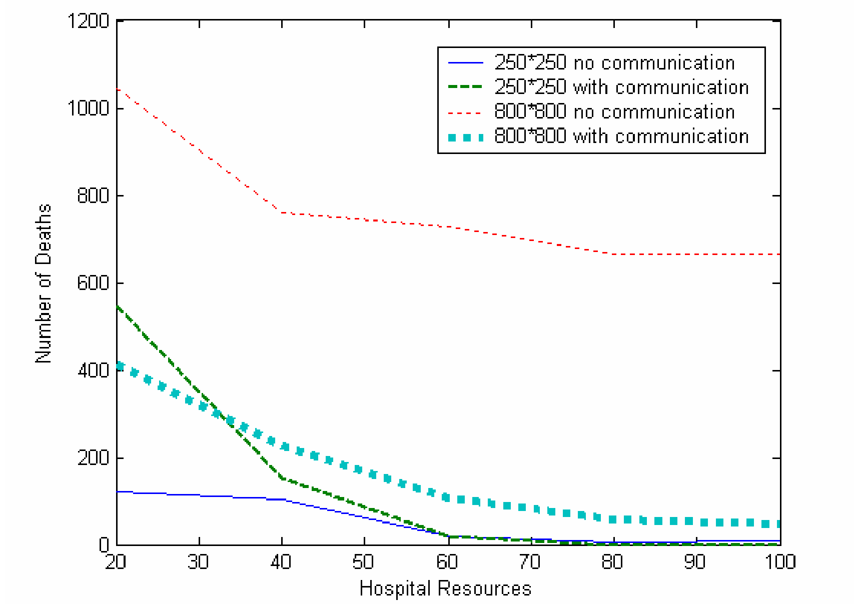

FIGURE 1 Effect of hospital resources, communication, and grid size on death rate

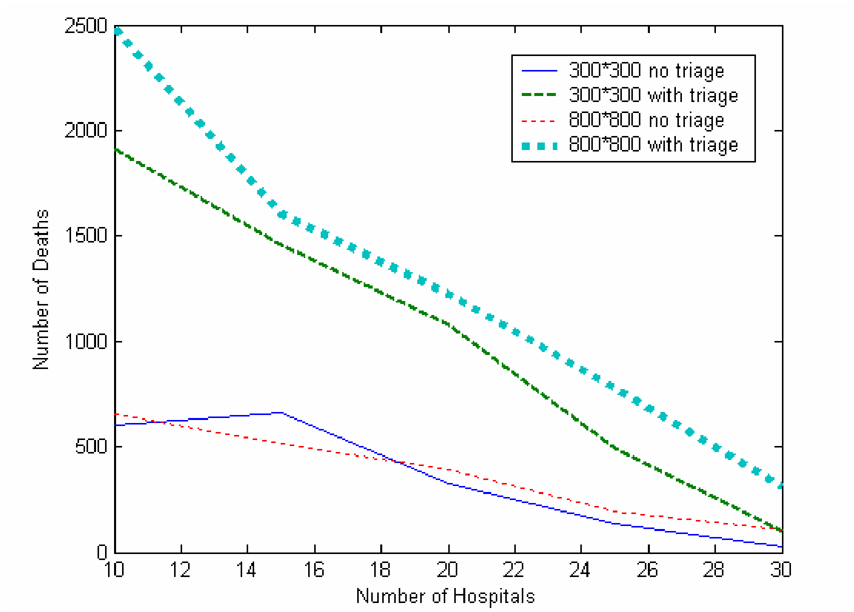

FIGURE 2 Effect of number of hospitals, triage, and grid size on death rate

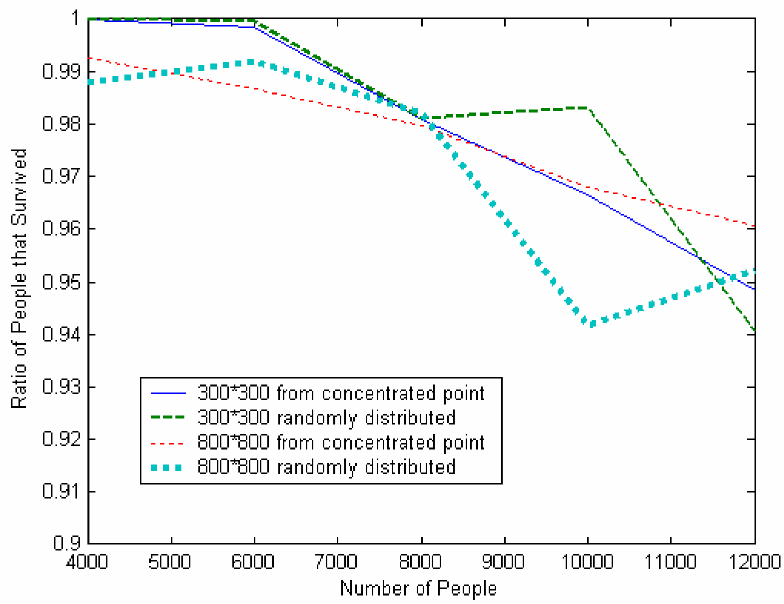

FIGURE 3 Effect of number of people, grid size, and initial pattern on survival rate

the same in both cases. However, the number of initial neighbors in the distributed case must be sufficient to supply the required information about the nearest hospitals.

Trace Analysis in XSSYS

The XSSYS temporal logic trace analysis system can answer linear temporal logic (LTL) queries about the time course behavior of a set of traces. It was developed originally as a part of Simpathica for simulating and analyzing biochemical pathways. XSSYS allows the user to formulate queries about multiple traces in temporal logic or English (via a natural language interface*)*. The person and hospital traces of Repast can be read by using XSSYS. These traces reveal very insightful aspects of the behavior of persons and hospitals and serve as a good starting point for coming up with new policies to be tested. Complex temporal queries linking different traces can help in discovering finer truths about the underlying dynamics of the system. In this section, we demonstrate the XSSYS trace analysis tool in some simple examples.

Time-trace of a Person

The variation of a person's health with time (in this case, *Person-78*) during the course of a simulation is plotted in Figure 4. XSSYS plots this curve by using data imported from Repast in the *btd* format by using the *PtPlot* tool. In addition to the health level (*HealthLevel)*, the person's current location *(x, y)* and destination *(destx, desty)* are plotted. To indicate when the person actually received treatment, a Boolean value *admitted* is also plotted.

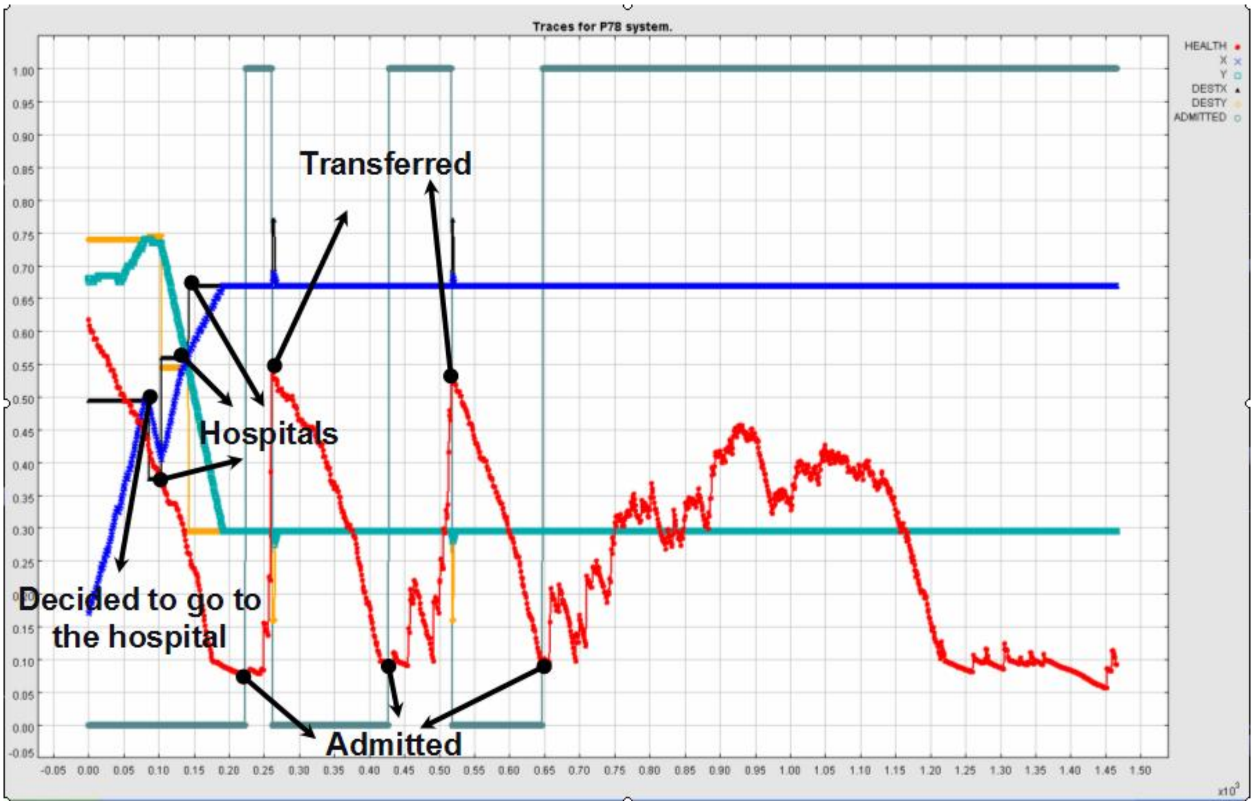

FIGURE 4 Time-trace of a person

Time-trace of a Hospital

In the case of a hospital, we plot the depletion of resources (*HospitalResources)* with time (Figure 5). The number of people admitted and the number of people waiting indicate the stress on the hospital (in this case, *Hospital-1*). The successful creation of vacancies by early discharge and their filling by critically ill persons awaiting treatment are also presented.

Temporal Logic Analysis

Temporal properties of these traces can be analyzed by formulating queries in linear temporal logic by using the operators *Eventually* (sometime in the future) and *Always* (henceforth in the future). In the specific case being demonstrated (Figure 6), the traces of *Person-13* and *Person-113* are being compared. *Person-113* is seen to have a consistently better *HealthLevel* than *Person-13*, although both their *HealthLevels* are dropping. *Person-113* is also seen to have reached the destination hospital, while *Person-13* has not.

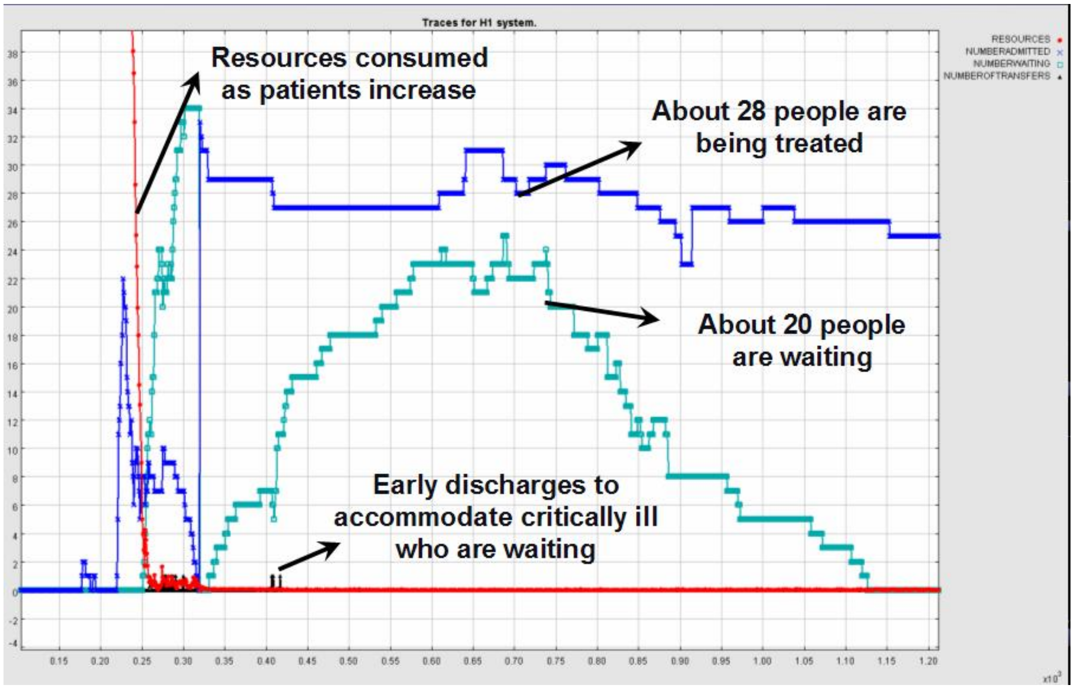

FIGURE 5 Time-trace of a hospital

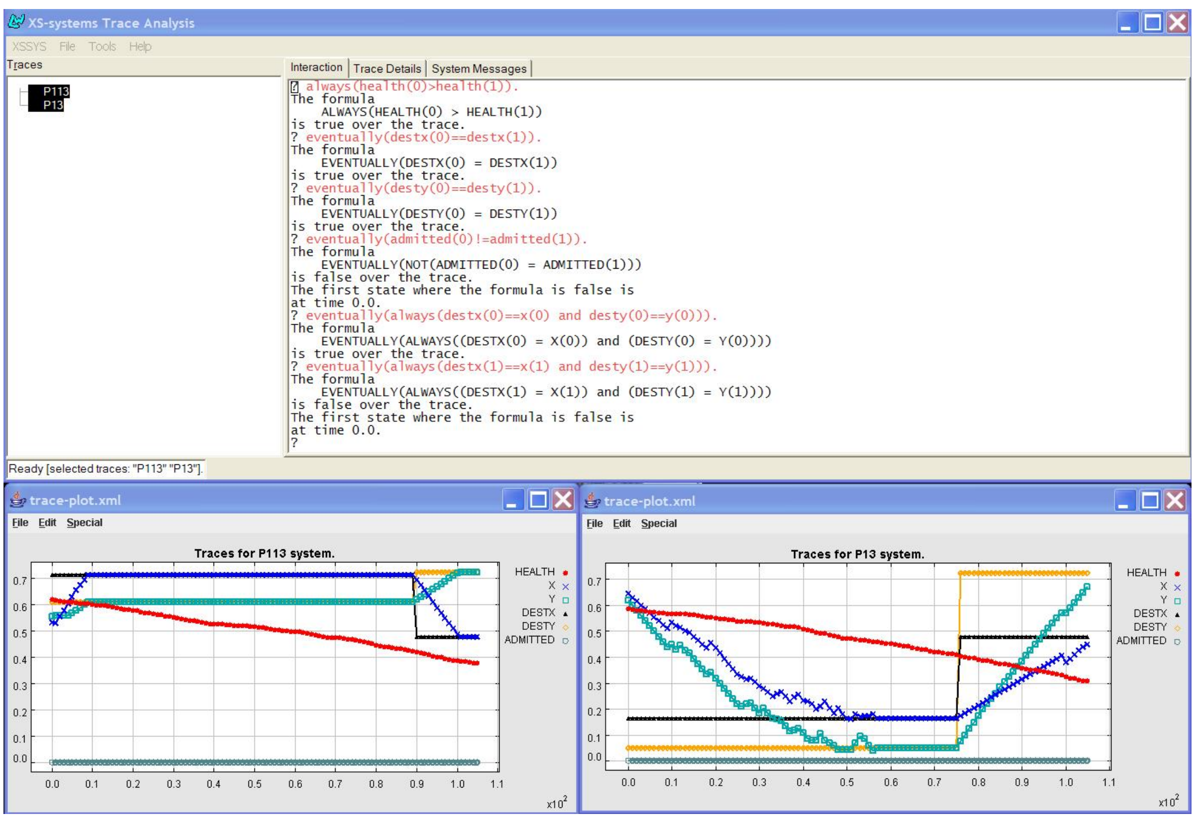

FIGURE 6 Temporal logic analysis in XSSYS

DISCUSSION

The Brazilian food-poisoning scenario proved to be a considerably complex problem, which had all the essential elements of a typical catastrophic scenario: a large number of agents (8,000 + 26), agents of different types (persons and hospitals), external factors governing the time evolution of the agents' features (effect of food poisoning on health), mobility (persons), mutual interaction (within persons, and between persons and hospitals), and multiple communication channels (talking, broadcast, and radios).

Repast proved to be sufficient to model and simulate the Brazilian food poisoning scenario. The analytical capabilities were enhanced by feeding its output to XSSYS. Despite the extreme parameter sensitivity of the model, we were able to explore the effectiveness of different emergency response strategies and catastrophe preparedness policies. The complexity and unpredictability of the model, because of the vast number of parameters, became apparent very quickly. Our model was able to capture the reported statistics to a reasonable extent, and it elucidated different conditions that could have led up to them. Factors that could have increased or decreased the number of fatalities also became evident. More specifically, the results showed that the distance the people need to travel to reach the hospital greatly determines how many people survive. We also observed that the survival rate increases when either the resources each hospital has or the number of hospitals increases, and that the survival rate decreases when the number of people increases. When the average distance to the nearest hospital is almost the same, there is almost no difference in survival rates between concentrated and random initial patterns. We found that communication among people about hospitals is beneficial when the difference in distances to hospitals is substantial, but it is harmful when all hospitals are close by and have few resources. We also found that our triage system harms the survival rate, since it is better to keep patients at a hospital, even if it has low resources, rather than have them transfer to another hospital and then having to treat a sicker person. The emergence of such interesting unanticipated behaviors already suggests a potential utility of such simulation-based analysis tools.

Many additional enhancements to the outbreak model to make it more realistic are possible. We might need to switch the environment to a real city. Transportation constraints and modes, roads, subways, and other geographical information might need to be incorporated. The moment these additional constraints emerge, we will need to model the agent's transportation choices. For example, Raney and Nagel (2004) describe a framework for running large-scale multi-agent simulations of travel behavior on the basis of each agent's "plan" of activities, times, and preferred modes of transport. However, as described by Sono and Ishibashi (2004), the change in the transportation choices after a disaster will need to be worked into the plan, with commuters and noncommuters having to be treated differently (a rather simple situation, which nonetheless seems to have had a major impact in the Katrina disaster). A somewhat complex model of this nature will endow each agent with a current-mobility variable, which decreases with a decrease in the agent's health, increases if the agent is being helped by a neighbor, and decreases if the agent is helping a neighbor.

We will need to add social networks at various levels (families, friends, etc.) and the social characteristics of subsets of the population to model the cultural differences in response behavior. A good example of the application of social judgment theory appears in the work on group attitude emergence via assimilation and contrast effects as described by Jager and Amblard (2004). The benefits of cooperation could be captured by increased mobility and information, while moving in groups. We could also add social infrastructure, like first responders, volunteer-based relief organizations, and law enforcement officers. Also, some of the people who consumed the contaminated food could belong to these groups, thus complicating the interaction dynamics even further.

We could also add more detailed models of communication and information exchange. For instance, the logic-based framework for handling messages and belief-state changes discussed by Perrussel and Thevenin (2004) could be combined with ideas from the work on the geographical divergence of knowledge via interactive-learning-based diffusion by Morone and Taylor (2004). This could prove useful in capturing the realistic transmission and accumulation of information during calamities. We could incorporate into the model long-distance 1-to-1 and 1-to-many communication channels, where 1-to-1 channels are between persons via cell-phone and 1-to-many are from authorized broadcasters to equipped receivers. We could model the ability to give instructions and the ability to receive instructions separately. Similarly, there could be a difference in the transmission of different kinds of information (e.g., the location of the nearest hospital, measures to use to slow down the progression of the sickness, instruction to proceed to a hospital). (See the work of Lawson and Butts [2004] on the propagation of rumors and information in crisis contexts.)

The food poisoning in itself could have been modeled differently. For instance, the spread of *Mycoplasma pneumoniae* via interaction between patients and caregivers is modeled by using network theory by Meyers et al. (2003). Similarly, Rahmandad and Sterman (2004) analyze the pros and cons of agent-based modeling versus differential equation modeling for contagion modeling. Although the work of Eidelson and Lustick (2004), who developed a stochastic agent-based model, VIR-POX, to explore the viability of available containment measures as defenses against the spread of smallpox, is similar to the Brazilian scenario analysis, it is different in its approach and goals.

On the pure computational side, the biggest challenge is in scaling up to a very large-scale simulation through parallelization, abstraction, hierarchy, and other strategies. We are working on enhancements to XSSYS to improve its expressivity and power. We also need to investigate the applicability of other formal reasoning techniques, such as probabilistic reasoning (Xiang 2002) and probabilistic argumentation systems and causal analysis (see the WIZER tool of Yahja and Carley 2004). We could treat the estimation of the triage policy parameters (e.g., the health level at which a person who is waiting gets deemed as critically ill, or the health level at which a recovering patient may be discharged to create a vacancy) as an optimum-value computation problem. From a practical utility point of view, we need to identify a way of describing the simulations in a manner that is formal and accurate enough to create a meaningful simulation but simple enough for a nonprogrammer to use. We are in the process of compiling a survey of approaches to model and analyze catastrophic scenarios. Our goal is to first extend this modeling and analysis approach from the Brazilian food-poisoning example to other scenarios. For example, the effects of several people independently consuming botulinum-contaminated milk at their homes (following the scenario investigated by Wein and Liu [2005]) could be modeled by a different health-modulation curve, and with people starting at their homes as opposed to congregating at a church. We would additionally need to model the transmission of the instruction to not consume any more contaminated milk. Eventually, we hope to develop and demonstrate the tools and technologies necessary for such simulation-based analysis to provide a rigorous yet user-friendly approach for exploring assumptions about public health policies in catastrophe preparedness and emergency response.

ACKNOWLEDGMENT

The research reported on in this paper was funded by a LaserMed grant from NYU's CCPR.

REFERENCES

Ancel Meyers, L., M.E. Newman, M. Martin, and S. Schrag, 2003, "Applying Network Theory to Epidemics: Control Measures for *Mycoplasma pneumoniae* Outbreaks," *Emerging Infectious Diseases* 9(2):204–210, Feb.

Antoniotti, M., F. Park, A. Policriti, N. Ugel, and B. Mishra, 2003, "Foundations of a Query and Simulation System for the Modeling of Biochemical and Biological Processes," in *The Pacific Symposium on Biocomputing: PSB 2003*, (Editors R.B. Altman, A.K. Dunker, L. Hunter, T.A. Jung, and T.E. Klein), pp. 116–127, World Scientific, Jan.

Chaturvedi, A., S. Mehta, and P.L. Drnevich, 2003, "Measured Response 2003," invited white paper for the Purdue e-Business Research Center (PERC) and Homeland Security Institute meeting.

Collier, N., T. Howe, R. Najlis, M. North, and J.R. Vos, 2005, "Repast 3.0: Recursive Porous Agent Simulation Toolkit;" available at http://repast.sourceforge.net/.

Do Carmo, L.S., C. Cummings, V.R. Linardi, R.S. Dias, J.M. De Souza, M.J. De Sena, D.A. Dos Santos, J.W. Shupp, R.K. Pereira, and M. Jett, 2004, "A Case Study of a Massive Staphylococcal Food Poisoning Incident," *Foodborne Pathogens and Disease* 1(4):241–246, winter.

Eidelson, B.M., and I. Lustick, 2004, VIR-POX: An Agent-based Analysis of Smallpox Preparedness and Response Policy," *JASSS* 7(3), June 30.

Jager, W., and F. Amblard, 2004, "A Dynamical Perspective on Attitude Change," NAACSOS Conference, Pittsburgh, PA, June.

Lasker, R.D., 2004, "Redefining Readiness: Terrorism Planning through the Eyes of the Public," The New York Academy of Medicine, New York, NY.

Lawson, G.C., and C.T. Butts, 2004, "Information Transmission through Human Informants: Simulation," *CASOS'04.*

Louie, M.A., and K.M. Carley, 2004, "VISualization of Threats and Attacks (VISTA): A Decision Support Tool for Urban Threat Environments," *CASOS'04.*

Mehrotra, S., C. Butts, D. Kalashnikov, N. Venkatasubramanian, R. Rao, G. Chockalingam, R. Eguchi, B. Adams, and C. Huyck, 2004, "Project RESCUE: Challenges in Responding to the Unexpected," *SPIE*, Vol. 5304.

Morone, P., and R. Taylor, 2004, “Small World Dynamics and the Process of Knowledge Diffusion: The Case of the Metropolitan Area of Greater Santiago De Chile,” *Journal of Artificial Societies and Social Simulation*, Vol. 8.

Perrussel, L., and J.M. Thevenin, 2004, “A Logical Approach for Describing (Dis)belief Change and Message Processing,” in *Proceedings of the Third International Joint Conference on Autonomous Agents and Mutli-Agent Systems.*

Rahmandad, H., and J. Sterman, 2004, “Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-based and Differential Equation Models,” *CASOS’04.*

Raney, B., and K. Nagel, 2004, “An Improved Framework for Large-scale Multi-agent Simulations of Travel Behavior,” in *Proceedings of Swiss Transport Research Conference (STRC),* Monte Verita, Switzerland; available at www.strc.ch.

Sono, K., and K. Ishibashi, 2004, “Constructing Commuters’ and Non-Commuters’ Action Choice Model after a Great Earthquake,” *CASOS’04.*

van Heerden, I.L., and A. Binselam, 2004, “Hurricane Pam Exercise — July 2004, Adcirc Storm Surge Animations and Related Information,” Center for the Study of Public Health Impacts of Hurricanes; available at http://hurricane.lsu.edu/floodprediction/PAM_Exercise04/Introduction/Introduction.htm.

Wein, L.M., and Y. Liu, 2005, “Analyzing a Bioterror Attack on the Food Supply: The Case of Botulinum Toxin in Milk,” *PNAS* 102(28):9984–9989, July 12.

Xiang, Y., 2002, *Probabilistic Reasoning in Multiagent Systems: A Graphical Model’s Approach*, Cambridge University Press.

Yahja, A., and K.M. Carley, 2004, “The Design of Automated Validation and Explanation for Large-scale Social Agent Systems,” *CASOS’04.*

POLITICIANS, BUSINESSMEN, WARRIORS, AND CIVILIANS: ANALYZING THE COMPLEXITY OF THE IRAQ CONFLICT

A. GELLER,* Military Academy at the ETH Zurich, Zurich, Switzerland
P. JERMANN, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

ABSTRACT

How can contemporary conflicts (e.g., like those experienced in Afghanistan, Chechnya, Colombia, Sierra Leone, or Iraq) be more adequately described? An intuitive WarSocietyModel (iWSM) ontology (which was subsequently case-study-validated) is suggested as a starting point. The implementation of iWSM into an agent-based model — WSM — should lead to a better understanding of contemporary conflicts in general and the Iraq conflict in particular. We find that political heterogeneity in a contemporary conflict setting may be more prone to violence than is a less fragmentized political landscape. Our results also suggest that neither greed nor violence contributes independently from the other to a contemporary conflict's outcome. Our data appear to be congruent with the common notion that the civilian population bears the brunt of contemporary conflicts. From a methodological perspective, our research suggests a qualitative but, at the same time, formalized procedure.

Keywords: Agent-based modeling, complex systems, social simulation, contemporary conflict, Iraq

INTRODUCTION

The armed hostilities in Afghanistan, Chechnya, Colombia, Sierra Leone, and Iraq have epitomized the nature of contemporary conflicts. While some scholars relate contemporary conflicts to concepts such as new barbarism (van Creveld 1992; Huntington 1993) or economization (Collier and Hoeffler 1998, 2004; Jean and Rufin 1999), others suggest that they are driven by more complex mechanisms (Bayart et al. 1999; Kaldor 1997; Reno 1998; Richards 1996). We tend to support the second argument and agree with the notion that research in this area must be based on the findings of field research and case studies. Gerring (2004, page 342) defines a case study as "an intensive study of a single unit for the purpose of understanding a larger class of (similar) units." We believe that a partial formalization of what is perceived of as contemporary conflicts into a model is possible and, from an epistemological point of view, is essential.

An agent-based modeling approach allows for both the implementation of qualitative data and the formalization of these data. Conservatively speaking, an inductive proceeding avoids the pitfall of prejudice, sheds light on the "veritable" character of the research subject, and highlights the necessity of interdisciplinary research. A deductive proceeding, on the other hand, allows us to test hypotheses and theory building. Furthermore, it assures the accurate embedding of the research subject into a broader context. However, in agent-based modeling, it is more suitable to

* *Corresponding author address*: Armando Geller, Military Academy at the Swiss Federal Institute of Technology (ETH) Zurich, Steinacherstr. 101B, 8804 Au (ZH), Switzerland; e-mail: armando.geller@milak.ethz.ch.

speak in Axelrod's terms (1997, page 24) of "the third way of doing science." We believe that the third way allows us to proceed creatively and exploratatively. This may solve some of the material problems of analyzing contemporary conflicts and enhance our understanding of them, both in general and specifically in the case of Iraq.

We first examine the notion of contemporary conflicts. Then we introduce the intuitive WarSocietyModel (iWSM)[1] in an attempt to validate it qualitatively in the case of Iraq.[2] Finally, we implement this into an agent-based architecture for further analysis.

CONTEMPORARY CONFLICTS, iWSM, AND THE IRAQ CONFLICT

Contemporary Conflicts

While van Creveld (1992) has postulated the end of the Clausewitzian wars and Jean and Rufin (1999) have compiled a seminal book on the economies of civil wars, it was Kaldor (1997) who introduced a vague but nevertheless insightful conceptualization of contemporary conflicts. Her notion — developed in a comprehensive case study of the conflict in Bosnia — has gained broad influence. We believe she has caught the main characteristics of the subject.[3]

Kaldor (1997, page 7) defines (contemporary) conflict "as conflict between politically organized groups involving large-scale violence." States do not matter anymore; conversely, their disintegration leads to the diminishing of state borders, interdependence of internal and external actors, and emergence of new nonstate actors. Furthermore, shadow economies develop where property rights cannot be guaranteed anymore — an anomic space evolves. While it would be misleading to underestimate the greed factor in contemporary conflicts (Collier and Hoeffler 2004), one of the main driving forces in a contemporary conflict setting is what Kaldor (1997) refers to as the "politics of identity." By this, she means the manipulation of groups by leading actors with a repertoire of "real" or invented cohesion-generating categories, such as ethnicity, religion, or tribe.

The intuitive WarSocietyModel (iWSM)

Kaldor's (1997) characterization of contemporary conflicts can, we believe, be further formalized into a preliminary and effectively intuitive model. A previously intuitive and subsequently qualitatively validating modeling approach is not only reconcilable with the

---

1 When we mean the intuitive model, we refer to it as iWSM. When we mean the agent-based model, we refer to it as WSM. Both iWSM and WSM are part of Geller's doctoral dissertation.

2 Note that because of a lack of space, we conduct this step only in an exemplified manner.

3 In her writings, she uses the term "new wars." For many, but especially for historical, reasons (cf. Kalyvas 2001), we reject that and instead apply the term "contemporary conflicts." Besides, we are well aware of the fact that the notion of "new wars" has been criticized upon various aspects we cannot discuss here further. However, we believe that it catches some of the main characteristics of contemporary conflicts. We would like to thank Lars-Erik Cederman for reminding us of this important point.

epistemology of critical realism[4] but also with recent tendencies in agent-based modeling. Outhwaite (1987, page 55) states that "we have intuitions about the structure of almost all the social processes we may care to think about; these may be right or wrong, but they at least give us an *entrée* into the subject matter." Moss and Edmonds (2005) highlight the importance of relying on qualitative data when designing agent-based models. For the moment, a viable theory of contemporary conflicts is nonexistent. Therefore, we use an intuitive model as a starting point. However, since reliable quantitative data on the subject are scarce (Barakat et al. 2002), we base our intuitive model on qualitative case studies.

Kaldor's (1997) comments bring forth — in congruence with other case studies (Bayart et al. 1999; Reno 1998; Richards 1996) — the depicted formalization (Figure 1), namely, the iWSM. The iWSM illustrates in a simple way the main constituents of a contemporary conflict setting: *politics*, *economy*, and the *military*.[5]

The political system is based on the mechanisms of neo-patrimonialism (Médard 1990) and redistribution (Reno 1998; Richards 1996). As reported by Weber ([1921] 1980), patrimonialism is based on authority, military power, and suppressed subjects.[6] The suffix "neo" deposes the notion patrimonialism from the Weberian connotation, for it must not be based on traditional grounds. However, it is not only military power that assures authority but also the capability to redistribute social and economic resources. Actors in the political system (i.e., politicians) are therefore resource redistributors who can rely on military authority. Hence, the political system penetrates the economic system and the military system. These two

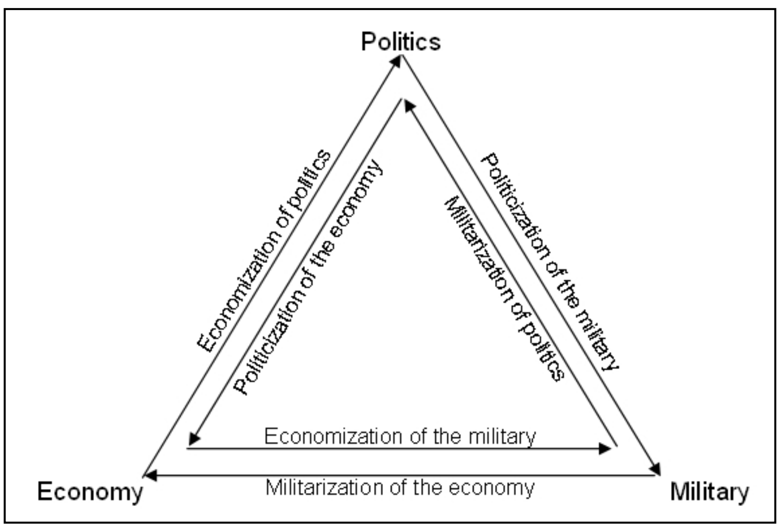

FIGURE 1 The intuitive WarSocietyModel (iWSM)

[4] The conceivability of emerging properties is, as it is in agent-based modeling and complexity theory, inherent to critical realism's ontology.

[5] We emphasize that the way we apply these three terms must not be confounded with their quotidian or general scientific usage.

[6] In our context, the suppressed subjects would be civilians. We do not consider them as being constitutive for the iWSM. Nevertheless, they will be introduced to the WSM.

processes are depicted in our model as the politicization of the economy and the military, respectively.

In the context of Sub-Saharan Africa, Hibou (1999, page 71) speaks of economies of plunder and defines them as "acquisition . . . by the representatives of public authority . . . of economic resources for private purposes." For contemporary conflicts, where public authority is either humble or has collapsed, it is, in our view, appropriate to speak of a double economy of plunder: "Officials" plunder a territory's resources in which there are also private economic actors. The idea of the shadow state can be used to link the political system with the economic system via the mechanism of neo-patrimonialism (Reno 2000). Economic actors need to play the game of neo-patrimonialism in order to open their own trade channels as a means either to get access to resources or to sell them. In Figure 1, this is referred to as the economization of politics. The economization of the military is about the self-interest of economic actors in sustaining the economic system in a highly insecure environment in order to ensure economic profit (Hibou 1999; Reno 2000).

Contemporary conflicts, notorious as they are, cannot be reduced to phenomena that erupt in violence as an end in itself. Ellis (1999) and Richards (1996) expose, in two carefully conducted case studies, the role of violence with respect to politics in Sierra Leone and religion in Liberia. The militarization of politics and the economy brings this particular aspect into the foreground. While the former tends toward the representation of political ideas with violent means, the latter identifies rent-seeking behavior.

The Iraq Conflict

However, how may an intuitive model like the iWSM be validated? How may the construct validity be assured? We propose to contrast our intuitive model with a real case — Iraq. This methodological step is a primary validation. A suitable description of the case of Iraq in the categories of our intuitive model could be considered as a successful primary validation. The secondary validation would be to contrast our data with real data, which are notoriously lacking in conflict-torn societies. However, this secondary validation is not conducted in this study.[7]

While the Iraq conflict used to be, in its first phase, a traditional Clausewitzian war, it has changed its nature in its second phase to what we described above as contemporary conflict. Although it appears on the surface to be a war between a resistance force and an occupying army, the current conflict in Iraq is, in our view, strife over political power among various actors. Political actors compete over scarce resources to build their own power bases and try to position themselves in a highly fluid situation. To achieve this aim, economic actors and combatants are being instrumentalized, with each one hoping to profit from the situation.

Political System in Iraq

During Saddam Hussein's tyranny (1979–2003), the Sunnis were favored over the Shias and Kurds. Following the overthrow of the dictator, the Sunnis, Shias, and Kurds became the

---

[7] See, for example, Cederman (2003).

major three actors in Iraq. Their current animosity toward each other is both a legacy of Saddam Hussein's rule and grounded in the traditional tribal and family structures of Iraqi society. Saddam Hussein's tyranny was based on five pillars, mainly taken by Sunnis: the Baath Party, the government, the military, and the security and intelligence services, in addition to the families and tribes. The coalition's invasion destroyed four of these pillars, leaving one intact – the families and tribes.[8] Party and government members, the soldiers, and the myrmidons of the security and intelligence services returned to their families and tribes (Baram 2003; Tilgner 2003). Nothing similar happened in the Shia or Kurdish community. However, the Shia and Kurds also seem fragmentized along tribal as well as religious cleavages. This is epitomized in the Shias' case, for example, by the tensions among the Ayatollah Ali Sistani, Ayatollah Mohammed Bakr Hakim (murdered), and the radical and militant cleric, Moqtada Sadr. In the Kurdish case, the rivalries between Massoud Barzani and Jalal Talabani (today's president) may be contemplated.

The anomic situation that resulted in the aftermath of Saddam Hussein's overthrow led to the emergence of new "old" political actors who are reorganizing themselves and are building new power bases. Hereunto they affiliate themselves with economic and military actors — and vice versa.

Economic System in Iraq

From an economic perspective, Iraq is undergoing radical economic privatization. Large-scale fraud and corruption are daily business. As reported by the British newspaper *The Guardian*, 100 Mio. $ of reconstruction money has disappeared (*The Guardian* 2005). In August 2003, the United Nations published a report in which it named organized crime (such as drug, oil, and copper smuggling, kidnapping, and ransom killing) as one of the most pressing problems in Iraq (UNODC 2003). Approximately 60% of the investment volume in Iraq is flowing into the private security apparatus. Subsequently, what has developed is a shadow economy that is plundering Iraq's resources. Notwithstanding this, it is difficult to investigate who is profiting from this hustle and bustle. A plausible conjecture is that local or regional political actors are financing themselves via these activities and that economic actors are profiting for their own sake.

Military System in Iraq

After the dissolution of the Iraqi army, most soldiers fled back to their ancestral places (Baram 2003; Tilgner 2003). Undisputedly, it is they who form the backbone of the Iraqi resistance — not international terrorist fighters, as has been claimed by some officials and parts of the media (Hottinger 2004). The big picture of the resistance is that Sunni and Shia armed groups are fighting coalition troops and collaborators. But a closer look reveals that they are also fighting each other, the Kurdish armed groups, and, most obviously, the civilian population. None of these entities are monolithic, which is disclosed by the multitude of armed groups currently active in Iraq. Equally heterogeneous are their aims, as some of them are pursuing political intentions, whereas others are after economic profits.

---

8 We would like to thank Albert A. Stahel for this insight.

WSM: EXPLORING THE IRAQ CONFLICT[9]

In a qualitative perspective, the iWSM is capable of shedding some structuralizing light on the somewhat deteriorating case of Iraq. This superimposition can be considered as a primary validation. The iWSM's ontology also enlightens the complex systems characteristics of contemporary conflicts in general and the Iraq conflict in specific. Two of the most important of them are sensitivity to initial conditions and emerging properties. Complex systems deprive themselves from being analyzed with standard methodological instruments (Moss and Edmonds 2005; Saperstein 1995). For this reason, and complementary to the aforementioned considerations on data availability, we developed an agent-based model, grounded on the iWSM, in order to deepen our understanding of the subject examined.[10] Cederman (2001, page 16) describes agent-based modeling as "a computational methodology which allows the analyst to create, analyze, and experiment with artificial worlds populated by agents that interact in non-trivial ways and that constitute their own environment."

Emergent properties depend on the rules given to the model's agents, and sensitivity to initial conditions may be tested by controlling the model's parameters. Our simulation model assumes three kinds of proactive agents, vicarious for the three aforementioned systems: *politicians*, *businessmen*, and *warriors*. The fourth sort of agents, *civilians*, are considered as being only re-active. The agents are placed randomly on a simple 50-by-50-cellwide torus grid and thus are rationally bounded (Simon 1955), since each type of agent has a specific vision (i.e., the range it can scan the grid [Moore neighborhood]). Businessmen, warriors, and civilians have default visions; politicians have a dynamic vision, representing also their power. It is assumed that only warriors and civilians can die. Table 1 specifies the agents' rules.

These rules are, of course, closely linked to the intuitive, meanwhile primarily validated, iWSM. Although not implemented as a rule, politicians are affiliates with businessmen and warriors.[11] Hence, in their behavior, they are following what has been described above as the politicization of the economy and the military, respectively. Affiliation with either businessmen or warriors renders politicians more powerful, but politicians are accountable for a larger power increase than warriors. The sensitivity to loss of military power — proxied by loss of warriors — is not implemented directly into the iWSM but is plausible for the WSM (Figure 2) insofar as military losses lead to defeat and, as a consequence, loss of territory.

Businessmen seek politicians, as the economization of politics suggests. However, for modeling reasons, they do not actively seek warriors,[12] yet they do emanate something like a money pheromone in the range of their vision that makes them attractive to warriors. Businessmen avoid competition insofar as they do not want to be close to other businessmen (Hotelling 1929).

---

9 The WSM is implemented in Repast 3.1.

10 Note that we are not modeling guerrilla warfare.

11 When affiliated, businessmen and warriors remain spatially bounded in the politician's power shade.

12 Appropriately, they would seek a leader that we have not modeled.

TABLE 1 The WSM Agent's Rules

Rule, According to Type of Agent			
Politician	Businessman	Warrior	Civilian
Changes position when losing too many warriors	Seeks politicians	Seeks politicians	Seeks politicians
	Avoids places where other businessmen are	Affiliated with politicians	
		Seeks businessmen	
		Affiliated with other warriors, forming a marauding horde	
		Fights only when affiliated with a businessman or if a member of a horde	
		Does not fight warriors of the same politician or horde nor civilians with identical affiliations	
		Recruits civilians	

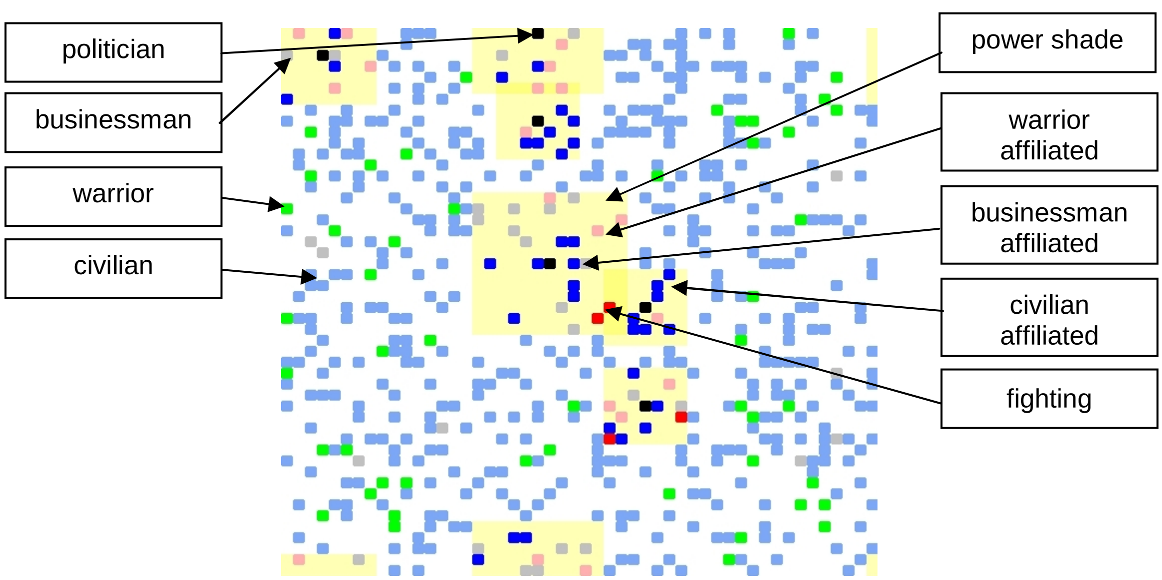

FIGURE 2 WSM snapshot at step 69 (default run; seed 1127491088036)

Warriors seek politicians, and may affiliate with them, as well as with businessmen, in their behavior (in accordance to the militarization of politics and the economy). Warriors obviously fight, but only when they are affiliated with a politician or when they are part of a horde. Fighting when warriors are affiliated with a politician is also vicarious for the militarization of politics and (since politicians attract more businessmen when they have more military power) for the militarization of the economy. Warriors can fight and kill each other as well as civilians. Warriors cannot fight warriors or civilians affiliated with the same politician or warriors who are part of the same horde. They can, however, recruit civilians. Thus, warriors are the only agents who are able to increase their total number. All other agents' numbers either remain stable or decrease.

Last but not least, civilians seek politicians as they are looking for protection. However, they do not provide the politicians with power, since the civilians are not constitutive to the iWSM. Figure 2 shows a typical WSM situation.[13]

Figure 3 gives a statistical impression of the WSM's dynamics measured in the number of times that fighting occurs per time-step. There is no obvious pattern after which the course of fighting develops. However, periods of intensified fighting cluster together, as do periods of relative peace. We have not introduced any kind of mechanism that would allow for reciprocal intensification of fighting. This suggests that the assigned reason for intensified fighting or expanded periods of peace may be found in the constellation of the WSM's agents. We would also like to point out that even after the number of times that fighting occurs starts decreasing (at around 400 ticks), fighting remains highly volatile, and it is still difficult, if not impossible, to predict new outbursts of violence (as, for example, around ticks 600 and 700).

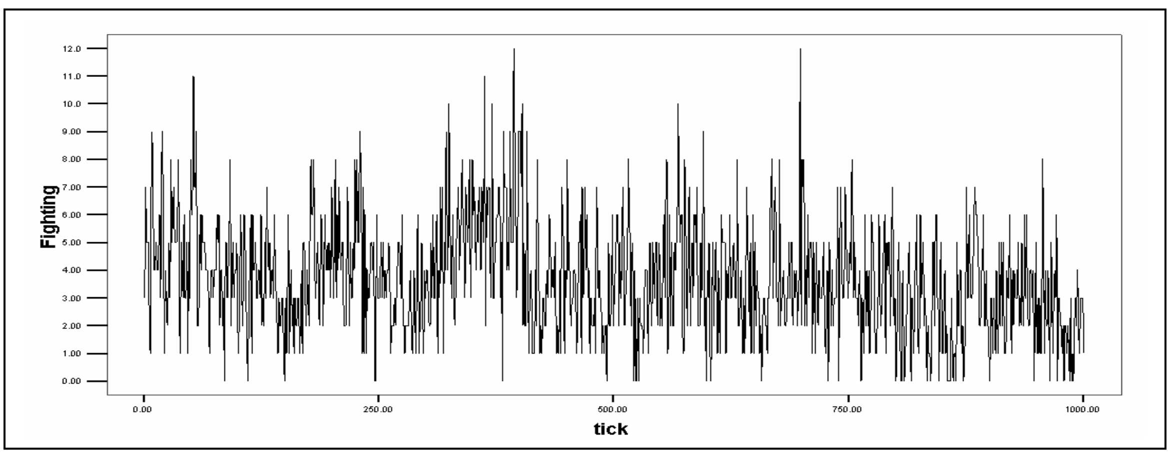

FIGURE 3 Number of times fighting occurs per time-step for one WSM run (default run; seed 1127332881837)

[13] If not otherwise indicated, the default parameters for all the runs are as follows: 6 politicians (vision 3), 30 businessmen (vision 2), 75 warriors (vision 1), and 500 civilians (vision 1). All simulation runs are stopped at step 1,000.

Figure 4 depicts the mean number of affiliated businessmen and warriors of six different politicians and their power, respectively. The two most successful politicians (politician 1 and 2) hold more than 7 businessmen and more than 10 warriors, on average. This could be a first allusion on the single agent level that neither greed nor violence stands for itself. In addition, it could be an indicator for a self-reinforcing process, whereby more powerful politicians become even more powerful in the long run.

What is the politicians' impact on the WSM? Or, to phrase it differently: Is a more fragmentized political landscape more prone to violence than a less fragmentized one? We examine this question by incrementing the number of politicians by 1 from 3 to 10. Violence is proxied by the number of times that fighting occurs per time-step, dead warriors, and dead civilians (Figure 5). The solid line indicates the mean out of 10 runs, and the error bars indicate one standard deviation (as they also do for Figures 6–8). We find that with an increasing number of politicians, the average number of times that fighting occurs per time-step increases as well (Figure 5). Bear in mind that warriors may fight only when they are affiliated with a politician or are members of a horde. In the WSM, hordes are rare and temporally constrained phenomena. Affiliation with a politician, however, is "lifelong." Thus, with an increasing number of politicians, the chance of organized violence increases. The number of dead warriors increases slightly. Bear in mind that warriors contribute to the politician's power. An exceeding number of politicians and exceedingly more powerful politicians make space scarce, resulting in more contact between politicians' power shades, with this, in turn, resulting in more fighting and slightly more dead warriors. Although we expected the number of dead civilians to increase (because of running into more affiliated warriors) or decrease (because of finding more places to hide), it remains fairly stable. This encourages us in our notion that, whatever the situation is, the civilian population suffers the most. However, the wide range of the standard deviation (also

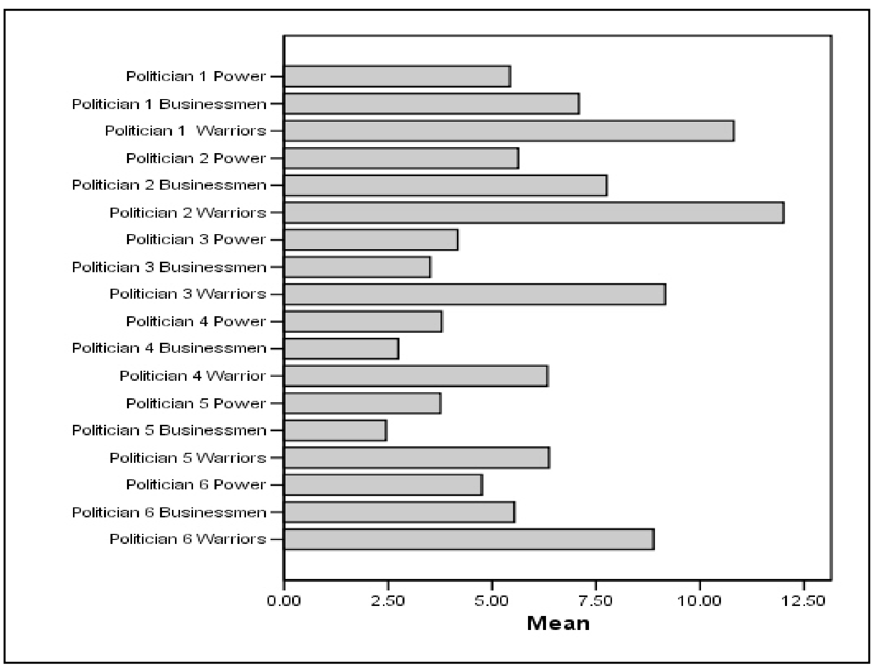

FIGURE 4 Politicians' power bases broken down to affiliated businessmen and warriors (default run; seed 1127332881837)

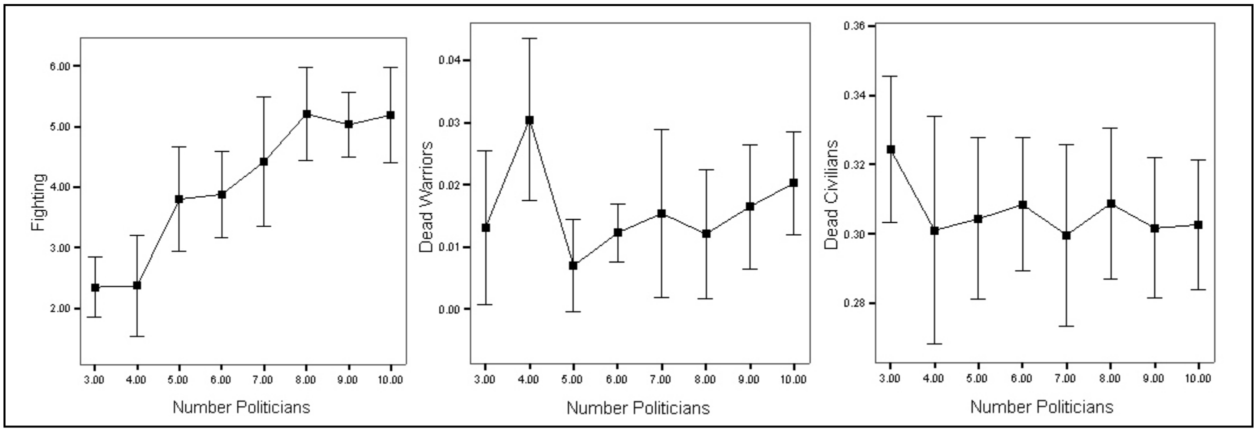

FIGURE 5 Impact of increasing number of politicians on WSM

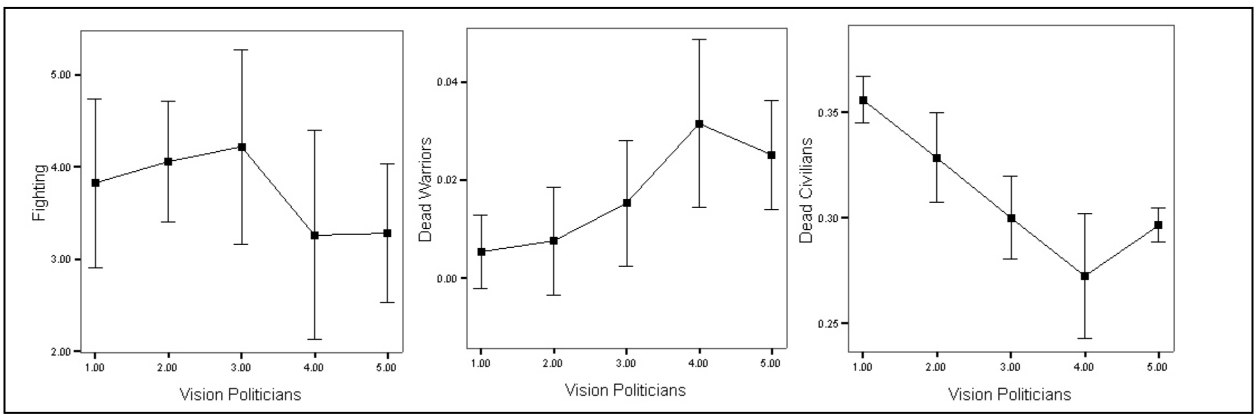

FIGURE 6 Impact of increasing vision of politicians on WSM

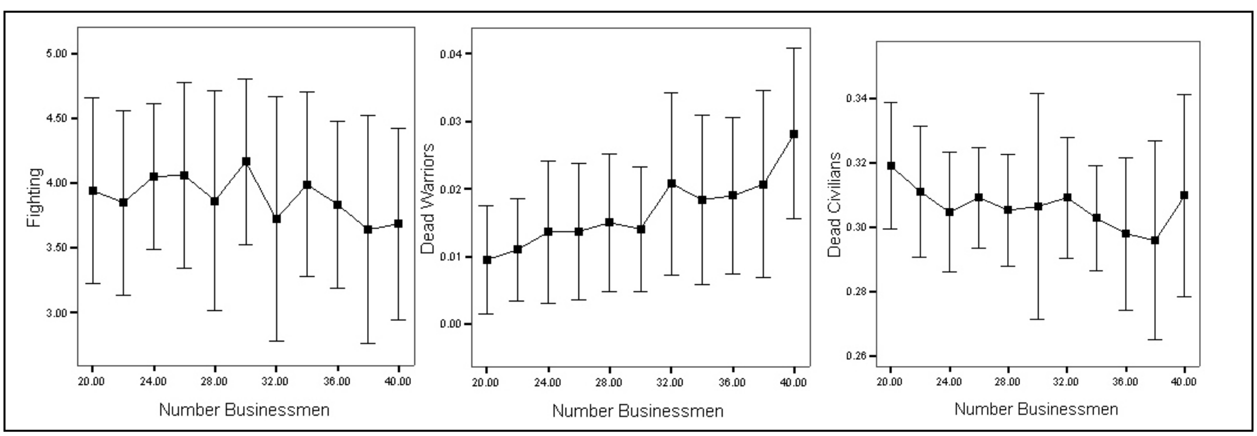

FIGURE 7 Impact of increasing number of businessmen on WSM

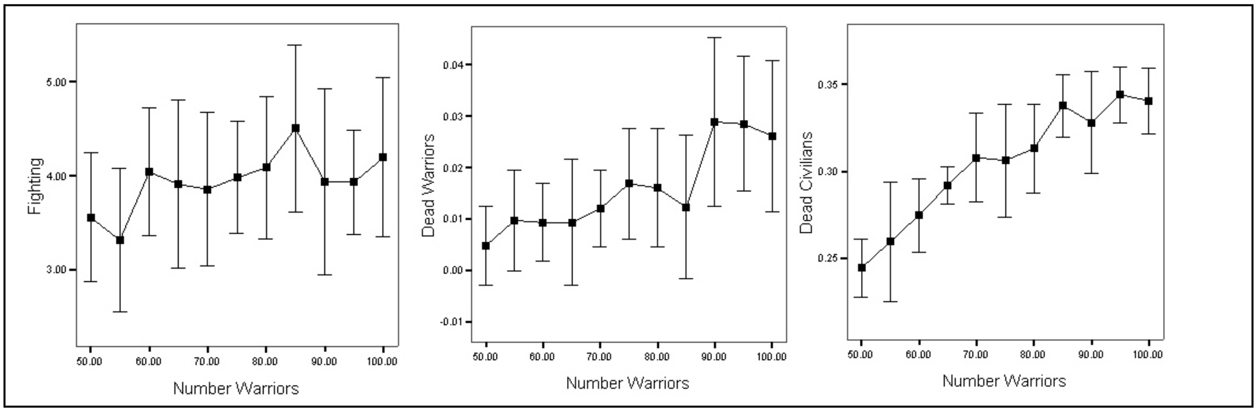

FIGURE 8 Impact of increasing number of warriors on WSM

shown in Figures 6–8) renders a precise interpretation of the results difficult and, at the same time, underlines the highly volatile character of the WSM.

The aforementioned alluded-to results instigated us to analyze whether more powerful politicians entail less violence. Vision, our proxy for power, is incremented by 1 from 1 to 5. More powerful politicians abate fighting but do not abate the number of dead warriors (Figure 6). This could be the result of major combat in the power shade fringe area. In contrast to this, the number of dead civilians is decreasing, suggesting that they are more likely to hide. However, constraints in the WSM were circumvented to simulate larger values for the politicians' visions and thus gain deeper insight into the matter.

Collier and Hoeffler (1998, 2004) report that greed is one of the main incentives for actors in (contemporary) conflict. We scrutinize this finding by incrementing the number of businessmen by 2 from 20 to 40. More businessmen should lead to more violence. Figure 7 shows that violence is decreasing only slightly. However, a wide standard deviation, when compared with an increase in politicians' power, suggests that fighting becomes more volatile and therefore less predictable. An explanation for this may be that more businessmen cause faster power changes of politicians, since their affiliation is not "lifelong" but instead depends on the politicians' power. The increasing number of dead warriors may be explained as follows: an increasing number of businessmen leads to more powerful politicians, which further increases the likelihood of warriors getting affiliated with politicians and therefore being able to fight. In addition, more powerful politicians increase the likelihood of fighting in fringe areas. Thus, the economization of conflict may lead to an increasing militarization. The number of civilian deaths is slightly decreasing. Again, more businessmen lead to more powerful politicians, under whom civilians may hide. Nevertheless, civilians bear the brunt of conflict.

If businessmen do not increase violence significantly on the entire scale, then what about warriors? The number of warriors in Figure 8 is incremented by 5 from 50 to 100. More warriors (compared to those in Figure 7) lead to more fighting, as could be expected. They also lead to more dead warriors. Bear in mind that one warrior is accountable for less of a power increase for a politician than is one businessman. Only a large number of warriors lead to a significant increase of the politicians' power shades and to more fighting among warriors in the fringe areas. If we compare fighting in Figure 8 and Figure 4, we see again that the number of times that fighting occurs per time-step depends mainly on the number of politicians. Yet an increasing number of warriors entails a significant increase in the number of civilian deaths. As more warriors in the beginning lead to more affiliated warriors, there is also a higher chance for a civilian to have contact with a warrior with the ability to fight. Figures 7 and 8 suggest that neither businessmen nor warriors contribute independently to the development of a contemporary conflict.

CONCLUSIONS

Encouraged by the paucity of research on the systemic character of contemporary conflicts, we have developed an intuitive model of a contemporary conflict setting: the iWSM. Three interacting systems — politics, economy, and the military — are constituents of the iWSM. Out of this creative process, a first desideratum accrues: How may an intuitive model that later becomes the basis for an agent-based model be validated?

The implementation of the iWSM in an agent-based model (WSM) unearthed some interesting results that may help us better understand contemporary conflicts in general and the Iraq conflict in particular. First, more political leaders entail more violence. In other words, a more fragmentized political landscape in a conflict-torn society seems to be more prone to violence than a less fragmentized landscape, as could be observed in Iraq after the topple of Saddam Hussein. Second, our findings suggest that an economization of conflict (i.e., an increasing number of businessmen) also leads to a militarization of conflict. Third, an increasing number of warriors results in an increasing number of civilian deaths — a development that was observed after the dissolution of the Iraqi Army. Fourth, neither businessmen nor warriors contribute independently of each other to a contemporary conflict's outcome. Fifth, civilians bear the major burden, as can be seen daily in Iraq. Last but not least, our data suggest that a contemporary conflict's system effects are complex and difficult, if not impossible, to predict. This leads us to a second desideratum: How may social simulation data can be validated against reality in the worst case (i.e., in contemporary conflicts)?

Rather than concluding that (1) the nongovernmental leaders in Iraq should be taken out, (2) the Iraqis should unite under one national leadership, (3) foreign private enterprises should leave the country, and (4) warriors should be disarmed, we would like to emphasize a third desideratum: How may the hermeneutical circle in agent-based modeling be further standardized? We hope our suggested procedure is a first step in the right direction and will help us gain a better understanding of the events in Iraq and perhaps of contemporary conflicts as a whole. For this, further investigation of the iWSM and the WSM, both qualitatively and quantitatively, is needed. From a qualitative perspective, plans are to design the iWSM as an integrated model (integrated intuitive WSM; iiWSM). By this, we mean that real actors should be consulted for developing the agent rules. While Iraq does not seem to be the place for this at the moment, Afghanistan may be. From a quantitative perspective, the WSM's output may be put under further investigation. It would be most interesting to see if some of our data, perhaps the aggregated data, are as power-law-distributed as well as Cederman (2003) or Johnson et al. (2005), for example, suggest. Another equally interesting question would be if our data also exhibit clustered volatility, as Moss and Edmonds (2005) found for their data. To examine our model output in a more dynamic fashion, a repeated measurement analysis of variance could be conducted. With this, we could perhaps identify typical development paths for different settings of the WSM.[14]

REFERENCES

Axelrod, R., 1997, "Advancing the Art of Social Simulation," in *Simulating Social Phenomena*, ed. by R. Conte, R. Hegselmann, and P. Terno, Berlin: Springer.

Barakat, S., M. Chard, T. Jacoby and W. Lume, 2002, "The Composite Approach: Research Design in the Context of War and Armed Conflict," *Third World Quarterly* **23**(5):991–1003.

Baram, A., 2003, "Saddam's Power Structure: The Tikritis before, during and after the War," in *Iraq at the Crossroads: State and Society in the Shadow of Regime Change*, ed. by T. Dodge and S. Simon, Adelphi Paper 354, London: International Institute for Strategic Studies, pp. 93–114.

Bayart, J.-F., S. Ellis, and B. Hibou (eds.), 1999, *The Criminalization of the State in Africa*, Bloomington/Indianapolis, IN: Indiana University Press.

Cederman, L.-E., 2003, "Modeling the Size of Wars: From Billiard Balls to Sandpiles," *American Political Science Review* **97**(1):135–150.

Cederman, L.-E., 2001, "Agent-Based Modeling in Political Science," *The Political Methodologist* **10**(1):16–22.

Collier, P., and A. Hoeffler, 1998, "On Economic Causes of Civil War," *Oxford Economic Papers* **50**(4):563–573.

Collier, P., and A. Hoeffler, 2004, "Greed and Grievance in Civil War," *Oxford Economic Papers* **56**(4):563–595.

Gerring, J., 2004, "What Is a Case Study and What Is It Good for?," *American Political Science Review* **98**(2):341–354.

Ellis, S., 1999, *The Mask of Anarchy, The Destruction of Liberia and the Religious Dimension of an African War*, New York, NY: New York University Press.

Hotelling, H., 1929, "Stability in Competition," *The Economic Journal* **39**(153):41–57.

Hottinger, A., 2004, *Islamische Welt, der Nahe Osten: Erfahrungen, Begegnungen, Analysen* (*Islamic World, The Middle East: Experiences, Encounters, Analyses*), Zurich: Verlag Neue Zürcher Zeitung.

[14] We would like to thank Wander Jager for this point.

Hibou, B., 1999, "The 'Social Capital' of the State as an Agent of Deception or the Ruses of Economic Intelligence," in *The Criminalization of the State in Africa*, ed. by J.-F. Bayart, S. Ellis, and B. Hibou, Bloomington/Indianapolis, IN: Indiana University Press, pp. 69–113.

Huntington, S.P., 1993, "The Clash of Civilizations?," *Foreign Affairs* **72**(3):22–49.

Jean, F., and J.C. Rufin (eds.), 1999, *Ökonomie der Bürgerkriege* [Economy of Civil Wars], Hamburg: Hamburger Edition.

Johnson, N., M. Spagat, J. Restrepo, J. Bohórquez, N. Suárez, E. Restrepo, and R. Zarama, 2005, "From Old Wars to New Wars and Global Terrorism"; available at http://arxiv.org/ftp/physics/papers/0506/0506213.pdf. Accessed Oct. 24, 2005.

Kaldor, M., 1997, "Introduction," in *New Wars, Restructuring the Global Military Sector, Volume I: New Wars*, ed. by M. Kaldor and B. Vashee on behalf of UNU World Institute for Development, Economics Research (UNU/WIDER), London/Washington: Pinter, pp. 3–33.

Kalyvas, S.N., 2001, " 'New' and 'Old' Civil Wars, A Valid Distinction?," *World Politics* **54**(1):99–118.

Médard, J.-F., 1990, "L'État Patrimonialisé [The Patrimonialized State]," *Politique Africaine* **39**:25–36.

Moss, S., and B. Edmonds, 2005, "Sociology and Simulation: Statistical and Qualitative Cross-validation," *American Journal of Sociology* **110**(4):1095–1131.

Outhwaite, W., 1987, *New Philosophies of Social Science, Realism, Hermeneutics and Critical Theory*, London: Macmillan Education.

Reno, W., 2000, "Shadow States and the Political Economy of Civil Wars," in *Greed and Grievance in Civil Wars*, ed. by M. Berdal and D.M. Malone, Boulder/London: Lynne Rienner Publishers, pp. 43–68.

Reno, W., 1998, *Warlord Politics and African States*, Boulder/London: Lynne Rienner Publishers.

Richards, P., 1996, *Fighting for the Rain Forest, War, Youth and Resources in Sierra Leone*, The International African Institute, Oxford: James Currey; Portsmouth: Heinemann.

Saperstein, A.M., 1995, "War and Chaos," *American Scientist* **83**(6):548–557.

Simon, H.-A., 1955, "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics* **69**(1):99–118.

The Guardian, 2005, "U.S. Can't Account for $100M Spent in Iraq," May 5; available at http://www.theinsider.org/mailing/article.asp?id=1158. Accessed June 26, 2005.

Tilgner, U., 2003, *Der Inszenierte Krieg, Täuschung und Wahrheit beim Sturz Saddam Husseins* [*The Staged War, Deception and Truth at the Overthrow of Saddam Hussein*], Berlin: Rowohlt.

UNODC (UN Office on Drugs and Crime), 2003, *Organized Crime to Be a Growing Problem in Iraq UNODC Fact-finding Mission Reports*, UNIS/CP/445, Aug. 27; available at http://www.unis.unvienna.org/unis/pressrels/2003/cp445.html. Accessed June 26, 2005.

van Creveld, M., 1992, *On Future War*, London: Brasseys.

Weber, M., [1921] 1980, *Wirtschaft und Gesellschaft, Grundriss der Verstehenden Soziologie* (*Economy and Society, An Outline of Interpretative Sociology*), 5th ed., Tübingen: J.C.B. Mohr (Paul Siebeck).

SIMULATING INITIAL CONDITIONS IN AGENT-BASED MODELING

W.M. BULLEIT* and M.W. DREWEK, Michigan Tech, Houghton, MI

ABSTRACT

Predicting where terrorists are most likely to strike concerns planners, law enforcement and government agencies at various levels, and engineers who must design facilities of all kinds. The present work is an effort to use agent-based modeling to examine the interaction of civilians, terrorists, and security to determine the types of facilities in a town or city that are most susceptible to attack. Agent modeling of civil violence has been performed in the past. The ultimate goal of our research is to be able to estimate the probability of attack for various types of facilities in a population center so that resources can be allocated for hardening or otherwise protecting those facilities. Because of the nature of resource-based agent modeling, the agents must be allowed to evolve in the town or city environment before the day-to-day behavior of the community is simulated. We have approached that problem by breaking the total simulation into two parts: (1) the incubation of the community, where the agent population evolves to live in the environment, and (2) the simulation of the behavior of the evolved agents in the community environment. Results from this work indicate that incubation can be ended at any desired time and still allow modified time-step simulation. This result allows modified time-step simulation of a population in any stage of its evolution. When transitioning from incubation to simulation, the behavior of the population must be allowed to stabilize in the early stages of the shortened time-step simulation.

Keywords: Agent-based modeling, artificial societies, simulation, terrorism

INTRODUCTION

Predicting where terrorists are most likely to strike concerns planners, law enforcement and government agencies at various levels, and engineers who must design facilities of all kinds. The present work is an effort to use agent-based modeling to examine the interaction of civilians, terrorists, and security to determine the types of facilities that are most susceptible to attack. Agent modeling of civil violence has been performed in the past (Epstein 2002). The ultimate goal of our research is to be able to estimate the probability of attack for various types of facilities in a population center so that resources can be allocated for hardening or otherwise protecting those facilities.

Agent models comprise a range of types, of which this one is an extension of the type used by Epstein and Axtell (1996) in which society evolves by using the basic concepts of resources in the environment, agent metabolism for those resources, and agent vision (knowledge of the environment). This model represents a community in which civilians evolve to become radicals (inactive terrorists) who may become active terrorists committing attacks on the community. The environment in which this community evolves consists of a rectangular grid on

* *Corresponding author address*: William M. Bulleit, Department of Civil and Environmental Engineering, Michigan Tech, 1400 Townsend Dr., Houghton, MI 49931; email: wmbullei@mtu.edu.

which a number of resources lie. Civilian agents evolve in this environment on the basis of their vision and metabolism for the various resources. The terrorist agents evolve from the civilian agents by using a tag-mediated procedure derived from that used by Axelrod (1997). After the agent becomes a terrorist, it remains an inactive terrorist until its age and wealth each reach a specific value that allows it to become active. An active terrorist agent stops looking for resources and begins to examine the agent wealth within its vision. When it finds a location of high local wealth, it moves to that location and becomes a suicide bomber that explodes, destroying the agents and the wealth on the surrounding grid points. The number and location of security agents are determined on the basis of the wealth, fear, and innate nervousness of the agents in the civilian population. The number of security agents in the community evolves as attacks occur. Security agents search for and arrest terrorists in regions of locally high populations.

Results from this basic model (Bulleit and Drewek 2005) show that the location of attacks is affected by the choice of the base level of security. Higher base levels of security shift many of the attacks away from the areas of high resources. In this work, a base level of security does not exist; security levels are endogenous. Thus, it appears that endogenous agent modeling of communities will require the use of an *incubation* period during which the community can evolve to allow the agents to acclimate to the environment and develop a set of initial conditions that are themselves endogenous. A limited use of incubation has been used by Cederman (2003). In that case, he merely allowed the simulation to run for a set number of time-steps before beginning data collection. We propose a more distinct incubation period. In the proposed incubation period, the time-step will be longer than what will later be used for the community simulation from which results will be gleaned. For instance, during the incubation period, the time-step might be representative of a year. The community will be allowed to evolve during the incubation to a user-chosen time. At that point, the time-step will be shortened (e.g., to a day), and the simulation will continue with the conditions at the end of incubation becoming the initial conditions for the short time-step simulation.

The objective of this paper is to describe the use of a simulation process that has an *incubator* in which the community evolves to a certain point and a *simulator* in which the day-to-day community simulation is performed.

MODEL DESCRIPTION

Community Environment

The environment in which this community evolves consists of a 50×50 rectangular grid on which lie a number of *piles* of resources. Each civilian agent requires a set amount of each different resource. The resource piles can be isolated in the sense that there may not be a resource gradient between the piles. This lack of gradient is important to the design of the civilian agents. For this study, the environment consists of four resource piles, each representing a different resource. All agents require each resource to live. A second aspect of the environment relates to the effects of a terrorist attack on the environment. A terrorist attack, modeled as a suicide bomber, results in the destruction of all resources on the grid point where the terrorist was at the time of the attack plus all agents, all their wealth, and all resources on the Moore neighborhood of that grid point. The resources at these nodes remain zero for 2 years before they begin to regenerate. The attack also makes agents *fear* the grid points where the attack occurred. The level of fear that agents feel for the attacked nodes dissipates with time and spreads to surrounding nodes. Figure 1 shows von Neumann and Moore neighborhoods.

For the environment that we discuss in this paper, the range of resource values for each resource is a maximum of 54.0 units and a minimum of 1.0 unit. The grow-back rates in the incubator are one-fourth of the maximum value allowed at each node. Hence, the maximum grow back rate is 13.5 units/year, and the minimum is 0.25 unit/year. Figure 2 shows the environment. The maximum resources are on the *peaks*, and the minimums are on the *plains*.

Civilian Agents

Civilian agents evolve in the environment. Each agent is assigned an initial metabolism for each different resource in the environment from a uniform distribution with a range of 1.5 to 3.0 or U(1.5, 3.0) for resources 1 and 3 and U(1.25, 2.50) for resources 2 and 4. The initial agent vision is an integer selected from U(3, 7). Vision is the number of grid points that an agent can see in the four cardinal directions from its current location. The agents are also randomly assigned an amount of each of the different resources in the environment from U(45, 90), their *wealth*. Thus, an agent's wealth is an agent's store of each of the various resources in resource units. Each agent's *initial endowment* is randomly selected from U(12, 24), in units of *generalized resources*. *A* generalized resource for an agent is one of its resources divided by the metabolism for that resource, thereby converting resource units into a time or, in other words, the amount of time an agent can live, assuming that it collects no more of the given resource. Initial endowment is discussed subsequently. The agents' death age is an integer selected randomly from U(40, 80), and the agent's nervousness factor is randomly selected from U(0, 1). Nervousness is a measure of how nervous an agent is in the presence of fear. Last, each agent in

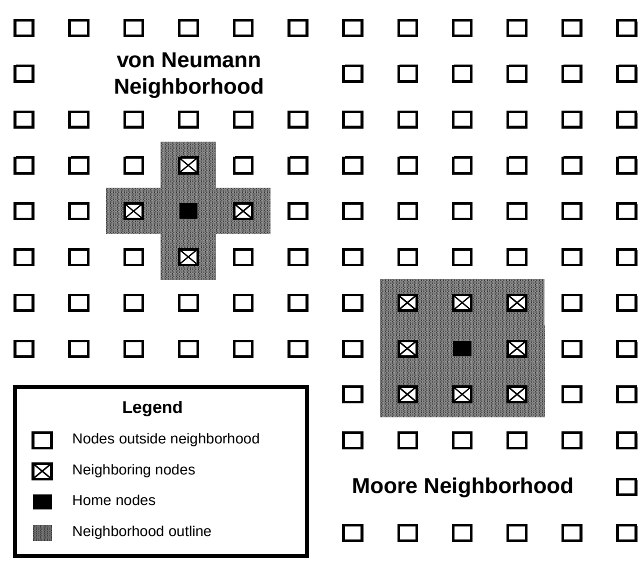

FIGURE 1 Von Neumann and Moore neighborhoods

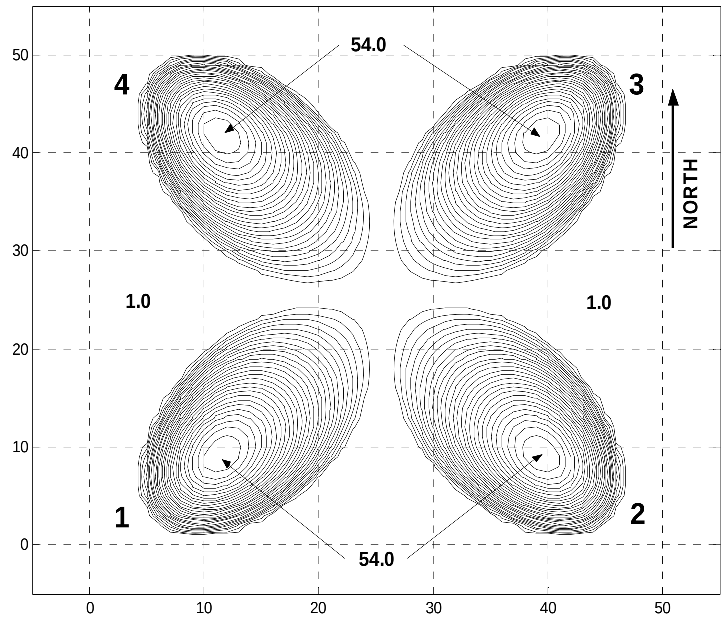

FIGURE 2 Community environment

the initial population is given a cultural tag in which each of the five tag integers is randomly selected from U(0, 9).

The initial population is made up of agents of an age between 10 and 50. Since there is some overlap in the initial age range and death age range, if the death age selected is less than the initial age, another death age is selected until the death age is greater than or equal to the initial age. Once the agents' initial wealth has been determined, the agents' initial *generalized wealth* is calculated. Generalized wealth is the length of time an agent can live assuming it collects no additional resources of any kind; i.e., the minimum of the generalized resources.

The agents move around the environment in search of resources that they need to live. The agents search for their *critical resource*. The critical resource is the resource currently limiting the agent's life span, assuming no additional resources are collected. The critical resource is the resource yielding the minimum generalized resource and is the resource used in determining an agent's generalized wealth. Because the environment does not have a resource gradient at all locations, the agents were given memory. Without this memory, it is difficult to evolve a stable population. The resource memory is simple: the agent remembers the grid point where the maximum of each of the different resources that it has encountered in its travels around the environment is located. Thus, in this case, since there are four different resources in the environment, the agent stores the location and amount of the maximum value of each of the four resources it has encountered. It updates these values as it finds a better source (larger value) of a specific resource. As well as allowing a stable population to evolve, this simple memory allows agents to evolve patterns of travel between resource locations; for instance, the path between two resources could represent travel between home and work in a real community.

Agents also have a memory of the maximum and minimum fear they have seen as they traveled around the environment. For the baseline case, discussed subsequently, fear memory includes the past 5 years. As an agent searches its local environment for resources, it considers the fear associated with the nodes it is examining. The maximum and minimum fears seen during a given time-step are recorded and will be remembered for the next 5 years. After 5 years have passed, maximum and minimum fears for the sixth year are forgotten. Fear and its use are described in detail below.

Agents have a gender, and when male and female agents meet they procreate if each of them has reached a fertile age and is wealthy. Procreation allows the agents' vision, metabolism, and nervousness to evolve. A potential parent is an agent that is fertile (i.e., has an age within the fertility age range) and possesses a generalized wealth equal to or greater than its own initial endowment. The minimum fertility age for males is an integer selected randomly from U(12, 15) with the maximum from U(50, 60). For females, the selection is made from U(12, 15) and U(40, 50), respectively. When an agent moves to a node, if that agent is a potential parent and one of its von Neumann neighbors is also a potential parent, and assuming that in one of their von Neumann neighborhoods there is an unoccupied node, then an child is born. If more than one potential parent of opposite gender is located in the agent's von Neumann neighborhood and an unoccupied node is still available, then the mate is selected at random. Potential parent agents who have a parent/child relationship or share a common parent are not allowed to procreate.

When a newborn agent is added to the population, its placement in the environment is selected randomly from all the unoccupied nodes in the parents' von Neumann neighborhoods. The newborn's vision is determined by taking the average of the vision of the parents (rounded to the nearest integer) with a mutation probability, P_{mv}, of 0.0025 that this value will be increased by 1.0 or decreased by 1.0. (Vision is limited to a minimum of zero and has no set maximum.) Infant agent metabolisms are determined in the same way, with mutation probability $P_{mm} = 0.0025$, except that the minimum metabolism cannot drop below the minimum of the range of the uniform distribution used in the selection of metabolisms for the initial agents (i.e., 1.25 or 1.5). Infant agent nervousness is determined the same way, with mutation probability $P_{mn} = 0.0025$, but the change is either +0.1 or –0.1. Nervousness is kept within the range of 0 to 1. The newborn's initial wealth is calculated by multiplying one-half of the father's initial endowment by his metabolism for each resource and adding to that the corresponding results of a similar calculation for the mother. The mother and father each donate the resources to their newborn; the resources donated are forfeited from the parents. This store of individual resources is used to determine the newborn agent's initial generalized wealth by dividing each resource level by the newborn's metabolism for each respective resource. The newborn agent's initial generalized wealth serves as its initial endowment. The newborn's gender is selected at random, with an equal chance of each. The newborn's fertility age range is selected from the ranges used by the initial population, as is the newborn's death age. The newborn's initial knowledge of where resources lie in the environment is taken from each parent's memory: The parents give the newborn the "best" locations of each resource in either of their memories. Note that the parents also exchange the best resource locations in their respective memories. The newborn's cultural tag is determined from its parents' tags; for each tag integer, there is an equal probability that the value will be taken from the mother or the father.

Terrorist Agents and Terrorist Attacks

Terrorist agents evolve from the civilian agent population. The evolution of a civilian agent to a terrorist is performed by using a tag-mediated process that is based on the approach used by Axelrod (1997). As described above, each agent is assigned a tag at the beginning of the simulation or at birth. The tag consists of a string of five integers in which each integer ranges from 0 to 9. As the agents move around, they interact with other agents. The interaction is controlled by the tags, and the evolution of a civilian to a terrorist is based on the tag values. First, consider interaction. When an agent moves to a grid point, it examines, at random, one of the grid points in its von Neumann neighborhood. If an agent is in that location, the agents compare the sum of the absolute value of the difference between each of the five integers in their tag:

$$S = \sum_{i=1}^{5} \left| I_{ij} - I_{ik} \right| , \tag{1}$$

where I is the value of the tag integer, i is the location of the tag integer, and j and k are the indexes of the interacting agents. The larger this sum is, the smaller the probability that the agents interact. If the sum is 45, then the probability is 0.0 that they interact. If the sum is 0, then the probability of interaction is 1.0. The probability of interaction is linear between these two end points. If the agents interact, then one of the integer locations on the tag is chosen at random — a 0.20 probability that any one of the five is chosen. Once one of the integer locations is chosen, the agents compare the integer they have at that location. If the integers are the same, nothing happens. If the integers are different, then one of two things occurs: (1) the agent that moved changes its integer to match the agent that it interacted with, or (2) the agent that moved has a radical change. The probability of a radical change is determined from using:

$$P_{rc} = P_b \frac{\left| I_{ij} - I_{ik} \right|}{9} , \tag{2}$$

where P_{rc} is the probability of radical change, P_b is the base probability (a P_b of 0.02 is used in all example simulations), and I_{ij} and I_{ik} have been defined previously. The direction of the radical change is determined by using the *changing* agent's current integer value. For example, if the current integer is 2, then there is a 2/9 probability that the agent will change to a 9, and a 7/9 probability that the value will change to 0. The agent that moved will be the changing agent. After the agents have interacted, whether or not an integer change has occurred in either of the above two ways, there is still a small, isolated change probability, P_{ic}, of 0.02 that one of the integers on its tag will change by –1 or +1. This ends the interaction. The agent that moves has the changes occur to it so that there is no possibility that an agent will be changed more than once during any time-step (Axelrod 1997).

An agent becomes a terrorist on the basis of the sum of the five integers in its tag (referred to as *cultural identity*). The probability that the agent becomes a terrorist is determined by using a U-shaped symmetrical polynomial function that passes through 1.0 at a sum of 0, passes through 0.0 at a sum of 22.5, and passes through 1.0 again at a sum of 45. Figure 3 shows the U-shaped curve. Thus, there is some probability that any agent can become a terrorist, but the probability is greatest near the end points of the sum of the tag integers. After the agent becomes a terrorist, it remains an *inactive* terrorist until its age and wealth each reach a specific value that allows it to become active. An inactive terrorist agent becomes *active* if it is 18 years old or older and its generalized wealth is equal to 5.0 or greater. Once active, the terrorist agent will remain active as long as its generalized wealth remains greater than 3.0. After every change to the tag,

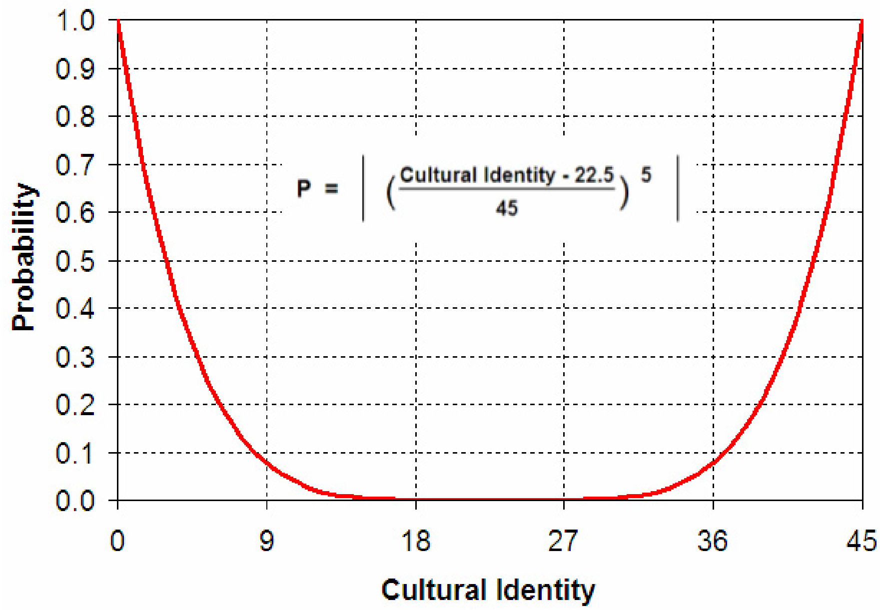

FIGURE 3 Probability of becoming a terrorist

the agent's new sum is used to determine the probability of becoming a terrorist (if the agent is already a terrorist, the probability is that of remaining a terrorist).

An active terrorist agent stops looking for resources and begins to examine the wealth on the von Neumann neighborhood of grid points within its vision and moves to the node with the largest surrounding wealth. (Note that even though the terrorist stops looking for resources, it continues to collect resources at the nodes it is visiting) This wealth information, referred to as surveillance data, consists of the present agent generalized wealth and the moving average of the agent generalized wealth on each grid point in the von Neumann neighborhood over the past 10 time-steps (referred to as *historical nodal wealth*).

This approach is used because terrorists do not strike just very wealthy locations but also locations where wealth passes through (e.g., airports). The active terrorist agent then keeps track of the mean and standard deviation of the largest five surveillance data values that it has seen in its travels. When it finds a grid point that has a surveillance data value that is greater than the mean plus some number of standard deviations (typically 1.0) *and* the coefficient of variation of its surveillance data is less than 0.25, it becomes a suicide bomber and explodes, destroying wealth on the Moore neighborhood as discussed above. These two criteria for detonating allow an active terrorist agent to attack when it finds a local region with a relatively consistent high level of wealth.

When a terrorist agent conducts an attack, all of the agents with their wealth and the nodal resources are destroyed in the terrorist agent's Moore neighborhood. Agents will fear the nodes in the destroyed area and, over time, in areas surrounding the destroyed area. The sum of all the agent generalized wealth destroyed becomes the fear at each of the nodes in that Moore neighborhood. If the level of fear is greater than the level that existed before the attack, then the portion of the fear that is greater than the existing fear level will diffuse outward over time, reducing the nodal intensity of the fear. Eventually, if enough time elapses without another terrorist attack, the fear level from the attack becomes uniform across the environment, thereby affecting every agent equally and thus having no effect on any agent's decision process. The spreading of fear has been modeled in the same way that Epstein and Axtell (1996) modeled the diffusion of pollution> Details on the diffusion of fear can be found in that reference. If the level of fear is less than the level that existed before the attack, then the fear level remains the same as it was before the attack for those nodes, that is, new fear is not summed on top of existing fear.

Agents evaluate the critical resource that they see by using the *adjusted critical resource*. The adjusted critical resource is used to evaluate the resources on a given node. The adjusted critical resource is the amount of the critical resource at a node adjusted to take into account the fear level at that node, the fear memory of the agent, the generalized wealth of the agent, and the innate nervousness of the agent. The resource values (for the critical resource) at each node are adjusted according to the following equation:

$$AR_i = R_i - \Delta R_i \ , \tag{3}$$

where

$$\Delta R_i = R_i n \left[1 - \exp\left(- \frac{(f_i - f_{min})}{(f_{max} - f_{min})} \times \frac{(f_i - f_{min})}{9GW} \right) \right], \tag{4}$$

where AR_i is the adjusted critical resources located at node i, R_i is the amount of the critical resource located at node i, ΔR_i is the amount of resource located at node i that an agent is willing to sacrifice for less fear, n is the agent's innate nervousness factor, f_i is the fear at node i, f_{max} is the largest fear seen in the last 5 years, f_{min} is the smallest fear seen in the last 5 years, and GW is the agent's generalized wealth. The two terms making up the exponent serve two purposes: the first term normalizes the fear level at node i to the range of fear seen in the recent past, and the second term normalizes the relative fear, $f_i - f_{min}$, to the agent's generalized wealth. The "9" appears in the second term because there are nine nodes in a Moore neighborhood. Once each node has been considered, the agent moves to the node with the largest adjusted critical resource.

Security Agents

The number of security agents in the community evolves as attacks occur. The number of security agents is based on characteristics of each agent in the population. These characteristics include the wealth of the agent, the resources that the agent collects at each time-step, the level of fear that the agent feels at that time-step, the maximum and minimum amount of fear that the agent has felt in the past, and the inherent nervousness of the agent. These characteristics are used to determine the amount of resources that the agent is willing to contribute to buying security. Note that the agents do not actually give up any resources. One method of putting a dollar value on a nonmarket good is to conduct a survey, essentially asking people how much they would be willing to pay for something to happen (Dorfman and Dorfman 1993). The responses are summed for the population affected, and this becomes an estimate for the value of that nonmarket good. This is called *contingent evaluation* and corresponds to the approach we are using to assign security to the environment.

At each time-step, adult civilian agents (agents that have reached their minimum fertility age) consider what portion of the resources being collected at that time-step they would be willing to contribute to purchase security. This willingness to contribute resources, without actually giving them up, is the agent's demand for security. Equation 5 is used to calculate the amount of each resource that each agent is willing to contribute:

$$C_j = R_{ij} n \left[1 - \exp\left(- \frac{(f_0 - f_{min})}{(f_{max} - f_{min})} \times \frac{(f_0 - f_{min})}{9GW} \right) \right], \qquad (5)$$

where C_j is the contribution of resource j, R_{ij} is the amount of resource j at node i, f_0 is the maximum fear that the agent has seen during that time-step, and the other variables are the same as in Equation 4. The first term making up the exponent in Equation 5 normalizes f_0 to the range of fear seen by the agent in the recent past, and the second term normalizes $f_0 - f_{min}$ to the agent's generalized wealth. During the time-step, the agent contributions are summed, and at the end of the time-step, there is a pool of each resource. The average metabolism for all nonsecurity agents for each of these resources is determined. Each resource pool is divided by the average metabolism for that resource. The minimum of these values becomes the number of security agents required at the end of the time-step (i.e., the number of security agents is the average number that can be supported by the contributed resources).

If the existing number of security agents needs to be increased to meet the calculated requirement, new security agents are introduced in the environment. The new security agents are given vision randomly selected from the range given by the absolute minimum and absolute maximum vision in the agent population. Nodes with higher historical nodal wealth have a higher probability of receiving these new security agents. Specifically, each unoccupied node is assigned a random number from U(0,1). Each of these random numbers is multiplied by the historical nodal wealth at the node divided by the maximum historical nodal wealth found in the environment. After all of the unoccupied nodes have been considered, the new agents are located on the nodes with the largest adjusted random number.

If the existing number of security agents needs to be decreased to meet the requirement, some existing security agents are removed from the environment. The security agents located on nodes with lower historical nodal wealth have a higher probability of being removed. The process for removing security agents is the same as for adding them, except that the security agents located on the nodes with the smallest adjusted random number are removed from the environment.

Security agents search for terrorists in regions of high population. Each security agent moves to the open grid point within its vision that has the most agents on its von Neumann neighborhood. Once on that grid point, the security agent examines its von Neumann neighborhood. It interacts with (investigates) each agent on the von Neumann neighborhood with a probability related to the number of agents in the neighborhood; for example, if there are three agents in the von Neumann neighborhood, then it interacts with each of those agents with a probability of one-third. If the security agent interacts with an agent, there are two possible outcomes: (1) it releases civilians or (2) it arrests terrorists (active or inactive), with a probability determined by using the U-shaped symmetrical polynomial function described above. When used for security agents arresting terrorist agents, the U-shaped symmetrical polynomial function is cubic. (When used for generating terrorist agents, the U-shaped symmetrical polynomial function is fifth order, as shown in Figure 3.) The probability of arresting a terrorist agent increases for more radical agents (sum of tag integers closer to 0 or 45). An arrested terrorist agent is permanently removed from the environment.

Incubation and Community Simulation

The day-to-day behavior of the community becomes apparent only after the agents, particularly the civilian agents, have learned to live in the environment. The *incubation period* is that time during which the agents are allowed to adjust to the environment. The portion of the simulation encompassing the incubation period is referred to as the *incubator.* The time-step of the incubator is representative of a *year*. The portion of the simulation following the incubator is referred to as the *simulator.* When switching from the incubator to the simulator, the model must be calibrated to the adjusted time-step. First, the number of time-steps that make up a year, T_y, must be chosen. At each time-step, the current time of the simulation is incremented by $1/T_y$. For ease of discussion, we refer to a time-step with a duration of $1/T_y$ as a *day.* (If the incubation time-step is representative of a year, then using $T_y = 365$ would produce the day that we are familiar with.) All time-related parameters must be adjusted. The agents' age, in years, is converted from an integer to a real number by adding a random number from U(0,1). In the simulator, each agent's age is incremented at each time-step by $1/T_y$. The agents' maximum and minimum fertility ages, as well as the agents' death ages, remain integers. To maintain consistent agent evaluations of resources, in order to maintain stable agent wealth when switching between the incubator and simulator, the resource concentrations in the environment are divided by T_y. The agents' resource metabolisms are also divided by T_y.

When a terrorist attack occurs, all resources are destroyed on the nodes involved, and the area remains devoid of all resources for 2 years. In the incubator, this time period is equal to two time-steps. In the simulator, the damaged area also remains devoid of all resources for 2 years, but the number of time-steps is $2 * T_y$. Fear is generated in the same way in the incubator and simulator, but the fear dispersion rate (on a per-time-step basis) must be adjusted. In the incubator, fear dispersion on a per-time-step basis is the rate at which fear disperses in a year. In the simulator, on the first day, the fear dispersion that will occur over the first year is calculated by using the procedure from the incubator. The resulting change in fear over the next year at each node is then divided by T_y, producing the change in fear per day at each node. After 1 year passes or when a terrorist attack occurs or when 1 year has passed since the last terrorist attack, the change in fear per day at each node is recalculated.

The occurrence probability for procreation and cultural exchange are reduced from 1.0 in the incubator to $1/T_y$ in the simulator. This modification means that in the incubator, a potential parent agent will consider procreating at each time-step, but in the simulator, this will happen at each time-step with a reduced probability. The same is true for cultural exchanges. In the incubator, an agent will consider culturally interacting with a neighbor at each time-step. In the simulator, at each time-step, an agent will consider culturally interacting with only one of its neighbors, with a probability of $1/T_y$.

Agent resource memory remains unchanged, except that the resource magnitudes held in memory are divided by T_y. Agent fear memory also changes in the transition from the incubator to the simulator. For the baseline case, fear memory consists of 5 years of memory of the largest and smallest nodal fears seen each year. In the incubator, this means keeping track of the largest and smallest nodal fears seen at each time-step. However, in the simulator, 1 year is composed of multiple time-steps. At the end of 1 year, an agent remembers the minimum and maximum fear seen in the past year, t_1, as well as the minimum and maximum fear seen in the previous four years, t_2, t_3, t_4, and t_5. At the beginning of the new year, minimum and maximum fears seen in t_5 are forgotten. The present year becomes t_1; what had been referred to as t_1 becomes t_2, t_2 becomes t_3, and so on. For the first day of the new year t_1, the minimum and maximum fear remembered are the minimum and maximum fear seen while searching the environment on that day. On the second day of year t_1, the minimum fear seen is compared to the minimum fear remembered on the first day, and the smallest value is remembered. A similar comparison is done for maximum fear. The procedure is repeated for each day in t_1. When the year is over, at the beginning of a new year, the minimum and maximum fear remembered in t_5 is forgotten, and the process is repeated.

The last change in the transition from the incubator to the simulator involves the determination of agent historical nodal wealth. For the baseline case, as well as all sensitivity studies done for this paper, the historical nodal wealth is the average agent generalized wealth that has been present on a node over the previous 10 years. In the incubator, this is easily calculated, since 10 years equal 10 time-steps. However, in the simulator, determining historical nodal wealth is not so easy. When an agent first moves into the simulator, it has 10 years' worth of data from the incubator and nothing from the simulator. On the first day of a year in the simulator, the agent generalized wealth present on each node that day is determined, and the value at each node is divided by T_y. Since the historical nodal wealth is calculated by using 10 years' worth of data, on the first day of a year in the simulator, the data would be taken from the first day of year t_1; all of the data collected for years t_2 through t_{10}; and 364 days out of 365 for year t_{11}. In this way, the historical nodal wealth is still the average nodal wealth over a 10-year period. In general, for any day in year t_1, the historical nodal wealth at a given node can be calculated from the equation:

$$HNW_{1,k} = \frac{1}{10}\left[\sum_{x=1}^{k} GW_{1,x} + GW_2 + \ldots + GW_{10} + \left(\frac{T_y - k}{T_k}\right) GW_{11}\right], \tag{6}$$

where $HNW_{1,k}$ is the historical nodal wealth calculated on day k of year t_1, $GW_{1,x}$ is the agent generalized wealth present on the node on day x of year t_1, GW_2 is the total agent generalized wealth present on the node over year t_2 (similar for GW_3 through GW_{11}), and T_y is the number of days in a year. On the last day in the first year of the simulator, year t_1's contribution to the historical nodal wealth is based on T_y days, or one full year, and year t_{11}'s contribution has shrunk to zero. On the first day of the next year in the simulator, all of the year subscripts are incremented by adding one, and year t_1 once again represents the current year, and year t_{12} is forgotten.

RESULTS AND DISCUSSION

The process described above was implemented by using MatLab (MathWorks 2002).

Generating Initial Conditions with Various Incubation Cut-off Times

The simulation of initial conditions involves a two-step process. First, pre-incubation conditions are formulated by using input parameters, where some input parameters define deterministic characteristics of either the environment or the agent population and others define ranges for uniform distributions. The pre-incubation conditions are then used to begin incubation where the time-step is analogous to 1 year. Upon termination of this incubation at a specific time, the ending conditions of both the environment and agent population are recorded. The post-incubation conditions eliminate much of the bias introduced by the user input parameters and the methods used to generate the pre-incubation conditions. More significantly, the post-incubation conditions will also typically represent an agent population acclimatized to its environment. The agents have had time to evolve and gain knowledge of their environment. The post-incubation conditions are then used to begin a simulation in which the time-step is much shorter. The occurrences during these simulations are the occurrences of interest.

In performing this process and analyzing the results, two scenarios were considered. First, a single set of pre-incubation conditions was generated. The input parameters are those defined throughout the previous sections of this paper. These conditions were then used to begin an incubation run. From this incubation run, post-incubation conditions were recorded after 200, 700, and 1,200 time-steps (years). Figure 4 shows the total nonsecurity agent population over the time this incubation run was performed.

In each case, when the post-incubation conditions were generated after 200, 700, and 1,200 years of incubation, they were used as the initial conditions for the simulator. Each simulator run used 365 time-steps per year. Figure 5 shows the total non-security agent population over a 5-year period for the cases where the initial conditions are based on 200-, 700-, and 1,200-year incubations, respectively.

Qualitatively, the simulated population histories continue on from the point at which the incubator left off. For example, when the post-incubation conditions were based on an incubation run of 200 years (see Figure 4), the incubator showed a relatively small population (less than 100 non-security agents) with a relatively small growth rate. In the corresponding simulator population history (see Figure 5), the population continues to be low and the growth rate continues to be small. The opposite is the case in which the post-incubation conditions were based on an incubation run of 700 years. Here the incubator had a moderate population (over 300 nonsecurity agents) and was experiencing rapid growth. In the corresponding simulator population history (see Figure 5), the population is moderate and the growth rate continues to be rather rapid. In the case in which the post-incubation conditions were based on an incubation run of 1,200 years (see Figure 5), the population is large (over 600 non-security agents) and rather stable but becomes somewhat cyclic. At the time that post-incubation conditions were generated, the population was climbing toward the upper cusp of one of those cycles. Not unexpectedly, the simulator shows a large population with a moderate growth rate.

At 200-, 700-, and 1,200-year incubation times, the growth rates shown in Figure 5 are superimposed on Figure 4. In each case, the growth rate from the simulator was significantly greater than the growth trend seen in the incubator in the same time frame. Although the causes of this phenomenon require further experimentation, several observations can be made at this time. When 1,200-year incubation was used, the corresponding simulator growth rate closely

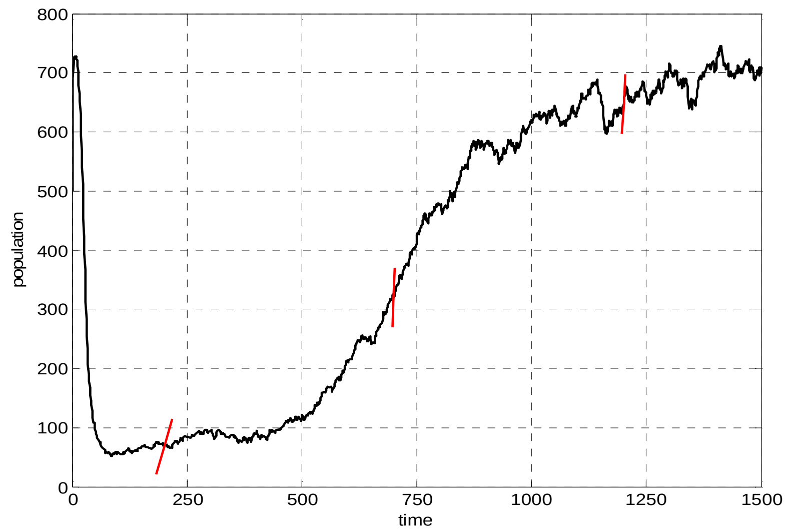

FIGURE 4 Incubator population history over 1,500 years

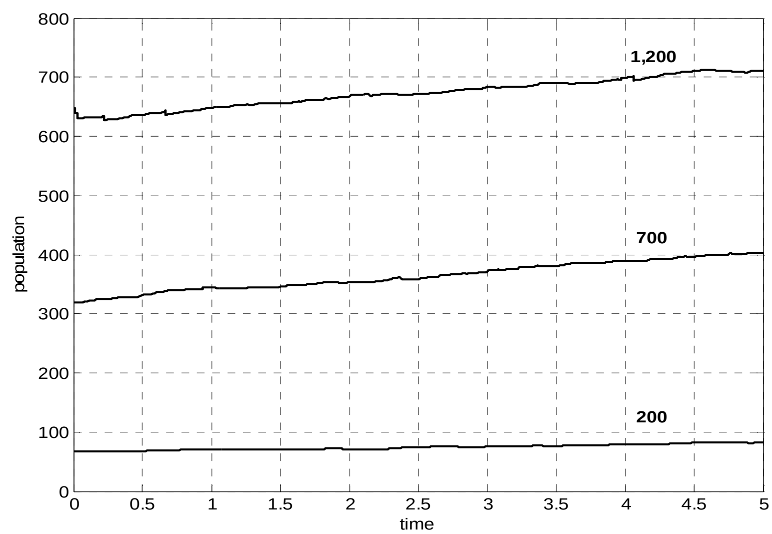

FIGURE 5 Simulator population history from using initial conditions from 200-, 700-, and 1,200-year incubations

matched the growth rate seen in the incubator at approximately 1,200 years over the 5-year period considered. Additional experiments using 5-year simulations in the simulator show that as T_y approaches 1.0, the growth rates in the simulator approach the growth trends in the incubator, as would be expected. Other experiments using 20-year simulations in the simulator with a T_y of 12 and 52 show that the growth rates in the simulator approach the respective growth trends exhibited in the incubator.

The behavior of the total nonsecurity agent population when shifting from incubator to simulator is important to consider. Major changes are neither expected nor desired. What happens in the incubator occurs at a particular rate per year; in the simulator, that particular rate per year is expected to fall within the same range as it does in the incubator, except one year is divided into multiple time-steps. In essence, the incubator speeds through history; the simulator moves slowly from the present into the future. With regard to population (and total population is indicative of many aspects of individual agent behavior), this transition is smooth and appears to be insensitive to the time in the incubation when the transition takes place.

Of course, the population history is not the only comparison of interest. A large amount of data are extracted from both the incubation and simulator runs. For example, the mean cultural identity can be plotted for both the incubation and simulator. The mean cultural identity coming out of the incubator remains relatively stable throughout the simulator. The same is the case for the variation in the cultural identity. The portions of the population with a cultural identity between 0–9, 10–18, 19–27, 28–36, and 37–45 were also examined. The population coming out of the incubator had cultural identity demographics very similar to those throughout the simulator. Other agent population characteristics where also examined, including: wealth, age, vision, resource metabolisms, innate nervousness, and initial endowment. The agent population characteristics at the end of the incubator were similar to the characteristics throughout the 5-year simulator run regardless of the incubation time. For example, when incubation was terminated after 1,200 time-steps, the average agent vision was 10 nodes in the four cardinal directions. The maximum vision in the population was 12; the minimum vision was 9. Throughout the simulation, the maximum and minimum remained the same, although the average vision showed a very slight increase. This behavior was typical of the other characteristics defining the population, including: wealth, resource metabolism, age, death age, and innate nervousness. For the 200-year incubation, more changes were observed in the simulator. The population was smaller; consequently, births and deaths had a larger impact. Similar behavior was observed for the 700-year incubation. The fairly limited adjustments between the incubator and the simulator do not drastically affect the agent population.

The next issue to be examined is the effect on agent behavior as the agents move around the environment. To examine this, consider the historical nodal wealth averaged over the last 10 years of the incubator — specifically the case in which the incubator was run for 1,200 years. Figure 6 is the historical nodal wealth contour plot, showing the average agent generalized wealth present on a node over the last 10 years of the incubation. After the 5-year simulation, another historical nodal wealth contour plot was generated. In this case, the historical nodal wealth comprises the last 5 years of the incubator plus the additional 5 years of the simulator (i.e., it is still based on 10 years of data). Figure 7 shows historical nodal wealth after simulation.

A comparison of Figures 6 and 7 shows no notable differences between the two. Some minor changes in magnitude and variations in contour shapes may exist, but, for the most part, a detailed comparison indicates agent behavior is unchanged between the incubator and simulator. This is especially true for areas where the historical nodal wealth is relatively high and less true for areas where the historical nodal wealth is relatively low. In the simulator, agents are maintaining their wealth and moving around the environment in much the same way as they did in the incubator. A similar comparison was done for population density contours, where the average nodal population over identical 10-year periods was considered. The results of that analysis showed that nodal population density also showed few significant changes between the incubator and simulator.

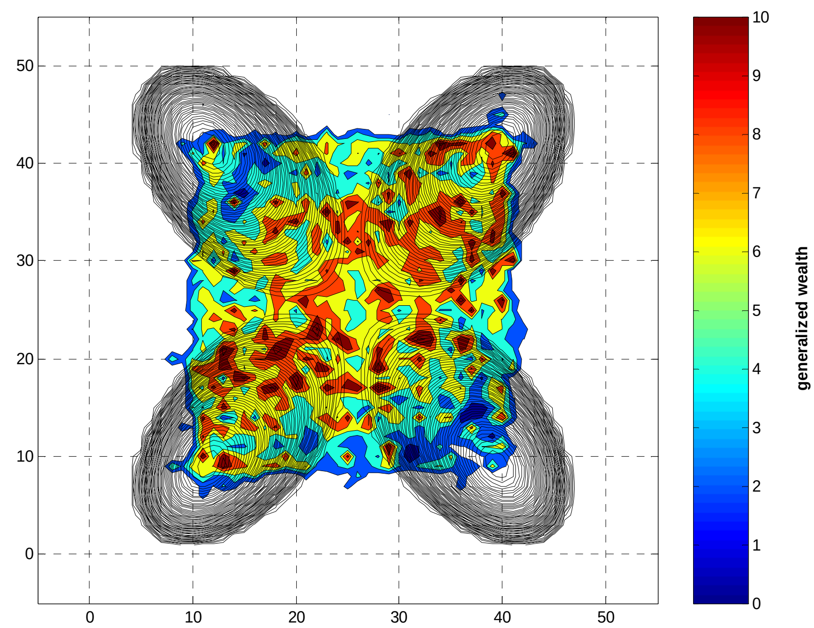

FIGURE 6 Historical nodal wealth over the last 10 years of a 1,200-year incubation

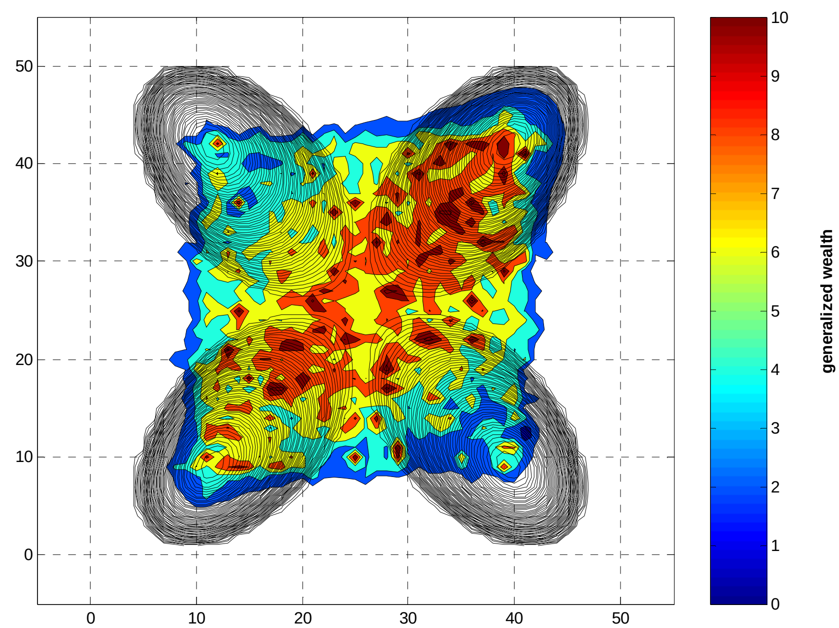

FIGURE 7 Historical nodal wealth after 5 years in the simulator

The formation and behavior of terrorist agents, as well as the generation of security agents within the population, varied significantly among the three scenarios considered here. When the incubation was stopped after 200 years and the simulation began, the population was small, and relatively little terrorist activity had occurred in that short period of instability. Low levels of terrorist activity led to only a small or nonexistent security population. After the transition from incubator to simulator, no additional terrorist activity occurred, and only minimal security was present over short periods of time (a result of fear created by the attacks in the incubator). When the incubator was run for 700 and 1,200 years, the populations were either growing rapidly or were relatively stable, respectively. Terrorist activity within the incubator resulted in a constant security presence by the end of the incubation runs. After the transition from incubation to simulation, the trends appeared to continue. In both cases, multiple terrorist attacks occurred in the 5-year simulations, and security levels jumped upward just after the attacks and slowly declined in the time following. When only the transition between incubation and simulation is considered, the behavior related to terrorism and security appears to be relatively constant even after the time-step definitions change.

However, the transition between incubation and simulation does not appear to be seamless. While the agent demographics, terrorist and security behavior, and aggregate behavior of the simulator remain fairly steady, there is a short time in the beginning of the simulator run in which the behaviors of agents seeking resources change. When an agent seeks its critical resource, it has a tendency to try to *equalize* its generalized resources. If one generalized resource drops below the others, the agent begins searching for the corresponding resource. Provided that the agent remembers where to find adequate concentrations of all the resources to ensure survival, the agent's generalized resources will reach an equilibrium in which each generalized resource is "close" to being equal. How close together the generalized resources can get depends on the magnitudes of the resources collected during a time-step. Since in order to maintain wealth, environmental resource concentrations have been factored by $1/T_y$ in the simulator, when the agent population moves into the simulator, there is a period of time during which an agent will attempt to equalize its generalized resources. Since an agent can collect only a small amount of a given resource at a time-step in the simulator (whereas in the incubator, it can collect 1 year's worth of that resource), the agent has a tendency to spend more time seeking its critical resource. In other words, the agent will spend more time on the pile of its critical resource. Since all agents are doing the same for their critical resources, a noticeable change in behavior can be observed in the early phase of a simulator. This behavior will continue to occur until the differences between each agent's generalized resources are reduced. Once the generalized resources have been equalized, the agent resource usage will once again stabilize. The agent's behavior when seeking resources will again look like it did in the incubator.

The time required for the simulator to stabilize varies with how stable the agent population is at the time of simulation. For example, when the incubator ended after 200 years, the simulator required approximately 1 year to stabilize; in this case, the population is small but there is little growth. When the incubator ended after 700 years, the simulator required approximately 2 years to stabilize; in this case, the population is established but experiencing rapid growth. And when the incubator ended after 1,200 years, the simulator required as little as one-half year to stabilize; in this case, the population is large and relatively stable. Thus, the change in the behavior of agents seeking resources always occurs, no matter how long the incubator is allowed to run. However, the more stable the population, the quicker the simulator stabilizes.

Incubation Runs Using Identical Input Parameters

The next issue to examine is the variability of post-incubation conditions generated by using one set of input parameters and their effects on 5-year simulator runs. For this analysis, four post-incubation conditions were generated, again by using the input parameters discussed previously, and the post-incubation conditions were based on 1,200 years of incubation. The only difference in each of the trial runs was produced by using a different random number seed to generate the initial population. After the post-incubation conditions were generated, each was used to begin a single 5-year simulation, where one year was subdivided into 365 time-steps.

Each of the four incubation runs was successful. A stable population evolved in the environment, and after 1,200 years, the post-incubation nonsecurity agent populations were all about 600 to 700. (Incubation run 4 had an ending population just below 600. This lower population was the result of an increased level of terrorist activity.) The population history for the first incubation run is shown in Figure 4. The behavior over time was similar for all four runs. While some differences did exist (e.g., in the second incubation run), the population dropped to approximately 25 agents around year 200, and the basic shape of the population history remained the same. In each case, the initial population of 500 agents grew rapidly for a brief period of time, then crashed, struggled to gain a foothold in the environment as the agents gained knowledge of their environment and evolved, and this was followed by rapid growth and finally a decreased growth rate as the population stabilized.

Agent characteristics (including, vision, resource metabolisms, wealth, and innate nervousness) evolved in a similar fashion for all four incubation runs. For example, by the end of 1,200 years, in each case, the average agent vision evolved to include approximately 10 nodes in the four cardinal directions. The average agent vision going into the incubator was 5 nodes in the four cardinal directions. The agent resource metabolisms at the end of the incubator were approximately the same. Agent age demographics were also steady between the four incubation runs. Agent generalized wealth (including the absolute maximum and minimum, as well as the average) was also similar after 1,200 years of incubation. In each case, agent innate nervousness (originally uniformly distributed between 0 and 1) averaged somewhere between 0.4 and 0.6, with the absolute maximum and minimum at +0.1 or –0.1 of the average. Thus, the agent characteristics defining the post-incubation agent population were very similar among the four incubation runs.

Even though most agent characteristics were similar among the four incubation runs, the agent cultural identities differed significantly, and this led to differences in the level of terrorist activity and the corresponding level of security present in the environment. Essentially the incubation runs fell into three categories: (1) runs 1 and 2 had a moderate level of terrorist activity over the 1,200-year incubation (for run 1, 20 terrorist attacks occurred, 11 terrorist agents were arrested, and the security level was approximately 3.6% of the nonsecurity population; for run 2, 28 terrorist attacks occurred, 23 terrorist agents were arrested, and the security level was approximately 5.0% of the nonsecurity population); (2) run 3 had significantly more terrorist activity over the 1,200-year incubation (70 terrorist attacks occurred, 121 terrorist agents were arrested, and the security level was approximately 10.0% of the nonsecurity population); and (3) run 4 showed significantly less terrorist activity over the 1,200-year incubation (11 terrorist attacks occurred, 1 terrorist agent was arrested, and the security level was approximately 4.0% of the nonsecurity population at the end of incubation).

Other information about the four incubation runs casts light on the reason for the differences in the conditions among the different incubation runs. First, the level of terrorist activity is directly related to the agents' cultural identity. For example, in incubation run 4, where the level of terrorist activity was relatively low, at one point during the incubation run (specifically, when the population reached its low point), the agent population was entirely made up of agents with a cultural identity between 19 and 27. While this anomaly was quickly eliminated by the population, by the end of the 1,200-year incubation, the group with a cultural identity between 19 and 27 was still dominant, and the fringe groups (those with cultural identities between 0 and 9 and between 37 and 45) were an extremely small portion of the total population. Likewise, in incubation run 3, the group with a cultural identity of 19–27 made up approximately 40% of the population, which led to significant increases in the groups with cultural identities between 10 and 18 and between 28 and 36 (approximately 20% and 25% of the population, respectively). This difference led to larger-than-usual fringe group populations.

Thus, the differences in the level of terrorist activity are directly attributed to the distribution of cultural identities in the population throughout the incubation run. The distribution of cultural identities is controlled not so much by the randomness of the pre-incubation population but by the randomness during the early "collapse" of that pre-incubation population (Figure 4). When the population crashes, the distribution of pre-incubation cultural identities can be drastically modified. Sometimes the distribution is flattened out, thereby adding to the fringe groups. At other times, the distribution is tightened up, and the population tends toward the median. Once the population begins to grow again, cultural interactions lead to further changes in the agents' cultural identities. It is reasonable to assume that given enough time, the distribution of cultural identity for the individual incubation runs would stabilize. However, since each of the four incubations was run for 1,200 years, different cultural identity distributions resulted.

Generally, when the level of terrorist activity was high at the end of an incubation run, the corresponding level of security was also high, and vice versa. But the correlation between these is not particularly strong. When the case with the least terrorist activity, run 4 (where 11 terrorist attacks and 1 arrest occurred) is considered, the security level of 4.0% coming out of the incubator was approximately the same as that of run 1: 3.6% (where 20 terrorist attacks and 11 arrests occurred). Keep in mind that the level of security is related to the level of fear seen in the environment, as well as the maximum and minimum fear seen over the past 5 years. Thus, sporadic terrorist attacks, which allow time for fear to dissipate, can lead to relatively low security levels. Such is the case for incubation runs 1, 2, and 4. In the case where the level of terrorist activity was substantially higher — incubation run 3 (where 70 terrorist attacks and 121 arrests occurred) — the attacks are no longer sporadic. The fear levels are high and consistently increasing; attacks are occurring throughout the environment and thereby effecting larger portions of the population.

One last observation is about the differences between these four different incubation runs and the resulting post-incubation conditions. Notice that only in incubation run 4 did the security get the better of the terrorists (121 terrorist agents arrested for 70 terrorist attacks). Only when the level of terrorist activity was high were the results of security really felt. Intuitively, this makes sense. When terrorist activity is low and attacks are sporadic, it is difficult to keep a sense of urgency in the population; consequently, the level of security is highly variable (increasing immediately after an attack, decreasing in the times when no attacks occur). This leads to a great disadvantage for the security agents. They are not present in substantial enough numbers to keep the threat level under control; therefore, the terrorist agents have the advantage. If the population is complacent about terrorism, the terrorists will gain an advantage, allowing them to conduct attacks with a much lower risk of being thwarted. However, as the level of terrorist activity increases, the level of security also increases, and the terrorist agents are subjected to a significantly greater probability of being arrested.

Post-incubation Initial Conditions in the Simulator

Each of the four individual sets of post-incubation conditions generated previously for 1,200 years of incubation will be used as the initial conditions for a single 5-year simulation where each year is subdivided into 365 time-steps. As was observed before, the nonsecurity agent population defined by the post-incubation conditions was stable throughout the 5-year simulation, with relatively slow growth. The agents' characteristics also remained stable, including vision, metabolism, age demographics (average, maximum, minimum, and average death age), wealth (average, maximum, and minimum), innate nervousness, and initial endowment (average, maximum, and minimum). The cultural demographics of the population were also stable over the 5-year simulation. The cultural demographics measured were the average cultural identity, variation in cultural identity, and the portions of the nonsecurity population made up of agents with a cultural identity of 0–9, 10–18, 19–27, 28–36, and 37–45.

Considering the level of terrorist activity during 5 years in the simulator, the simulation associated with incubation run 1 had 5 terrorist attacks and 4 terrorist arrests. The level of security ranged between 1.5% and 6.5%, with a ballpark average more than 3% over the 5-year simulation. For the simulation associated with incubation run 2 (see Figure 8a-d), 7 terrorist attacks occurred, and 12 terrorist agents were arrested. The level of security ranged between 1.5% and 6.5% (Figure 8c), with an average of about 3% over the 5-year simulation. For the simulation associated with incubation run 3, 12 terrorist attacks occurred, and 13 terrorist agents were arrested. The level of security started the simulation around 15% and steadily declined to approximately 5%. Attacks at the very end of the incubation and early in the 5-year simulation caused a significant increase in security at the incubation/simulation interface. For the simulation associated with incubation run 4, 8 terrorist attacks occurred, and 5 terrorist agents were arrested. The level of security at the beginning of the simulator was about 4%, fell quickly down to 1%, steadily increased to an average of about 7%, and then declined, ending the simulation with an average of 6.5%.

Comparing the level of terrorist activity in the simulator to the results from the incubation reveals some trends. Incubation runs 1 and 2 exhibited moderate levels of terrorist activity, incubation run 3 showed a high level of terrorist activity, and incubation run 4 showed a minimal level of terrorist activity. Note that when these post-incubation conditions were used to begin a 5-year simulation, the level of terrorist activity roughly corresponded to that in the incubator. The level of terrorist activity experienced in simulation runs 1 and 2 was moderate (with 7 and 5 attacks, respectively; averaging 6 attacks in 5 years). Simulation run 3 experienced twice the number of attacks (12 attacks in 5 years). Simulation 4 had 8 attacks in 5 years. At the end of incubation run 4, the number of terrorist attacks was increasing, and the level of security was relatively low. Therefore, the post-incubation conditions depicted a population in a very different phase of its development than in the other three simulations.

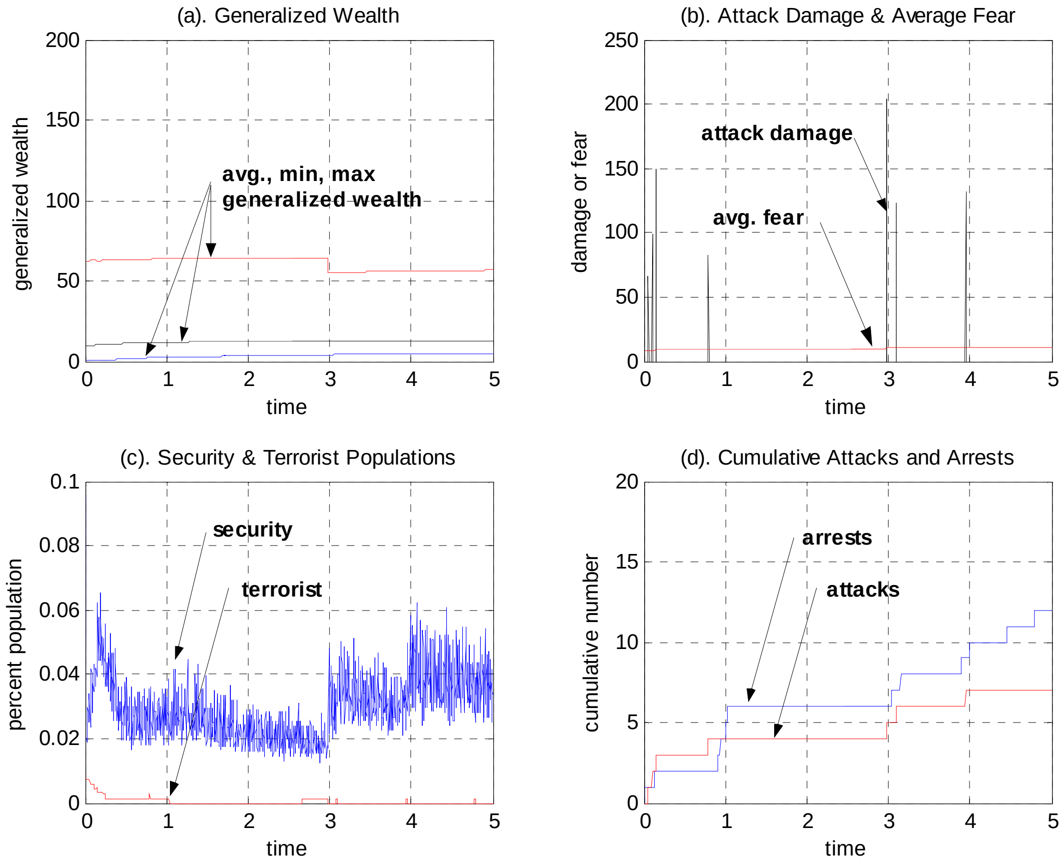

FIGURE 8 Terrorism and security activity plots for simulation 2

Further examination of Figure 8, which shows behavior typical of all four simulations, allows for some additional observations. First, all seven terrorist attacks are shown separately in Figure 8b; the attack damage represents the total agent generalized wealth destroyed by terrorists during that time-step. By comparing Figure 8b to 8c, the responsiveness of the security demand algorithm can be seen clearly. Early in the first year of the simulator, three terrorist attacks occur, and the security level spikes upward. Just prior to the end of year 1, a small terrorist attack occurs, and this attack causes a slight increase in the security level. Right around year 3, two terrorist attacks occur in quick succession; one very large one is immediately followed by another smaller one. The security level exhibits a significant increase before gradually declining. However, before the gradual decline in security can reach its pre-attack level, another terrorist attack occurs at the end of year 4. This attack causes another significant increase in security, higher than the previous jump, even though the attack was not nearly as large. The generalized wealth time history plot (average, minimum, and maximum generalized wealth for the nonsecurity population; see Figure 8a) provides a scale for the magnitude of each of the terrorist attacks relative to the agent population.

In all four simulations, there seemed to be an inordinate number of terrorist attacks occurring early in the first year. For example, for simulation run 2, Figure 8b shows that three terrorist attacks occurred in the first quarter of year one. The other simulation runs show similar scenarios. This phenomenon is primarily caused by the brief change in the behavior of civilian agents as they seek resources, which causes changes to the way the civilian agents are congregating. The terrorist agents see these increases in agent wealth concentrations and decide to conduct attacks. Once the population has settled down, terrorist behavior also settles down. This observation merely supports the previous assessment: the simulator requires a period of time to stabilize.

Effects of Changes in the Fear Memory

The sensitivity of changes in fear memory to terrorist activity and security levels was also examined. Although not directly related to simulating initial conditions, fear memory appears to be an important component in the behavior of the system. Three levels of fear memory were considered: 3-, 5-, and 8-year durations (the previous fear memory of 5 years is considered the baseline). A single incubation was run to 1,200 years for each of the fear memory cases. On the basis of the resulting post-incubation conditions, a single 5-year simulation was run for each case. All other parameters were set as previously discussed. The following discussion focuses on results from the 1,200 years of incubations. The results from the 5 years of post-incubation simulation exhibited similar trends.

First, consider the levels of terrorist and security activities. In Figures 9–11, side (a) shows the damage and average fear time history plots for the 1,200-year incubation phase for each of the three cases, and side (b) shows the cumulative number of terrorist attacks and terrorist agent arrests. A comparison of these figures indicates that the changes in fear memory have varied effects. For example, when the agents have a 5-year fear memory, 20 terrorist attacks and 11 terrorist arrests occur. When the fear memory is decreased and set at 3 years, 32 terrorist attacks and 35 terrorist arrests occur. Decreased fear memory appears to lead to an increase in terrorist and security activities. Likewise, when the fear memory is increased and set at 8 years, 82 terrorist attacks and 104 terrorist arrests occur. In other words, increased fear memory also leads to an increase in terrorist and security activities. Similar differences were observed between the corresponding 5-year simulator runs.

The relationship between fear memory and terrorist and security activity levels is the result of complex interactions within the model. Consequently, the reasons for the observed differences in the terrorist and security activity levels for the 3-, 5-, and 8-year fear memory cases can only be gleaned from a more thorough analysis of the results. For this reason, the overall historical nodal wealth should be considered. The overall historical nodal wealth is simply the average agent generalized wealth that has been present on each node over the entire incubation period of 1,200 years. Figures 12, 13, and 14 show the overall historical nodal wealth for the 3-, 5-, and 8-year fear memory cases, respectively.

Close examination of these figures will show that for the 3-year and 8-year fear memory cases, the concentration of historical nodal wealth is greater than for the 5-year fear memory case. In fact, in the 8-year fear memory case, where the terrorist and security activity levels were the highest, the concentration of historical nodal wealth was the greatest. In all three cases, the maximum, minimum, and average generalized wealth time histories were very similar; in other words, the agent populations were of similar overall wealth. These changes in the concentration of historical nodal wealth can be directly related to the terrorist agents' attack-triggering mechanism, since the value of a site includes the agent generalized wealth and historical nodal

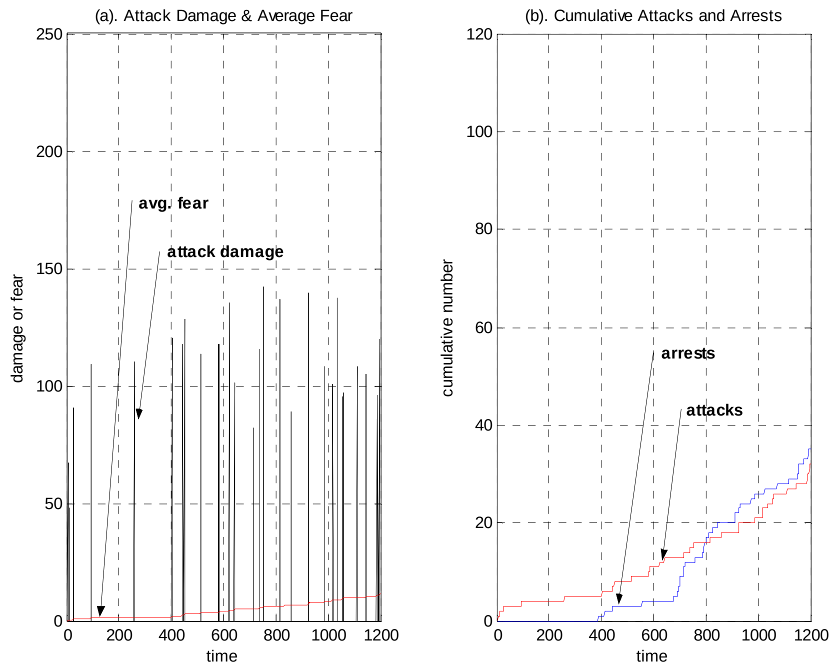

FIGURE 9 Terrorism and security statistics, 3-year fear memory

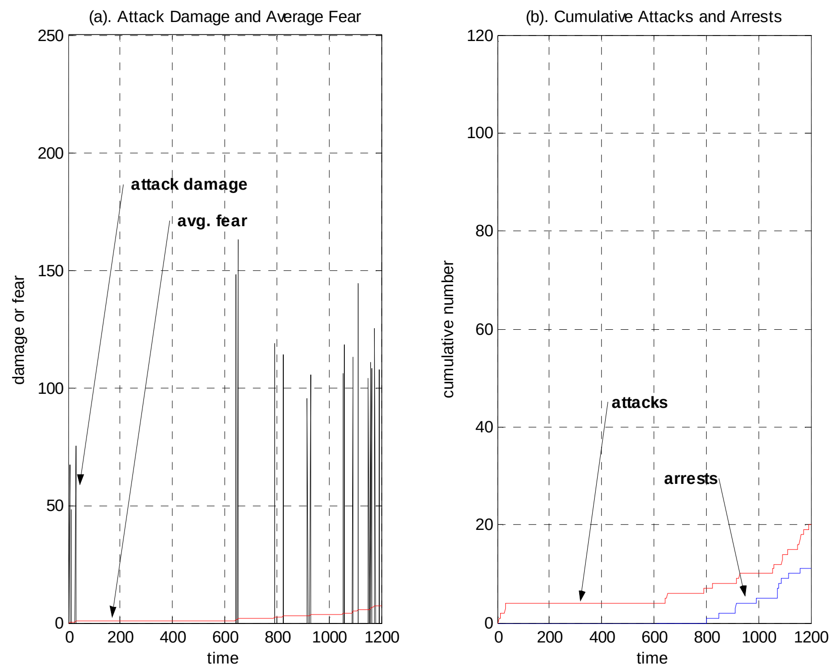

FIGURE 10 Terrorism and security statistics, 5-year fear memory

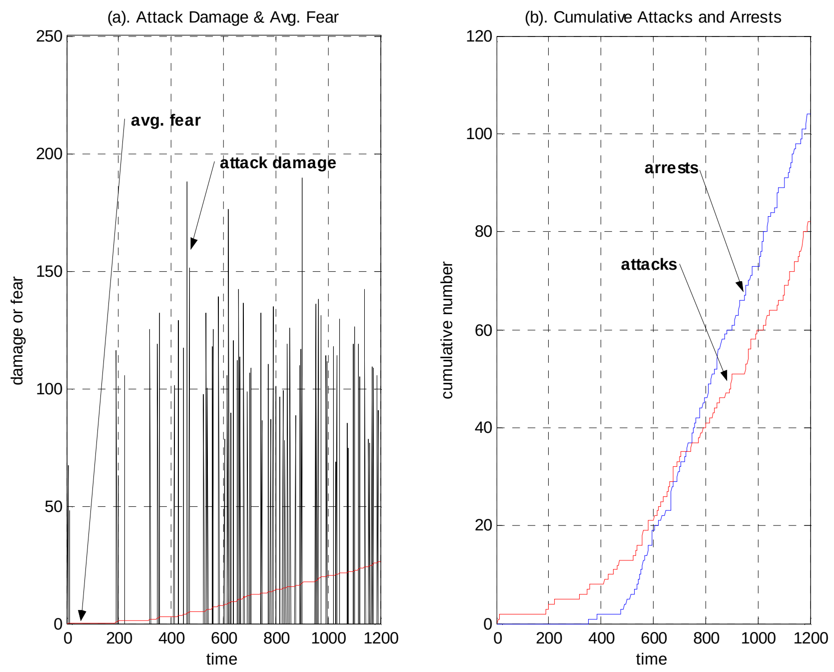

FIGURE 11 Terrorism and security statistics, 8-year fear memory

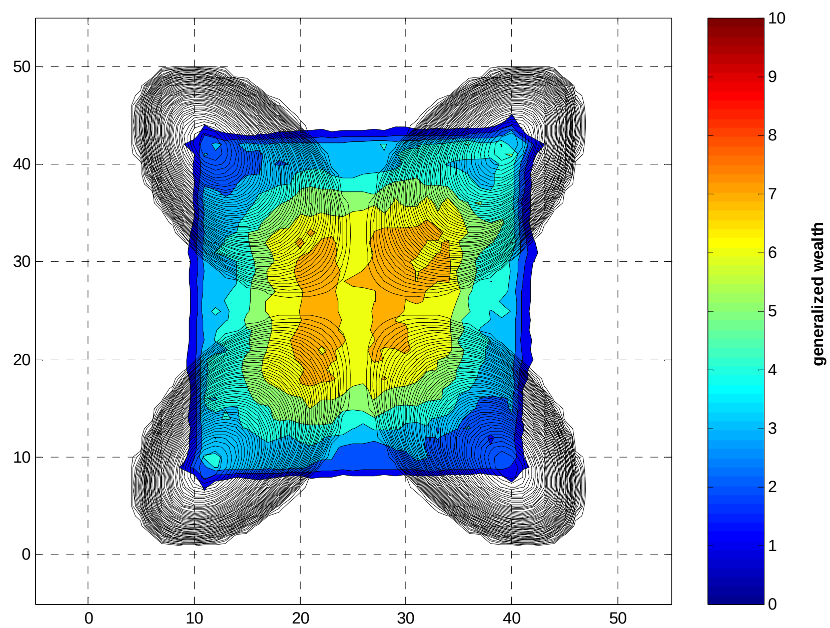

FIGURE 12 Overall historical nodal wealth, 3-year fear memory

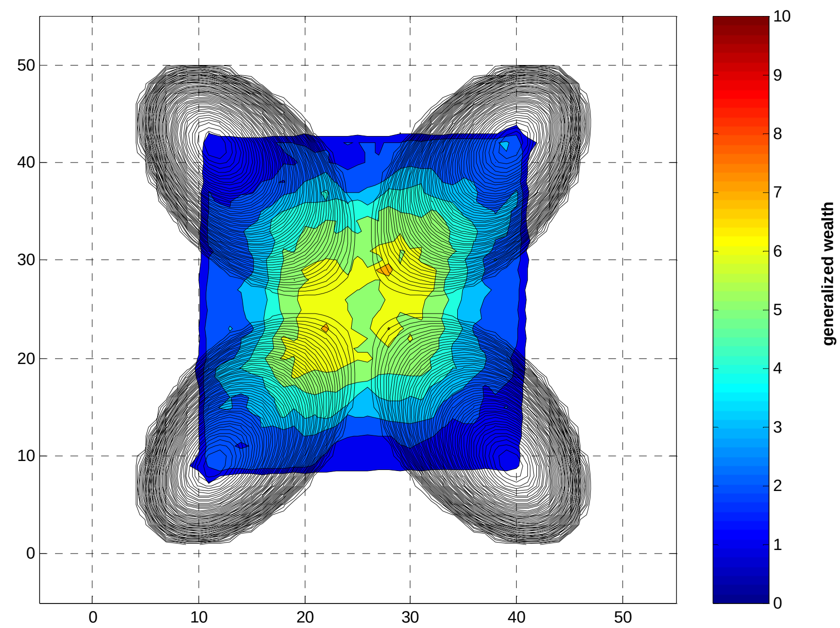

FIGURE 13 Overall historical nodal wealth, 5-year fear memory

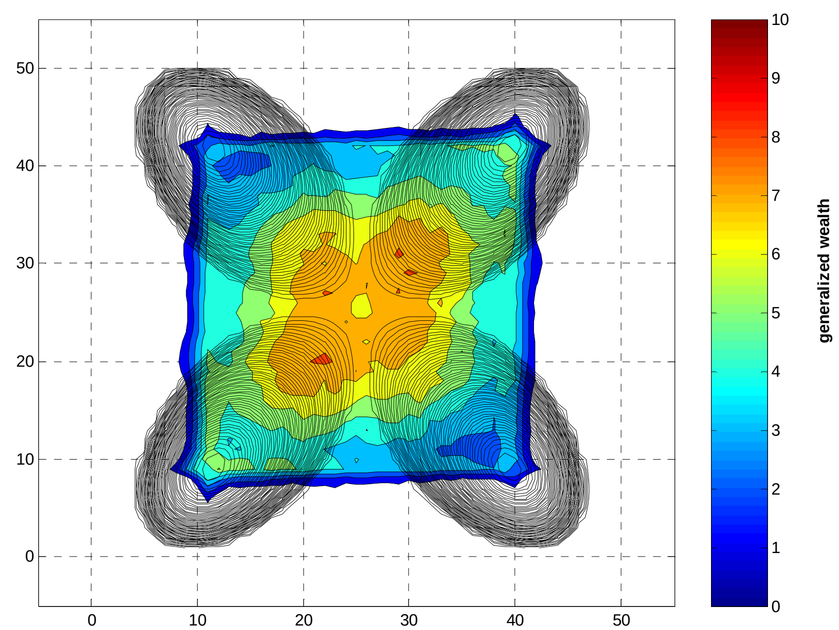

FIGURE 14 Overall historical nodal wealth, 8-year fear memory

wealth on that site's von Neumann neighborhood. The more concentrated the wealth, the more likely a terrorist attack will occur. Furthermore, if the wealth is more concentrated, the population is likely to be more concentrated too. (Overall population density contour plots verify this assumption.) The terrorist agents are more likely to be interspersed among this group and will therefore have the chance for more cultural interactions, helping to generate more agents with radical cultural identities. And finally, if more agents are in this area and it is an area where more terrorist attacks are occurring, then this will increase the security demand even further, resulting in more security agents. This is an important observation, since in the cases where the historical nodal wealth concentrations were relatively high, the security tended be more successful in arresting terrorists than terrorists were successful in carrying out attacks.

But none of this explains why different fear memory caused this increased concentration of historical nodal wealth. Recall that when civilian agents are searching for their critical resource, they are attempting to balance the reward a node has to offer (the amount of the critical resource present) with the fear to which they are subjected. Fear memory is an integral part of this balancing. The agents are looking to balance resources with fear so that they can perhaps avoid becoming a victim of a terrorist attack. However, there is apparently an optimal memory. Too great a fear memory leads agents to behave in a way that makes them become a prime target for a terrorist attack. Too little of a fear memory has agents moving around, oblivious to the risk; therefore, agents are again behaving in a way that makes them become a prime target for a terrorist attack. In other words, fear memory is a tool that can either help or hurt, depending on its scope. Further work is required to confirm that this effect of fear memory is a general trend and not just an artifact of this particular case.

CONCLUSIONS

A resource-based agent model has been developed to model terrorist activity. Endogenous terrorist agents are formed from within the civilian agent population by using a tag-mediated cultural identity. The terrorist agents conduct surveillance and commit terrorist attacks. When a terrorist attack occurs, fear is generated in the area subjected to the attack. Over time, this fear spreads out to the surrounding area. When searching for resources, civilian agents attempt to balance the rewards of visiting certain nodes with their fear of those nodes and with their innate nervousness. Civilian agents demand security on the basis of their wealth, fear, and nervousness.

When the model is run by using a time-step analogous to 1 year, the model is said to be acting as an incubator. After the incubator is run for a period of time, the agent population evolves from that set by the user to a population in tune with its environment. When the incubation phase ends, the post-incubation conditions become the initial conditions for a relatively short time-step simulation: the simulator.

A number of incubator and simulator runs were conducted. First the incubator was run and terminated at various points, thereby generating post-incubation conditions at various stages of the population's development. Second, multiple post-incubation conditions were generated by using identical input parameters. On the basis of these results, the following conclusions could be drawn. (1) In the transition between incubation and simulation, the characteristics defining the agent population remain stable. (2) The level of terrorist activity and, hence, the level of security in the environment remain consistent between the incubator and simulator. (3) The incubator and simulator are robust, requiring no special criteria to be satisfied in order to generate initial conditions. (4) The behavior of civilian agents when searching for resources is altered when they enter the simulator but stabilizes after a period of time (the length of which depends on the level of evolution attained in the incubator).

A combined incubator/simulator simulation was performed for a 3-, 5-, and 8-year fear memory. The 5-year fear memory case was considered as the baseline. The following qualitative conclusions resulted. (1) An increase in the duration of the fear memory resulted in significant increases in the level of terrorist activity and consequently the security level. (2) A decrease in the duration of fear memory also resulted in an increase in terrorism and a larger security presence. (3) The increased terrorism seems to result from changes in civilian agent behavior — changes that lead to increased concentrations of wealthy agents. (4) The changes in the behavior of civilian agents may be allowing increased numbers of radical agents to form (increased population densities lead to more cultural interactions) and may be making it easier for terrorists to conduct attacks. Thus, the duration of fear memory (i.e., the amount of time that an agent remembers fear that it has seen) can have significant effects on the level of terrorist activity. The effect is nonlinear; too great or too little fear memory works against the civilian agents, promoting increases in terrorism. This effect must be examined more carefully.

ACKNOWLEDGMENTS

We would like to thank both the National Science Foundation Graduate Research Fellowship Program and the Michigan Technological University Department of Civil and Environmental Engineering. Their support has made this project possible. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the writers and do not necessarily reflect the view of the National Science Foundation.

REFERENCES

Axelrod, R., 1997, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton, NJ: Princeton University Press.

Bulleit, W.M., and M.W. Drewek, 2005, “An Agent-based Model of Terrorist Activity,” in *Proceedings of the North American Association for Computational Social and Organizational Science (NAACSOS) Conference,* on CD, Notre Dame, IN, June 26–28.

Cederman, L.-E., 2003, “Modeling the Size of Wars: From Billiard Balls to Sandpiles,” *American Political Science Review* 97(1):135–150.

Dorfman, R., and N.S. Dorfman, 1993, *Economics of the Environment: Selected Readings,* New York, NY: W.W. Norton & Co.

Epstein, J.M., 2002, “Modeling Civil Violence: An Agent-based Computational Approach,” in *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 3; available at http://www.pnas.org/cgi/doi/10.1073/pnas.092080199.

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom up*, Washington, DC: The Brookings Institution.

MathWorks, 2002, *MatLab Student Version with Simulink*, Version 6.5, Release 13, Natick, MA.

DISCUSSION

Social Simulation Applications

(National Security and Emergency Management
Friday, October 14, 2005, 3:30–5:30 p.m.)

Chair and Discussant: *Jonathan Ozik, Argonne National Laboratory*

Jonathan Ozik: The next paper was originally scheduled for the Public Policy Session this morning, but due to time constraints, Li agreed to move his paper to this session. So we will have four speakers instead of three. Let me introduce Zhian Li from Argonne.

Agent-based Model for Simulation of West Nile Virus Transmission

Zhian Li: After the invited speaker gave his talk, I wondered whether I should still give my talk. If his model is an elephant, mine is just a small ant. Even though he has an elephant there, my ant has more legs. So let me start by giving you a little background information.

[Presentation]

Ozik: Are there any questions? Yes.

John Sullivan: John Sullivan, Ford Motor. You had a map that showed the diagnosed cases of West Nile. I think it was for this year. It looked like it dipped to around 2,000 or something in that neighborhood. Of those cases, what percent would result in a fatality?

Li: The article says 6%; that would be 6% for all human-infected cases, which is high.

Sullivan: Yes. It's my understanding, if I've interpreted this correctly, that perhaps half the people in this room already have antibodies for West Nile; most of us have been exposed already.

Li: I think it's probably much lower than this rate because we show the probability of infection for humans is small. This is because most of the time we stay in the office. When you get home, you may go to a yard. The comparable rate of infection for the elderly, however, is much higher.

Unidentified Speaker: I'd like to build off of the previous question. When we looked at the maps that showed the spread of the disease, it seemed to go across the nation in a wave. Where had it been previously? There were fewer cases than the year before. Do you know the reason for that? Could it be a buildup of antibodies among those who've been infected but didn't know it?

Li: I do not know the reason. It's a very interesting phenomenon. Actually, you could ask an epidemiologist. It's showing a wave or sweep kind of propagation, but it's pretty much established all over the place. If we look at the patterns, though, it's progressing from east to west. It's a very distinguished pattern; why, I do not know. One of the purposes of this model is to create an environment to see why it's progressing. Maybe the bird population dropped significantly because of such a high death rate so that the number of dead birds dropped in the consecutive year.. It's reported that there is a very high casualty rate for the blue jay or American crow. Very high. But this is the primary, very primary vector, not of the host, but then you have a huge drop. How did it progress?

The other interesting point has to do with how the birds do over winter. This is a chicken-and-egg situation. Some people argue that it was the host that carried the birds over winter. Maybe the mosquito was the carrier. It's subject to study.

Ozik: Zhian, if you don't mind, I would like to comment here relative to the somewhat interesting dynamics of this situation, in particular, this year in Illinois. It was a drier year, and the number of cases initially went down compared to previous years. The weather has such an important effect on the mosquito population because of moisture. There's a large variability, so it's unclear what the actual trend is, whether it's increasing or actually stabilized at this point.

Li: Actually, there are more questions than answers.

Neil Silbert: Neil Silbert. I also noticed that there was a decline in some places. How would you represent abatement efforts? In Illinois, especially in the northwestern suburbs, there were significant spring abatement efforts. Can we see the impact of these things?

Li: At this stage, this model does not include that, but, in the future, we may have a parameter go through the mosquito reproduction cycle or the death rate. If you just get a step function, makes the model kill, mostly the adult mosquitoes, and you may see the population drop.

Silbert: Could you introduce an agent to represent abatement?

Li: Oh, certainly, you can introduce the government agents, so all the governments are doing something good for us.

Ozik: We have one more question.

Reginald Tucker-Seeley: Reginald Tucker-Seeley from the Harvard School of Public Health. My primary interests are in chronic disease, so I don't know much about infectious disease. I'm curious, however, about the wave. Was it a function of the development of better surveillance methods within the offices of public health across the regions that could then account for the type of turn that we saw with increasing prevalence of the disease across the states?

Li: I think that's certainly a big factor, but it is not the only factor. If that's the only factor, you would not see the wave. It is just a very distinct pattern.

Ozik: Thank you very much. Our next speaker is Venkatesh Mysore.

Multi-agent Modeling and Analysis of the Brazilian Food Poisoning Scenario

Venkatesh Mysore: Good evening. I'm Venkatesh Mysore from New York University, and I'm going to be presenting the paper, "Multi-agent Modeling and Analysis of the Brazilian Food-poisoning Scenario." My co-authors for this paper are O. Gill, who is a graduate student in my lab; Professor Mishra, who is our adviser; and R. Daruwala and Marco Antoniotti, who are also in our lab. Professor Saraswat was overseeing some of the work. Are any of these names familiar to you? We are very new to this community.

[Presentation]

Mysore [in response to unintelligible question]: That is bacteria, so what happened with it probably happened lots of places. People gathered for a very big event, and lunch was served. Because they had to plan for 8,000 people, they started cooking two days in advance; the food was improperly stored and it was summer. So bacteria were present in almost everybody's nails and several other places; it became infected, entered the food, and the weather conditions made it flourish. Hence, it gave out what is called a toxin, an Enterotoxin, which causes food poisoning.

[Presentation Continues]

Unidentified Speaker: When you give people information, the death rate increases. Does anybody have an explanation for this?

Mysore: When you give them information, they all go to the nearby hospital. The number of people overwhelms the hospital, and by the time they are transferred to another hospital, they are already too sick. So this transfer effectively makes them sicker. This is, again, the nature of the way the disease has been modeled.

Unidentified Speaker: How many simulation runs did you do?

Mysore: Actually, the effect of the random variables and the variations is tremendous, so this is, again, I think, an average about four or five, but in the paper, we've explained what we have done and why — these points about absolute and substantial variation. Still, it is meaningful to use these curves. You're welcome to read the paper, and I can explain after the talk.

Unidentified Speaker: Is this an average or is this a single run?

Mysore: No, this is an average. In most cases, it is an average. Again, what I would like to point out is not the absolute values, but the fact that consistently one trend prevails when you compare to cost. Those are the things that can help us evaluate our policy. We are never going to use any of these absolute numbers. We're going to say if this is what typically happens and if this particular policy consistently outperforms another policy. Hence, that is what we should choose. That would be the way somebody would use a tool like this.

[Presentation Continues]

Unidentified Speaker: I believe you have performed this simulation many, many times.

Mysore: That’s correct.

Unidentified Speaker: Have you performed any kind of a statistical analysis?

Mysore: No, not at all. Our aim was not to get absolute figures but to develop a bigger tool and to understand the sources of complexity in modeling catastrophic scenarios. There are mobile agents with complex behaviors. In our model, there is an external function that will look at some of their attributes and see how these things change with time.

[Presentation Continues]

Mysore: [demonstration of model] … milk is being supplied by the same place, same source. They all consume this contaminated milk one morning, and let’s see what happens. So this is Manhattan. We just picked a specific region.

Unidentified Speaker: This is hypothetical, right?

Mysore: This is hypothetical.

[Presentation Continues]

Mysore: So here the improvement over the previous thing, of course, is in the incorporation of realistic transportation constraints of a city. In addition, some people are constrained to move along the streets, and so this introduces new effects.

Unidentified Speaker: The thing is that they’re actually moving. In New York City, I would think that it would all be gridlocked.

Mysore: That’s true.

Unidentified Speaker: Are they dying on the way?

Mysore: This is what would happen as they spread out and reached the hospitals and populated them and so on.

[Presentation Concludes]

Ozik: We have time for a few questions.

Robert Reynolds: Bob Reynolds, Wayne State University. A couple things. One is that certainly, you can get general statistical trends that summarize population’s actions. On the other hand, as you suggested, you can get, for example, a particular scenario that emerges that you hadn’t expected, and often that’s information you didn’t expect because it’s paradoxical.

For example, a classic one in terms of flow is Braess’s paradox, where you have an additional route to remove congestion, and in fact you make things worse. You have the issue of too much information. You give people too much information, and in fact they all rush to the same place. In other words, a certain amount of disinformation would reduce that flow. That’s a paradox. We often think that everybody needs to know everything. In fact, you’re suggesting that some sort of timed or phased information may be more useful in terms of doing things. Also in your traffic pattern, I noticed that when people leave this event and rush to the hospital, there's no background traffic. If you throw in background traffic, you're going to get some interesting paradoxes as well that will produce some interesting results.

Mysore: Yes, I completely agree with you. There are several ways of expanding the model and making it more realistic, so we could add more transportation constraints, like you said. The effect of people who do not attend the church, or how their movement affects things; there are too many avenues for expanding this initial idea. Actually, in the paper we've summarized the different avenues that we plan to explore. Hopefully, in a year or so, we will have addressed some of those avenues.

Gabriel Istrate: Gabriel Istrate, Los Alamos. Do you plan to do a model checking and those kinds of things for your system?

Mysore: That's correct. The idea is, and the example may not have been very clear, but there is this ability to perform traces, so you can ask a query for all the traces. Is it true over all traces that if you reach a hospital when your health is about 0.4, you are guaranteed to survive?

Istrate: Have you done anything about verification?

Mysore: That's exactly what we want to do eventually.

Joanna Bryson: My question might tie in with the last one. First, though, I want to say that I really liked your presentation, particularly the way you presented the agents. We were just talking about how to characterize an agent, and you started out saying what can it do and then you showed the situations where it does something else. I think those are the two parts of what you need to describe an agent model.

Mysore: That's correct….

Bryson: However, I'm a little worried about your quote that we don't care about exact numbers. On the one hand, I can understand that in some contexts, but on the other hand, it looks like you had many more deaths than were evident in the case you reported, and so you do have to wonder about how that would affect the dynamics of your model if you have some piece like that so severely different.

Mysore: That's absolutely correct. Actually, in the paper, we've shown that we've considered different scenarios that produce the same numbers. However, only 16 people died out of 8,000, so that's not visually trackable. I wanted to show how people are transferred from one place to another. That means a lot of people have to be very sick for the transfer to be seen.

Bryson: [Unintelligible]

Mysore: Exactly, we cruelly kill people so you can understand very quickly what happens in a model. Otherwise, it's not clear at all what happens, and we had to start with fewer numbers and fewer hospitals and understand the behavior.

Bryson: Okay. I retract my reservations then.

Mysore: As for the numbers, the reason I made that remark is that at this point, we don't care about the numbers. At this point, we are trying to understand the complexity of modeling such scenarios. Next, we will go into specific scenarios. This was just a test case. We are actually working on various possible scenarios in New York City, which is why we had open maps. We're definitely getting there.

Ozik: I have one question. Did you learn something policywise from analyzing the Brazilian scenario?

Mysore: Yes, there are two conclusions. One is about the information we have about the triage policy. There are the two nonconclusions, if you will. You can cook up a triage policy, but you can find that in all cases it works exactly against your model. When you press into this research regarding the doctors, they immediately said, "Well, I know what's going on. You never move sick people." But that's what is typically done. People move in ambulances, and the doctors fly to the site of trouble.

Our next step is to add first responders. These are people who are equipped to rush to the scene of action and attend to the people so that the people themselves don't have to move around. When you're moving around when you're sick, you're going to die. You will be much sicker when you reach the hospital. That is obvious. The fact is that the triage policy is a three-state model, and people are sent away when the hospital is full; that is one conclusion.

The other conclusion is about information that has to be provided in the correct way. Actually, in Brazil, one of the reasons that only 16 people died was that they all returned to their homes then went to the nearby hospital. Notice that all went between 26 [hospitals]. They could have gone to only five hospitals if they had rushed to the hospital right away. These are insights that were not obvious when we started modeling.

Ozik: Thank you again. That concludes the question and answer session.

Politicians, Businessmen, Warriors, and Civilians: Analyzing the Complexity of the Iraq Complex

Ozik: Our next speaker is Armando Geller. He will be speaking about "Politicians, Businessmen, Warriors, and Civilians: Analyzing the Complexity of the Iraq Conflict."

Armando Geller: First, I would like to say what we are not doing. We are not modeling a Clausewitzian war, like a traditional war, and we are not modeling guerilla warfare. We are also not modeling networkcentric warfare or effects-based warfare. We're doing no consulting. Some people tell me that it's better that way.

[Presentation]

Unidentified Speaker: Where do clerics fit in here? Clerics, religious leaders? Are they part of the political …

Geller: They're part of the political sphere because the politician or the political subsystem was, in a certain way, created by an idea called the 'identity of politics,' which was formulated by Mary Kaldor [Centre for the Study of Global Governance, London School of Economics and Political Science]. The identity of politics is basically a manipulation strategy, and it doesn't matter if the manipulation is religious or political.

[Presentation Continues]

Geller: Warriors and businessmen affiliate with the politician and render the politician more powerful.

Unidentified Speaker: You say that businessmen avoid places where other businessmen are. Wouldn't you say that they avoid places where competing businesses are? There's always a tendency to aggregate with complementary…. In other words, there's a power in aggregation if you're doing complementary business.

Geller: Yes, okay.

Unidentified Speaker: That's an important thing, too, because aggregation then allows a group of businessmen to have more clout than a single individual.

Geller: Yes, but they will avoid places. Of course, they have a certain range of vision, and the vision is included in our model, so they can only see other businessmen within their range. That means that in a field of three by three, they already meet maybe two or more businessmen because there's bounded rationality. They don't see the whole field, and that's what emerges. That's a default run. A default run always has six politicians, the black dots. For the businessman, the gray ones, it has 75 warriors and 500 civilians. If somebody starts to count the warriors, there are more because they were recruited. What you see as a yellow shape is the politicians' power shape. That corresponds to his power, but it's also his vision. His vision is dynamic.

[Presentation Continues]

Geller: The most important aspect, I guess, is that fighting is highly unpredictable in our model, and this is in a certain way congruent to what we see in everyday life when we read about these conflicts. It is difficult to predict where and when fighting erupts again. But, again, that's just one rung.

Unidentified Speaker: Are you suggesting fighting is more opportunistic than strategic in this case?

Geller: Yes, absolutely, because that's also a reason. Sorry I didn't say that. We implemented this in Repast, but that's the reason why we did it as an agent-based model. We wanted to have it decentralized, and especially these contemporary conflicts seem to be very much decentralized. So that's why I also said we're not modeling guerilla warfare, which usually is quite centralized, as Tito has shown. We didn't model any kind of networkcentric warfare, where even the word implies a certain centralization.

These results are a little bit sketchy though, especially when you see the standard deviation. It's huge, and, well, the mean doesn't really say too much, but it gives us a certain point where we can stick to it to analyze the data.

[Presentation Concludes]

Ozik: We have time for a few questions.

Mysore: I'm very intrigued by your three-pronged model of social … in war. Maybe I missed it, but where are institutions like the police department? Where do they appear?

Geller: They do not exist. That's the whole story. As I said, it's an anomic space. I give you the same example, so we have two questions answered with one: how do you say that in English? It's the same thing with statistical analysis. We do not have statistical analysis because we cannot go there or because we don't. It's dangerous to go there. But all kinds of statistical or institutions that collect statistical data don't exist anymore because those are usually war-torn societies and they are absolutely in a condition where every kind of institution we know usually is destroyed. That's why the police do not exist, the military in a westernized way does not exist, and so on.

Mysore: One very related question. There's this notion of fighting, but there is also the notion of good fighting and bad fighting; that is, you can be defending your country and be quelling an uprising as opposed to being a terrorist. Again, I couldn't understand how that would fit in. Is it assumed that such a thing does not exist in the societies that you are studying?

Geller: Yes. Well, we do not really model a very complicated fighting. Our fighting is basically a random generator. I mean, they need to find each other, of course, but we didn't invest too much time on the fighting mechanism. There is no good or bad in our model. There's just opportunism. That's it.

Jager Wander: Jager Wander, University of Groningen, the Netherlands. I really enjoyed your presentation, and I also was impressed by the large variability and the trends you discovered, for example, the increase in fighting and victims. Have you thought about conducting statistical analysis on these data?

Geller: Yes, we're working on that. For example, we're wondering if our civilians are power law related, if the events are actually power law related. We had some first results. We also had some results on clustered volatility. At the same time, though, we were somewhat unsure, and we decided to not present it here because there still might be some problems.

Lars-Erik Cederman: I'd like to make two comments. The first concerns power laws because it's very interesting what you referred to. This research was actually written up in *The Economist* about power laws in civil wars. I'm somewhat torn about this. Obviously, I worked on power laws in wars, but I'm not so sure that this applies to all civil wars. At least I did some initial work on this using data from Uppsala — conflict data, gone beyond single civil wars. If you do it for a larger number of countries, the casualty numbers seem to be log normally distributed. This could depend on simply the population size of all these countries that went into that investigation.

Now, the other point is about Kaldor's concept of new wars. I'm quite nervous about all of that literature because at the end of the day, you really have to ask, "When was the first new war?" Almost all of the characteristics that are mentioned in Kaldor's work antedated the more recent period. So I just don't think there is a very clear category to speak of here, and that most of what is ostensibly new is actually quite old.

Geller: Absolutely. Thank you very much. Concerning power laws, we just thought we were giving it a try, and it's not comparable, for example, with your article, "Modeling the Size of Wars: From Billiard Balls to Sandpiles," but we wanted to give it a try. Maybe we won't do it so hard anymore.

As for your second question, you're absolutely right. I absolutely reject the term "new wars." There is, for example, a great paper by Stathis Kalyvas [Yale University Program on Order, Conflict, and Violence] on this topic, and he shows quite well why we should not term these kinds of conflicts new wars. That's why I was looking for a term less definitional, like "contemporary conflict." I also tried to identify them by naming these conflicts, but I absolutely agree, yes.

Ozik: Okay. Thank you again.

Simulating Initial Conditions in Agent-based Modeling

Ozik: Next is William Bulleit, and he'll be speaking to us about simulating initial conditions in agent-based modeling.

William Bulleit: I'm William Bulleit. Ph.D. student Matt Drewek is the one who's doing a lot of the programming work. We thank the National Science Foundation (NSF) Graduate Fellowship Program for funding him. He's a year and three months so far, so he's got over 18 months to go. And then of course the disclaimer that NSF would like, that they disavow any knowledge of what we're doing.

[Presentation]

Ozik: Are there any questions?

Mike North: I have a quick comment about the incubation.

Bulleit: Right.

North: Commonly, you see an issue of what we call 'spin-up' in simulations, you know, where you basically, but not necessarily, have to get things to a steady state, but you're compensating for the fact that the input data usually are very far from whatever equilibrium or standard pattern is found. There's actually been a reasonable amount that's been written about spin-up and how to deal with it. That might be something that's useful.

Bulleit: Spin-up is the term?

North: Spin-up is the common term, yes. Exactly.

Bulleit: Any more questions?

Kostas Alexandridis: I was wondering if you're using the probability of an attack as a public event for agents. In others words, does the probability of attack and the probability of not attack add up to one, or there is some kind of uncertainty that either of the parties for not knowing that probability?

Bulleit: The probability is what we would hope to get out of large numbers of simulations from the model, just like the probability of an earthquake. The probability of an earthquake is some number; the probability of not an earthquake is the rest of the set.

So this would be a similar thing, but it would be given a resource type. For instance, given a railway station, what is the probability that it will be attacked? If we can get it down to resources, resource types, then it becomes a conditional probability for that particular type of facility. I don't know if I answered your question, but if you think of it just as the probability of attack and then not attack, it's just that.

Alexandridis: Well, theoretically it makes sense, but in a practical sense, often there is the actual probability of an attack and the probability of non-attacks, and there is a region that is a kind of gray, which is that the probabilities of both as a strict event are kind of … we're not sure, or they are both large or both small.

Bulleit: It might have been better to have said the distribution of the probability. In a sense, you could say it's the distribution of the probability of attack, because, I agree, it's not going to be a single number. In a design case, though, you search for a high-end probability or one that you think you want to use for design. For instance, the probability of a magnitude 8 earthquake in LA is different than the probability of a magnitude 2 earthquake in LA. The same idea might be used here: you have the probability of a major attack that does significant damage, or the probability of a smaller attack that does less damage, or you hope it does less damage. The goal is to design so that does happen. Does that answer the question?

Alexandridis: Yes, thank you.

David Sallach: I was wondering, doesn't it worry you that you're in the business of reinventing social science by picking out what factors you think are going to represent and more or less arbitrarily assigning parameters? It seems like the most that you can hope for is kind of a common-sense framework. I wonder if that isn't worrisome. I mean, you've set it in motion and got something coherent out of it. It seems like a big gamble.

Bulleit: I would say, yes, it's potential — yes, it's worrisome in a way because, you're right, you could end up with nothing. But the other option is that this is a proof of concept and that even if we don't end up with the result we hope for, people more knowledgeable about the social behavioral models can do something more with it. I don't expect to reach the final answer. I don't expect to reach the final answer even at the end of three years. I hope we'll be on a path that would indicate that, yes, this is do-able, and maybe it needs to be less obviously wrong.

Sallach: I guess my question is more of a design methodology. That is, wouldn't it make sense to involve social science experts from the outside and try to incorporate those insights?

Bulleit: Perhaps. Like I say, we went at it from the very base level and have been trying to find those things that we think affect it based on the results. And you're right, we're not at a point where we can say, “Yes, we’ve screwed up, or no, we’re there.” But if it looks like we screwed up, we do intend to do exactly that. That would probably be future work. Remember, this is an NSF graduate fellowship for the graduate student, so we have to keep it in the balance of that funding as well.

Ozik: Are there any more questions? If not, let’s thank our speaker and all the speakers of this session.

Saturday, October 15, 2005

Computational Social Theory

Invited Speaker

Lars-Erik Cederman

GROWING SOVEREIGNTY: ORGANIZATIONAL SHIFTS IN STATE SYSTEMS

L.-E. CEDERMAN* and L. GIRARDIN,
Center for Comparative and International Studies,
Swiss Federal Institute of Technology, Zurich, Switzerland

ABSTRACT

Drawing on prominent theories of historical sociology, we model the emergence of the territorial state in early modern Europe. Our modeling effort focuses on systems change with respect to the shift from indirect to direct rule. Inspired by the historical logic, we take a first step toward formalization by introducing a one-dimensional model that helps us fix our thoughts about the tradeoff between organizational and geographic distances. To test our initial deductive findings, we also present an agent-based model that features states with a varying number of organizational levels. This model explicitly represents causal mechanisms of conquest and internal state-building through organizational bypass processes. The computational findings confirm our hypothesis that technological change is sufficient to trigger the emergence of modern, direct state hierarchies.

Keywords: Territorial states, agent-based model, international relations theory, systemic change, geopolitical model

INTRODUCTION

Conventional theories of political science assume not only that actors are fixed and given but also that fundamental actor types remain constant. For example, international relations theory postulates the existence of a system of states and then goes on to explore interactions among such actors. However, the world is not made up of only territorial states; moreover, these entities have not always been, and will not necessarily always remain, the most important actors in world politics.

The terrorist attacks of September 11, 2001, highlighted the subversive effect on territorial states from transnational, covert networks of terrorists. Some of these entities have been associated with civilizations (seen as large-scale identities following mostly religious lines). Furthermore, analysts have also drawn the conclusion that the European Union (EU) represents a new type of political actor that violates the traditional norms of territorial sovereignty. The EU, being neither a state nor an intergovernmental organization, represents a *sui generis* social form.

Given the fundamental importance of these developments, our theories need to catch up with the real world. We do so by exploring the emergence of the most important organizational form in world politics, namely, the modern territorial state. Focusing on changes in the internal structure of state organizations, we introduce a series of models that show how the territorial

* *Corresponding author address:* Lars-Erik Cederman, Center for Comparative and International Studies (CIS), Swiss Federal Institute of Technology, Zurich (ETH), Seilergraben 49, 8032 Zurich, Switzerland; e-mail: http://www.icr.ethz.ch/.

modern state came to dominate the Westphalian system. Setting aside other aspects of this transition, we model the shift from “indirect rule” to “direct rule.” This change entailed a flattening of medieval hierarchies, which allowed state leaders to govern their territories without intermediaries. Rather than assuming sovereignty, we attempt to grow it.

We start our investigation by classifying different types of change in international systems. This exercise in conceptual ground-clearing indicates that our analysis must confront the problem of systems change, of which the emergence of the territorial state is an important special case. Highlighting the shift from indirect to direct rule, we survey some of the best writing in historical sociology for ideas about the mechanisms that brought about this case of systems change. Inspired by this historical logic, we take a first step toward formalization by introducing a very simple, one-dimensional model that helps us fix our thoughts about the tradeoff between organizational and geographic distance. To test our initial deductive findings, we present an agent-based model that features states with a varying number of organizational levels. The model explicitly represents causal mechanisms of conquest and internal state-building through organizational bypass processes. The computational findings confirm our hypothesis that technological change, captured by a sliding logistical function, is sufficient to trigger the emergence of modern, direct state hierarchies.

CAPTURING CHANGE IN WORLD POLITICS

While many unique aspects characterize the changes that the international system is currently undergoing, world history has always been in flux, albeit at varying levels of structural depth. This point is nicely illustrated by Gilpin’s (1981) taxonomy, which introduces three types of change in world politics:

- *Systems change* is the most fundamental type because it concerns the very nature of the units (i.e., the actor types).
- *Systemic change* is the next level, and it relates primarily to the emergence and disappearance of specific units, as well as to shifts in their outer boundaries.
- *Process change* is the least profound type since it relates to the dynamics and behavioral interactions among given units, such as cooperation, conflict, and alliance formation.

Despite the changes associated with the end of the Cold War, *process change* is still the bread and butter of international relations (IR) theory. Given a fixed set of actors, researchers typically study interaction patterns among them. Such analyses tend to focus on the balance between conflict and cooperation. Taking issue with structural theories that expect little change at all (e.g., Waltz 1979), neoliberal theorists explain the emergence of cooperation under anarchy. Drawing on Axelrod’s (1984) and other game theorists’ important result that cooperative strategies may thrive in iterated games, a whole generation of scholars has applied this logic to the development of cooperative regimes in various issue areas (Oye 1986; Keohane 1984). To a large extent, constructivist IR scholars account for the same phenomenon while relying on ideational factors, including processes of identity formation (Adler and Barnett 1998; Wendt 1999).

The next level of Gilpin's taxonomy shifts the focus to changes in specific actors and actor constellations. Viewed as an endogenous development, *systemic change* features the rise and fall of states. The so-called hegemonic theories of Gilpin (1981) and Kennedy (1989) explain the shifting fortunes of great powers as instances of systemic change caused by differential growth within the system. For various technological and organizational reasons, individual states are sometimes able to pull far ahead of the competition and assert themselves as hegemons. Over time, however, competitors tend to catch up with the hegemon, which ultimately loses its leading role, typically by losing a hegemonic war.

All these theories of systemic change assume that there is only one actor type in the system: the territorial state. Variation merely concerns the scale of the actors, without any reference to changes in their underlying structure. *Systems change*, on the other hand, implies a more radical transformational logic that introduces entirely novel organizational forms. This emphasis on novelty requires a much more flexible ontology than standard theories assume (Cederman 1997). This type of transformation is truly configurative rather than merely purposive because it presupposes "change in the organising principle governing the distribution of power and authority. It involves a change in the principles that determine the spatial configuration of politics, as well as the moral language used to justify the spatial order" (Reus-Smit 2002, page 137).

Analysis of systemic and systems change, as opposed to process change, requires explicit attention to the actors' corporate identities. A corporate identity denotes an actor's spatio-temporal extension, including the mechanisms that mark its spatial boundaries and that reproduce it over time (Wendt 1999, Chapter 4). Conventional constructivist theory is unable to deal with this type of change because of its exclusive preoccupation with social identities as opposed to corporate identities. Social identities merely specify the role repertoires of actors, such as friendship or enmity.

As suggested by Georg Simmel's process theory, a complete explanation of systems change needs to theorize changes both in terms of space and time (Cederman and Daase 2003). The spatial dimension highlights the actor's internal structure, its external environment, and its boundary to the environment. It is also necessary to specify a set of dynamic mechanisms that provoke changes in spatial structures over time. Evolutionary theory tells us that such processes rely on either natural selection or social evolution (Kahler 1999). Whereas the former involves a Darwinian winnowing of the best-adapted organizational forms, the latter features at least some adaptation of the actors' "internal models." It is also possible to imagine hybrid theories that combine natural selection and learning processes.

THEORIZING SYSTEMS CHANGE: THE EMERGENCE OF THE TERRITORIAL STATE

Any attempt to explain the emergence of the territorial state has to start by considering its constitutive principle: sovereignty. According to Bull's (1977, page 8) classical definition, states can be seen as

> "…independent political communities each of which possesses a government and asserts sovereignty in relation to a particular portion of the earth's surface and a particular segment of the human population."

Bull (1977, page 8) proceeds by dividing the notion of sovereignty into an internal and an external component:

> "On the one hand, states assert, in relation to this territory and population, what may be called internal sovereignty, which means supremacy over all other authorities within that territory and population. On the other hand, they assert what may be called external sovereignty, by which is meant not supremacy but independence of outside authorities."

As a corollary, it has to be inferred that the distinction between internal and external sovereignty presupposes sharp inter-state boundaries.

How did this configuration emerge over time? An answer to this question calls for a dynamic account of the three dimensions of sovereignty. Internally, sovereign rulers ridded themselves of internal competition within their territories. At the same time, they expanded their territories in the face of external competition. Together, these two processes generated increasingly thin and clearly defined borders.

Before searching for the mechanisms that brought about this complex process, we must first consider what preceded the modern territorial state of modern Europe. Indeed, sovereignty, as it is understood today, did not exist in the Middle Ages (Strayer 1970). Although territorial states had started to emerge as organizational cores, these were characterized by fading central control that occurred as distance from the capital increased. Moreover, centralized political control was limited by feudalism's indirect arrangement for broadcasting power (Poggi 1978, page 28). This was reinforced by the primitive state of the transportation infrastructure, which made it practically impossible to govern large units directly. Tilly (1990, page 104) explains that

> "…city-state, autonomous bishoprics, petty principalities, and other microstates ruled in a relatively direct way. Agents who were immediately responsible to the crown and served at the monarch's pleasure collected taxes, administered courts, tended crown property, and maintained day-to-day contact with local communities falling under the crown's jurisdiction. Larger states, however, invariably opted for some form of indirect rule, co-opting local powerholders and confirming their privileges without incorporating them directly into the state apparatus."

This situation in these large organizations thus stood in stark contrast to that of the smaller, directly ruled units (Tilly 1990, page 104):

> "Before the seventeenth century, every large European state ruled its subjects through powerful intermediaries who enjoyed significant autonomy, hindered state demands that were not to their own interest, and profited on their own accounts from the delegated exercise of power. The intermediaries were often privileged members of the subordinate populations, and made their way by assuring rulers of tribute and acquiescence from those populations."

Under feudalism, "war lords" ruled their own fiefs while offering the state core military services in exchange for the right to extract resources within their own territories. This created a hierarchical organization that rested on several layers of semi-autonomous control.

In terms of internal organization, we can summarize the process of state formation as a shift from "indirect rule" to "direct" rule. Feudalism's limits on the central state's power were finally swept aside by the ascendance of the modern, territorial state — a process that culminated with the nationalization of state power after the French Revolution.

This process had both an internal and an external dimension, the most obvious being the external one. Thanks to the revolution in military technology (Downing 1992), power centers amassed resources that triggered a snowball process of conquest (Gilpin 1981). War led to conquest, and conquest led to increased resources to wage future wars. Only the largest units managed to survive this cutthroat competition.

The internal facet of this process of geopolitical consolidation was equally important, because as we have seen, sovereignty not only implies freedom from external challenges but also effective subjugation of the emerging state's internal enemies. Again, Tilly (1990, page 69) eloquently describes the process:

> "Since the seventeenth century ... rulers have managed to shift the balance decisively against both individual citizens and rival powerholders within their own states. They have made it criminal, unpopular, and impractical for most of their citizens to bear arms, have outlawed private armies, and have made it seem normal for armed agents of the state to confront unarmed civilians."

Thus, if conquest was the key mechanism that transformed sovereignty's external dimension, a process of *organizational bypass* operated within the emerging state's territory. Organizational bypass occurs when the central ruler manages to supplant the authority of the inferior subunits with direct rule. As a result, the bypass process connects subjects directly to the capital, thus depriving the regional power center of its capacity to extract resources from below.

In reality, this process happened through a prolonged series of conflicts between the center and the subunits and through the consolidating state's gradual penetration of the provinces. As the state expanded its power, "struggles arose between center and periphery over the new 'right' to tax" (Finer 1974, page 98). For example, "Louis XIII, the seventeenth-century monarch who with the air or Richelieu and Mazarin rebuilt the armed forces of the French state, probably tore down more fortresses than he constructed. But he built at the frontiers, and destroyed in the interior" (Tilly 1990, page 99).

Such activities were followed by the replacement of feudal lords with tax collectors and governors who were placed under the direct control of the sovereign (Finer 1974; Ardant 1975). In France, organizational geniuses such as Colbert and Richelieu were the chief architects of this transformation, but other states in Western Europe carried out similar campaigns of centralization. Under the leadership of Louis XIV, "administration, tax collection, the levying, command and disposition of troops had *all* been fully incorporated into the master system as its *functions*, not any longer as so many demarcated sub-systems which were linked in to the center by one or a few prestigious individuals" (Finer 1974, page 114). Throughout the state's territory, weapons were seized, militias demilitarized, and private armies suppressed. Organizational bypass secured the modern territorial state's monopoly on political power.

Finally, the external and internal reconfiguration of sovereignty entailed the gradual crystallization of territorial borders. To be sure, the pre-modern world was characterized by a blurred distinction between domestic and international politics. In particular, empires had porous borders constituting “gray”, semi-anarchic areas between the power centers. In some cases, even the extent to which there was any difference between domestic and international politics is questionable (Ruggie 1993; Kratochwil 1986).

Although a complete account of state formation would describe this transformation of borders, we will assume the existence of sharply demarcated borders and focus entirely on the internal and external aspects of sovereignty. We assume that these two processes created the modern state. Although repeated acts of conquest consolidated the external dimension of sovereignty, internally, organizational bypass led to a gradual shift from indirect to direct rule.

PREVIOUS ATTEMPTS TO MODEL SYSTEMS CHANGE

Historical sociologists have analyzed systems change along all three dimensions in great detail. However, because of the processes’ overwhelming complexity, it is not surprising that few have attempted to model them formally. Nevertheless, several attempts have been made to model the external dimension, particularly through computational models of conquest that allow boundaries to evolve endogenously over time. Already in 1977, Bremer and Mihalka (1977) proposed a model of featuring conquest in a hexagonal grid, which was later extended and further explored by Cusack and Stoll (1990).

Building on the same principles, Cederman (1997) introduced a new generation of models called Geosim. These models share a common architecture that starts with a territorial grid of fixed and indivisible primitive agents that can be thought of as villages or counties (Cederman 2002). Those states that survive grow, and their boundaries expand endogenously through a repeated process of conquest. The resulting states, which are organized in a dynamic network, are hierarchical organizations whose capitals are linked to their respective provinces through direct, asymmetric relations of domination.

Nevertheless, because all these models hard-wire direct rule into their foundational ontology, they are not designed to generate new actor types along the internal dimension of sovereignty. In fact, the only dimension that is explicitly captured by these models is the external aspect of sovereignty. The repeated process of conquest does illustrate the growth of territorial size and leads to organizations that resemble modern territorial states, at least in terms of size. But this process says little about the organizational aspect of sovereign rule. In other words, the models in this research tradition feature only two-level organizations, in which provinces are subordinate to capitals. As we have seen, however, modeling the transition from direct to indirect rule requires an explicit representation of intermediate layers of organization. Without such ontological flexibility, it is impossible to explore the “flattening” of hierarchies that characterize the shift from indirect to direct rule.

Moreover, the notion of sharp territorial borders is hard-wired into the model specification. While secession is explicitly modeled in some versions, and culturally bounded units (such as nations) are modeled in others (cf. Cederman 2002, 2004), state borders remain sharply delineated, thus excluding the possibility of modeling fuzzy frontiers and competing internal or external sovereignties. In this sense, the models endogenize systemic change, while ignoring systems change.[1]

In this paper, we focus on systems change with respect to internal structure. The endogenization of boundary types, however, is an interesting topic that needs to be addressed in future modeling efforts. To our knowledge, there are very few models that capture systems change as the emergence of novel social forms. However, computational organization theory and agent-based modeling open new avenues of research in this respect (Lomi and Larsen 2001; Cederman 2005). Organization theorists have typically studied how structural change affects overall performance, but their efforts typically treat the organizational topology exogenously (e.g., Morel and Ramanujam 1999 — though also see Prietula et al. 1998).

Before considering the possibilities of computational modeling, we turn the focus to a simpler, deductive approach. In a series of general conceptual papers, Kochen and Deutsch (1969, 1974) propose a simple mathematical framework for the analysis of organizational decentralization within private or public organizations, including territorial states. Whereas most preceding models had focused on the location of control, Kochen and Deutsch explicitly study organizations' performance as a function of their centralization.

One of their key arguments is that "a system should be decentralized if the additional cost of communication for coordination that a centralized system must have exceeds the difference between the higher expected profit of the centralized form and the lower expected profit of the decentralization" (Kochen and Deutsch 1974, page 107). Relying on this logic, they specify a mathematical model that allows them to draw inferences about the optimal number of organizational levels. The key bottleneck that forces delegation to deeper levels of organization is the cost of coordinating the provision of services at any specific level of organization. Their reasoning is so abstract that it is not obvious what it implies in terms of state formation. In the next section, we therefore take a first step toward modeling the shift from indirect to direct rule. This transformation hinges on a tradeoff between geographic and organizational distances.

ONE-DIMENSIONAL MODEL OF HIERARCHICAL STATE FORMATION

In feudal systems, there were formidable obstacles to geographic mobility that could only be overcome by the power centers' improved geographical reach. The introduction of intermediate instances of taxation and control reduced the geographic distances that had to be overcome. We illustrate this principle with a simple mathematical model that highlights the logic of resource extraction. Following Dacey (1974), the model is represented in one dimension (see also Cederman 1995).

Although we make the assumption of a linear system for tractability reasons, there are some historical cases of state formation that did unfold in a linear fashion. Perhaps the best-known examples are river valleys surrounded by impassable terrain, such as the Nile (Carneiro

---

[1] Axelrod (1997, Chapter 6) proposes a "tribute" model of new political actors that may be at least a partial exception to this limitation. According to Axelrod's algorithm, collective actors emerge if there is a pattern of interactions that confirms a number of properties that are seen to be constitutive of agency. These include effective control over subordinates, collective action, and recognition by third parties that an actor has been formed.

1978). Relying on computational methodology, the following section generalizes the mechanism to two dimensions and puts it on a dynamic footing. The purpose of the present discussion is to illustrate the main logic of resource extraction in multi-level state hierarchies.

We start by demonstrating how the tax mechanism operates in a "line state" with three levels of hierarchy (Figure 1). The nodes of the tree are equidistantly spread out along the horizontal axis, starting at location 1 and ending with location 5, and are labeled accordingly. The logic of resource extraction can now be summarized. Each node produces a resource unit on its own in addition to extracting taxes from its inferior units. All of this income is subject to taxation by the next-higher unit at a fixed tax rate k that is discounted δ per distance unit. While k determines the organizational cost of traversing from one level to another, the parameter δ penalizes attempts to reach over large geographic distances. The subunits retain whatever resources are not passed up the tree.

Let us now assume that the tax rate k is 0.5 and the geographic discount rate δ is 0.8. This means that the "leaves" of the state tree pay $0.5 \times 0.8 = 0.4$ of their income in tax to the next-higher instance, which in every case is located only one step from their locations. Nodes 1 and 3 pay node 2, and node 5 sends its taxes directly to the capital, node 4. This means that these subunits can keep $1 - k\delta = 0.6$. The intermediate node 2 receives tax revenue 0.4 + 0.4 from the inferior nodes 1 and 3 and from one resource unit from its own territory, yielding a total of 1.8 units of "taxable income." Taxation to the next-higher instance proceeds in analogy to the first level, but, in this case, the distance-discounting will be more severe because the distance between node 2 and the "capital" at node 4 is two distance units. This means that the effective tax rate amounts to $k\delta^2 = 0.5 \times 0.8 \times 0.8 = 0.32$ and that node 2 has to pass on 32% of its 1.8 resource units (i.e., 0.576) to the capital, while keeping the remaining resources (1.224). We can now compute the total resources controlled by the capital. The tax revenue is 0.576 + 0.4 from nodes 2 and 5. In addition, the resource unit of the capital province has to be taxed at $k = 0.5$, although without distance discounting, which means that the total resources amount to 1.576.

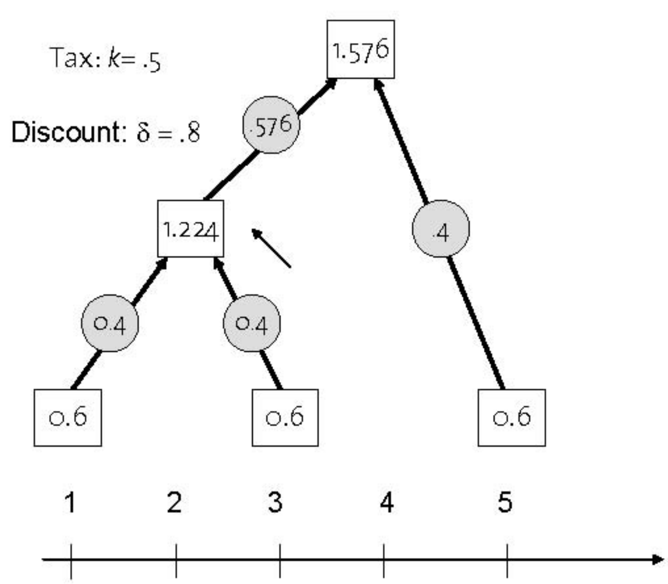

FIGURE 1 A simple example of a linear state with three levels of hierarchy

In sum, each level of the hierarchy taxes its subordinates while passing on a share of the revenue to the next superior instance. In this simple example, we assumed that the distance-discounting followed a simple exponential decay function $\delta(x) = \delta^x$. Functions of this type are often referred to as loss-of-strength gradients (Boulding 1963), and the functional form is often assumed to be exponential. By using the same discounting as in our first example, $\delta = 0.8$, Figure 2 illustrates how resource extraction declines as the distance from the capital increases. (The step function is discussed shortly.)

We are now ready to compare stylized versions of indirect and direct rule. In both cases, we represent the states' internal structure symmetrically. Figure 3 introduces a "chainlike" state of maximum depth, where the capital is located at zero. On each side of the capital, subordinate nodes are organized such that for each step away from the capital, a new, inferior level is added. Here distance is overcome through repeated acts of delegation, producing a very deep structure.

Thanks to the symmetry of the organizational form, it is sufficient to compute the resources extracted from one of the two branches $f_{IDR}(n)$. The total resources can be obtained as $p_{IDR}(n) = f_{IDR}(n) + k$, where k adds a taxed resource unit for the capital itself.

$$f_{IDR}(n) = k\delta(1)\{1 + k\delta(1)[1 + k\delta(1)(...)]\} = \sum_{i=1}^{n} k^i \delta(1)^i \quad ,$$

where $\delta(1)$ is the discounting for one distance unit. This is a more general functional form than the exponential decay function we relied on in Figure 1.

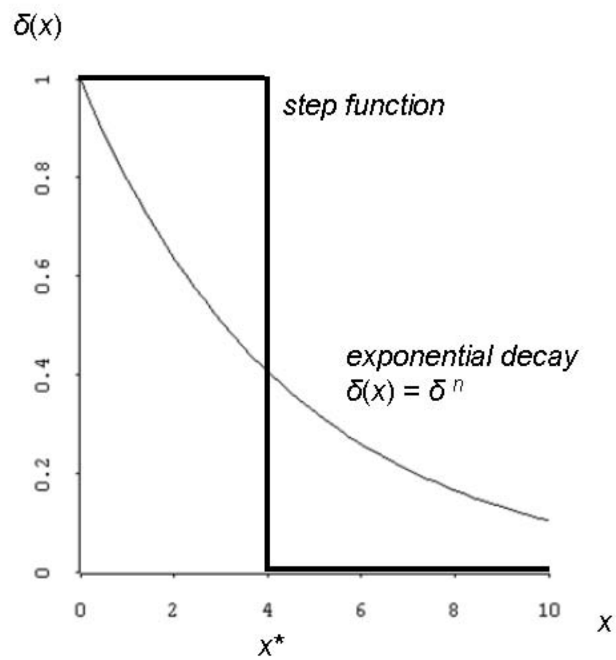

FIGURE 2 Two types of loss-of-strength gradients

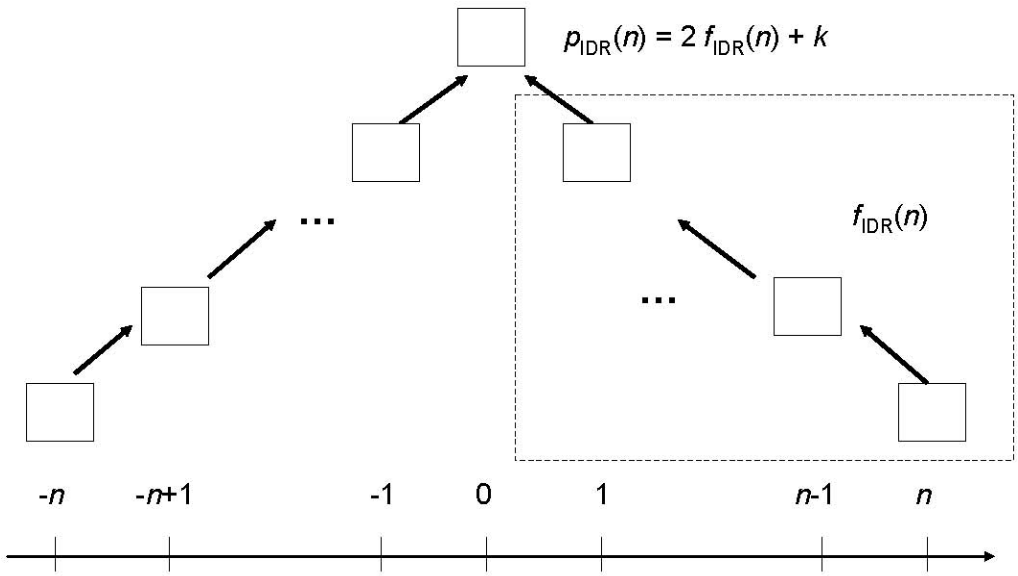

FIGURE 3 Symmetric, indirectly ruled state with a hierarchical depth of *n*

In stark contrast to the logic of indirect rule, Figure 4 illustrates the flattest possible hierarchy featuring a superior capital that rules all other nodes directly. Here it is straightforward to compute the total resources for direct rule as $p_{\text{DR}}(n) = f_{\text{DR}}(n) + k$, where

$$f_{DR}(n) = k\delta(1) + k\delta(2) + ... + k\delta(n) = \sum_{i=1}^{n} k\delta(i) \ .$$

In this case, the maximum discounting amounts to $\delta(n)$ rather than $\delta(1)$. Clearly, direct rule needs to overcome much more significant distances than does indirect rule.

We can now compare these two organizational forms for the exponentially decaying loss-of-strength gradient $\delta(x) = \delta^x$. This comparison yields our first proposition (see the appendix at the end of this paper for proof):

> *Proposition 1.* For an exponentially decaying loss-of-strength gradient $\delta(x) = \delta^x$, direct rule is always more efficient; that is, $p_{\text{DR}}(n) > p_{IDR}(n)$ for all $n > 1$.

We have just shown that, at least in our simplified, one-dimensional world, exponential distance-dependence cannot be responsible for a shift from indirect to indirect rule. Regardless of the value of δ, direct rule is superior. Thus, even if the loss-of-strength function shifts outward over the course of history, directly ruled states are always more efficient organizational forms. Under these conditions, feudalism could never have taken root!

What type of distance dependence has to apply for the shift from indirect to direct rule to occur? It is necessary to postulate a radically different loss-of-strength gradient in order to simulate Tilly's transformation. Let us instead assume a step function that drops from 100% resource extraction to zero at some distance x^* from the capital. We are now ready to state our second proposition (again, see the appendix for proof).

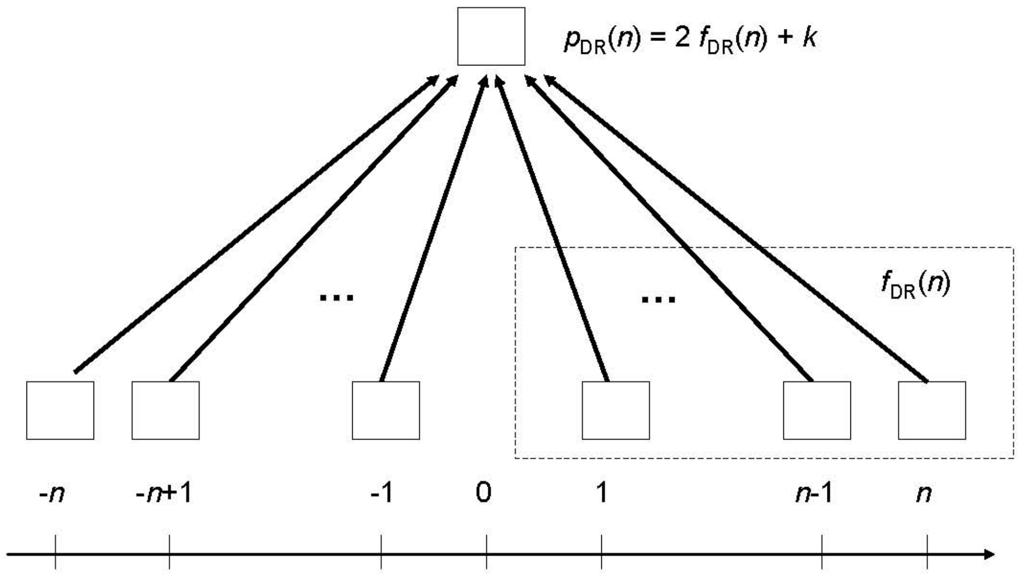

FIGURE 4 Symmetric, directly ruled state with the flattest possible structure

Proposition 2. Assuming that the loss-of-strength gradient is defined by a step function $\delta(x) = 1$, for $x \leq x^*$, and for $\delta(x) = 0$ for $x > x^*$, there are two cases:

Case 1. If $x^* < 1/(1 - k)$, direct rule is more efficient if

$$n < x' = \frac{\log[1-(1-k)x^*]}{\log k}$$

For all other values $n > x'$, indirect rule is superior.

Case 2. If $x^* > 1/(1 - k)$, direct rule is always more efficient.

Figure 5 provides a graphical illustration of the two curves. Direct rule grows as a linear function of n up to x^* and remains flat after that. In contrast, indirect rule yields a slowly growing increase in power that flattens out asymptotically at $k(3 - k)/(1 - k)$. For low cutoff points x^* below $1/(1 - k)$, direct rule is more efficient for low values of n up to x'. After this point, indirect rule surpasses its direct counterpart. However, for high values of x^* (i.e., wherever distance plays less of a role), direct rule is always the best choice.

We can now imagine a shift from indirect to direct rule by sliding the step function's cutoff point x^* from low to high values. In the former case, loss-of-strength gradients with little reach approximate the conditions of the Middle Ages. Here only the smallest political organizations can afford direct rule, as suggested by Strayer (1970) and Tilly (1990). This explains why the pre-modern world contained a combination of large, indirectly ruled empires and very small units, such as city states and other principalities characterized by immediate relations between rulers and the ruled. However, as the logistical conditions improve with increasing values of x^*, we enter the modern era. Now direct rule offers the superior logic, regardless of scale.

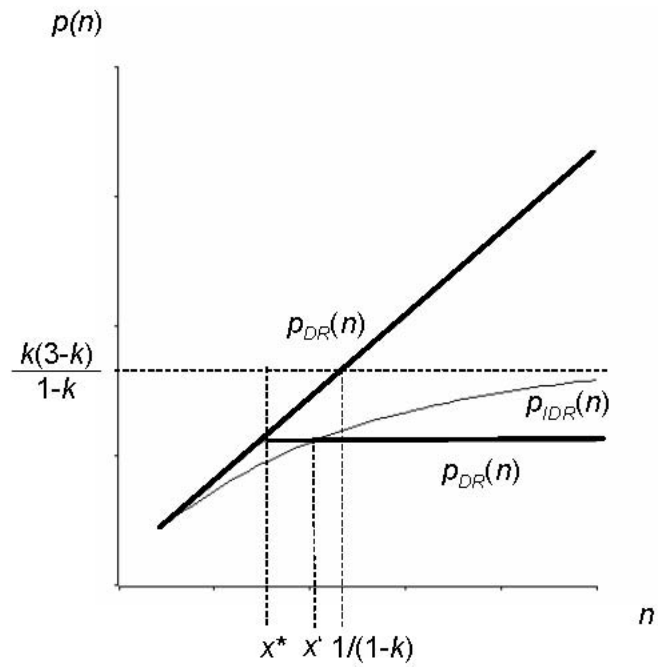

FIGURE 5 Resource extraction under direct and indirect rule based on a step function

Although some caution is appropriate given the starkly simplified situation, it is theoretically significant that exponential loss-of-strength gradients are incapable of triggering a change from indirect to direct governance. If Proposition 1 can be generalized, it represents a powerful argument against exponential decay functions, at least in early modern Europe. It seems that distance-dependence needs to feature some drop, albeit not necessarily as radical as that suggested by the step function.

COMPUTATIONAL MODEL OF TERRITORIAL SYSTEM CHANGE

The preceding section provides important insights into the tradeoff between organizational and geographic distances. However, our exceedingly simple mathematical model leaves many questions open. In particular, the deductive framework forces us to make many strong assumptions for tractability reasons. Therefore, this section introduces a computational model that enables us to relax these assumptions. First, we extend the linear setup to a more general, two-dimensional space. Second, we go beyond static comparisons of organizational efficiency by studying the operational causal mechanisms in a competitive, dynamic setting. Third, we generalize the results of the mathematical model in terms of the loss-of-strength gradient's functional form.

We proceed by building on the principles of Geosim (Cederman 1997, 2002, 2003) while extending the framework to an arbitrary number of hierarchical levels. The new model, which we call OrgForms, has been created from scratch by using the Java-based toolkit Repast. Like Geosim, OrgForms constitutes a dynamic network of endogenous hierarchical organizations that engage in constant geopolitical competition. These actors reside in a two-dimensional spatial grid and interact only with their immediate neighbors.

As opposed to the states in Geosim, which feature only two-level hierarchies, the actors in OrgForms resemble the hierarchical line states of the preceding section. We adopt exactly the same recursive taxation mechanism as in the linear case. This means that each actor, whether sovereign or not, produces exactly one resource unit locally. Each superior instance taxes the next inferior level by a constant, globally fixed tax rate of 0.2 that is discounted according to the Euclidian distance between the taxing and taxed units. Obviously, the efficiency of large states' resource extraction tends to decline in the periphery as their territories expand.

In our experiments, we use a 30×30 square lattice that is initially populated with 200 compound, state-like actors with merely two levels. Figure 6 describes such a system. The black lines denote state borders, and the red dots mark the capitals. The figure also indicates the resource level of each state below the capital. This setup is very similar to Geosim's initial configuration (cf. Cederman 2003).

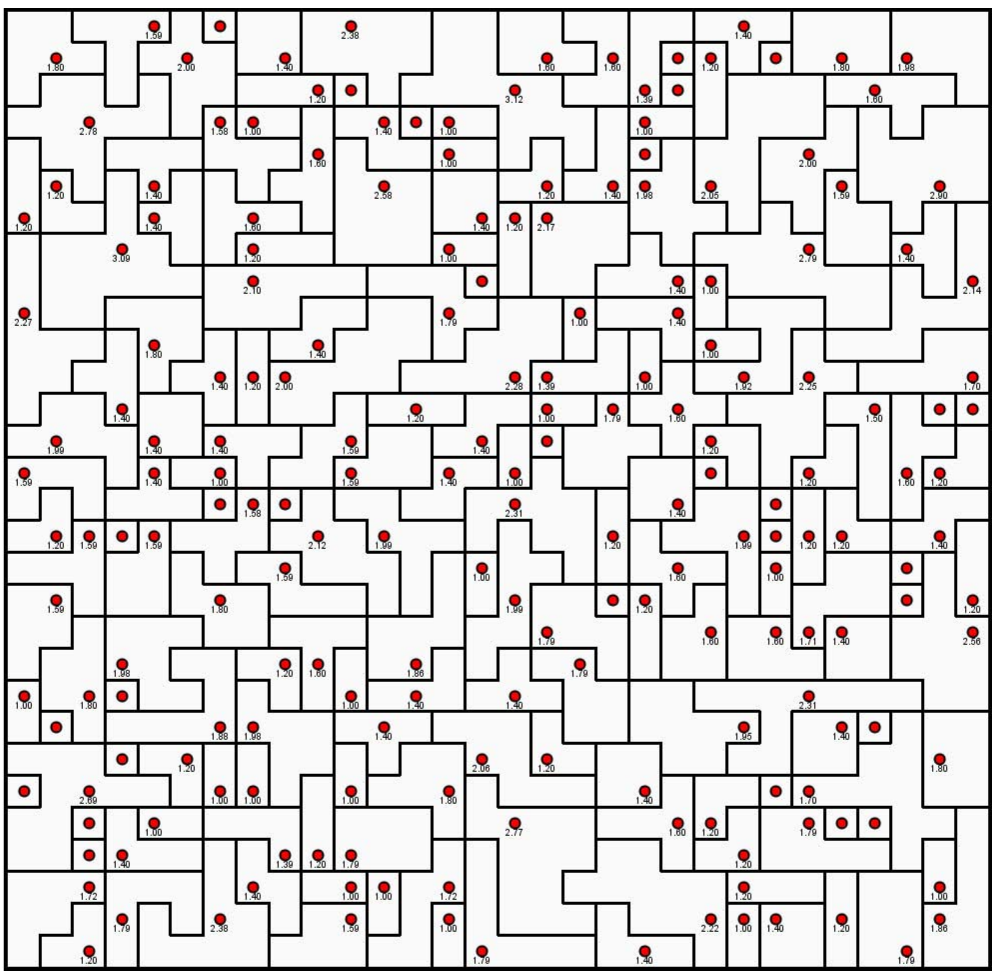

FIGURE 6 Initial state of the OrgForms model (*t* = 0)

After having created an initial state system, we now turn to the dynamics of the model. Technological progress is the master process that drives the system forward. In contrast to the perfectly sharp step function used in the previous section, the loss-of-strength gradient in the current, computational model follows a smoother logistical functional form:

$$\delta(x,t) = \frac{1}{1+\{x/x^*(t)\}^{c(t)}} \ ,$$

where x is the distance from the capital; $x^*(t)$ is a time-dependent threshold value for which resource extraction reaches 50%; and $c(t)$ is a dynamic, tunable parameter that controls the curve's steepness. While $c(t)$ = 0 yields a totally flat function, a perfect step function forms as $c(t) \to \infty$.

Technological change is modeled by sliding the threshold $x^*(t)$ from lower to higher values and by reducing the slope $c(t)$ over time t (cf. Cederman 2003 for a similar way to model technological change, which focuses on the threshold only). Figure 7 illustrates how the shape of the loss-of-strength gradient $\delta(x,t)$ changes as a function of time t. At the beginning of the simulation, the logistical obstacles are overwhelming. As the curve slides from left to right in the diagram, however, the capital's ability to tax at large distances improves dramatically. More precisely, throughout the course of a simulation run, we slide the threshold $x^*(t)$ from 3 to 10 at time period t = 1,000. At the same time, the slope is reduced from $c(t)$ = 10 to 5.

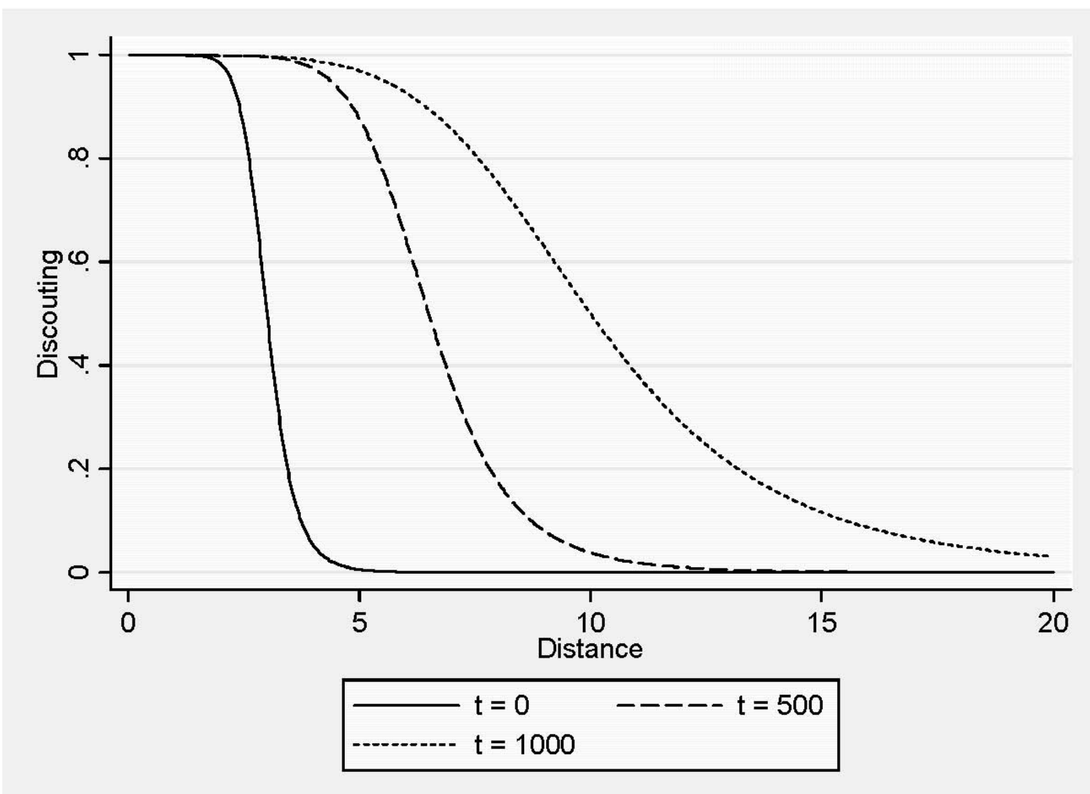

FIGURE 7 Loss-of-strength gradients represented as sliding logistical functions

How does this technological transformation influence states' organizational forms? First we need to specify explicit geopolitical mechanisms that have the potential to alter the actors' hierarchical structure. Drawing on the qualitative, historical theories discussed above, we postulate two main processes: conquest and organizational bypass. Whereas the former mechanism figures prominently in Geosim and other similar geopolitical models, organizational bypass is an entirely new mechanism that has, to our knowledge, never been modeled formally.

Because there is an arbitrary number of organizational levels in OrgForms, conquest is a more complicated process there than in Geosim. In order to specify the mechanism, three questions need to be answered: (1) When does conquest occur? (2) What is being fought over? (3) How is the conquered organization going to be inserted into the conquering state's multi-level hierarchy?

1. *Conditions of conquest*. Only sovereign states can conquer another neighboring sovereign state. The decision criterion is a deterministic threshold, as in the original Geosim model (Cederman 1997, Chapter 4). A state i with total resources R_i decides to conquer a neighbor j with resources R_j if its share of the dyadic resources exceeds the predetermined threshold (i.e., if $R_i/[R_i + R_j] > 2.5$). Note that in its basic version, the OrgForms model does not attempt any resource allocation to separate fronts. Nor is there any separate battle mechanism. Thus, we assume that conquest is always successful and costless.

2. *The nature of the conquered organization.* While conquest proceeds locally ("nibble by nibble") in Geosim, OrgForms resembles Bremer and Mihalka's (1977) original model in that each conquered state is swallowed as a whole. This rule makes the geopolitical changes more radical, but it also allows for "inheritance" of organizational forms when entire "victims" are incorporated inside the conquering states. This is how the states gain organizational depth.

3. *The point of insertion of the conquered organization.* For our present purposes, the most important aspect of conquest is where the subjugated state is to be inserted. We assume that this decision is resolved efficiently. The conquering state locates all of its own provinces that border on the conquered state. For each such neighboring province, it follows its own organizational tree all the way up to the capital. Repeating this operation for all possible paths from the border to the capital defines a set of potential insertion points, including the capital itself. The point of insertion is the location out of this set that maximizes the resource extraction of the resulting expanded state, given the distance-discounting at the time of conquest.

The second geopolitical mechanism to be specified is organizational bypass. In this case, we need to address two questions: (1) When is bypass conducted? (2) What are the consequences of bypass?

1. *Conditions of bypass*. As the case with conquest, only the capitals are allowed to initiate action, which proceeds within the state's territory. Moving one level down the organizational hierarchy, the capital randomly selects a subordinate unit, which, in turn, controls subordinate provinces itself. Two conditions

must be fulfilled for bypass to occur. First, the capital must be powerful enough to bypass the intermediate unit. The resource level is given by the taxation algorithm described in the previous section. The decision to bypass is defined by exactly the same deterministic threshold as conquest, here with the capital as state *i* and the subordinate unit to be bypassed as actor *j*. Second, organizational bypass requires that the resulting, flattened organization be more efficient than the status quo.

2. *Consequences of bypass*. Once bypass has been decided, the capital establishes direct organizational links with all sub-trees of the bypassed unit, which itself loses its connection to these provinces. The result is a reduction in organizational depth.

In order to avoid conflicts between the two mechanisms, one or the other is randomly chosen per time period.

We can now summarize the model's logic. Figure 8 depicts how the process of technological change drives the entire system. The two micro-level mechanisms — conquest and bypass — channel the effect of this process on the evolving organizational forms. The main output variable to be measured is the territorial share of indirect rule in the system.

Because the goal is to explore the shift from indirect to direct rule, we first have to "grow" an indirectly ruled, "medieval" system with deep state hierarchies. Given the flat organization of the initial system, this deepening hinges on repeated acts of conquest. At some point, however, technological change makes organizational bypass both possible and desirable, thanks to improved communications. A flattening of the deep state hierarchies follows, which produces a modern state system characterized by direct rule.

We are now ready to test if this logic does materialize from the assumptions made so far. Figure 9 shows a snapshot of the system at time period 136. It is apparent that plenty of conquest has already taken place. Four conquest centers have managed to subjugate their surrounding hinterlands. The picture illustrates the organizational dependencies with thick lines radiating from the capitals to the inferior units. For each level further down the tree, the lines become thinner. To increase the readability of the picture, we suppress lines from the capital to the unitary provinces within the white, directly ruled areas. All other provinces that are indirectly controlled by the capital are shown in grey shading with their organizational connections. Thus,

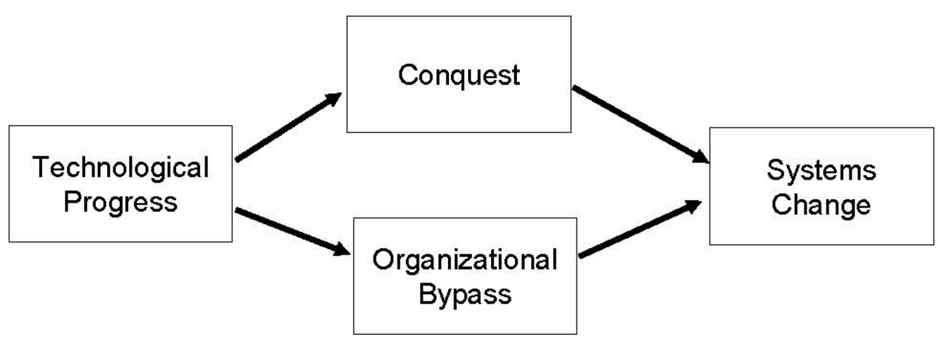

FIGURE 8 Main causal logic of the OrgForms model

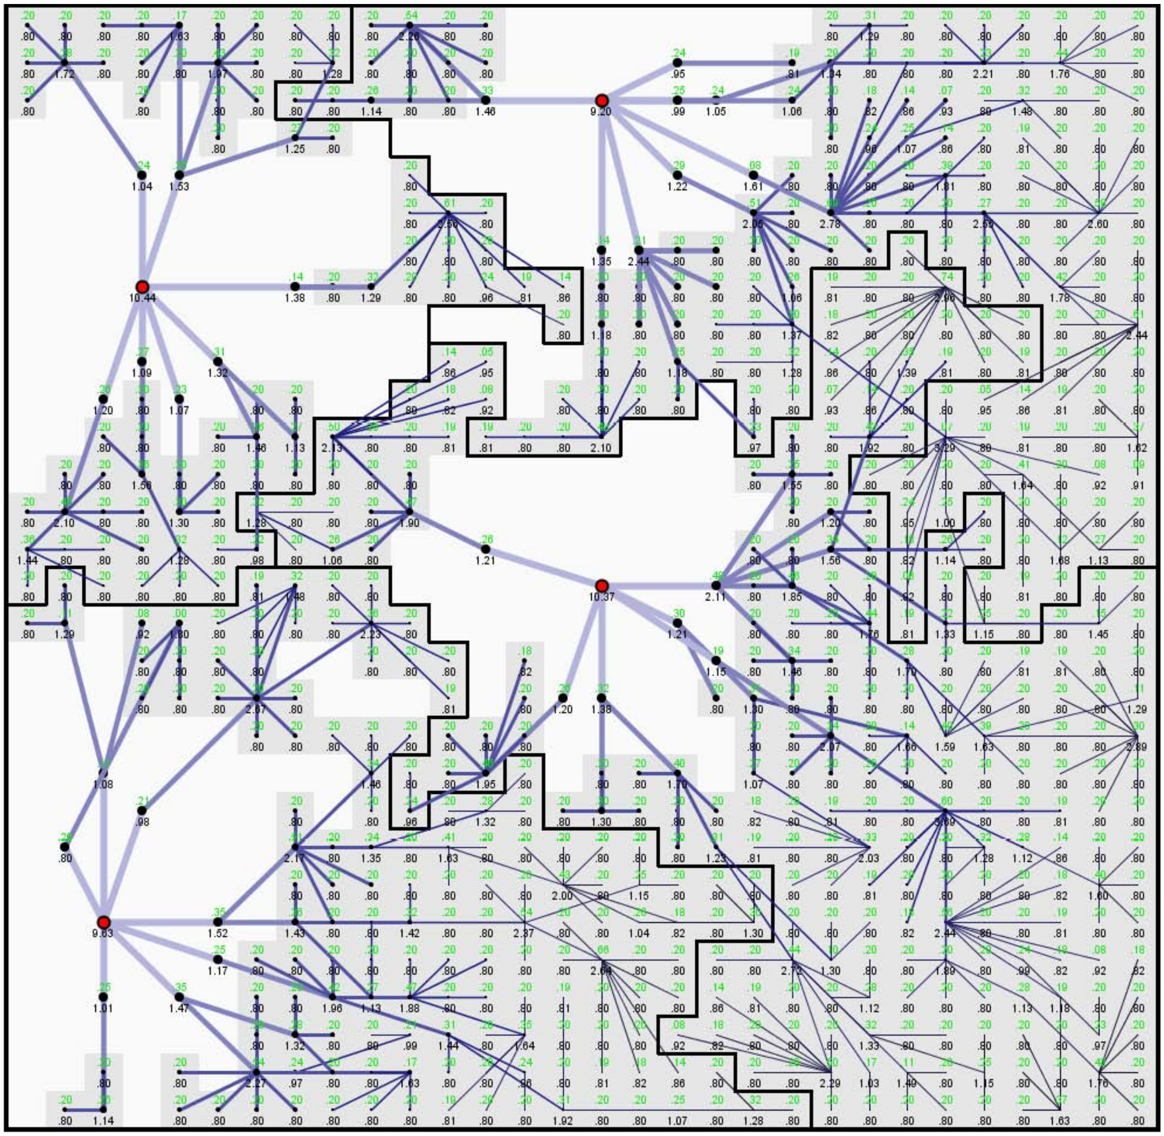

FIGURE 9 Evolution of deep hierarchies in the "Middle Ages" of OrgForms (t = 137)

the extent to which a system is shaded reveals how far indirect rule has been established. In this "medieval" situation, only the basins around the capitals are directly ruled.

The system in Figure 9 is relatively stable. The threshold of the decision function maintains a meta-stable equilibrium amongst the deep hierarchies. However, as technological progress makes long-distance interactions more effective, the conditions of conquest and bypass start to change. Figure 10 depicts the situation at time period 1,000. In contrast to the deep state organizations that characterize the Middle Ages, this modern state system features entirely flat organizations. Organizational bypass has managed to eliminate the internal competition for resources, thus flattening the organizational structure in each of the surviving states. In fact, the final system resembles a modern state system that is completely dominated by direct rule.

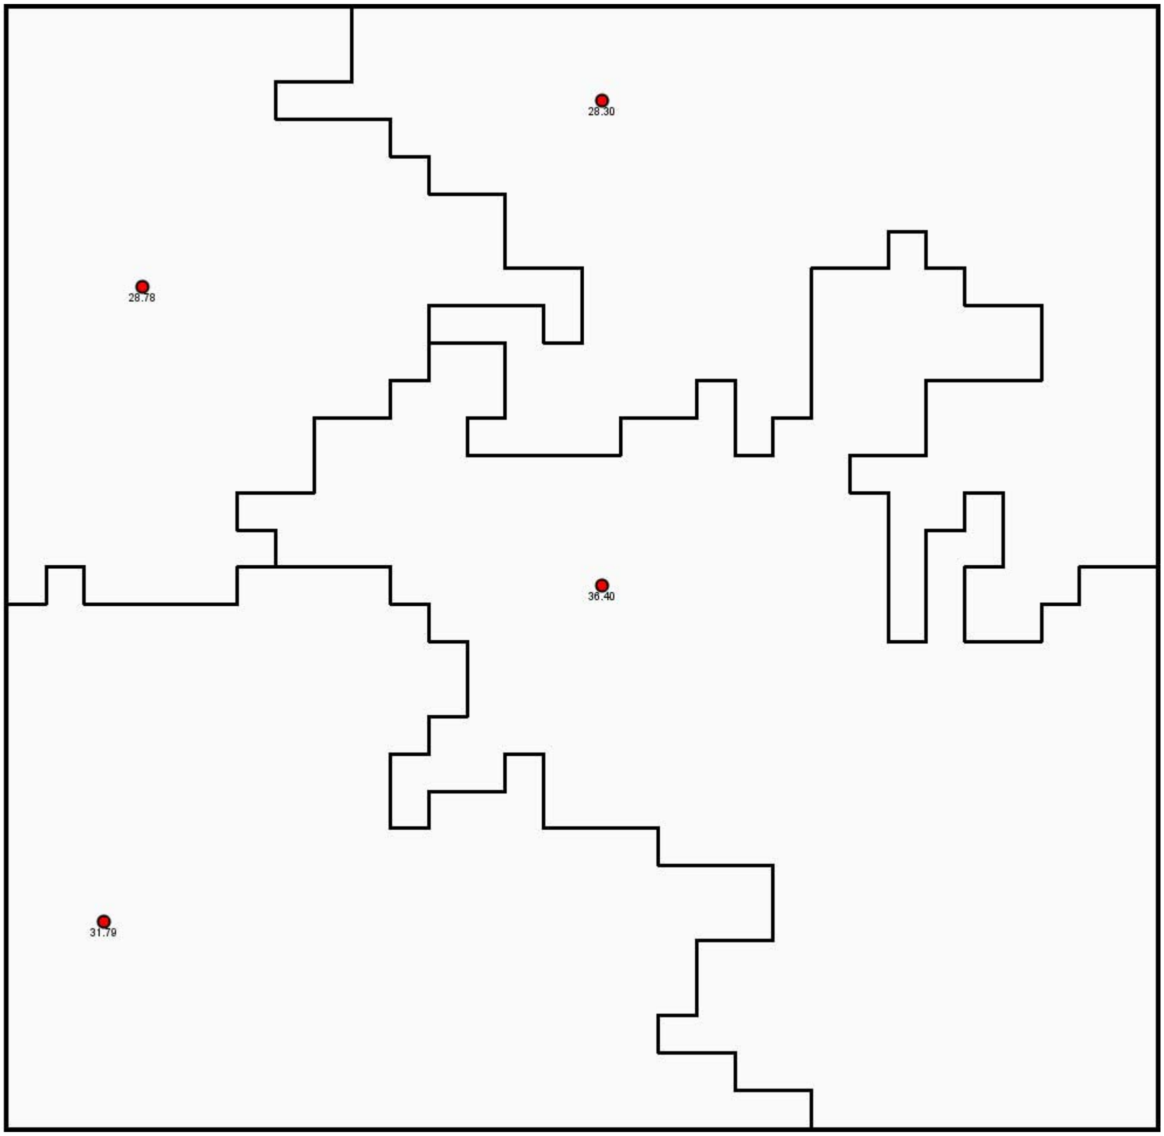

FIGURE 10 Emergence of direct rule in the OrgForms model (t = 1,000)

The snapshots tell us a good deal about how the system develops over time. However, it is desirable to trace its behavior more continuously. Figure 11 charts three important systemic properties over time. First, it shows the system's maximum depth (i.e., the number of hierarchical levels of the deepest organization). The figure reveals that this measure grows very quickly at the beginning of the simulation run, where it briefly reaches a maximum depth of 11, before collapsing abruptly around time period 600. Second, the average maximum depth follows a similar trend line, even though the changes are somewhat less dramatic for statistical reasons. Finally, the average depth of the system is much lower. After a relatively steep ascent, this variable slowly declines down to two, which represents the flattened, modern system.

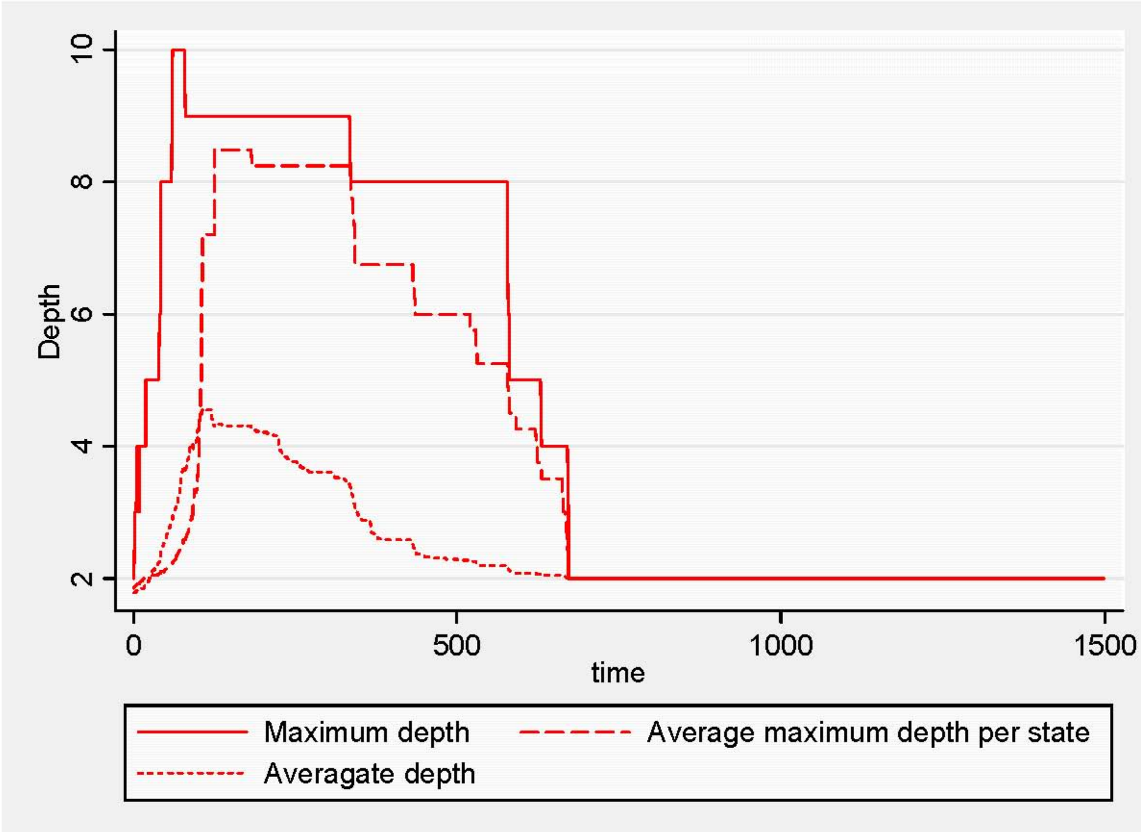

FIGURE 11 Tracing the sample run over the course of the simulation

REPLICATION RESULTS AND POSSIBLE EXTENSIONS

So far, we have only considered a single sample run. In order to check whether the shift from indirect to direct rule is more than a historical accident of the computational specification, we reran the system with 30 randomly generated initial systems, using different random seeds for each configuration. In all cases, the simulations started with 200 states in random locations. Summarizing the behavior of the entire ensemble of replications, Figure 12 traces the share of the system that is governed by indirect rule.

As expected, conquest triggers a rapid, initial increase in indirect rule. The curves peak shortly after time period 100 and then start a steady decline in response to the organizational bypass mechanism. The red curve represents the sample run shown above. Despite plenty of variation among the trajectories, a rising and falling trend is clearly observable. After time period 1,000, the last vestiges of indirect rule disappear. After this point, indirect rule materializes only temporarily in conjunction with conquest. Such situations do not last for long, however, because organizational bypass immediately follows.

In all runs up to this point, we have postulated a logistical loss-of-strength gradient. This choice of functional form is justified by the deductive findings of the previous section, which suggests that there has to be a threshold in the curve for systems change to emerge. To verify if

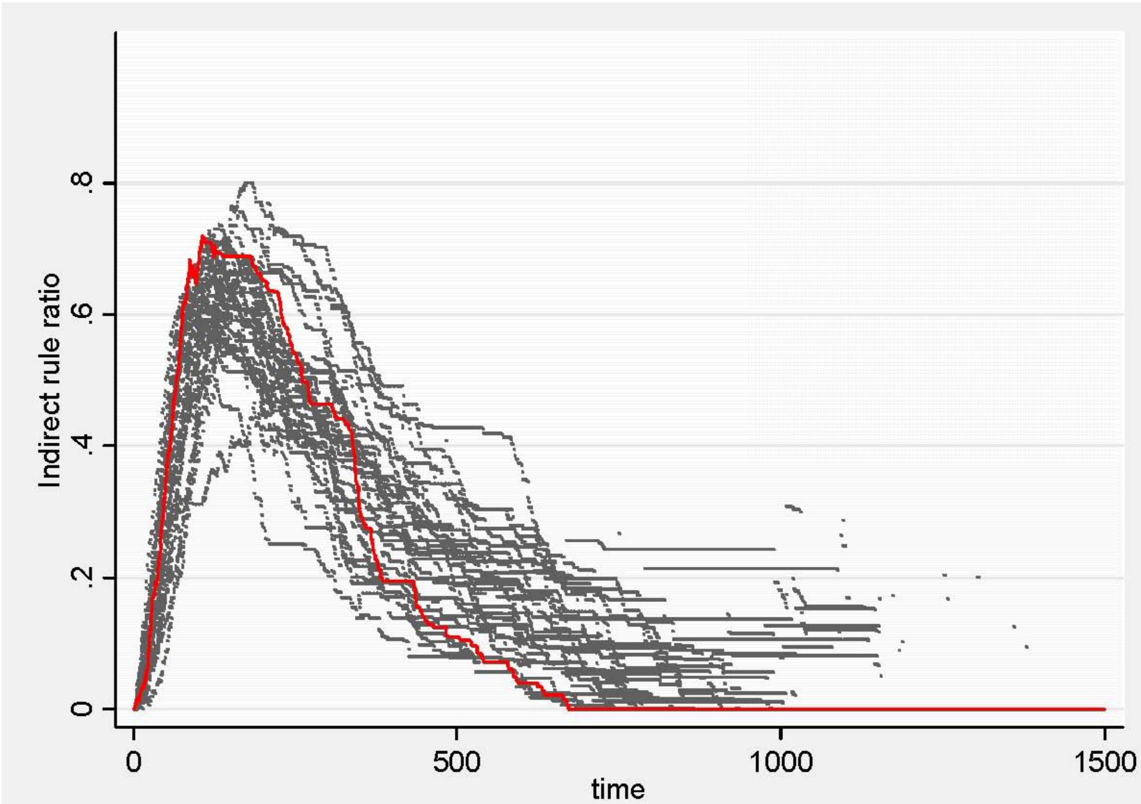

FIGURE 12 Share of indirect rule over time with sliding threshold and slope

this observation applies to the computational case, we reran the 30 runs with an exponential decay function that was shifted from $\delta = 0.8$ to 0.9 at $t = 1{,}000$. Here technological change was simulated through a stretching of the decay function, which improves geographic reach over time.

Figure 13 displays a radically different picture from the trend in Figure 12. In this case, indirect rule fails to take hold. In fact, the maximum share of indirect rule never exceeds 0.25, which is well below the peak in the previous set of runs. The majority of the trajectories with exponential decay do not even reach that level. These findings confirm the general conclusion of Proposition 1 in a more realistic setting.

Our final experiment investigates the pace of the organizational shift. Figure 11 traces a rather slow decline of indirect rule. What would it take for this transformation to unfold more suddenly? Figure 14 presents the evolution of indirect rule assuming a logistical distance function that changes its slope $c(t)$ over time from 10 to 0.95 while keeping the threshold value $x^*(t)$ constant at 3.

Under these conditions, the distance curve gradually develops into a decay function. The results indicate that while the modern world may well be best-described by an exponential loss-of-strength gradient, a more abrupt distance-dependence is initially necessary in order for the shift from feudal, multi-level hierarchies to occur in the first place.

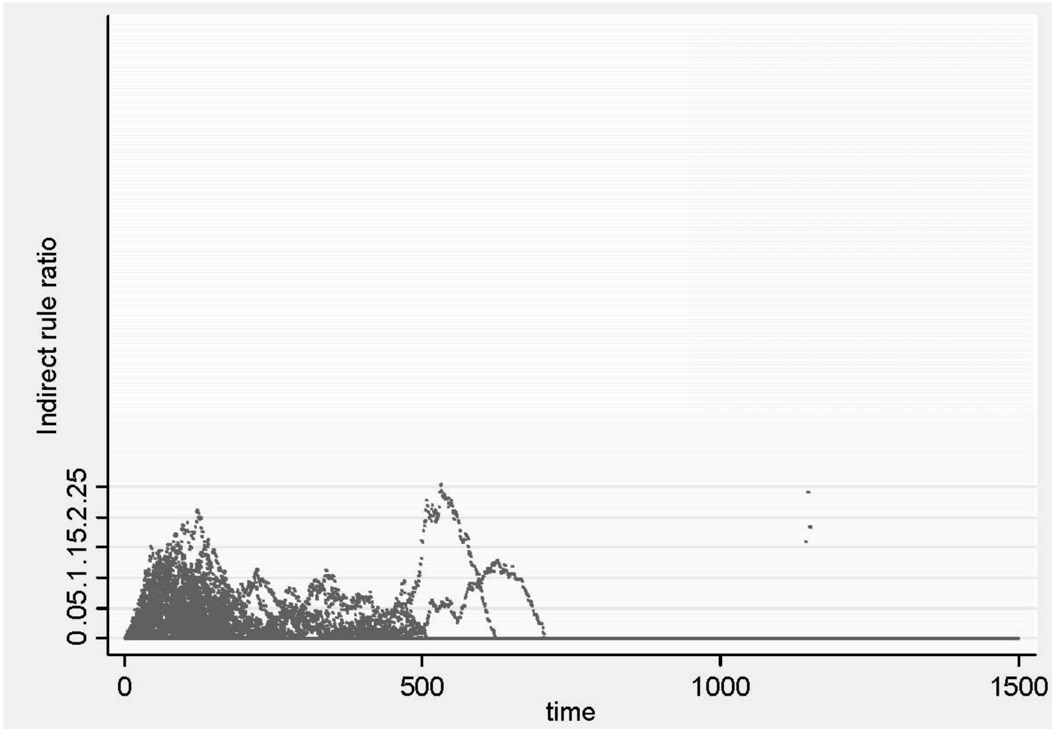

FIGURE 13 Share of indirect rule over time with an exponential decay function

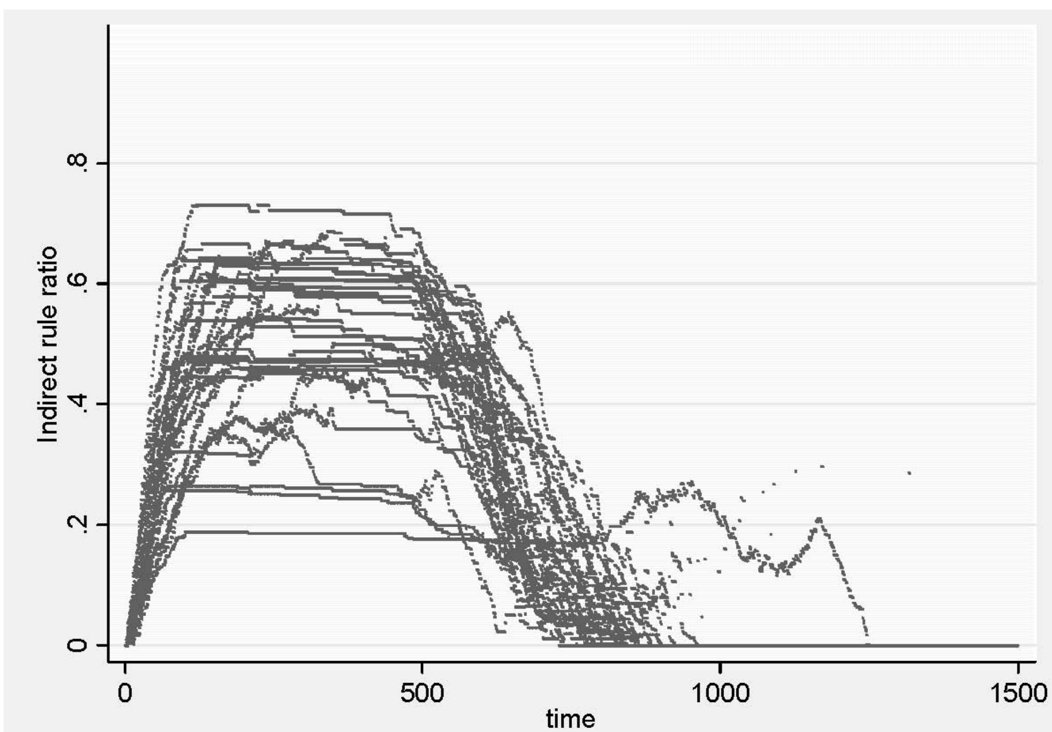

FIGURE 14 Share of indirect rule over time with a sliding slope only

Obviously, the empirical validation of these trends remains a wide-open question. Still, in the absence of comprehensive data, our models offer some important hints that could help guide more careful, theoretically guided calibrations of functional forms. Previous work has focused entirely on exponential decay functions (e.g., Lemke 2002). To the extent that criticism has been directed at Boulding's (1963) initial formulation of the loss-of-strength gradient, it typically relates to the deviations imposed by weapons systems or geography (Wohlstetter 1968). Our findings add support to such critiques, suggesting that the reach of power centers may not always decline smoothly with increasing distance. The limitations of specific infrastructural technologies may thus explain deviations from the assumption of exponential decay.

Another empirical question concerns the speed of systems change. Our initial set of runs, summarized in Figure 12, exhibits a rather slow decline of indirect rule. With a slightly different assumption of technological change, we were able to generate a much more abrupt transformation (see Figure 14).

Moreover, abrupt change could also result from organizational imitation. World history contains many examples of technological transfer, which suggests that infrastructural and military innovations diffuse from successful states to other members of the system (Gilpin 1981). While the current version of our computational model relies entirely on natural selection, it could be extended to encompass strategic adaptation (cf. Cederman 1997, Chapter 5; Cederman and Gleditsch 2004). Such an extension would need to feature some transferable representation of organizational structures (i.e., an "organizational code"). On the basis of such a formalization, it would be possible to analyze the co-evolution of technological change and the ideological spread of sovereignty.

CONCLUSION

Traditional theories of IR assume that sovereignty operates as a constant. In contrast, we have shown how this assumption can be relaxed with respect to the internal dimension of sovereign governance. Building on the insights of historical sociologists, we have posited a coherent explanation for how technological change triggered a shift from indirect to direct rule. Our modeling framework allows us to specify how this dynamic process influences the organizational topology of the system.

It should be reiterated that this internal reorganization does not exhaust all aspects of territorial systems change. The emergence of sharp boundaries is at least as important, but it was set aside in this study in order to simplify the modeling effort. However, the Middle Ages exhibited many examples of overlapping and conflicting rules that deviated from the perfect hierarchies assumed in this paper. In future work, we hope to be able to endogenize the crystallization of sharp borders that exclude competing sovereignties, rather than postulating such a modern order from the outset.

Despite this simplification, we believe that the current paper represents significant theoretical progress. As far as we can tell, our computational framework is the first geopolitical model that incorporates variable-depth hierarchies. Our model of systems change demonstrates that structural changes and power extraction are two sides of the same coin. This is a general insight that applies to contemporary examples of power competition involving radically different actor types, such as networks of insurgents and terrorists.

Thanks to its flexible ontology, computational modeling is ideally placed to assist historical research on such macro-transformations. With the current study of state formation in early modern Europe, we hope to inspire other researchers to rely on similar tools to explore complex transformations. Such a modeling effort promises to yield profound insights into the constantly changing realm of world politics.

APPENDIX

Proof of Proposition 1. Because the symmetric functional form of $p_{DR}(n) = 2\, f_{DR}(n) + k$ and $p_{IDR}(n) = 2\, f_{IDR}(n) + k$, it is enough to prove that $f_{DR}(n) > f_{IDR}(n)$. This is the same as proving that $f_{DR}(n) - f_{IDR}(n) > 0$.

$$f_{DR}(n) - f_{IDR}(n) = \sum_{i=1}^{n} k\delta^i - \sum_{i=1}^{n} k^i \delta^i = k\sum_{i=1}^{n}(1-k^{i-1})\delta^i \ .$$

However, it is immediately clear that each term of the sum $(1 - k^{i-1})\delta^i$ is positive for $n > 1$ because $0 < k < 0$ and $0 < \delta < 1$, which means that the entire sum has to be positive too. Hence $p_{\mathrm{DR}}(n) > p_{IDR}(n)$ for all $n > 1$.

QED.

Proof of Proposition 2. All comparisons can be made between $f_{DR}(n)$ and $f_{IDR}(n)$. We start by computing the latter, which holds for all values:

$$f_{IDR}(n) = \sum_{i=1}^{n} k^i \delta(i)^i = \sum_{i=1}^{n} k^i = \sum_{i=1}^{n} k^i = \frac{k}{1-k}(1-k^n) \ .$$

For direct rule, it must be that for $n \leq x^*$,

$$f_{DR}(n) = \sum_{i=1}^{n} k\delta(i) = \sum_{i=1}^{n} k = nk \ .$$

It can be shown that for $n \leq x^*$, direct rule is superior. This applies for $f_{DR}(n) - f_{IDR}(n) > 0$, which can be written

$$f_{DR}(n) - f_{IDR}(n) = \sum_{i=1}^{n} k - k^i = k\sum_{i=1}^{n} 1 - k^{i-1} .$$

However, we know that this expression has to be positive because the terms of the sum are positive for $(1 - k^{\mathrm{i}-1}) > 0$ for $n > 1$.

For $n > x^*$, we compute

$$f_{DR}(n) = \sum_{i=1}^{x^*} k\delta(i) = x^* k .$$

This quantity will always be greater than $f_{IDR}(n)$ where $x^* > 1/(1-k)$ (Case 2):

$$f_{IDR}(n) \to \frac{k}{1-k} \text{ as } n \to \infty .$$

Clearly $f_{IDR}(n)$ cannot exceed this value.

Otherwise (Case 2), we need to find the point x' where indirect rule surpasses direct rule as $n \to \infty$: $f_{DR}(x') = f_{IDR}(x')$, or

$$x^* k = \frac{k}{1-k}(1-k^{x'}) \ ,$$

which implies that $n < x' = \dfrac{\log[1-(1-k)x^*]}{\log k}$.

QED.

ACKNOWLEDGMENTS

We would like to thank Alessandro Lomi for organizing the workshop “Forms, Functions and Organization,” in Bologna, Italy, on September 9 and 10, 2004, which stimulated the early phase of this project. We are also grateful to Kimberley Sims for her helpful linguistic advice.

REFERENCES

Adler, E., and M. Barnett, eds. 1998. *Security Communities*. Cambridge: Cambridge University Press.

Ardant, G., 1975, “Financial Policy and Economic Infrastructure of Modern States and Nations,” in *The Formation of National States in Western Europe*, C. Tilly, ed., Princeton, NJ: Princeton University Press.

Axelrod, R., 1984, *The Evolution of Cooperation*, New York, NY: Basic Books.

Axelrod, R., 1997, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton, NJ: Princeton University Press.

Boulding, K.E., 1963, *Conflict and Defense,* New York, NY: Harper and Row.

Bremer, S.A., and M. Mihalka, 1977, “Machiavelli in Machina: Or Politics among Hexagons,” in *Problems of World Modeling*, K.W. Deutsch, ed., Boston, MA: Ballinger.

Bull, H., 1977, *The Anarchical Society*, London: Macmillan.

Carneiro, R.L., 1978, "Political Expansion as an Expression of the Principle of Competitive Exclusion," in *Origins of the State: The Anthropology of Political Evolution*, R. Cohen and E.R. Service, eds., Philadelphia, PA: Institute of for the Study of Human Issues.

Cederman, L.-E., 1995, "Expansion or Unity? Placing the European Union in a Historical Perspective," in *Towards a New Europe: Stops and Starts in Regional Integration*, G. Schneider, P. Weitsman, and T. Bernauer, eds., Westport, CT: Praeger Greenwood.

Cederman, L.-E., 1997, *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*, Princeton, NJ: Princeton University Press.

Cederman, L.-E., 2002, "Endogenizing Geopolitical Boundaries with Agent-based Modeling," *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 3, 7296–7303.

Cederman, L.-E., 2003, "Modeling the Size of Wars: From Billiard Balls to Sandpiles," *American Political Science Review* 97:135–150, Feb.

Cederman, L.-E., 2004, "Nationalism and the Puzzle of Changing State Sizes," Paper presented at the Annual Convention of the American Political Science Association, Chicago, IL, Sept. 2–5.

Cederman, L.-E., 2005, "Computational Models of Social Forms: Advancing Process Theory," *American Journal of Sociology* (Forthcoming).

Cederman, L.-E., and C. Daase, 2003, "Endogenizing Corporate Identities: The Next Step in Constructivist IR Theory," *European Journal of International Relations* 9(1):5–35.

Cederman, L.-E., and K.S. Gleditsch, 2004, "Conquest and Regime Change: An Evolutionary Model of the Spread of Democracy and Peace," *International Studies Quarterly* 48:603–629, Sept.

Cusack, T.R., and R. Stoll, 1990, *Exploring Realpolitik: Probing International Relations Theory with Computer Simulation*, Boulder, CO: Lynnie Rienner.

Dacey, M.F., 1974, "A Model of Political Integration and Its Use in Reconstruction of a Historical Situation," in *Locational Approaches to Power and Conflict*, K.R. Cox, D.R. Reynolds, and S. Wiley, eds., New York, NY: Wiley.

Downing, B.M., 1992, *The Military Revolution and Political Change: Origins of Democracy and Autocracy in Early Modern Europe*, Princeton, NJ: Princeton University Press.

Finer, S.E., 1974, "State Building, State Boundaries and Boundary Control," *Social Science Information* 13(4/5):89–126.

Gilpin, R., 1981, *War and Change in World Politics*, Cambridge: Cambridge University Press.

Kahler, M., 1999, "Evolution, Choice, and International Change," in *Strategic Change and International Relations*, D.A. Lake and R. Powell, eds., Princeton, NJ: Princeton University Press.

Kennedy, P., 1989, *The Rise and Fall of the Great Powers*, New York, NY: Vintage Books.

Keohane, R.S., 1984, *After Hegemony: Cooperation and Discord in the World Political Economy*, Princeton, NJ: Princeton University Press.

Kochen, M., and K.W. Deutsch, 1969, "Toward a Rational Theory of Decentralization: Some Implications of a Mathematical Approach," *American Political Science Review* 63:734–749, Sept.

Kochen, M., and K.W. Deutsch, 1974, "A Note on Hierarchy and Coordination: An Aspect of Decentralization," *Management Science* 21:106–114, Sept.

Kratochwil, F., 1986, "Of Systems, Boundaries and Territoriality," *World Politics* 34:27–52.

Lemke, D., 2002, *Regions of War and Peace*, Cambridge: Cambridge University Press.

Lomi, A., and E.R. Larsen, eds., 2001, *Dynamics of Organizations: Computational Modeling and Organizational Theories*, Cambridge, MA: MIT Press.

Morel, B., and R. Ramanujam, 1999, "Through the Looking Glass of Complexity: The Dynamics of Organizations as Adaptive and Evolving Systems," *Organization Science* 10:278–293 (May–June).

Oye, K., ed., 1986, *Cooperation under Anarchy*, Princeton, NJ: Princeton University Press.

Poggi, G., 1978, *The Development of the Modern State: A Sociological Introduction*, Palo Alto, CA: Stanford University Press.

Prietula, M.J., K.M. Carley, and L. Gasser, eds., 1998, *Simulating Organizations: Computational Models of Institutions and Groups*, Menlo Park, CA: AAAI Press and Cambridge, MA: MIT Press.

Reus-Smit, C., 2002, "The Idea of History and History with Ideas," in *Historical Sociology of International Relations*, S. Hobden and J.M. Hobson, eds., Cambridge: Cambridge University Press.

Ruggie, J.G., 1993, "Territoriality and Beyond: Problematizing Modernity in International Relations," *International Organization* 47:139–174.

Strayer, J.R., 1970, *On the Medieval Origins of the Modern State*, Princeton, NJ: Princeton University Press.

Tilly, C., 1990, *Coercion, Capital, and European States, AD 990–1990*, Oxford: Basil Blackwell.

Waltz, K.N., 1979, *Theory of International Politics*, New York, NY: McGraw-Hill.

Wendt, A.E., 1987, "The Agent-structure Problem in International Relations Theory," *International Organization* 41:335–370.

Wohlstetter, A., 1968, "Theory and Opposed-system Design," *Journal of Conflict Resolution* 12(3):301–331.

DISCUSSION

Computational Social Theory

(Invited Speaker, Lars-Erik Cederman, Saturday, October 15, 2005, 8:30–9:30 a.m.)

Chair and Discussant: *Doowan Lee, The University of Chicago*

Growing Sovereignty: Organizational Shifts in State Systems

David Sallach: Good morning. I'd like to welcome you back to Agent 2005 for a day devoted to computational social theory. Our invited speaker today is Lars-Erik Cederman, and we're very happy to have him. He's been with us before and always makes a signal contribution. I should say that he is also one of that small group of people who has done work that serves as an exemplar for the field. The discussant for the session is Doowan Lee. So it's a good way to start the day, and we'll just move forward. Thank you for coming, Lars.

Lars-Erik Cederman: Well, thank you very much, David, for the invitation. As always, I'm delighted to be here at an Agent conference, and it certainly isn't my first Agent conference. I am going to be presenting work today that I coauthored together with Luc Girardin, who's sitting right here. I want to make sure you don't overlook that because I think the name slipped out of the program. We are very happy Repast users in Zurich at the Swiss Federal Institute of Technology at ETH. In fact, this presentation and this paper build on something I did here at Agent 2002. I gave a talk on computational models of social forms, work that subsequently was published in the *American Journal of Sociology*, this year in fact. Today, together with Luc, I'm going to follow up on some of the theoretical work that was done in that paper. This is now a much more complete work.

I'm going to be talking about the types of social forms that we may want to generate. There is a conceptual transition from the theme of generative methods to whatever it is that we're trying to grow. After that introduction, I'm going to say something about the specifics. I'm going to hone in on emergent systems change in world politics, which is a very tricky topic, and talk specifically about state formation in Europe, a favorite topic of mine. As you will see, it's a slightly different take on the subject compared to my previous work. I'm going to present a deductive model, which is, of course, a sacrilege at this gathering on computational modeling, before moving into the computational material.

[Presentation]

Doowan Lee: My name is Doowan Lee from The University of Chicago. I'm very happy to comment on this excellent paper. In fact, the basic theoretical setup of the paper shows the way to follow subsequent logics because, in many cases, if you take a look at a number of models in international relations or comparative politics, they are mostly about processes or systems changes at best, as opposed to systemic changes.

To use Walt's terminology, most people talk about changes *in* systems as opposed to changes *of* the systems. In that sense, I think the paper identifies a huge void in the field of international relations and comparative politics. I think this paper is going in the right direction.

To take a few examples of different models in the field, I would say the evolution of strategists using the Prisoner's dilemma settings are good examples of process-oriented models. On the other hand, your own emergent polarity models are a good example of systemic changes as opposed to systems change. In that sense, our org forms is a good extension of your previous work as well.

I have about four comments. The first two are internal to the model, and the last two are my suggestions for extensions. My first question is about the unit homogeneity of the model. In case of a state with multiple layers of internal resource extraction mechanisms, and they're all paying successive portions of their wealth to the tactile, the power-holding defense, it is difficult to differentiate this political form from a federal system or even from an empire. I'm guessing that this might have something to do with the fact that the model does not explicitly implement the consolidation or crystallization of external borders. You emphasized the distinction between internal and external sovereignty, and I think this internal homogeneity issue will be easily solved if you differentiate internal and external political boundaries and let them play out over time to see how these internal organizational forms become differentiated from external space structures like federal systems and so forth.

That was my first concern about the paper because the basic unit is sometimes a little unclear to me. The second question is basically curiosity on my part. For example, one of the basic hypotheses of the paper is that when the loss of strength gradient is defined by a step function, indirect rule is much more likely to prevail over time. I think this is very counterintuitive to the common sense we are used to in the field because, you know, it is very common to use an exponential distribution to limit power projection capabilities of space, etc. So it is very counterintuitive and actually sheds a lot of light on what we understand about power projection. At the same time, I think it will help us understand the uniqueness of your findings a lot more if you could give us a simple anecdotal story so that we could see where exponential distributions fail and step functions are much more likely to prevail in reality. I think that will help us understand the logic much better.

The third question is about division of technology and mimicking. When I was reading the paper, I thought, "Wait a minute. Cederman obviously talks about the division of technology, and he was implementing technological division in his paper. How come I didn't see it here?" Toward the end of paper, though, I noticed that you acknowledged this issue and then suggested an extension based on division of technology.

My conjecture at this point is that it is actually going to slow down the process because no one additional form will be able to take over the whole system very quickly because everybody will be competing and then once something is successful, everybody will try to emulate that one additional form. On the other hand, just the opposite will be possible, too; that is, once something becomes very popular, everybody follows the rule and then the whole system changes very quickly. I think this is much more consistent with what really happens in world politics when it comes to military technology, doctrines, grand strategies, and etc. In this sense, it is not clear whether division of technology will make the system go much faster or slower. That means this is unclear, and I think it would be great if you could comment on this aspect.

My last comment is about resource extraction. On one hand, of course, if you take a look at Charles Tillian's work, taxation and warfare are the two driving forces behind state building processes, especially in Europe, and they make perfect sense. At the same time, you should take a look at how internal units react to resource extractions imposed by a central authority. There is a very different story because a lot of warfare in Europe was triggered by this conflict between the peripheries and the center, and these internal units actually pooled their resources to get a new form of organizational innovation or a new form of direct resource extraction. Thus, we had the typical conflict between a feudalist system and a capitalist system. They're actually competing with different imperatives of resource extraction. All generations are feudalistic, with system laws, and they pool their resources and rebel against the central authority from time to time. Since warfare is such a key notion in your paper, I think just the opposite corollary is perfectly possible in the same scenario. I think it would be great if you could comment on this as well.

I have one very small comment to finish my response. I could be wrong about this, and I'm going to be conflicted to actually say this because I don't know whether it's an important aspect of the paper or not, but I'll let you be the judge of that. I'm only concerned about the source of technological innovation because there are two key things in your paper. The first one is an additional bypass, and the second one is technological innovation or progress. You seem to place a lot of emphasis on the state as the carrier of technological innovation. Since you talk about the feudal system in comparison to the modern space system, a little bit of Marx wouldn't hurt here; that is, the source of innovation comes from the economic system, not from the political system. Charles Tillian would agree with me that the actual source of technological innovations comes from the economics and then it drives the political forces to reorganize and create an additional form and resource extraction methods.

So I'm a little troubled by your heavy emphasis of technological innovation on the state as opposed to the economic level. Those are my comments. I think this is a great paper and sets the standards for a lot of people, and one last methodological comment: I think that it is really great to see another paper that starts with a closed-form solution and then builds up complexity gradually. I think that is the only way to go whenever you're building an agent-based model.

Cederman: Do you want me to respond directly to your comments? Okay, I'll try to be brief because I'd like to hear from other people, too, but these are absolutely excellent comments.

Let's start from the last point. I wouldn't, by the way, go so far as to say that you had to have a deductive model. I would say that wherever it's possible, we should try to tighten things up with deductive inference if we can. If we can't, well, you can still do very important work. In any case, the point about the economy and capitalism is well-taken. This is very reductionist work. It is focusing on logistical and military processes and geopolitical transformations only in that dimension. That's a conscious choice. In future extensions, it will be wonderful to factor in economic processes that are much more sophisticated, but we just can't do everything at the same time.

Exactly how to do that is difficult to tell because I'm not an expert on political economy of state formation. But it's a fascinating topic, and I hope we'll be able to do work on that. I'm not sure, however, that we could treat innovation as being completely driven by the economy. I think that the state and the market, or whatever economic systems, are co-evolving here.

Certainly if you look at Russia after the fall of the Berlin Wall, the state does have a certain importance for the economy in terms of setting up the rules of the game. That's often underestimated and ignored by neoclassical scholars today.

Now, obviously the resource extraction is completely simplified here also with respect to legitimacy. I mean the neoconservatives were expecting that they would be able to siphon off the oil and be welcome in Iraq without any problems. In today's world, nationalism is a very important factor, and anyone who ignores that will have to pay for it, as we have seen. We have other models that look into nationalism, but we haven't, as I say, grown nationalism as a new antitype yet. That will be another difficult problem to grapple with, and it's certainly on our research agenda.

Very briefly, on diffusion of technology — it's extremely difficult to tell exactly what that would lead to. Perhaps it's going to lead to a stickier system, as you say, but when change happens, it's going to be more rapid; it's going to flip more quickly. Whether that is realistic or not, is another question. How fast was this transition in the real world here? You know, we can only guess, and much more work would have to be done to figure out exactly how that happened.

Going back through your list, your second point, about the loss of strength gradient and the functional shape of that empirical question of how we came up with some examples. In the first write-up of this paper, we overlooked Michael Hechter's very important book on containing nationalism. It has an excellent chapter with details and at least examples. We were certainly going to go through another write-up, and hopefully we'll be able to say more about that in that connection.

Finally, your first point about unit homogeneity is also important — and I think for precisely the reason you mentioned: without different configurations of boundaries, it's virtually impossible to differentiate states from empires. One of the core features of empires is precisely this fading, as I say, control in the borderlands, and that is by definition ruled out here. That would be very, very interesting to model, but there's more than that, of course. There is also fluid, as I say, border between the state, the territorial state, and the nation state. Where do you become a nation state? That involves things like democracy, an active citizenry in the state, so direct rule takes on a very different meaning in that context. That also involves things like federalism. We are not modeling federalism per se, but we think that modeling federalism explicitly, including democratic forms of it, would be very, very important in the context of, for instance, civil war studies. A lot of our activity — I would say the main thrust of our research in Zurich — is not along this line; it's not primarily on the social theory agenda like what I've shown today, but it's actually much more down-to-earth empirical work on civil wars and trying to use computational models to figure out what mechanisms are driving these horrible conflicts that afflict our world.

Sallach: We have time for a few questions.

Joanna Bryson: Joanna Bryson, Bath. I loved your paper, and I've been trying to think about this. In a way, I went from worrying about modular single agents into multi-agent systems, but using totally different representations, and I've been interested in this distinction.

Anyway, I was confused about your second rule. You went quickly toward the end, and so I can see what's going on with the conquering, but you talk about that when do you spread out the organization stuff. It's something about it's a greater power, but then also something about efficiency. Are you maintaining something of the old structure? It wasn't showing up in your graphics, but is there a bit of the old structure still around the infrastructure or something that you have to say, "Okay, are the paths shorter?" What did you mean by simpler efficiency there? And then, just as an aside, when you do go into your regular sim, will those chunks be like Alsace and keep flipping back and forth or whatever?

Cederman: Yes. In fact, the efficiency requirement is just to make sure that when bypass happens it should be sustainable, so it's a proxy for the possibilities. We would have wanted to have a more complete model of the resources on both sides and a more evolutionary model where bypass is attempted, but it creates completely unsustainable organizations that would then collapse. In order to keep things under control and not completely chaotic, we had to introduce these slightly unrealistic — not only slightly, actually, quite radically unrealistic — assumptions that move away a little bit from, as I say, bound in rationality. They are still bound in rationality in many senses, but we inject a healthy dose of efficiency into the agent's rules just to make the outcomes more manageable to study.

Bryson: So what's the definition of efficiency?

Cederman: It's just how much you extract. If the bypass doesn't make you better off, why would you do it? That's the basic question. In reality, politicians do many things that don't make them better off, obviously, and you have a more Darwinian process that weeds out the bad choices, although in some cases you really wonder if that process is operating. No more comments on that.

Larry Kuznar: Yes, Larry Kuznar, from Indiana Purdue University of Fort Wayne. I appreciate Dr. Lee's comment about economic causation. I too am an unabashed materialist in a lot of ways. However, I was thinking, one of my regional areas of expertise — I'm an anthropologist — is the Andean region. I was thinking about the evolution of the Incan Empire, and in general of empires constructing roads. Now, this is a technological innovation that in a very emergent way can only be done by a higher unit, in this case, empire, to put down such a system. I was thinking about the issue of exponential decay versus step functions, and I was wondering if both of you might want to comment on empires constructing infrastructures so as to force a new form of power decay so that they can better rule their domains.

Cederman: That's a difficult question, since I'm certainly no expert on the Incan systems. I do believe that they form, that they constitute, excellent examples where these types of models could be applied. I mean, there were models. I think Douglas Dacy came up with a deductive model of state formation, but those river valleys, I'm not sure whether they were Incan or somewhere else, but there are attempts to do this, to apply these models.

I think the general point is about the technologies, if there is some kind of drop in efficiency beyond a certain point due to the technology in question. I mean, how much gas can you fit into your fuel tank, as it were, into the fighter jet? When do you have to refuel, and what about the social organization? There may be some kind of boundary beyond which your efficiency falls because it is extremely hard to get more than 20 people into a meeting and make sense. There are different rules of that type, so I think there are inflection points. We don't know very much about them, empirically speaking. I would be fascinated to talk to experts in your area, including yourself, to see if one could, as I say, tease out these empirical examples that our discussant asked for.

Steven Bankes: Steve Bankes, Evolving Logic. Ideally, one would like to find a case where one can get data on this efficiency curve. In fact, one would like to find a case where the efficiency curve has changed for some reason so you can establish the tracking — as the curve changes the organization form changes. That's going to be hard looking backward a few hundred years, and there may be cases that we can find evidence of, but it's going to be highly uncertain at best. I'm wondering, in spite of the motivation for this, whether in fact the nature of the model is very economic, in a sense. Form emerges from efficiency considerations. That logic can't be applied in an organizational setting, perhaps.

I'm thinking a merger and acquisition followed by bureaucratic reorganization is not so different than some of the things we see going on in this model, in a way; if one could make that mapping a little more firm than my hand wave, one would be in a position to look back only 10 years with the deployment of the Internet, and a lot of organizational theorists that were strongly asserting that organization would get flat, and there must be some data sources out there that one could begin to try and make an alignment between data that is actually pretty solid, and perhaps find some level of support for this model of organization coming out of essentially efficiency in extracting resources. I was asking to get your comments on whether that's a completely insane idea or a viable possibility.

Cederman: Well, the political scientists who have tried to do that are nearly insane, but I will be very happy to give it a shot. In fact, the Hechter book contains an example of Cornwall. There were actually serious innovations happening in early modern Europe, and it's not all about the Industrial Revolution. I mean, there was the standardization of coinage, the abolition of tolls, and so forth. I remember that when I was working in Oxford, there was a little bridge just a couple of miles outside Oxford where you have to pay 10 pence, so there are some remainders of that, and that just slows you down.

I would not completely exclude the possibility of at least providing some preliminary evidence, but I think it would be a major data project if you want to get the functional form right, and certainly how it moves over time. I'm not sure our group should be doing that kind of work, but at least we have come up with a hypothesis that is counterintuitive, and we'll do what we can, you know, scouring the texts without doing any original research ourselves. Whether the Internet can tell us about this — yes, perhaps, although I would hate to include any reference to Tom Friedman's book on the flat world, which I think is one of the most misplaced metaphors I've ever seen.

Robert Reynolds: Bob Reynolds, Wayne State University. I think you make a really nice point for the generative sufficiency of your approach, and, having worked with origins of the state, I'm very pleased to see that. One of the things I was thinking about, and not something you have to do right now, but you're assuming a closed system. One of the interesting things is that if you allow perturbations by tweaking something within this stable equilibrium or semi, what happens to the system? In other words, what kinds of reorganizations? Are they going back to the same organization, or do the changes ripple through the configuration to make slight adjustments? And how does the system deal with the resilience of the system or respond to external perturbations?

The other comment is that, often when people look at the emergence, well, archeologically, for example, in terms of the origin of the state, they have this notion of the rank size rule, where you can look at the relative size of cities associated with a state organization. There tends to be a concavity, which means that the state tends to attract resources into its area and away from the surrounding area. It sounds like that's something you could measure in terms of the current parameters of your model to see whether those types of observational things are re-created in your simulation.

Cederman: Yes, these are very good comments. In fact, the last point about trying to re-create distributions would be a very natural step, given our previous work in the group. I think that's an excellent idea. In fact, one observation that we're not really replicating here at all, but that's very important, it is that you have a bimodal size distribution. You have huge empires and lots of small units. That's very unrealistic. That's not what we have here. So I would want to get that right, and that requires us to get empires right. It's really about the empires, and this is not a very realistic model of that kind of entity.

The other point was about perturbations, and one has to be a bit careful, you know. You don't want to have the outcome be exogenously driven too much. You want to be as indulgent as possible to come up with true emergence. However, you can, as I say, introduce a loss of strength gradient that is snapping into place that is not very smooth, where you have jerky evolution of progress. That's where diffusion also fits in. So you could have these ideas travel suddenly and then transforming states very quickly. That would probably lead to different outcomes, but it's very hard to tell exactly what it would lead to and why we would like to do it, but it's certainly something that's on our agenda.

Charles Macal: I have one comment or possibly a question, but it's strictly methodological in nature. Your results are obviously critically dependent upon the functional form that you're assuming, and you have two distinctly different types of results. I'm aware of a function that is more general than either of the two functions that you've assumed, which could be varied by two continuous parameters to get either at one extreme an exponential or at the other extreme a step function, albeit continuous. I think it would be interesting to range through that function in parameter space to see where the qualitatively different regimes begin to make a transition.

Cederman: That would be very elegant because the robustness tests have not been included that will be needed for this paper to be complete. I'm always in favor of parametric sweeps, rather than having, as I say, disparate functional forms or a grab bag of different possible configurations. If you can capture it like you suggest, we are all ears and we would like to know more about how to get this down on paper.

Macal: Okay, there is a function called the Fermi function that allows you to do that, and I can give you that since I'm a function hobbyist. When you work with David Sallach and others, they are continually requesting functions with certain properties that I come up with.

Cederman: Is it included in Repast already?

Macal: No, no.

Cederman: It should be.

Macal: Well, it's included in the *Mathematica* Repast linkage, you see, but not the Repast part alone.

Cederman: Thank you.

Sallach: I think that's all the time that we have for this session.

Social and Conflict Dynamics

EMERGENT AGENTS AND THE SIMULATION OF POLITICAL UNREST: APPLICATION TO PALESTINIAN POLITICAL COALITIONS

L.A. KUZNAR,* J. TOOLE, and N. KOBELJA,
Indiana University–Purdue University, Fort Wayne, IN

ABSTRACT

Computational social science uses agent-based computer models for simulating social phenomena from the "bottom up," allowing social phenomena to emerge from the interactions of individual agents. This approach grounds social phenomena in entities that have a clearer and more concrete existence. Yet higher order social phenomena, such as alliances, kin groups, villages, and states, at times act on goals that transcend those of their constituent members. Programming special rules for higher-order phenomena fails to capture their emergent nature. We model the emergence of radical political coalitions in Palestinian society and explore the possibility of their evolving group-level actions.

We have developed a theoretical tool, expo-sigmoid utility theory, to model the risk-sensitive behavior of agents. Expo-sigmoid utility theory is derived from the empirical observation that wealth is typically distributed from poorest to wealthiest, with quasi-periodic fluctuations around an overall exponential increase. Such distributions are typical in chiefdoms, ancient states, modern states, and the global economy. Individuals whose social rank places them on a convex (concave upward) section of a wealth curve behave in risk-prone ways; they take chances for social advancement that most individuals would avoid. Our empirical research shows that the risk proneness of individuals in a coalition is a necessary, although not necessarily sufficient, condition for collectively violent behavior.

In this paper, we utilize data on wealth distribution in the Palestinian Authority to model the formation and evolution of political coalitions within Palestinian society. We consider two primary influences on coalition-joining behavior: risk sensitivity and communication proximity. Agents play a coordination game with their neighbors, and their probabilities of taking a risk (in this case joining a coalition) are altered by their risk sensitivities. Coalitions form as emergent phenomena, exhibit collective measures of risk sensitivity, and change by the same rules as individuals, producing new attributes for these emergent agents. Risk-prone coalitions socially isolate their individual agents, in accordance with social psychology theories, altering the situational definition that agents use for decision making. Preliminary validation of the model is done by a comparison with historical developments in Palestinian political alliances.

Keywords: Agent-based modeling, emergence, adaptive agents, terrorism, Palestine

* *Corresponding author address:* Lawrence A. Kuznar, Decision Sciences and Theory Institute and Department of Anthropology, Indiana University–Purdue University, Fort Wayne, IN 46805; kuznar@ipfw.edu.

INTRODUCTION

Dynamically adaptive agents, emergent phenomena, and the agent-like behavior that is exhibited by emergent phenomena characterize social processes (Barabasi 2003; Watts 2003). Realistic modeling of social phenomena therefore requires researchers to enable simulations to produce adaptive agents and emergent phenomena, and it requires emergent structures to behave dynamically in a manner that is greater than the sum of their parts (Holland 1998; Axelrod and Cohen 2000).

Nowhere is this need more apparent than in the simulation of radical political behavior and terrorist activity. Terrorist organizations seemingly materialize out of nothing and adapt dynamically to changing political circumstances and military confrontation (Stern 2003; Arquilla et al. 1999). In this paper, we describe a simulation of political activism and terrorist activity in which we strive to incorporate realistic individual- and group-level dynamics. We apply our results to a consideration of political activity and terrorist recruitment within contemporary Palestinian society.

TERRORISM AS AN ADAPTIVE SYSTEM

Analysts argue that terrorism has entered a new era characterized by a decidedly international scope, a horizontal organization, and an ability to change and adapt to rapidly changing circumstances (Ellis 2004; Hoffman 1999). Terrorist organizations are now made up of increasingly diverse individual agents who coalesce into relatively independent cells that change with shifting social conditions. For instance, al Qaeda began with a relatively top-down command-and-control organization and recruited Jihadis from Islamic countries such as Egypt, Saudi Arabia, Yemen, and Pakistan. Increasingly, al Qaeda recruits in Western nations, it lacks explicit command and control from its founders, and its members appear to act independently on the basis of a general notion of waging a Jihad against the West. Terrorist organizations and their attendant political manifestations exhibit the classic features of an adaptive system of agents who coalesce into emergent social forms that behave by their own rules.

SIGMOID UTILITY THEORY, POLITICAL ACTIVISM, AND TERRORISM

Sigmoid utility theory provides an explanation and a method of prediction for risk-taking behavior. The basic concept is that individuals would accept gambles with negative expected value if the potential gains of such gambles exceeded the potential losses (Friedman and Savage 1948). Kuznar (2001, 2002) operationalized this concept and demonstrated its relevance in a variety of cultural settings. Kuznar and Frederick (2003) developed the concept further and demonstrated its applicability to political activist and rebellious behavior. They found that ranking individuals in a society on a wealth scale (x-axis) from poorest to wealthiest and then measuring each individual's wealth on the y-axis produces an oscillating, S-shaped or sigmoidal curve of wealth differences (Figure 1). This S-shaped oscillation occurs along a generally exponential increase in wealth in complex societies (chiefdoms, states, global economies), leading the authors to call such distributions expo-sigmoid (Kuznar and Frederick 2005). Convex (concave upward) wealth distribution curves imply that potential gains from a gamble typically exceed losses, leading to risk-prone behavior (acceptance of negative expected utility). Because

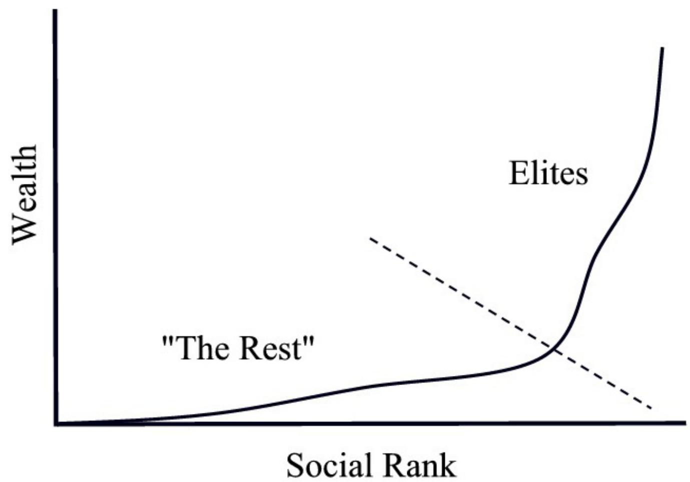

FIGURE 1 Expo-sigmoid wealth distribution

political activism and terrorism are risky, sigmoid utility theory is applicable to these phenomena. Our preliminary research indicates that other dimensions along which humans are valued (social status in terrorist groups, regard in religious organizations) also distribute expo-sigmoidally, and so the method is applicable beyond strict material economic applications. However, our research also shows that concerns over material wealth often lurk behind many revolutionary activities. The advantage of the sigmoid utility approach over other economic theories is that it gets us beyond simplistic poverty-based explanations of rebellion and terrorism.

The close study of the demographics of any revolution demonstrates that few poor people ever rebel and that comparatively wealthy individuals typically lead rebellions (Brinton 1964). Recent studies of terrorist recruitment demonstrate that there is a near-zero correlation between wealth and terrorist activity (Krueger and Maleckova 2002, 2003; Palmer Harik 1996). Even comparatively wealthy individuals can find themselves in a convex distribution of wealth or status and correspondingly feel much aggrieved that they are not as well off or as well regarded as their near superiors. As aggrieved individuals seek one another out, rebellious coalitions can emerge spontaneously from across a social spectrum, often under the banner of nationalism, ethnic identity, or religious affiliation (Kuznar et al. 2005).

We use the Arrow-Pratt measure of risk aversion (Pratt 1964), $r(x)$, to measure an individual's sensitivity to risk by fitting an expo-sigmoid curve, $W(x)$, to a wealth distribution over some wealth or status, x. Negative values of the measure indicate risk proneness, positive values indicate risk aversion, and zero values indicate risk neutrality.

SOCIAL PSYCHOLOGY OF TERRORIST AND EXTREMIST GROUPS

Studies of the social psychology of terrorist groups demonstrate that group members actively work to isolate recruits from their families and the wider society (Stahelski 2004; Hudson 1999; Post 1990). Such isolation reinforces group goals and an individual's attachment to and dependence on a terrorist group (Soibelman 2004). It also reinforces the group's situational definition, the classification of others, and the emotional attachment of an individual to a group or a cause (Sallach and Mellarkod 2005; Jackson, 2002).

There exists a continuum — from groups that are so socially isolated that they remain small, suffer from shared distortions of reality, and lack networked contacts that can fuel funds and expertise, to groups that are so open that they are easily infiltrated by new ideas and personnel and consequently never achieve a stable identity or goal (Axelrod and Cohen 2000, pp. 50-52). In this project, members of risk-prone groups avoid joining with nonmembers in proportion to the overall risk proneness of their current group. Agents in risk-neutral to risk-averse coalitions join or defect with others according to the rules applied to individual risk sensitivity.

This approach imbues a risk-prone coalition with the social dynamic described in social psychological studies of terrorist groups: as individuals join radical groups, their individual perceptions of reality are brought more in line with the group's view, while the groups both psychologically and physically isolate members from the outside world. Such groups begin to take on a life of their own, and their members behave in a coordinated fashion so that the emergent group itself will begin to behave as though it was an individual agent, in accordance with rules that originally applied to individuals. Our approach also recognizes that while resentment over economic inequalities may be an initial factor in an individual's decision to join a radical group, these psychological factors may increase in importance with a group's evolution.

SIMULATING COALITION FORMATION IN PALESTINE

Axelrod (1997) advocates the use of simple, abstract simulations for exploring the basic relationships between parameters. On the other hand, if simulations are going to capture the complexity and dynamism of real social systems, then researchers need to develop more complicated and realistic models (Kuznar 2005). However, the complexity of realistic simulations makes establishing the causal relationships between variables nearly as difficult as in simulations as it is in the real world. Our intention in this exercise is to provide a simulation in which the relationships between parameters are apparent but in which there is enough verisimilitude to make a reasonable connection to real social phenomena.

The first step in establishing a model's realism is to use real data as inputs. We use data on the average incomes of Palestinians provided by the Palestinian Authority (Palestinian National Authority–Palestinian Central Bureau of Statistics 2003) to initialize the wealth distribution in the model. The Palestinian data have the classic expo-sigmoid distribution of a complex society. To fit a curve to these data, we first take the natural log of the data, and then we produce a periodogram to establish the dominant frequencies that create the oscillations in the curve (Lomb 1976). Once the dominant frequencies are identified, we produce a trigonometric polynomial, plus a linear term and a constant, by multiplying sine and cosine terms by the coefficients of the dominant frequencies from the periodogram. That equation is then used as the argument for an exponential function, yielding an equation of the general form:

$$W(x) = e^{f(x)},$$

where $f(x) = k + ax + \Sigma\, c_i\, sin(x) + d_i cos(x)$ for i dominant frequencies. Our simulation uses 256 agents assigned wealth according to the initialized function and ranked according to their wealth.

We model agent coalition formation and its payoffs with a variety of game variously referred to as the stag hunt or coordination game (Battalio et al. 2001). This game has two equilibria — either both players defect or both players cooperate — and the Nash optimum is to play a mixed strategy of join and defect (Table 1).

This game is especially useful for modeling the payoffs associated with joining a rebellious coalition. Many people never join rebellions, simply continuing on with their lives and earning the payoffs they normally expect from life. However, if rebels were to succeed, rebel payoffs would be much higher. Refusing to join a rebellion usually has no particularly severe punishments, although there can be exceptions. The risk-averse choice in this game is to defect (not join or cooperate) and continue to earn moderate payoffs. The risk-prone choice is to join (or cooperate with the high payoff) because of the risk that one's partner will defect. We randomly assign agents for game play within their Moore neighborhood, mimicking small-scale, face-to-face societies.

Once each player has played the game, the agents are ranked according to their wealth levels. We then fit a curve to these data and use divided difference numerical techniques to estimate the first and second derivatives for each agent over the wealth function. These estimates are then used to produce the Arrow-Pratt risk sensitivity measure for each agent. The Arrow-Pratt measures alter the probability of joining. We alter the probabilities in proportion to the degree to which an agent is either risk prone or risk averse. The most risk-prone agent always joins, and join probabilities are altered continuously up to the most risk-averse agent, who never joins. Coalitions form when agents join. If an activated agent solicits another agent and it joins, then the activated agent joins the other agent's coalition.

Once coalitions form, average risk sensitivities for these coalitions are calculated. Then the agents' probability of joining with nonmembers is altered in the reverse fashion from individual agent probabilities. The probability of joining with a nonmember is varied

TABLE 1 Payoff matrix for coordination game[a]

	Column Player	
Row Player	Join	Defect
Join	(R, R)	(S, T)
Defect	(T, S)	(P, P)

[a] Payoffs where R > T = P > S reflect the relative benefits and costs of joining and defecting from a rebellious coalition.

continuously, from zero for members of the most risk prone (most insular) coalition, to the Nash optimum for risk-neutral agents.

RESULTS

Relatively stable coalitions emerge by iteration 75, making comparisons to empirical data useful. In iteration 75, there are 22 coalitions and 2 independent agents. The coalitions range in size from 2 to 29 members, with a mean of 12. The risk sensitivity of coalitions is an average of –0.00064 and ranges from –0.017 to 0.026. Since we are interested in political radicalism and terrorism, we concentrate on the most risk-prone coalitions.

Figure 2 shows the distribution of agents from the three most risk-prone coalitions. Two results may be particularly important for understanding the emergence of radical groups, terrorist recruitment, and the long-run behavior of such groups. First, all groups contain agents from a range of wealth levels. Coalition 3 is the most striking because it is simultaneously the most risk-prone group and composed entirely of agents from the upper half of the wealth distribution. Once again, wealth or poverty per se is not important, only the relative wealth differences measured by the convexity of wealth distribution curves. The other two coalitions (20 and 63) contain a mix of impoverished and relatively wealthy agents. Second, coalitions 20 and 63 contain a minority of risk-averse agents. This is possible because each agent plays a mixed strategy, so there is usually a probability that an agent can either join or defect from a coalition. While a Nash mixed strategy does not model the precise mechanism by which otherwise risk-averse individuals would join a radical group, it nonetheless provides for this real-world possibility. Both of these basic results may help explain the curious lack of correlation between wealth and terrorist activity, despite the facts that terrorist groups often cite economic grievances and that terrorists recruit in impoverished neighborhoods.

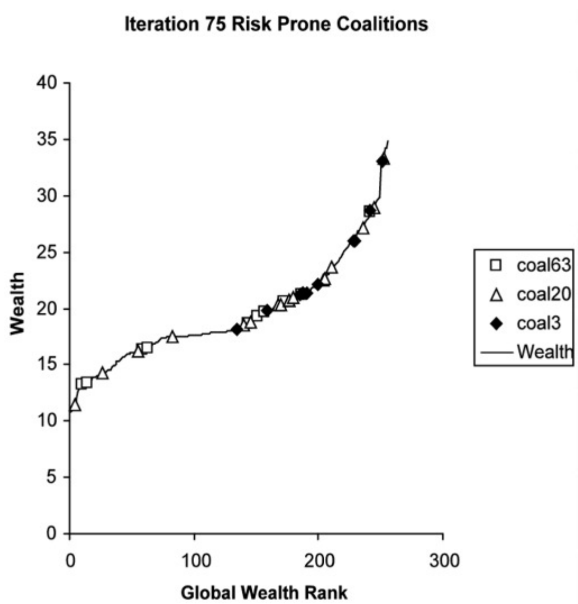

FIGURE 2 Agents of the three most risk-prone coalitions in iteration 75

The risk sensitivity of all coalitions changed dynamically through time with changes in the collective fortunes and risk sensitivities of group members (Figure 3). Coalition 3 was the shortest-lived of the three coalitions and remained risk-prone throughout its history. Coalition 20 was the most dynamic, swinging from risk prone to risk averse, but fluctuations were drastic. Coalition 63 follows a secular trend in the simulation, toward general risk neutrality but, once again, with unpredictable fluctuations between risk proneness and risk aversion. These fluctuations, along with the variable lives of coalitions, mirror the dynamic lives of real Palestinian political parties and radial groups (Rubin 1999).

DEMOGRAPHICS OF PALESTINIAN RADICAL GROUPS

Most studies of Palestinian political and terrorist groups are qualitative, and the few quantitative studies present aggregate statistics that mask the details of wealth and status distribution. Therefore, complete validation of our model is not possible. However, some comparisons can be made, and they indicate the potential of our approach.

The data on socioeconomic status and political involvement/terrorism are mixed for Palestine. Some studies find no correlation (Krueger and Maleckova 2003), some find a positive correlation between wealth and radicalism (Inbar and Yuchtman-Yaar 1989, Table 3), some document that the political elite originate from elites and upwardly mobile middle-class families (Brynen 1995, page 39), some note that suicide bombers tend to come from poor communities (Weinberg et al. 2003, page 143; Pedahzur et al. 2003, page 418), and others argue that economic deprivation is an insufficient explanation of radicalism (Moghadam 2003, page 76; Soibelman 2004, page 185).

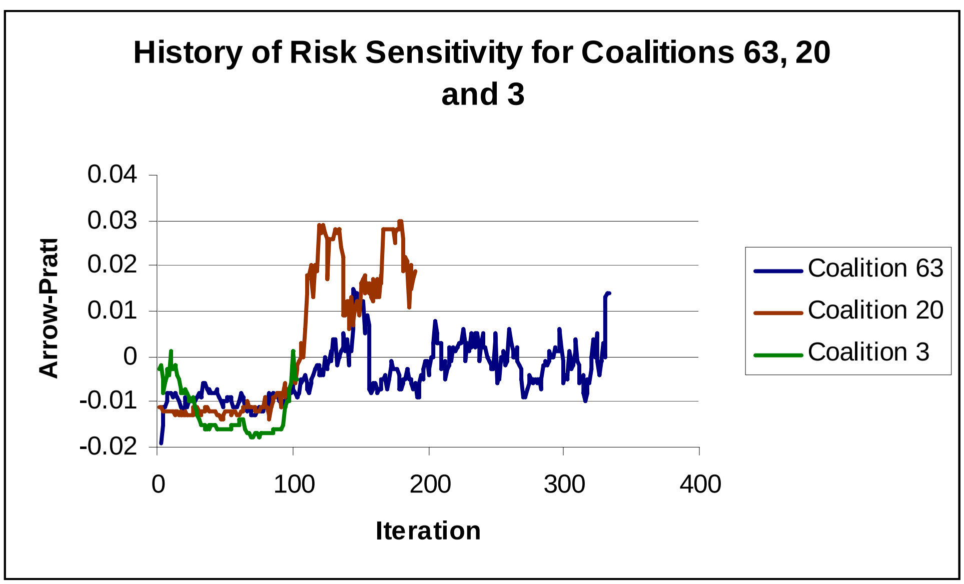

FIGURE 3 History of risk sensitivity for coalitions 63, 20, and 3

The simulation outputs, however, provide an explanation for the mixed results of empirical studies. First, since risk proneness occurs at the low and high ends of the wealth distribution, any correlation between wealth and risk is bound to be near zero. The fact that the most risk-prone coalition in iteration 75 (coalition 3) was statistically significantly wealthier than the whole population ($p = 0.01$) and found among the wealthiest half of the agents mirrors the origin of some radical groups from the middle-class and university community (Paz 2003, p. 35) and the derivation of the political elite from upper- and middle-class families (Brynen 1995). On the other hand, the mean wealth of the other two coalitions in iteration 75 was not statistically significantly different from the population, with one coalition being slightly higher and the other being slightly lower than the mean. These coalition results resemble the data found by other researchers (Krueger and Maleckova 2003). The problem with standard approaches is that they ignore the important variability in wealth distributions that mask relevant variation and lead to mixed results. Our approach utilizes a fresh approach by considering the significance of relative wealth differences and using the flexibility of simulation to produce these varied results.

CONCLUSION

Our simulation of political group formation based on economic risk sensitivity focused on a few key factors that could explain the mixed demographics of radical groups within Palestinian society. Our programming approach allowed agent's preferences to evolve through time. Agent coalitions were not hard-coded but instead emerged according to autonomous agent decisions. We further altered agent behavior to make agents more insular once risk-prone coalitions emerged. The social isolate algorithm we used allowed coalitions to behave dynamically and in accordance with individual behavioral rules as though they themselves were agents. This approach thoroughly grounds the dynamic behavior of emergent phenomena in the interactions among their constituent elements (Holland 1998). The result was a simulation that contains essential elements for realistic modeling of complex phenomena: autonomous agents, emergent phenomena, and emergent behaviors for these phenomena. The model produced results that were consistent with qualitative analyses of Palestinian radical groups.

In order to improve upon our preliminary approach, future work will include: alternative decision rules, more boundedly rational agents, alternate games and alternate payoffs for sensitivity analysis, changed random seeds to produce a range of possible runs, and more detailed empirical data for better validation.

REFERENCES

Arquilla, J., D. Ronfeldt, and M. Zanini, 1999, "Networks, Netwar, and Information-age Terrorism," in *Countering the New Terrorism*, I.O. Lesser, B. Hoffman, J. Arquilla, D. Ronfeldt, and M. Zanini (Eds.), pp. 39–84, Santa Monica, CA: RAND.

Axelrod, R., 1997, *The Complexity of Cooperation: Agent-based Models of Conflict and Cooperation*, Princeton, NJ: Princeton University Press.

Axelrod, R., and M.D. Cohen, 2000, *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, New York, NY: Basic Books.

Barabasi, A.-L., 2003, *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*, New York, NY: Plume.

Battalio, R.C., L. Samuelson, and J. van Huyck, 2001, "Optimization Incentives and Coordination Failure in Laboratory Stag Hunt Games," *Econometrica* **69**(3):749–764.

Brinton, C., 1964, *The Anatomy of Revolution, Revised and Expanded Edition*, New York, NY: Vintage Books.

Brynen, R., 1995, "The Dynamics of Palestinian Elite Formation," *Journal of Palestine Studies* **24**(3):31–43.

Ellis, B., 2004, "Countering Complexity: An Analytical Framework to Guide Counter-terrorism Policy-making," in *Terrorism and Counterterrorism: Understanding the New Security Environment*, R.D. Howard and R.L. Sawyer (Eds.), pp. 109–122, Guilford, CT: McGraw-Hill/Dushkin.

Friedman, M., and L.J. Savage, 1948, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy* **4**:279–304.

Hoffman, B., 1999, "Terrorism Trends and Prospects," in *Countering the New Terrorism*, I.O. Lesser, B. Hoffman, J. Arquilla, D. Ronfeldt, and M. Zanini (Eds.), pp. 7–38, Santa Monica, CA: RAND.

Holland, J. H., 1998, *Emergence: From Chaos to Order*, New York, NY: Basic Books.

Hudson, R.A., 1999, *The Sociology and Psychology of Terrorism: Who Becomes a Terrorist and Why?*, Washington, DC: Federal Research Division, Library of Congress.

Inbar, M., and E. Yuchtman-Yaar, 1989, "The People's Image of Conflict Resolution: Israelis and Palestinians," *Journal of Conflict Resolution* **33**(1):37–66.

Jackson, J.W., 2002, "Intergroup Attitudes as a Function of Different Dimensions of Group Identification and Perceived Intergroup Conflict," *Self and Identity* **1**:11–33.

Krueger, A.B., and J. Maleckova, 2002, "Does Poverty Cause Terrorism?," *New Republic*, pp. 27–33, June 24.

Krueger, A.B., and J. Maleckova, 2003, "Education, Poverty and Terrorism: Is There a Causal Connection?," *Journal of Economic Perspectives* **17**(4):119–144.

Kuznar, L.A., 2001, "Risk Sensitivity and Value among Andean Pastoralists: Measures, Models and Empirical Tests," *Current Anthropology* **42**(3):432–440.

Kuznar, L.A., 2002, "Evolutionary Applications of Risk Sensitivity Models to Socially Stratified Species: Comparison of Sigmoid, Concave and Linear Functions," *Evolution and Human Behavior* **23**(4):265–280.

Kuznar, L. A., 2005, “High Fidelity Computational Social Science in Anthropology: Prospects for Developing a Comparative Framework,” *Social Science Computer Review,* in press.

Kuznar, L.A., and W.G. Frederick, 2003, “Environmental Constraints and Sigmoid Utility: Implications for Value, Risk Sensitivity, and Social Status,” *Ecological Economics* **46**:293–306.

Kuznar, L.A., and W.G. Frederick, 2005, “The Effect of Nepotism on Political Risk Taking and Social Unrest,” in *Proceedings of the 2005 Annual Conference of NAACSOS* (North American Association for Computational Social and Organizational Science).

Kuznar, L.A., R.L. Sedlmeyer, and W.G. Frederick, 2005, “Agent Based Models of Risk Sensitivity: Applications to Social Unrest, Collective Violence and Terrorism,” paper presented at the Third Lake Arrowhead Conference on Human Complex Systems, Lake Arrowhead, CA, May 18–22.

Lomb, N.R., 1976, “Least-squares Frequency Analysis of Unequally Spaced Data,” *Astrophysics and Space Science* **39**:447–462.

Moghadam, A., 2003, “Palestinian Suicide Terrorism in the Second Intifada: Motivations and Organizational Aspects,” *Studies in Conflict & Terrorism* **26**:65–92.

Palestinian National Authority–Palestinian Central Bureau of Statistics, 2003, *Number of Enterprises and Persons Engaged and Main Indicators by Economic Activity in Remaining West Bank and Gaza Strip — 2003;* available at http://www.pcbs.org/economic/table1b.aspx.

Palmer Harik, J., 1996, “Between Islam and the System,” *Journal of Conflict Resolution* **40**(1):41–67.

Paz, R., 2003, “The Development of Palestinian Islamic Groups,” in *Revolutionaries and Reformers: Contemporary Islamist Movements in the Middle East*, B. Rubin (Ed.), pp. 23-40, Albany, NY: State University of New York Press.

Pedahzur, A., A. Perlinger, and L. Weinberg, 2003, “Altruism and Fatalism: The Characteristics of Palestinian Suicide Terrorists,” *Deviant Behavior* **24**:405–423.

Post, J.M., 1990, “Terrorist Psycho-logic: Terrorist Behavior as a Product of Psychological Forces,” in *Origins of Terrorism: Psychologies, Ideologies, Theologies, States of Mind*, W. Reich (Ed.), pp. 25–40, Baltimore, MD: Johns Hopkins University Press.

Pratt, J.W., 1964, “Risk Aversion in the Small and in the Large,” *Econometrica* **32**:122–136.

Rubin, B., 1999, *The Transformation of Palestinian Politics: From Revolution to State-building*, Cambridge, MA: Harvard University Press.

Sallach, D.L., and V. Mellarkod, 2005, “Typifying Situations: Contextualizing & Cohering Action,” in *Proceedings of the 2005 Annual Conference of NAACSOS* (North American Association for Computational Social and Organizational Science), pp. 1–6.

Soibelman, M., 2004, “Palestinian Suicide Bombers,” *Journal of Investigative Psychology and Offender Profiling* **1**:175–190.

Stahelski, A., 2004, “Terrorists Are Made, Not Born: Creating Terrorists Using Social Psychological Conditioning,” *Journal of Homeland Security,* pp. 1–7, March.

Stern, J., 2003, “The Protean Enemy,” *Foreign Affairs* **82**(4):27–40.

Watts, D.J., 2003, *Six Degrees: The Science of a Connected Age*, New York, NY: W.W. Norton and Company.

Weinberg, L., A. Pedahzur, and D. Canetti-Nisim, 2003, “The Social and Religious Characteristics of Suicide Bombers and Their Victims,” *Terrorism and Political Violence* **15**(3):139–153.

MULTIPLE ATTITUDE DYNAMICS IN LARGE POPULATIONS

W. JAGER,* University of Groningen, Groningen, The Netherlands
F. AMBLARD, IRIT-UT1, University of Social Sciences, Toulouse, France

ABSTRACT

A multi-agent simulation for studying opinion dynamics in the context of two opinion dimensions is presented. The agent rules that describe changes in opinion are based on a theory about persuasion and distinguish between central and peripheral processing. Central processing is formalized as assimilating or contrasting the opinion of a contacted agent, depending on the initial (dis)agreement with the other agent. Peripheral processing is formalized as a source effect: If an agent agrees with another agent about one issue, it will also assimilate the position of this agent about another unrelated issue, regardless of the initial difference. Experiments show that the correlation between opinions on the two dimensions increases if agents engage in peripheral processing on one dimension. In addition, some experiments are performed with a meta-actor influencing the whole population.

Keywords: Social simulation, agent-based simulation, opinion dynamics, social judgment theory, elaboration likelihood model

INTRODUCTION

The recent rejection of the European Constitution by the voters in France and The Netherlands (2005) instigated a debate on how this could happen in countries having a basically pro-European-Union attitude amongst the population. We hypothesize that the complexity of the constitution, along with the limited information on its potential effects, caused many people not to process the arguments in defining their vote but instead to use the position of other people, and in particular that of major politicians, to determine their position. Especially the fact that the unpopular leaders of the government strongly campaigned in favor of the constitution may have resulted in a contrasting effect on this topic, despite the population's initial pro-European attitude.

Experimenting with the dynamics of attitude or opinion dynamics is not possible by using laboratory studies. Field data on the contrary are too complex to identify the causalities of observed dynamical processes. Multi-agent simulation provides a tool allowing experimentation with these dynamics, because large series of experiments can be performed systematically by varying assumptions on how people change their opinion and on conditions of the initial opinions of the population. This has resulted in an increasing body of research on opinion dynamics from using multi-agent simulation. Several researchers have worked on simulating how opinions, attitudes, or voting behavior in groups emerges from locally interacting people. Some work on binary opinions (e.g., Latane and Nowak 1997; Galam 1999), and some use

* *Corresponding author address*: Wander Jager, Faculty of Management and Organization, University of Groningen, Groningen, The Netherlands; email: w.jager@bdk.rug.nl.

continuous opinions, where influence depends on distance (using a threshold, e.g., Deffuant et al. 2001, 2002; Weisbuch et al. 2002; Hegselmann and Krause 2002).

These studies mainly used attraction of opinions as a mechanism to generate opinion dynamics, and hence did not use existing behavioral theory on attitude change to formalize agent rules. More recently, researchers have started to use behavioral theory in formalizing these rules. For example, they have used social judgment theory (SJT) as a formalization of both assimilation and rejection effects (Jager and Amblard 2005) and self-categorization theory in studying meta-contrast effects (Salzarulo 2004).

To study the dynamics involved in attitude change, we formalize relevant social psychological theories in the architecture of agents. The field of persuasion, social influence, and attitude change provided us with a rich theoretical perspective on how people change their attitudes and on the factors determining the degree and stability of these changes. In particular, SJT (Sherif and Hovland 1961) is relevant in understanding how people assimilate or contrast their opinion after being confronted with another position. The basic idea of this theory is that a change of a person's attitude depends on the position of the persuasive message that is being received. If the advocated position is close to the initial position of the receiver, it is assumed that this position falls within the *latitude of acceptance* of the receiver. As a result, the receiver is likely to shift in the direction of the advocated position *(assimilation)*. If the advocated position is distant to the initial position of the receiver, it is assumed that this position falls within the *latitude of rejectance* of the receiver. As a result, the receiver is likely to shift away from the advocated position *(contrast)*. If the advocated position falls outside the border of the latitude of acceptance but is not that distant that it crosses the border of the latitude of rejectance, it will fall within the *latitude of noncommitment*, and the receiver will not shift its initial position. Formalizing this SJT in an agent-based model allowed us to study the conditions under which the attitudes in populations tend to polarize, converge, or display pluriformity (Jager and Amblard 2005). One main result was that when the latitude of noncommitment gets small, which has been found to happen in crisis situations (O'Keefe 1990), our model produces polarization effects.

However, both the experimentally based laboratory studies and the social simulations addressed processes where only a single attitude is taken into consideration. Yet the example of the vote on the constitution indicates that often more than one attitude is taken into consideration. Many people reported to have voted against this constitution not because of their negative attitude toward this constitution but because of their negative attitude toward the political leaders advocating a positive vote.[1] In this paper, we study to what extent processes such as congruity affect attitude dynamics in large populations. In the work that we present in this paper, we focus on (1) two attitude dimensions rather than one, (2) cognitive effort in processing information, and (3) possible effects of mass-media performances of popular versus unpopular leaders.

People may spend more or less cognitive effort in elaborating on the attitude position of another person. This is captured in the *elaboration likelihood model (*ELM; Petty and Cacioppo 1986), which discerns a *central* and a *peripheral route* to attitude change. The central route pertains to the elaboration of pure arguments in a persuasive message and/or new information. Here people are motivated and capable of processing the arguments of the message,

---

[1] Actually, also a lot of people reported to have voted in favor of the constitution because they disliked the political position of politicians advocating to vote against the constitution.

whereas peripheral processing is more likely when people's motivation to elaborate is low and/or their cognitive processing ability is limited (i.e., complex issues). The peripheral route is concerned with the elaboration of form aspects or cues of a message, such as the number of arguments and the credibility and attractiveness of the source. The attractiveness of the source is related to similarity of attitudes. Generally, people like to have opinions similar to those of people with whom they interact (Festinger 1954). This implies that when engaging in peripheral processing, people may compare on one attitude dimension how similar they are, and depending on that observed (dis)similarity, either accept or reject the information of the other attitude dimension.

In the following text, we outline the formalization of this theory in rules that apply to the agents we use.

THE MODEL

For the formalization of the SJT, which refers to central processing, we follow the model as used by Jager and Amblard (2005). This formalization implies that we have a population with N individuals. Each individual i has an opinion (an attitude) x_i, a threshold determining the latitude of acceptance u_i, and a threshold determining the latitude of rejection t_i, with $t_i > u_i$. Varying the values of t_i and u_i allows for modeling agents having different attitude structures. For example, an agent having a high ego-involvement can be formalized as an agent where t_i is slightly larger or equal to u_i: The agents are scheduled to communicate on a random basis by scheduling random pairs for each time-step of the simulation. During the interaction between individual i and individual j, the following rules are applied:

If $|x_i - x_j| < u_i$, then $dx_i = \mu.(x_j - x_i)$.

If $|x_i - x_j| > t_i$, then $dx_i = \mu.(x_i - x_j)$.

where the parameter μ controls for the strength of influence. The same rules are applied for the update of the opinion of the individual j.

For the formalization of peripheral processing, we formalize two attitude dimensions that agents discuss. After encountering another agent, the attitudinal shift on one dimension will affect the shift in the other dimension, thus indicating peripheral source effects. A assimilation or contrast effect on the first attitude dimension will also translate in a similar assimilation or contrast effect in the second dimension. Here agents select attitude A for the interaction process, and depending on the outcome (assimilation, noncommitment, or contrast), they will also apply this outcome to dimension B. The rule describing peripheral processing is:

If $|xA_i - xA_j| < u_i$, then $dAx_i = \mu.(xA_j - xA_i)$ and $dBx_i = \mu.(xB_j - xB_i)$.

If $|xA_i - xA_j| > u_i$, then $dAx_i = \mu.(xA_j - xA_i)$ and $dBx_i = \mu.(xB_i - xB_j)$.

RESULTS

In experimenting with the model, we use a research design that uses three basic experimental conditions. In experiment 1, we replicate the experiments of Jager and Amblard (2005), only here we formalize two attitude dimensions instead of one. Three conditions are tested, which lead in the original single dimension experiment to polarization, convergence, and pluriformity. In experiment 2, we introduce peripheral processing on attitude dimension B. The same three conditions are run. Finally, in experiment 3, we explore how a meta-actor that is capable of addressing all agents simultaneously affects the attitude dynamics. Also here we explore these effects for the three conditions, and we explore the effects of extreme versus average positions of the meta-actor on the two attitude dimensions.

Experiment 1: Central Processing on Two Dimensions

In the first experiment, agents engage exclusively in central processing on both dimensions according to the principles of SJT. Sixteen hundred (1,600) agents are positioned on regular lattice and randomly contact one of their four neighbors, either south, east, north, or west (Von Neumann neighborhood). The contact implies a comparison and resulting shift first on attitude dimension A, and subsequently on dimension B.

If $|xA_i - xA_j| < u_i$, then $dxA_i = \mu.(xA_j - xA_i)$.

If $|xA_i - xA_j| > t_i$, then $dxA_i = \mu.(xA_j - xA_j)$.

If $|xB_i - xB_j| < u_i$, then $dxB_i = \mu.(xB_j - xB_i)$.

If $|xB_i - xB_j| > t_i$, then $dxB_i = \mu.(xB_j - xB_j)$.

Conditions for Experiment 1

In this experiment, we create a condition where the latitude of acceptance is high and the noncommitment is high, by setting U at 1.0 and T at 1.5. In the single attitude condition (Jager and Amblard 2005), this condition stimulated convergence to a single attitude position.

Results of Experiment 1

Figure 1 presents the developments on both attitudes for different time-steps of the simulation. In every time-step, a single agent is randomly selected. This agent randomly interacts with one of its four neighbors. Hence, in 1,600 time-steps, each agent on average had two interaction contacts, one because it was selected to engage in an interaction, and one because it was selected by another agent. On each grid, the color figures the opinion of the agent between –1 (red) and +1 (green) coding for opinions near 0. The right-hand figure positions agents on the basis of their attitude position on A (horizontal axis) and on B (vertical axis), thus indicating the relation between positions on A and B. The blue lines here indicate the social network (i.e., the links between the agents).

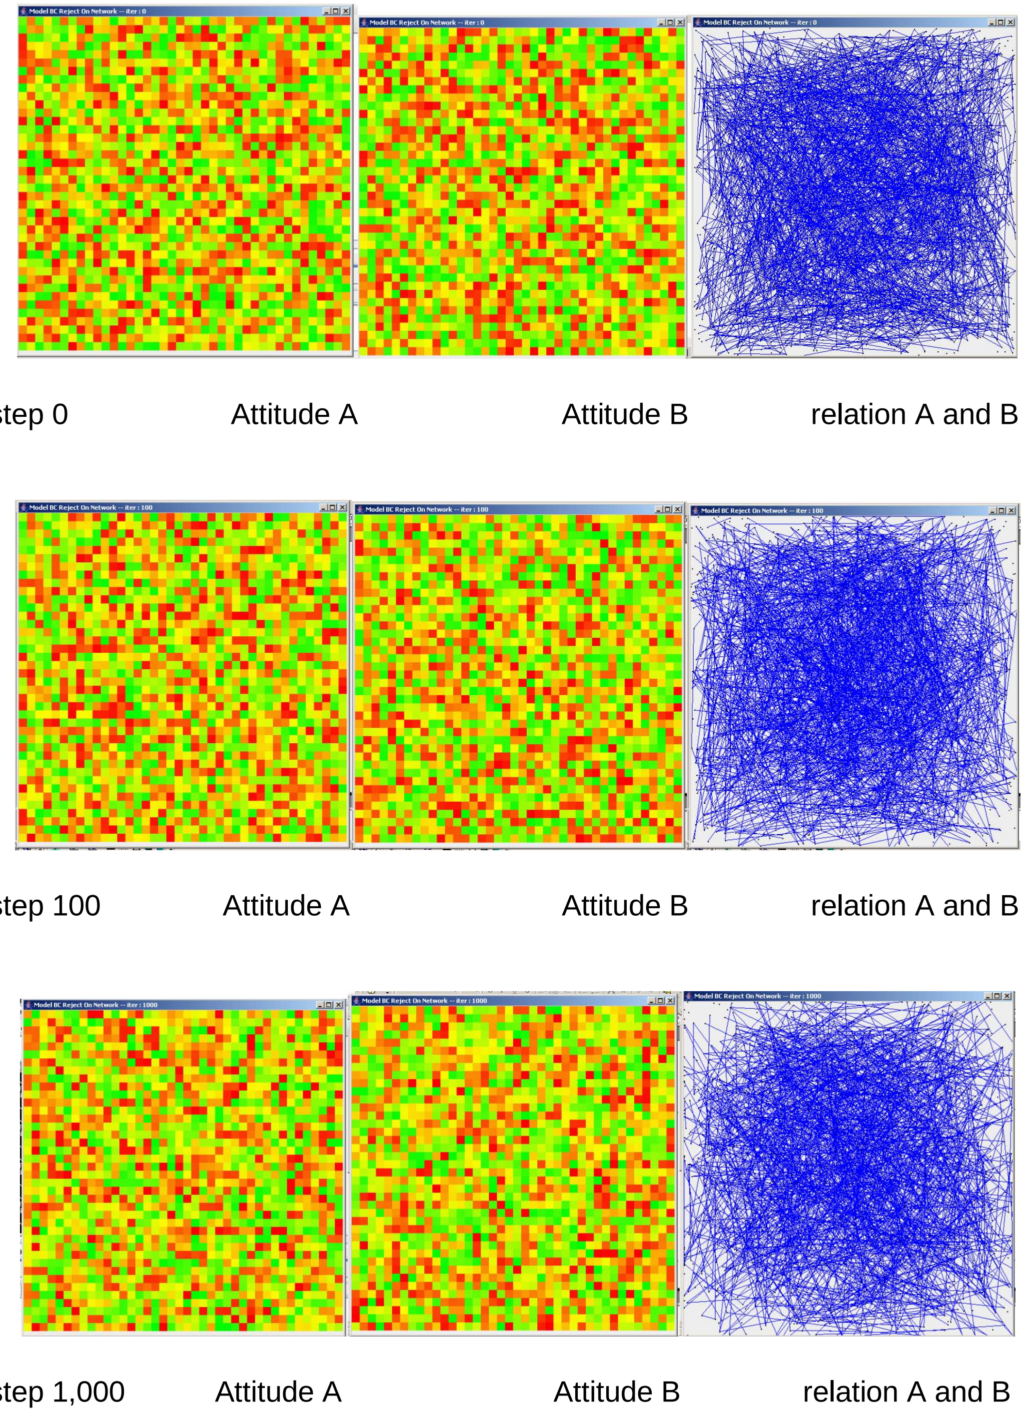

FIGURE 1 Attitude position on A and B over time

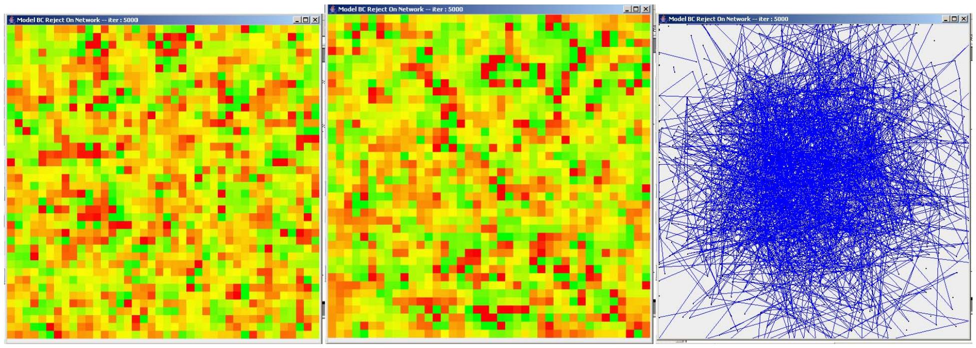

Time-step 5,000 Attitude A Attitude B relation A and B

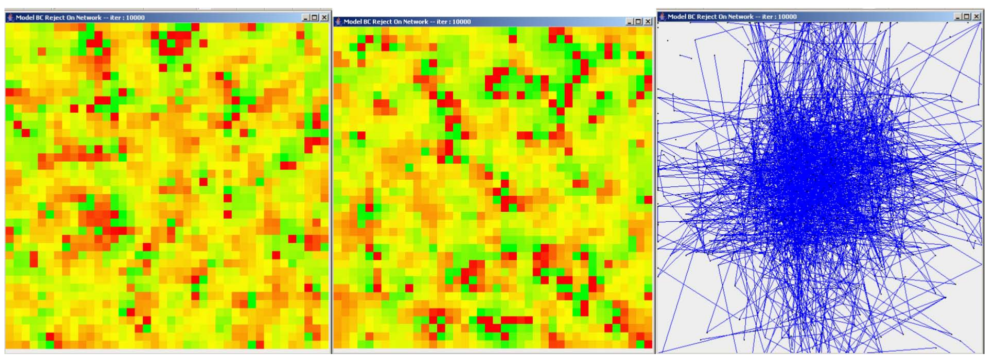

Time-step 10,000 Attitude A Attitude B relation A and B

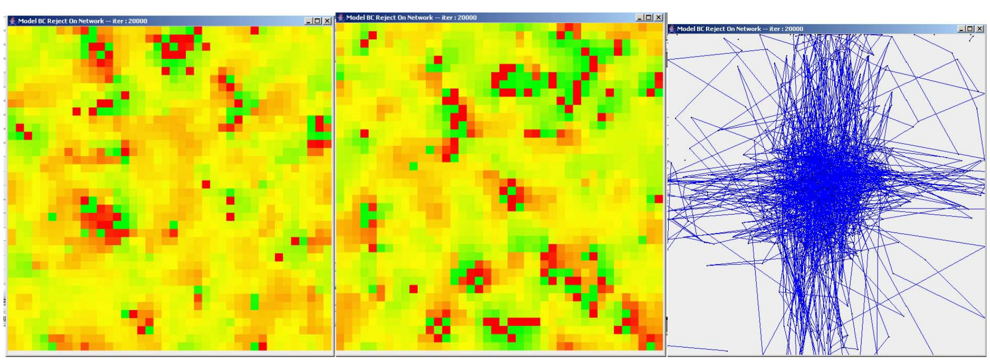

Time-step 20,000 Attitude A Attitude B relation A and B

FIGURE 1 Cont.

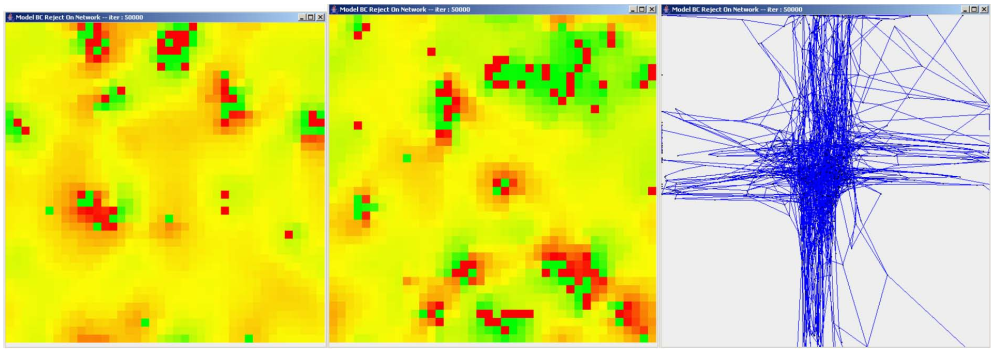

Time-step 50,000 Attitude A Attitude B relation A and B

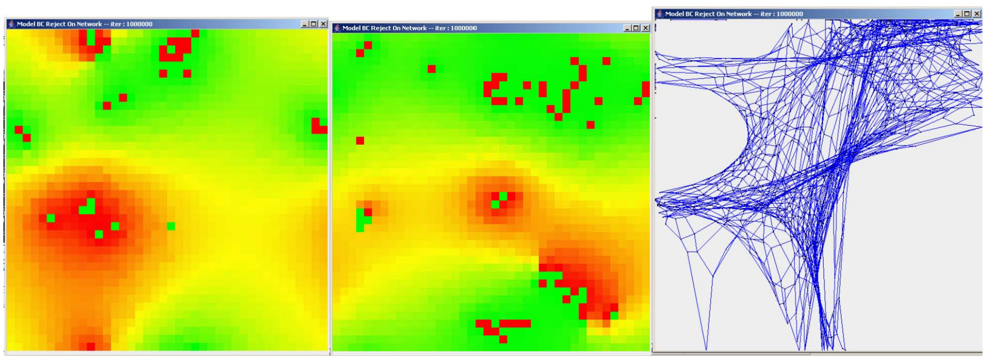

Time-step 1,000,000 Attitude A Attitude B relation A and B

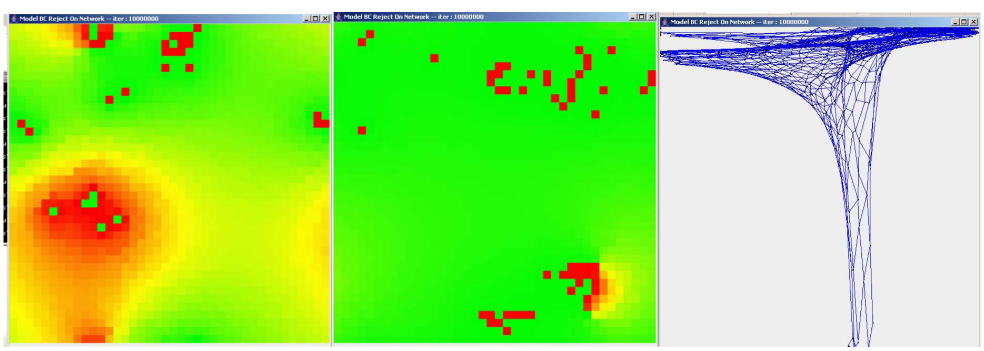

Time-step 10,000,000 Attitude A Attitude B relation A and B

FIGURE 1 Cont.

What can be observed in this experiment is that after a large number of time-steps, attitudes appear to converge on both attitude dimensions. There are still a minority of agents having an extreme position. However, an agent having an extreme position on one dimension is most likely to have a mid position on the other dimension, resulting in the emergence of the cross-like figure in the relational graph. One has to be aware that this cross-like figure is not a systematic outcome of this condition. Sometime, the population converges quickly toward an extreme on the first attitude A, and then the second dimension B stays quite uniformly distributed between –1 and +1. Instead of a cross-like figure, convergence to an extreme on attitude A results in a vertical line either on the left (A = –1) or the right (A = +1) of the figure. Looking at Figure 1, we observe that whereas at $t = 50{,}000$, it appears that the attitude dimensions tend to grow toward a convergence, the number of extremists is still large enough to generate large attitude shifts, as the results of $t = 1{,}000{,}000$ and 10,000,000 indicate. Here we observe that despite the initial tendency toward convergence, a polarization on dimension B emerges, with a large majority adhering to the green position. Also it can be observed that in the most extreme attitude areas (red or green), small numbers of dissidents show up. Here a sharp polarization effect emerges on the very local level.

In addition, the results do not indicate a string correlation between the attitude position on A and B. To get a better view of the relation between A and B, we calculated the correlation between A and B over time for 10 simulation experiments (see Figure 2).

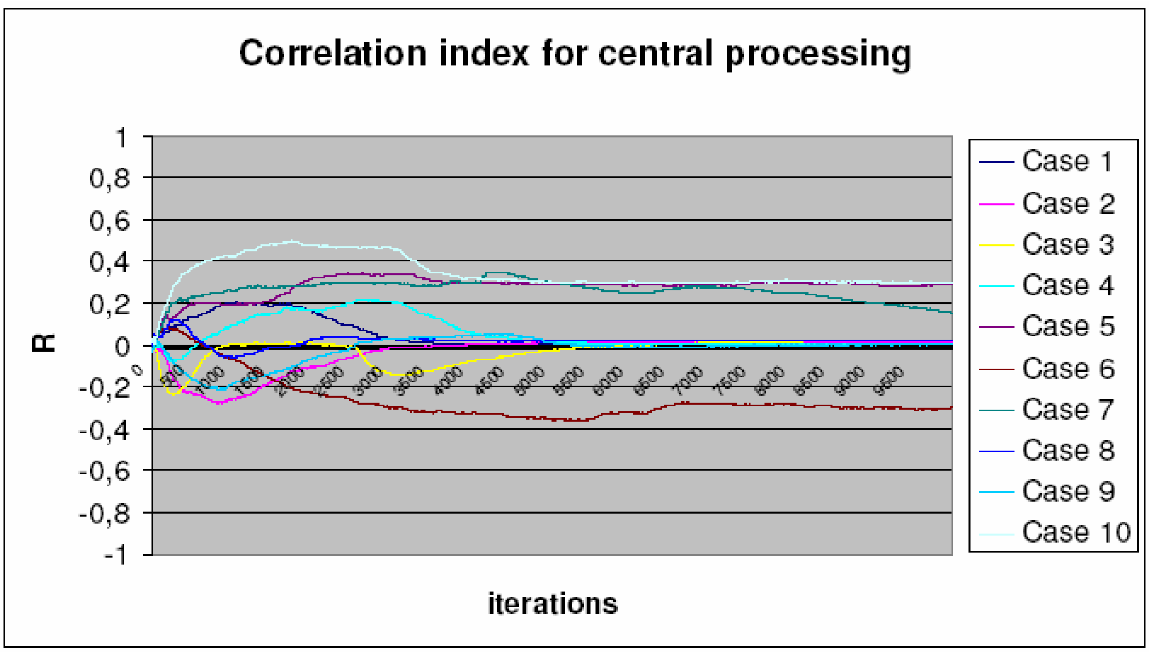

FIGURE 2 Correlation between A and B over time for 10 simulation runs

Experiment 2: Central and Peripheral Processing

In the second experiment, we implement central processing on dimension A according to the SJT, and peripheral processing on dimension B. Here we select at random existing relationships on the social network, and we let the agents interact on dimensions A and B. They apply the central processing rule for attitude A. For attitude B, they apply the peripheral rule as follows:

If $|xA_i - xA_j| < u_i$, then $dAx_i = \mu.(xA_j - xA_i)$ and $dBx_i = \mu.(xB_j - xB_i)$.

If $|xA_i - xA_j| > t_i$, then $dAx_i = \mu.(xA_i - xA_j)$ and $dBx_i = \mu.(xB_i - xB_j)$.

Conditions for Experiment 2

Experiment 2 replicates experiment 1 by setting the latitude of acceptance and the noncommitment high (U at 1.0 and T at 1.5).

Results of Experiment 2

Figure 3 presents the developments on both attitudes for different time-steps of the simulation. The figures again represent the position on attitude dimension A, attitude dimension B, and the relation between positions on A and B, respectively.

This experiment shows that when agents engage in peripheral processing on dimension B, the attitude positions on A and B are becoming related. Whereas most agents tend to converge toward a mid position, we observe especially during time steps 5,000 to 20,000 that a proportion of agents having an extreme position on attitude dimension A also develop an extreme position on dimension B. This is the result of the peripheral processing on B, where contrast and assimilation effects on A translate to the same effects on B. Initially there appears to be no strong correlation, as having an extreme positive position on A may coincide with a extreme positive or negative position on B, as indicated by the X-shaped relation graph. However, developments in later time-steps show that a virtually perfect (in this case negative) correlation between the attitude positions emerges. This can be seen in the color distribution on dimensions A and B, where the B figure is almost a perfect negative of the A figure (red is green and vice versa). Whereas here we observe that a positive position on A is coupled with a negative position on B, for other simulation runs, we may find an equally strong positive correlation. Therefore, we conducted 10 experiments and recorded the correlation over time (Figure 4).

Figure 4 indeed shows that the correlations between A and B are much more prominent than in the condition of only central processing. Moreover, it can also be observed that this correlation may be positive or negative. It can also be observed that the correlations are not stable over time, indicating that attitude dynamics are continuous. In Figure 3, this can be seen at t = 10,000,000, where a number of agents have an extreme negative position on both A and B (left bottom corner of the relation graph), thus indicating a positive correlation between both dimensions for these agents, which originally was negative. This is being explained by the

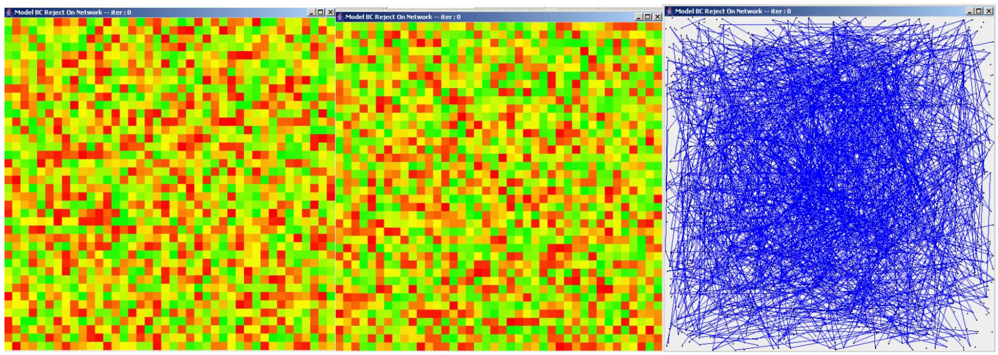

Time-step 0 Attitude A Attitude B relation A and B

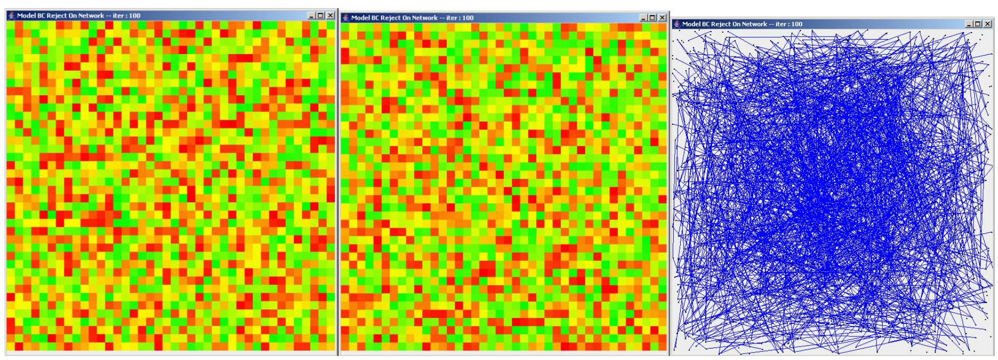

Time-step 100 Attitude A Attitude B relation A and B

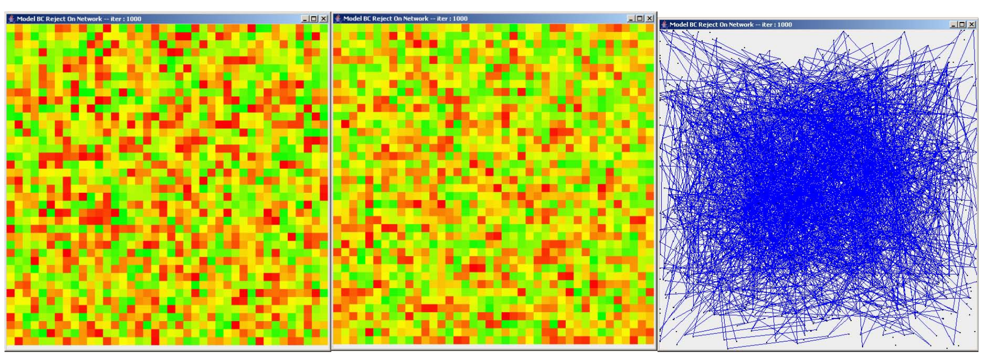

Time-step 1000 Attitude A Attitude B relation A and B

FIGURE 3 Attitude position on A and B over time

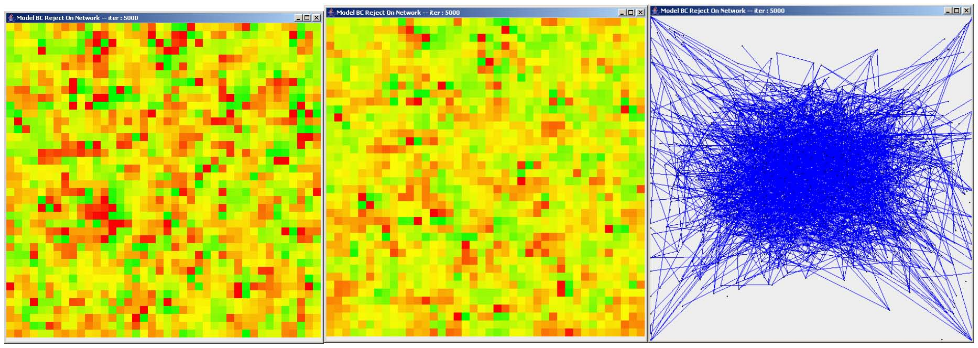

Time-step 5,000 Attitude A Attitude B relation A and B

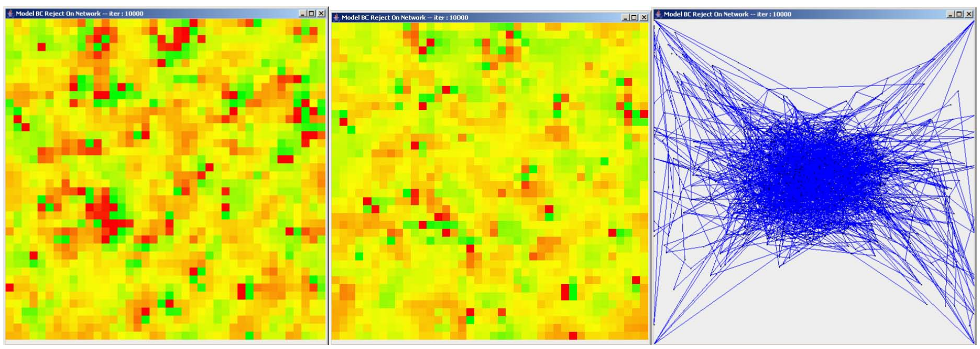

Time-step 10,000 Attitude A Attitude B relation A and B

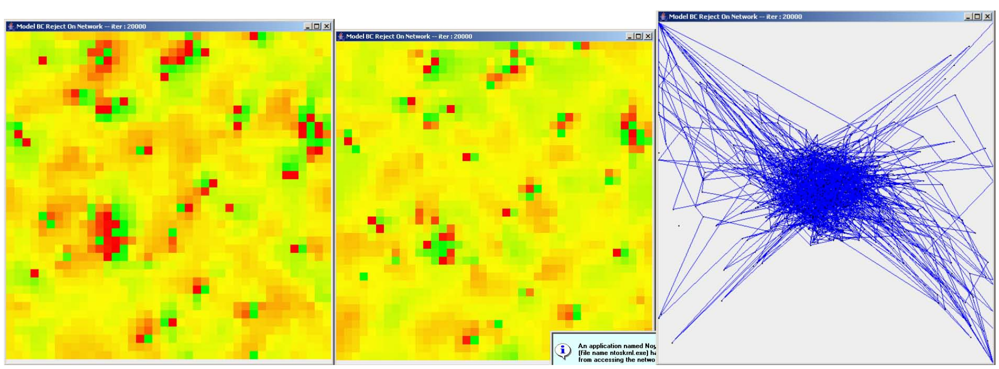

Time-step 20,000 Attitude A Attitude B relation A and B

FIGURE 3 Cont.

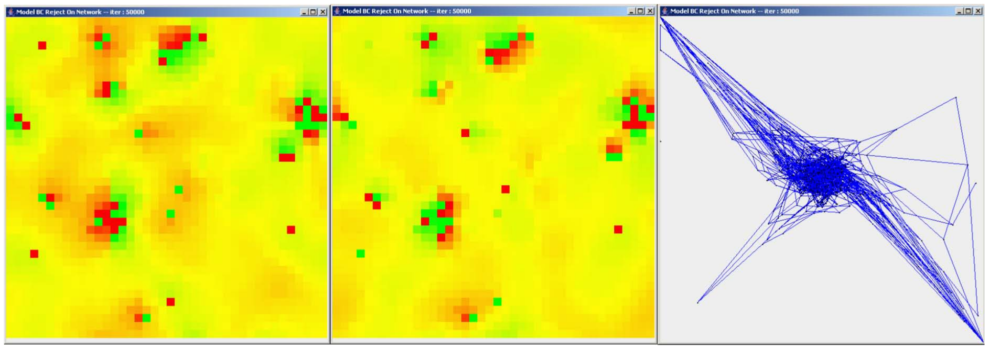

Time-step 50,000 Attitude A Attitude B relation A and B

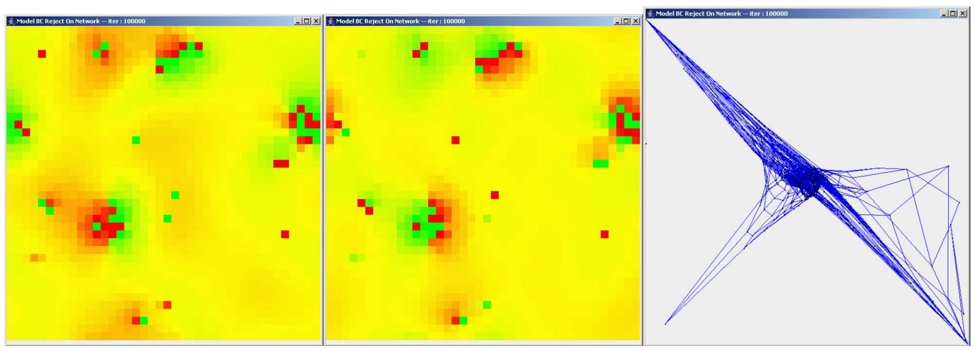

Time-step 100,000 Attitude A Attitude B relation A and B

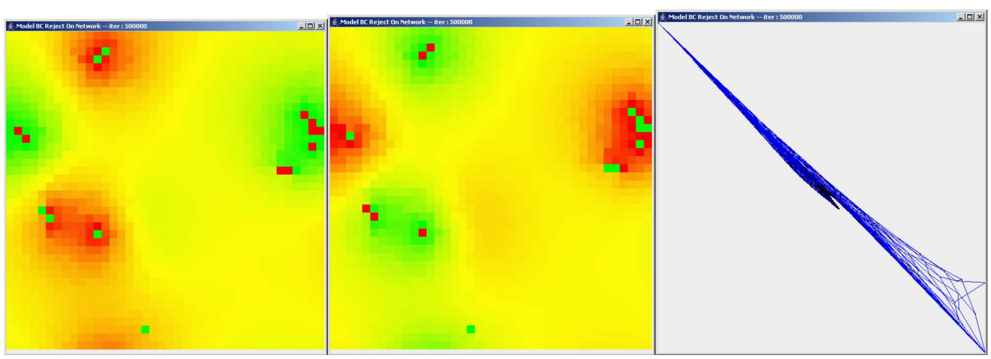

Time-step 500,000 Attitude A Attitude B relation A and B

FIGURE 3 Cont.

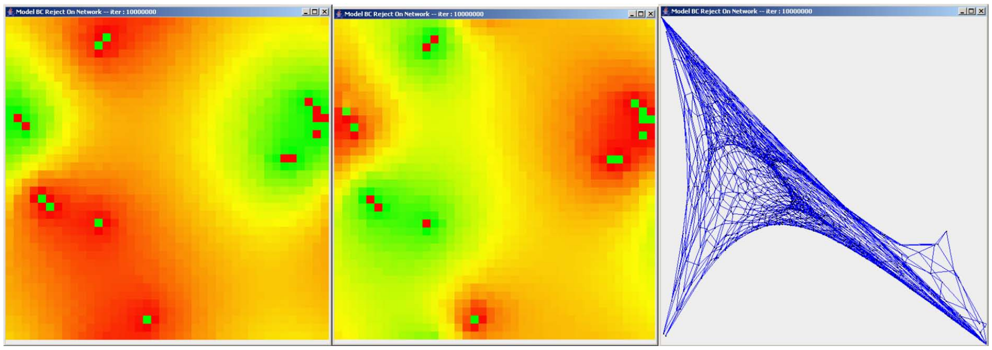

Time-step 10,000,000 Attitude A Attitude B relation A and B

FIGURE 3 Cont.

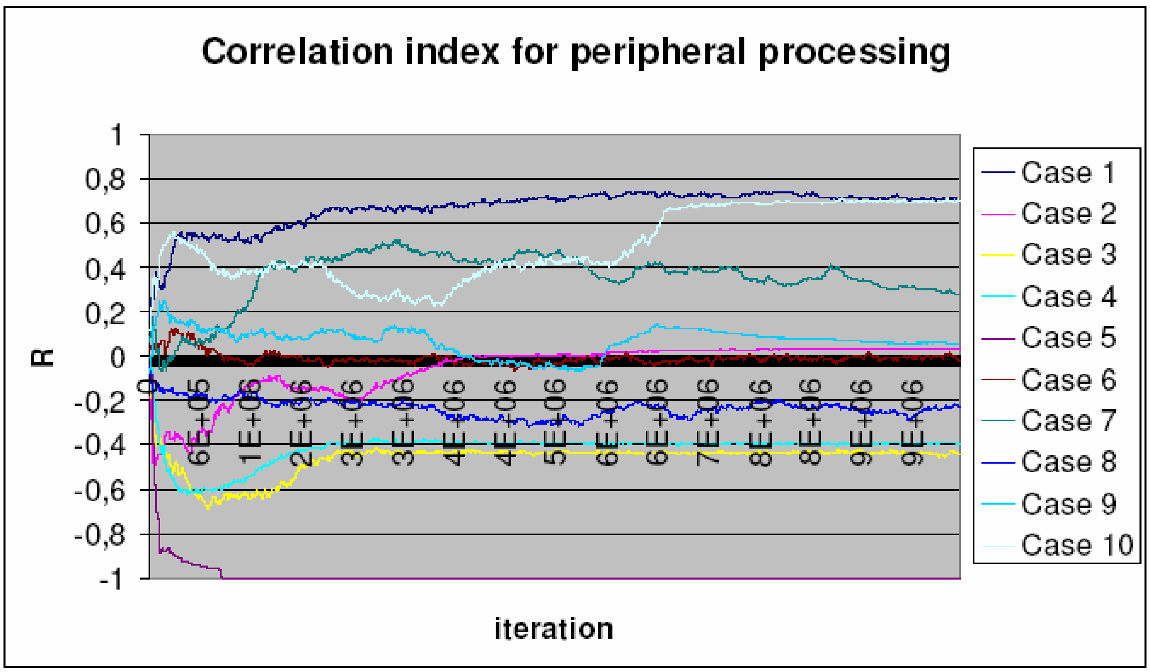

FIGURE 4 Correlation between A and B over time for 10 simulation runs

contrast effect, as elicited by the green position on dimension A of the single (green) agent located at the center bottom. It can be observed that agents in the neighborhood of this agent respond with reactance — in this case turning to red. Because this reactance effect translated to dimension B according to the peripheral processing, we also observe this reactance effect on dimension B, where the neighboring agents also turn to red. These results indicate that the dynamics on the second attitude are quite unstable, as singularities (like the green dot) tend to get amplified depending on the dynamics on the first attitude. In other words, the dynamics on the first attitude control the dynamics on the second one, but in a different context. This may lead to situations where in one region, the correlation between A and B is positive, whereas in another region, this correlation is negative. Agents that are located in a transitional zone between these two contrasting situations are thus experiencing instability concerning the direction of the peripheral processing on attitude B, and thus may move hence and forth on this dimension.

Experiment 3: Influence of a Meta-actor

In the previous experiments, the agents only interacted with their direct neighbors. However, often politicians or other spokesmen have a large audience they address on a frequent basis. Hence, before elections or votes, people not only discuss issues with their local peers but are also influenced by what we call “meta-actors.” In the model, we formalize a meta-actor as an agent having a fixed position; hence, it is not susceptible to influences of the opinion of others. In selecting an interaction partner, each agent randomly contacts either one of the four neighbors or the meta-actor. Hence the meta-actor has a chance of 20% of being contacted every time-step. In the experiments, the agents process centrally on attitude A, and peripherally on attitude B, thus replicating the conditions of experiment 2.

Conditions for Experiment 3

For the meta-actor, we formalize an extreme position (−1, or red) on dimension A (central processing) and a neutral position (0 or yellow) on dimension B (peripheral processing). We use different settings for the agents. In experiment 3A, the population is rather accepting by setting U at 1.5 and T at 1.7. In experiment 3B, the population is less accepting by setting U at 1.0 and T at 1.2. Furthermore, the population is set at 10,000 agents. Concerning the interaction structure, we connect the meta-actor to all agents in the population. Each individual agent is now connected with five agents: North, South, East, West, and Meta-Actor.

Results of Experiment 3a, An Accepting Population

Figure 5 presents the developments on both attitudes for different time-steps of the simulation. The figures represent the position on attitude dimension A (left) and attitude dimension B (right). The black dot in the middle represents the meta-actor.

These results show that under conditions of an acceptable population, the vast majority of the population accepts the attitude position of the meta-actor. Only a few agents contrast with the meta-actor on attitude A (the green dots), and because their neighbors contrast themselves with these particular agents on dimension A, they also contrast on dimension B, resulting in the more red position of the neighbors on dimension B.

Results of Experiment 3b, A Less-accepting Population

Figure 6 presents the developments on both attitudes for different time-steps of the simulation. The population is less accepting by setting U at 1.0 and T at 1.2. The figures represent the position on attitude dimension A (left) and attitude dimension B (right). The black dot in the middle represents the meta-actor.

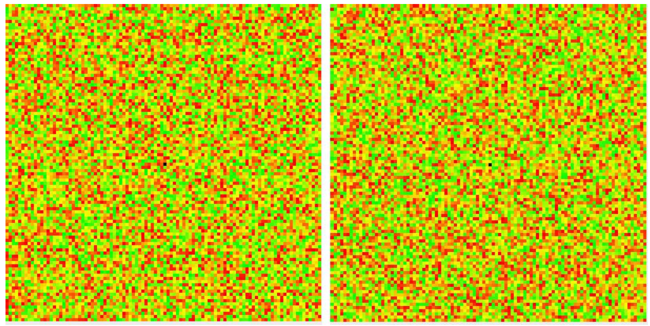

Time-step 0 Attitude A Attitude B

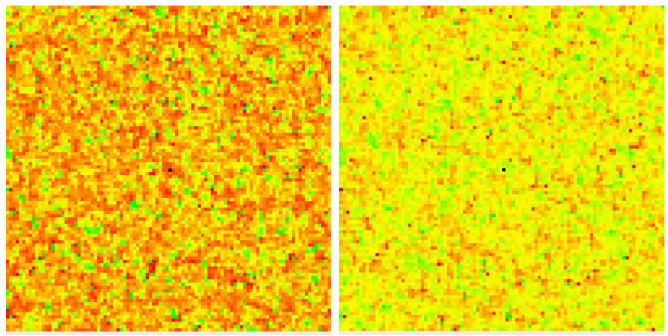

Time-step 50,000 Attitude A Attitude B

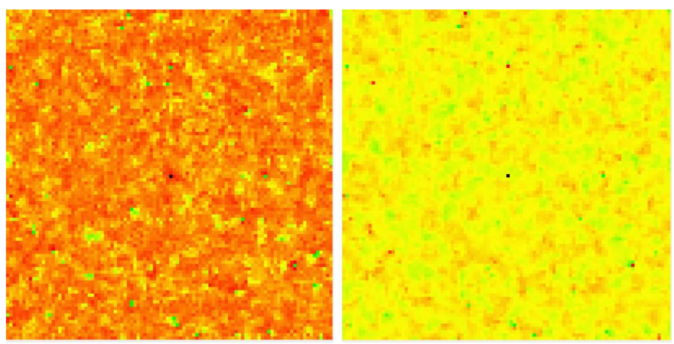

Time-step 100,000 Attitude A Attitude B

FIGURE 5 Experiment 3a developments on both attitudes for different time-steps

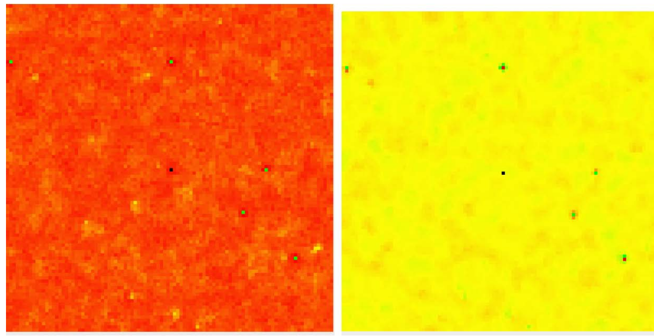

Time-step 200,000 Attitude A Attitude B

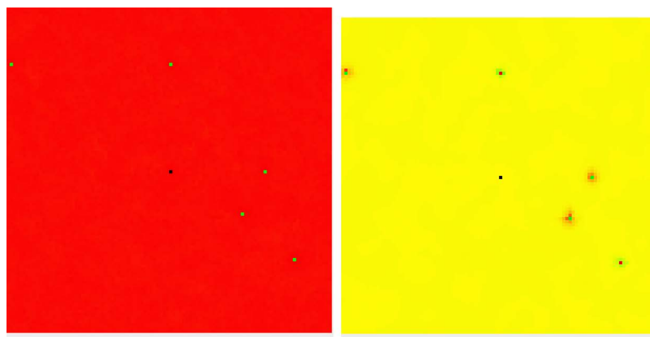

Time-step 500,000 Attitude A Attitude B

FIGURE 5 Cont.

It can be observed that when the attitudes on dimension A polarize, agents either become red or green. Apparently, the reds are having a slight majority because of the systematic influence of the meta-actor. Concerning dimension B, we see heterogeneity. This is due to the fact that when agents contact the meta-actor and assimilate his position, they also assimilate the meta-actor's position on B. Close observation indeed reveals that the agents contrasting with the meta-actor on A (the green ones) also have an extreme position on attitude B, whereas for many actors, being red on A holds that they are yellow on B, showing the systematic effect of the meta-actor. A particular case concerns those agents having a red position on both A and B. Interacting on dimension A with a green agent results in a contrast effect on both A and B, thus also stimulating a red position on B. However, interacting with the meta-actor results in an assimilation effect, which draws them to the yellow position on attitude B. The dynamics are then stable on both attitudes for the opponents of A (the greens) but rather unstable for the followers of the meta-actor on A (the reds), resulting in alternating positions between red and yellow on attitude B. Hence the meta-actor succeeds only in drawing people to his position on B for the agents that agree with him on A.

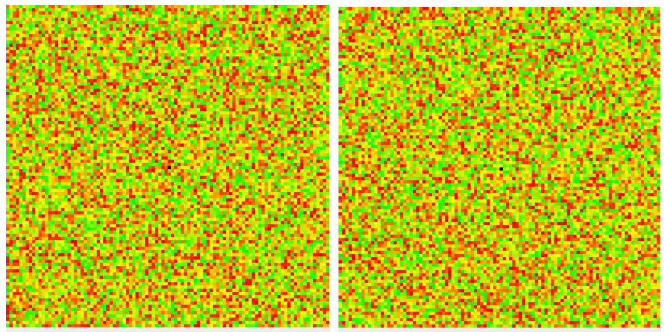

Time-step 0 Attitude A Attitude B

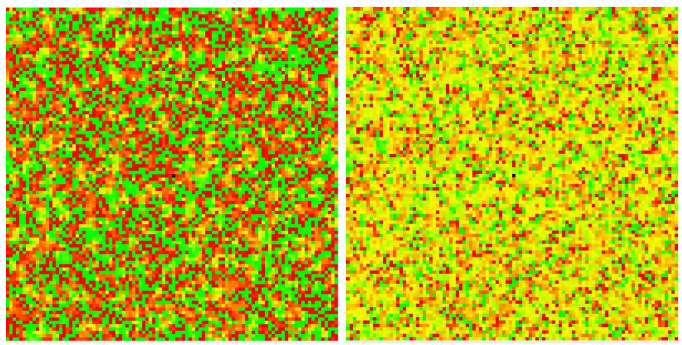

Time-step 50,000 Attitude A Attitude B

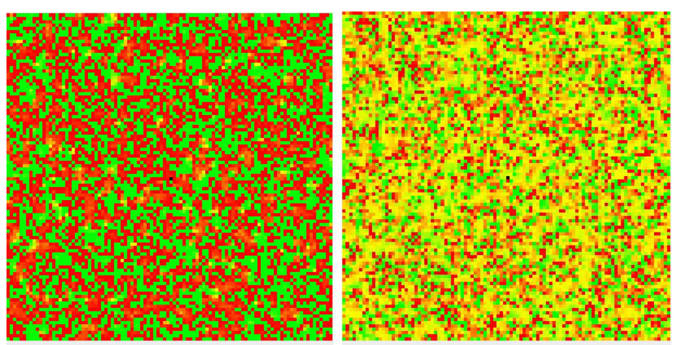

Time-step 100,000 Attitude A Attitude B

FIGURE 6 Experiment 3b developments on both attitudes for different time-steps

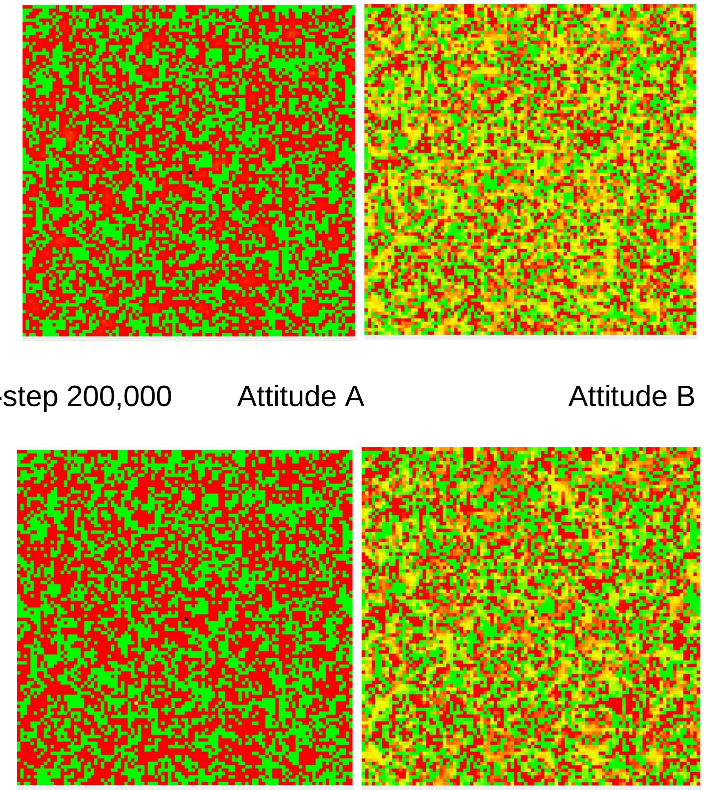

Time-step 200,000 Attitude A Attitude B

Time-step 500,000 Attitude A Attitude B

FIGURE 6 Cont.

GENERAL DISCUSSION

Although an increasing number of scientists study attitude or opinion dynamics by using multi-agent models, up until now, there has hardly been any attention on multiple attitude dynamics. Both in experimental laboratory settings and in simulation studies, researchers have focused on single attitudes/opinions. Yet observations from the field indicate that many people use a position on one attitude as a determinant for selecting a position on another, often unrelated, dimension. These effects may pertain to simple consumer preferences, where people may have a tendency to have the same preference for a variety of unrelated consumer goods, thus generating subcultures where people have about the same preferences on basically unrelated issues. Especially when people have to select a position on an issue that is complex and/or less personally important, they may engage in simple processing, taking the behavior of their peers to select a position. In the experiments as presented in this paper, it can be observed that such decision strategies — here formalized as peripheral processing — have major impacts on the attitude dynamics. Basically, we observe that peripheral processing is often responsible for the emergence of a correlation between originally unrelated issues. Hence the assimilation or rejectance of other people's attitudes on the basis of a perceived (dis)agreement on another, more important issue causes attitudes on different issues to become correlated. Because people are interacting with other people on a multitude of issues, it is expected that this relatedness of attitude dynamics may be important in understanding why certain clusters of people having the same opinion on various issues emerge, and how these clusters change over time (as formalized in the culture dynamics model of Axelrod, with discrete tags on each dimension).

In addition, the first experiments with the meta-actor demonstrated that under conditions of high acceptability of the population for deviant opinions, the meta-actor was capable of attracting virtually all agents in the populations to its own position on both attitude dimensions. The situation changed, however, when the population was less accepting. Here we observed that a polarization emerged on the dimension on which agents processed centrally, whereas heterogeneity emerged on the dimension where agents processed peripherally. These results differ from situations where no meta-actor was available, showing that such an actor may have a considerable impact on the attitude/opinion dynamics that emerge.

These first experiments reveal the importance of including several attitude/opinions simultaneously in understanding these dynamics and the effect a meta-actor has on these dynamics. Many experiments have to be conducted to get a better understanding of these multi-attitude dynamics and the critical factors that determine clustering effects. Some issues that remain to be studied are (1) the differences and heterogeneity between agents with regard to their tendency to assimilate, contrast, and firmness of opinions; (2) the effects of the connectivity between agents (social network effects); and (3) strategies that can be employed by meta-actors in affecting these dynamics.

REFERENCES

Deffuant, G., D. Neau, F. Amblard, and G. Weisbuch, 2001, “Mixing beliefs among interacting agents,” *Advances in Complex Systems* 3, 87–98.

Deffuant, G., F. Amblard, G. Weisbuch, and T. Faure, 2002, “How can extremism prevail? A study based on the relative agreement interaction model,” *Journal of Artificial Societies and Social Simulation* 5(4); available at http://jasss.soc.surrey.ac.uk/5/4/1.html.

Festinger, L., 1954, “A theory of social comparison processes,” *Human Relations* 7, 117–140.

Galam, S., 1999, “Application of statistical physics to politics,” *Physica A* 274, 132–139.

Hegselmann, R., and U. Krause, 2002, “Opinion dynamics and bounded confidence models, analysis and simulation,” *Journal of Artificial Societies and Social Simulation* 5(3); available at http://jasss.soc.surrey.ac.uk/5/3/2.html.

Jager, W., and F. Amblard, “Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change,” *Computational & Mathematical Organization Theory* 10(4), 295–303(9).

Latane, B., and A. Nowak, 1997, “Self-organizing social systems: Necessary and sufficient conditions for the emergence of clustering, consolidation, and continuing diversity,” in G.A. Barnett and F.J. Boster (Eds.), *Progress in Communication Sciences,* Ablex Publishing Corporation.

O’Keefe, D.J., 1990, *Persuasion Theory and Research*, Sage Publishing, Newbury Park, CA.

Petty, R.E., and J.T. Cacioppo, 1986, *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*, Springer Verlag, NY.

Salzarulo, L., 2004, "Formalization of self-categorization theory to simulate the formation of social groups," paper presented at the Second European Social Simulation Conference, September 16–19, Vallodalid, Spain.

Sherif, M., and C.I. Hovland, 1961, *Social Judgment*, Yale University Press, New Haven, CT.

CONTESTING HEARTS AND MINDS: A BASELINE MODEL OF SOLIDARITY DYNAMICS IN MILITARY OCCUPATIONS

K. RUBY, The University of Chicago, Chicago, IL
V. MELLARKOD, Texas Tech University, Lubbock, TX, and Argonne National Laboratory, Argonne, IL
D. SALLACH,* Argonne National Laboratory, Argonne, IL, and The University of Chicago, Chicago, IL
C. MACAL, Argonne National Laboratory, Argonne, IL

ABSTRACT

In this paper, we discuss an Occupation Dynamics Model and provide an initial verification. We describe the parts of the model — the mechanisms and parameters that translate authority strategies into agent outcomes. The baseline model tests reveal that strategic interaction represents a key element in determining authority outcomes; even a materially weaker authority could gain the advantage in specific instances of strategic interactions. A series of tests designed to evaluate mechanisms and their impact on the model is then presented.

Keywords: Multi-mechanism model, strategic interaction, sanction effects, preference falsification, model verification

INTRODUCTION

As the U.S. experiences in Vietnam and Iraq attest, wartime and post-war occupations can generate counterproductive conflict dynamics with local resistance movements. The conventional wisdom is that “winning hearts and minds” is the key to the long-term success of an occupation, with both occupation forces and resistance movements employing a mix of coercive and noncoercive strategies (punishments and rewards) in a competitive effort to shape support within the occupied public (Edelstein 2004).

However, the relationship between material incentives and hearts and minds is complicated by at least two factors. First, individuals do not make decisions on the basis of purely economic considerations. Emotional, social, and psychological factors mediate the impacts of sanctions and play a significant role in influencing individual decisions. Second, individuals may not be representing their true feelings in public (Kuran 1991, 1995), suggesting that sanction strategies may be affecting hearts differently than minds, with unclear effects on the degree and stability of public support. The possibility of preference falsification means that the appearance of widespread public support may rest on fragile foundations vulnerable to tipping. In short, finding the optimal strategy for winning hearts *and* minds in order to build robust public support is complicated by competing claims to authority and cross-cutting incentives and pressures, often leading to unintended or surprising outcomes.

---

* *Corresponding author address*: David L. Sallach, Decision and Information Sciences Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439; e-mail: sallach@anl.gov.

This paper contributes to a growing interest in the computational modeling of political conflict, rebellion, and insurgency (Epstein 2002; Lustick et al. 2004; Cederman 2004), complementing a well established literature that uses formal models to understand and explain conflict and order dynamics (e.g., Tilly 1985; Grossman 1991; Olson 1993; Skaperdas 2001; Tákacs 2001). We present a baseline model of occupation dynamics in military occupation settings that uses agent-based simulation to represent the complex interaction of occupation and resistance strategies with private emotions (valence) and public support (alignment) of individuals in a simulated occupied public.[1] The model builds on previous models of social order and preference falsification to investigate the tension between hearts and minds and the effects of preference falsification on social order. Agent-based modeling represents a useful methodology for investigating complex social systems and processes because it effectively captures interactions between competing actors with conflicting incentives in a dynamic social environment characterized by nonlinearity. By creating, in effect, an artificial society in which the operation of parameters and mechanisms can be explicitly controlled and monitored, agent-based modeling affords the researcher the opportunity to test assumptions and explore how outcomes are arrived at (Cederman 1997, 2001; Pepinsky 2005). The paper presents a thorough analysis of model parameters and mechanisms. The intent of the model testing component of the paper is to contribute to an ongoing discussion on methods and principles of verification for agent-based simulation in the social sciences.

The paper proceeds as follows. In the next section, we introduce the model and its key variables, agents, and mechanisms and how they interact. Then in the third and fourth sections, we vary key parameters and mechanisms to investigate their effects on how the model operates. In the fifth section, we report preliminary model results. In the sixth section, we discuss the benefits of rigorous testing and experimentation with respect to the Occupation Dynamics Model and agent-based modeling more generally, and we present future directions.

THE OCCUPATION DYNAMICS MODEL

The baseline Occupation Dynamics Model is constructed in the J programming language[2] (Thomson 2001) by using relatively simple assumptions. As its name suggests, the baseline model is intended to define a basic frame of reference relative to which more complex models can be assessed. The model includes two basic actors: (1) authorities (occupation authority [OA] and counter authority [CA]), who apply rewards and punishments (sanctions) to individuals in order to shape public support, and (2) individuals (agents), who choose their level of public support relative to material, emotional, social, and psychological variables (Figure 1).

This section continues as follows. First, we specify authorities and the strategies by which they compete to shape public support. Next, we present the agents who make up the individuals in our notional occupied public, followed by the four mechanisms that translate the experience of sanctions on agents into degree and direction of public support for authorities.

---

1 It should be pointed out that the problem of occupation presented in this paper is a specific case of the more general problem of creating and maintaining order in political systems, suggesting that the model presented here may be applicable to a wider range of political phenomena.

2 J can be downloaded and used for free at www.jsoftware.com.

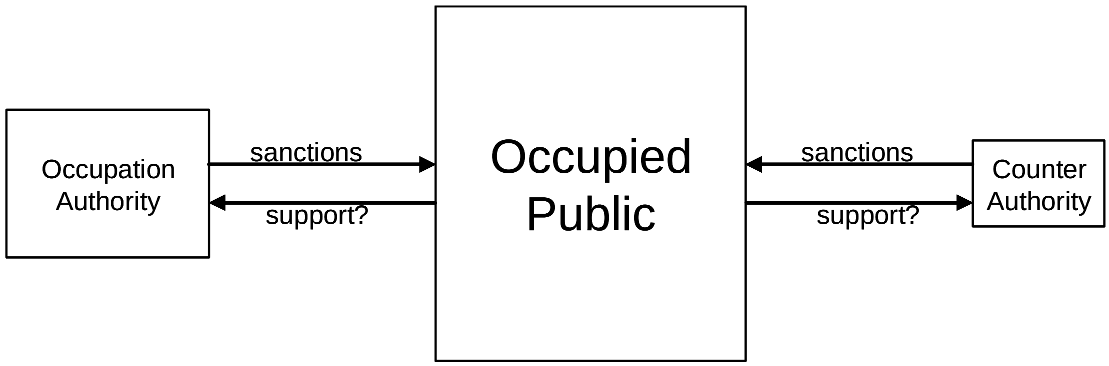

FIGURE 1 Model overview diagram

After specifying model flow, we conclude with some thoughts on how the model can be used as a research platform.

Occupation and Counter Authorities

The baseline model presumes the existence of two authorities: (1) an OA and (2) a CA who represents an organized resistance group. Both the OA and CA select strategies by which members of the population are rewarded or punished. Rewards may be thought of as various material benefits, such as contracts, jobs, and schools. Punishments may be thought of as detention, injury, or the destruction of resources. What typically distinguishes an OA from a CA in the baseline model is the availability of resources that can be applied to sanctions: The OA starts with a greater resource endowment than the CA, building asymmetry into the model. However, the model allows for the resource balance to change. While the OA receives a fixed budget from an exogenous "home government," the CA's operating budget depends on indigenous support and therefore increases or decreases in proportion to its success.

Both the OA and CA choose a strategy composed of two components. First, each authority selects the location of a breakpoint between agents receiving rewards and punishment (the authority's *sanction strategy*) on the basis of individuals' publicly expressed support (alignment), which ranges from –1 (support for CA) to 1 (support for OA). While, in principle, the breakpoint can fall anywhere on the alignment continuum, the baseline model limits the breakpoint to three settings (–0.5, 0, 0.5) to simplify analysis. Once the breakpoint has been set, the second aspect of strategy selection concerns the distribution of resources for carrying out sanctions (authority's *resource strategy*). Thus the OA and CA must choose where along the alignment spectrum to concentrate the impact of sanctions. Authorities choose among the following resource allocation policies: (1) *focus on neutrals* (i.e., concentrate resources on agents in the center of the alignment spectrum); (2) *focus on friends/enemies* (i.e., concentrate resources to reward the strongest supporters and punish most vocal opponents); and (3) *focus evenly* (i.e., uniformly distribute resources across the alignment spectrum). *Focus on neutrals* may be thought of as trying to win over the middle by focusing rewards and punishments on the uncommitted, while the *focus on friends/enemies* resource allocation option may be thought of as rewarding an authority's closest supporters and punishing its strongest adversaries. The nine resource/sanction strategy combinations are depicted graphically in Figure 2, with colors denoting the breakpoint between reward and punish.

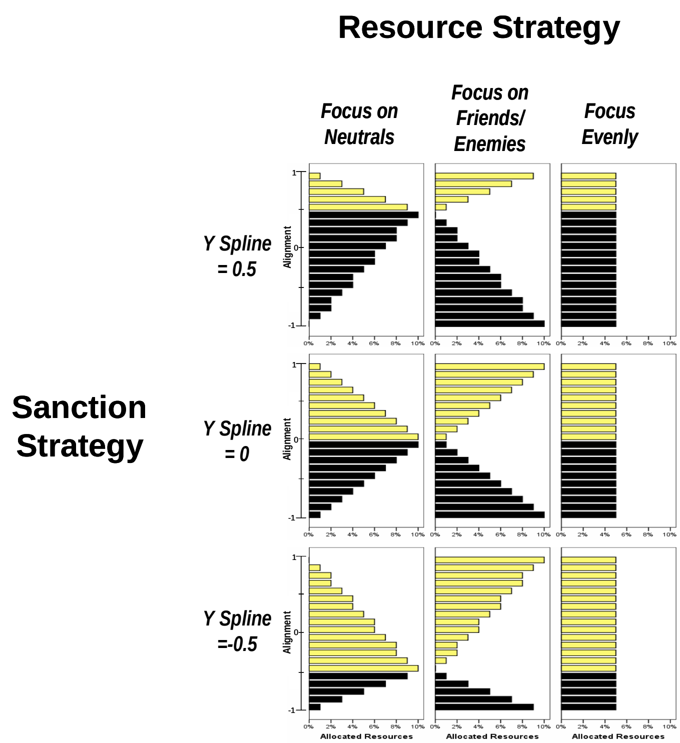

FIGURE 2 Authority strategy combinations

The nine possible sanction/resource combinations then combine for a total of 81 strategic interaction possibilities between the OA and CA (Table 1).

TABLE 1 Resource/sanction strategy interactions

OA Resource Strategy	Sanction Strategy	CA Resource Strategy								
		Neutrals			Friends/Enemies			Even		
		0.5	0	–0.5	0.5	0	–0.5	0.5	0	–0.5
Neutrals	0.5	1	2	3	4	5	6	7	8	9
	0	10	11	12	13	14	15	16	17	18
	–0.5	19	20	21	22	23	24	25	26	27
Friends/ enemies	0.5	28	29	30	31	32	33	34	35	36
	0	37	38	39	40	41	42	43	44	45
	–0.5	46	47	48	49	50	51	52	53	54
Even	0.5	55	56	57	58	59	60	61	62	63
	0	64	65	66	67	68	69	70	71	72
	–0.5	73	74	75	76	77	78	79	80	81

Each authority is allocated resources with which to administer sanctions, distributed among the occupied public on the basis of the selected resource strategy. For each resource point available to an authority, that authority is able to sanction one agent. Authorities have two sources for resources. First, authorities can be assigned a renewable but fixed endowment of resources supplied at every tick. Second, authorities can receive resources from supporters. When the public support option is activated for an authority, that authority receives a set number of resource points for every agent showing high levels of public support for that authority (e.g., alignment greater than 0.5 for OA and less than –0.5 for CA). The size of the renewable fixed endowment, the alignment cutoff for receiving public support, and the number of resource points for each supporter are parameters that can be set at initialization.

Individuals in the Occupied Public

How agents respond to sanctions is determined by the interaction of emotional, economic, social, and psychological factors and authority sanctions. Agents have two basic properties: valence and alignment. *Valence* is meant to capture an emotional state and ranges from extreme dislike (–1) to extreme like (1). Valence is distinguished from *alignment* in that alignment is a public position of support discernable by neighbors and authorities, whereas valence is an internal response that may not be fully represented by the agent's public position and thus can only be inferred. Each agent has valence and alignment toward the OA and CA. As is valence, alignment is calculated along a single scale, with strong support of OA equaling 1 and strong support of CA equaling –1. The interplay of alignment and valence includes the ambivalence that sometimes occurs in public discourse and provides a consideration outside of purely material calculations.

In addition to responses to direct sanctioning by the authorities, agents also respond to the impact of sanctions on the evolving social environment. Agents are embedded in a social network representing family, friends, and/or associates and have a valence toward each individual network neighbor (NN). The network is generated at initialization by randomly selecting an agent and assigning that agent up to the specified maximum number of randomly selected neighbors and then removing the new neighbor from the list of available neighbors (selection without replacement). Network links are reciprocal. In the baseline model, NNs are fixed throughout the simulation, as are valences toward each individual neighbor. In future iterations, networks will be generated and updated dynamically in response to agent-agent and agent-structure interactions. The range of possible agent parameters is summarized in Table 2.

TABLE 2 Agent (occupied public) parameters

Parameter	Maximum Range
Valence toward authority	From –1 to 1
Alignment toward authority	From –1 to 1
Network neighbors	Limited by population size
Valence toward neighbor	From –1 to 1
Resources	No limits

Mechanisms

Four computational mechanisms govern agent response to authority sanctions: (1) emotional response to sanction (valence mechanism [VM]), (2) expected cost of future sanction (cost-benefit mechanism [CBM]); (3) social conformity (imitation mechanism [IM]); and (4) dissonance constraint (dissonance reduction mechanism [DRM]).

Valence Mechanism

An emotional reaction to authority sanctions makes up the first mechanism of population response, with VM intended to capture the interaction between sanctions and local interests. Agents adjust how they feel about an authority on the basis of whether that authority rewarded or punished them. In addition to the direct impact of sanctions on valence toward an authority, agent valence is influenced by the sanctions visited upon members of the agent's social network (NNs). If an authority rewards an NN j that the agent likes (i.e., valence toward the NN, or v_{ij}, is ≥ 0) or punishes an NN that the agent dislikes (v_{ij} is <0), the agent adjusts its valence toward the authority in a positive direction. Conversely, if an authority rewards an NN j that the agent dislikes, or punishes an NN j that the agent likes, the agent adjusts the valence toward that authority in a negative direction. In relation to the agent's social network, then, VM follows the general logic contained in the statement, "the enemy of my enemy is my friend." VM adjusts valence toward authority (v_{ia}) as follows. For each NN j in an agent's network,

If j is a friend (v_{ij} is ≥ 0), and j was

(1) rewarded: $\Delta\, v_{ia} = I * v_{ia} + v_{ij} * (1 - v_{ia})$,
(2) punished: $\Delta\, v_{ia} = I * v_{ia} + v_{ia} - v_{ij} * (1 + v_{ia})$.

If j is an enemy (v_{ij} is <0), and j was

(3) rewarded: $\Delta\, v_{ia} = I * v_{ia} + v_{ij} * (1 + v_{ia})$,
(4) punished: $\Delta\, v_{ia} = I * v_{ia} + v_{ia} - v_{ij} * (1 - v_{ia})$.

In the equations above, I serves as an indicator function representing the authority responsible for sanctioning. For the OA, I is set to 1; for the CA, $I = -1$. VM calculates the new valence toward authority v_{ia}' by summing all adjustments to valence resulting from sanctions to NNs ($\Delta\, v_{ia}$) and self (with valence to self currently set to 0.8), then dividing by the total number of NNs + 1. In general terms, then, the results of the VM can be expressed as:

$$v_{ia}' = v_{ia} + \frac{\sum(\Delta v_{ia})}{\sum(j+1)} \, .$$

Cost-Benefit Mechanism

CBM represents an economic calculation to evaluate expected future costs and benefits of maintaining or changing a given level of public support (alignment). The agent looks at those whose alignment levels are similar and how they have fared with regard to punishments and rewards, then decides whether to increase, decrease, or remain unchanged in its alignment to achieve the best economic outcome. Specifically, an agent determines the average level of reward and punishment in three bands relative to the agent's own location on the alignment continuum and chooses to move toward a band on the basis of the expected benefits of moving. The size of the three bands is set at initialization by the parameter *alignment band percent* (*AB%*), representing a fixed percentage of the total alignment space. For example, at *AB%* = 5, the alignment continuum of –1 to 1 is divided into 20 bands, each with a width of 0.1. Given *AB%* = 5, an agent with alignment a_{ia} = 0.13 surveys the average degree of sanction for his own alignment band (a_{ia} is 0.1 to 0.2), the band below (a_{ia} is 0.0 to 0.1), and the band above (a_{ia} is 0.2 to 0.3). On the basis of the agent's calculation of the most beneficial band, determined by taking the average of all punishments and rewards for each of the three bands, the agent shifts his alignment in the direction of relative safety (which may mean not moving at all if the present alignment band is deemed the safest of the three). If CBM indicates a shift in alignment is warranted, v_{ia} is changed in the indicated direction, with the degree of change a random distance chosen from a uniform distribution between a_{ia} and ±*AB%* of a_{ia}. In the case of our agent *i* at a_{ia} = 0.13 and AB% = 5, if the upper band is most desirable, a_{ia} adjusts randomly to an alignment between 0.13 (starting position) and 0.1365 (0.13 + 5% of 0.13).

The Imitation Mechanism

The agent's social network provides a secondary basis of comparison that influences the degree to which it publicly supports an authority. The Imitation Mechanism is meant to capture social conformity pressures, adjusting agents' alignment decisions on the basis of the alignment decisions of members of their social reference group (agents' NNs). By basing an agent's alignment decision on not only the material costs (as calculated by the CBM) but also the social costs of assuming a given alignment, IM generates a process by which agents' alignment choices are imitated, propagated, and/or disseminated through a network.

Agents prefer to maximize similarity with friends and minimize similarity with enemies. Thus there are two factors that influence the operation of the imitation mechanisms: (1) valence toward neighbors (a_{ij}), and (2) absolute distance in alignment space to each neighbor (da_{ij}), calculated as $|a_{ja} - a_{ia}|$. With regard to friends (v_{ij} is ≥0), the more they are esteemed (i.e., the higher that v_{ij} is), the stronger is the attraction/pull, the greater is the distance in alignment (da_{ij}), the stronger is the need to imitate, and, therefore, the greater is the change in alignment with respect to that neighbor. With regard to enemies (v_{ij} is <0), the more they are disliked (i.e., the lower that v_{ij} is), the more the agent wishes to demonstrate dissimilarity in alignment, with close alignment proximity (av_{ij}) generating a stronger repelling effect. The function governing the IM and its affect on alignment ($\Delta\, a_{ia}$) is presented here:

$$\Delta\, a_{ia} = \sum \frac{\{IF * sda_{ij} * v_{ij} * [(da_{if} * I) + (2 - da_{ij}) * (I - I)]\}}{2 * (\sum j)} \ .$$

In the formula above, *IF* (imitation factor) moderates the degree of alignment change (set to 0.5 in the model) and sda_{ij} is the sign (positive or negative) of the difference in alignments, with sda_{ij} equal to 1 when da_{ij} is ≥0, and equal to –1 otherwise. *I* is an indicator function denoting whether a given NN *j* is a friend (for v_{ij} of ≥0, *I* = 1) or an enemy (for v_{ij} of <0, *I* = 0). Figure 3 illustrates the change in alignment relative to the alignment distance between a single agent *i* and a single NN *j* alternatively cast as a friend (v_{ij} = 1) and an enemy (v_{ij} = –1).

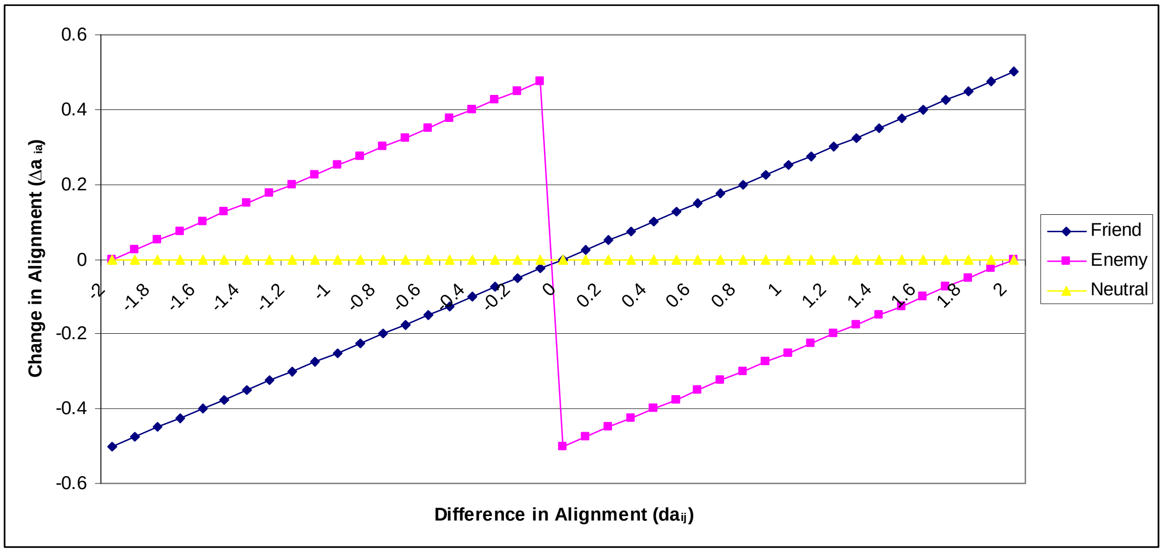

FIGURE 3 Effects of imitation mechanism on alignment by alignment difference and neighbor status (friend, enemy or neutral) for a sample agent

Dissonance Reduction Mechanism

The final mechanism imposes a constraint on the extent to which an agent tolerates the dissonance resulting from preference falsification. Dissonance refers to psychological stress experienced by agents whose public alignment diverges from private valence. Both theory and historical patterns suggest that there are limits on the extent to which public positions and private feelings can come into conflict.

Cognitive dissonance theory (Festinger 1962; Kuran 1998; Epstein 2002) argues that, ceteris paribus, individuals prefer to act publicly (alignment) as they feel privately (valence). Because valence and alignment are influenced by competing and often contradictory mechanisms, it is highly likely that discrepancies between valence and alignment will emerge. Preference falsification occurs when an agent prefers to publicly misrepresent its private preferences. Because cognitive dissonance is associated with psychological discomfort, agents are motivated to reduce it by adapting beliefs to behavior or vice versa.

The Dissonance Reduction Mechanism reduces the distance between alignment and valence by adjusting alignment to valence or vice versa, depending on the actor's location on the alignment spectrum. DRM adapts valence to alignment when alignment is greater than zero, with positive public alignment with authority reorienting privately held emotions. When alignment is less than zero (i.e., the agent shows support for the CA), alignment is adjusted to valence because deeply held negative emotions may ultimately restrict the public positions that an agent can comfortably take. Figure 4 illustrates the axes and directions of DRM's conditional effects on reducing discrepancies between valence and alignment (with all points on the diagonal line where $v_{ia} = a_{ia}$ representing the preferred positions of agents).

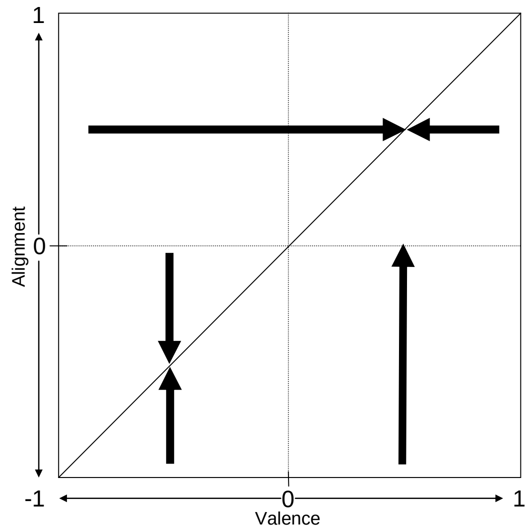

FIGURE 4 Axes and directions of valence and alignment change, dissonance reduction mechanism

DRM is currently modeled as a quadratic function where the degree of adjustment is proportional to the square of the perpendicular distance between the agent's current location in alignment and valence coordinates (v_{ia}, a_{ia}) and the point on the diagonal where alignment and valence are equal ($v_{ia} – a_{ia} = dva$). If the discrepancy between valence and alignment is low, the adjustment is minimal; increasing the level of dissonance results in stronger corrections. A preference constraint factor (*PCF*), set at initialization, further moderates the degree of adjustment due to dissonance. Adjustments are made according to the following formulas:

$$\text{if } a_{ia} \text{ is } \geq 0,\ \Delta\, v_{ia} = sdva * I * (PCF/2) * (dva)^2, \text{ and}$$

$$\text{if } a_{ia} \text{ is } <0,\ \Delta\, a_{ia} = sdva * I * (PCF/2) * (dva)^2.$$

In the formulas above, *I* is an indicator function equal to 1 when a_{ia} is ≥0 and –1 otherwise. The mathematical sign from taking the difference between valence and alignment (*sdva*) sets the direction of adjustment on the relevant axis and is set to 1 when *dva* is ≤0 and –1 otherwise.

Taken together, the four mechanisms represent a cluster of salient response patterns found in history and society. In a situated setting, rational calculation, social influence, emotional reaction, and psychic strain interact to produce a complex, nonlinear response to authority sanction policies.

Model Flow

The order of operations for the Occupation Dynamics Model is presented in the following flow diagram (Figure 5). The order of execution of authority strategies is random, as is agent activation. Mechanisms, however, are activated in the order indicated.

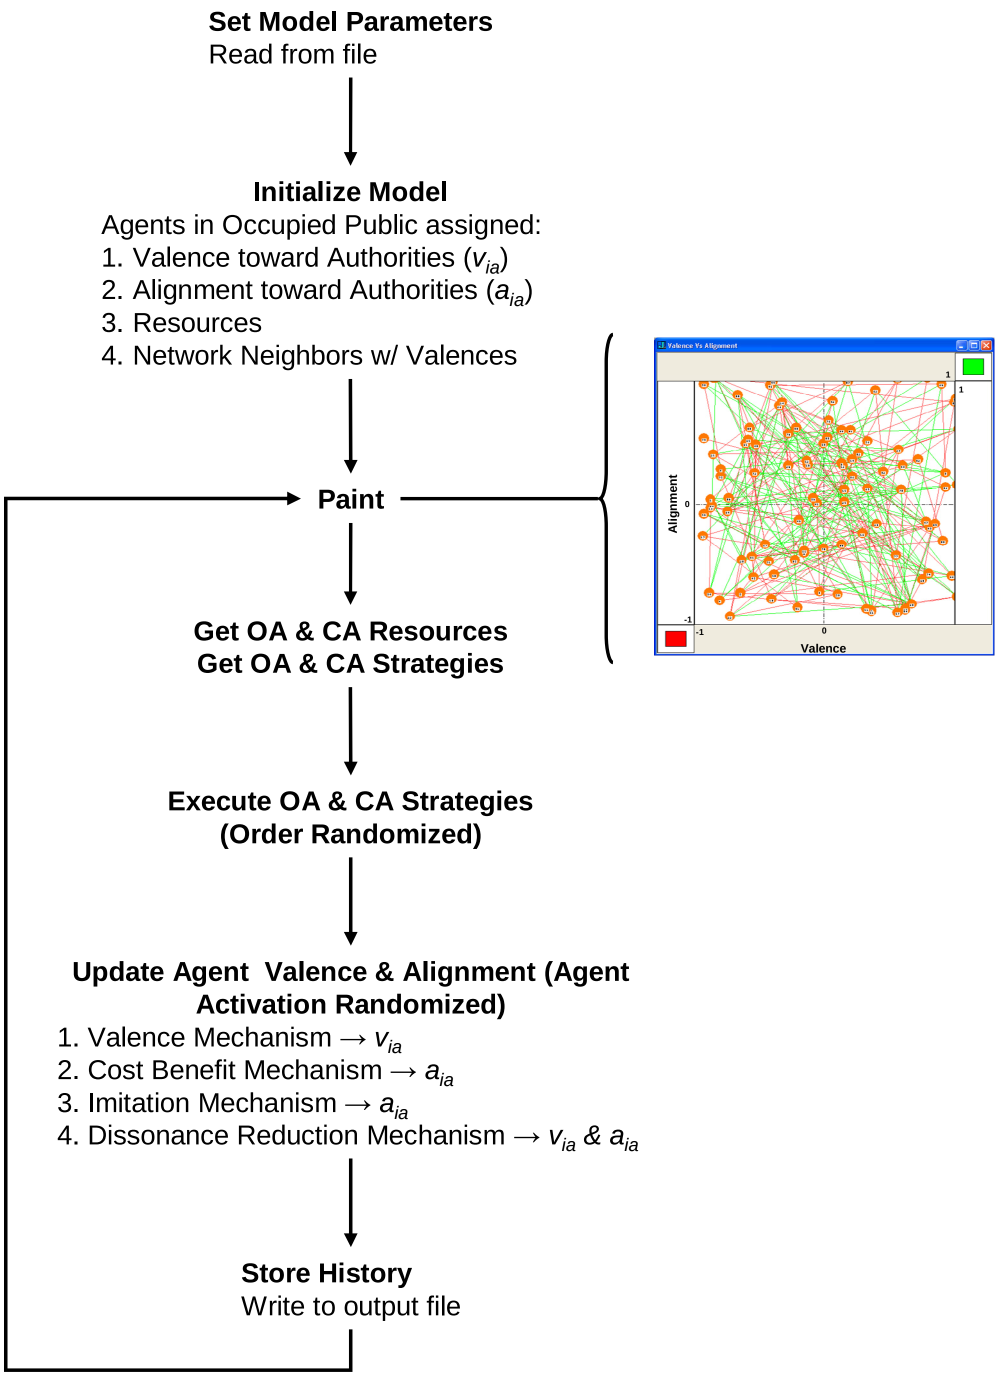

FIGURE 5 Model flow diagram

TESTING MODEL MECHANISMS

The four mechanisms described previously (valence [VM], cost-benefit [CBM], imitation [IM], and dissonance reduction [DRM]) mediate the interaction between authority strategy and agent orientation in valence and alignment space. Understanding how each mechanism — both singly and in concert with others — impacts how the model functions is therefore an important step in verification and validation of the model. In this section, we present findings from a series of tests designed to evaluate how the four mechanisms function, both alone and in combination. Each mechanism is subjected to two basic tests. First, we investigate *mechanism main effects* to ensure that each mechanism operates as expected.[3] Second, we investigate *mechanism interaction effects* to determine how the inclusion of each mechanism impacts outcomes in the full model.

We conducted mechanism tests by varying selected mechanisms while controlling for strategies and resources and by comparing outcomes to a baseline case with all mechanisms active. For the main and interaction effects tests, OA and CA are initialized with three different but symmetrical strategies, meaning that in each case, OA and CA employ identical strategies. The three strategies are composed of the *focus evenly* resource strategy and its three possible sanction strategy settings (0.5, 0, –0.5). Each authority receives equal resources, and none receives public support. Test parameters are summarized in Table 3.

For each test condition, results are averaged across the 50 runs by tick. The individual and combined effects of mechanisms are compared to a baseline that is in all ways identical, except that all mechanisms are simultaneously on. This section concludes with a test of model symmetry and the effects of variations in network and population size on model outcomes.

TABLE 3 General parameter settings for mechanism tests

Parameter		Setting
Authorities	OA and CA strategy	Focus evenly (0.5, 0, –0.5)
	OA and CA resources	100
Agents	Population	100
	Valence toward authority (v_{ia})	Uniform dist., [–1,1]
	Alignment toward authority (a_{ia})	Uniform dist., [–1,1]
	Resources	Normal dist., mean 50, st. dev. 25
	Network neighbors (NNs)	1 to 5
	Valence toward neighbors (v_{ij})	Uniform dist., [–1,1]
Mechanisms	Alignment grain (*AG*)	5%
	Imitation factor (*IF*)	50%
	Preference constraint factor (*PCF*)	0.15
Model	Runs	50
	Ticks	200

[3] For CBM, IM, and DRM, we also vary mechanism components (e.g., dampening factors) to explore how changing the settings internal to each mechanism affects how each mechanism performs. Results are reported on in the appendix.

Mechanism Main Effects

This section presents the main effects tests in which mechanisms are turned on, one at a time, while all others are kept off, next to a comparison case in which all mechanisms are turned on. Table 4 summarizes the mechanism main effects test settings.

Figure 6 presents the results of the main effects test on average valence and alignment over time. As expected, VM (orange line) has no independent effect on alignment (it operates on valence only), and neither CBM (green) nor IM (purple) have an effect on valence (both operate on alignment only). Also as expected, DRM (red) has effects on both valence and alignment, since it operates on both. Because both authorities are evenly distributing identical resources for sanctions across the alignment spectrum and varying only the point at which agents are rewarded or punished, and because the scope of the cost-benefit analysis at the heart of the CBM is local instead of global, it is not surprising that the CBM's independent effect on alignment is minimal. IM affects alignment only with respect to the distribution of alignment among friends and enemies in an agent's social network; thus, the independent effect of IM is strategy-independent. Because valence toward neighbors is uniformly distributed at initialization, IM's main effect is also minimal, reflecting the influence of very minor deviations from zero in average valence toward neighbors. The effect of different sanction strategies is very visible with respect to VM. VM reacts to three factors: (1) distribution and source of rewards and punishments, (2) distribution of friends and enemies in an agent's social network, and (3) distribution of the social network across the alignment spectrum. Because friends and enemies in the social network and their alignments are uniformly distributed at model start, the distribution and source of sanction should have the dominant effect, as is clearly evident on the right side of Figure 6. When the sanction strategy is set to 0.5, 75% of the agents are being punished by OA and rewarded by CA. Thus, VM shifts valence in favor of CA. When the sanction strategy is set to 0, the effect is also 0, while –0.5 mirrors the effect of the 0.5 setting. Most surprising is the independent effect of DRM. DRM, like IM, is strategy-independent, so the fact that the effects are the same across strategy conditions is expected. If DRM were to function as expected, the uniformly distributed agents would gravitate toward the diagonal where valence = alignment, for an average alignment and valence of 0. However, DRM shows a consistent positive bias (for OA). It is clear from Figure 6 that DRM, ceteris paribus, has the most significant impact on the full model (dotted blue line). The impact of DRM is also dominant in the interaction effects tests discussed in the following section.

TABLE 4 Configurations for mechanism main effects test

Test	VM	CBM	IM	DRM
VM	**On**	Off	Off	Off
CBM	Off	**On**	Off	Off
IM	Off	Off	**On**	Off
DRM	Off	Off	Off	**On**
Full	**On**	**On**	**On**	**On**

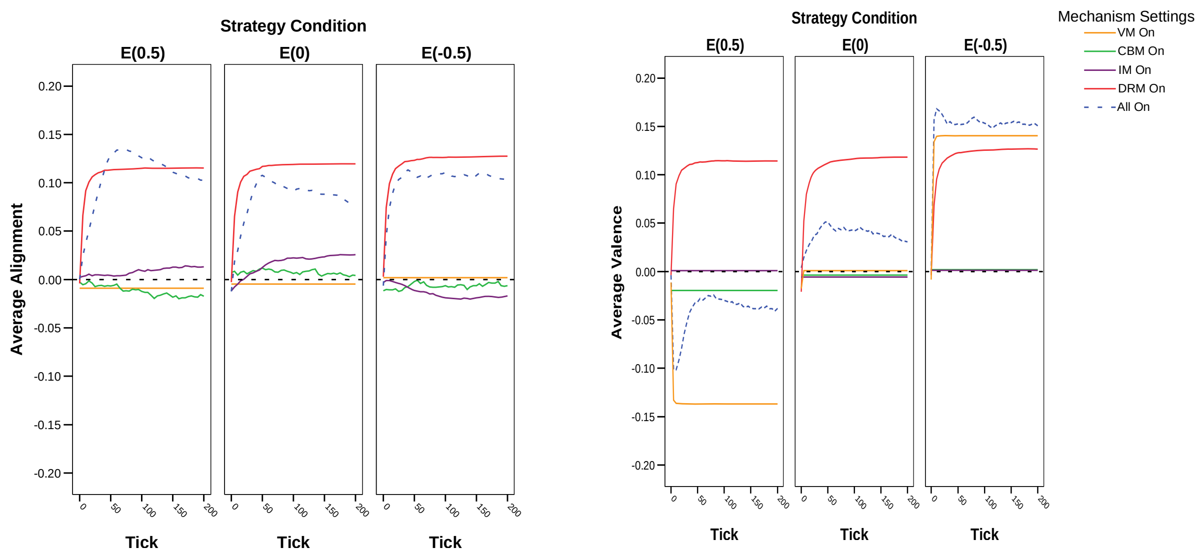

FIGURE 6 Mechanism main effects on average valence and alignment over time

Mechanism Interaction Effects

To test the impact of mechanisms in relation to the full model, we ran a series of tests in which each mechanism was individually turned off while the remaining mechanisms were kept on. By individually subtracting each mechanism from the fully specified model and controlling for strategy and resources, we have gotten insights into how the mechanisms interact to affect valence and alignment. The interaction effects test conditions are summarized in Table 5.

Because CBM and IM are the only mechanisms operating on alignment alone, we expect to find that removing them from the model will limit alignment adjustment to DRM, which operates on both valence and alignment. Similarly, VM operates only on valence; thus, removing it is expected to limit shifts in valence to the workings of DRM, since DRM keeps valence in synch with alignment. Results of the interaction effects tests are presented in Figure 7.

TABLE 5 Configurations for mechanism interaction effects test

Test	VM	CBM	IM	DRM
VM	**Off**	On	On	On
CBM	On	**Off**	On	On
IM	On	On	**Off**	On
DRM	On	On	On	**Off**
Full	**On**	**On**	**On**	**On**

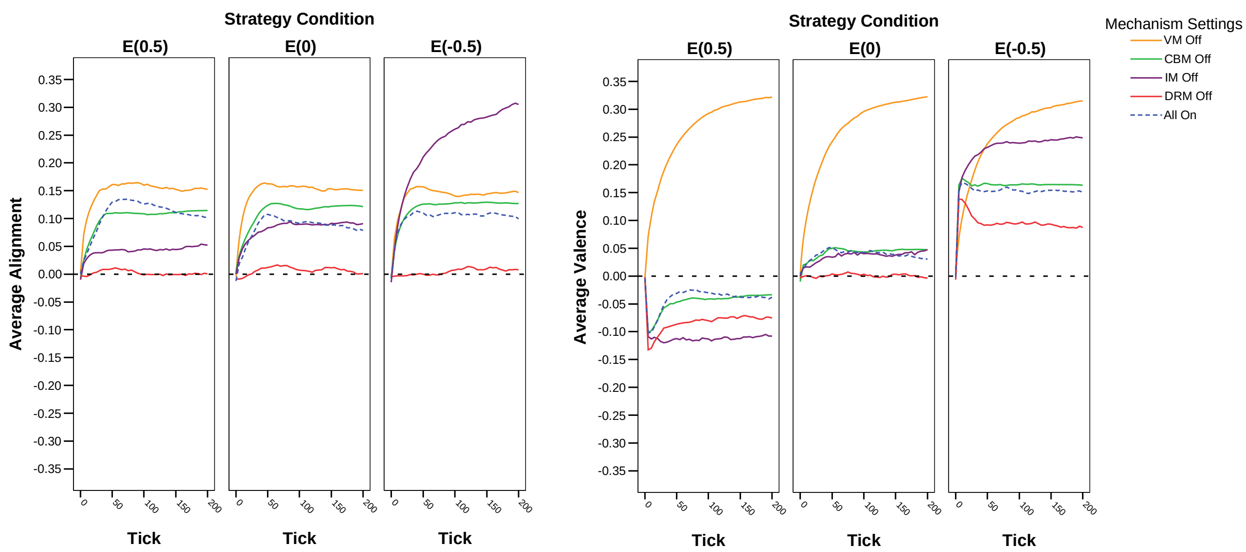

FIGURE 7 Mechanism main (**or interaction, not main?**) effects on average valence and alignment over time

The strong effect of DRM on model outcomes is clearly visible, as is the degree to which the operations of other mechanisms serve as a counterweight. When DRM is turned off (red line), alignment remains close to 0, while valence responds to the distribution of rewards and punishments resulting from variation in the sanction strategies (resource strategy remains fixed at *focus evenly*). When VM is removed from the model, DRM remains the only mechanism impacting valence directly, consistently causing a dramatic shift in average valence that favors OA (the positive bias noted earlier). Alignment is also strongly affected by the positive bias of DRM, with the interactions among VM, CBM, and IM acting as constraints.

Mechanisms and Model Symmetry

In the tests of mechanism main and interaction effects, authorities were initialized with identical resources. These tests revealed that DRM has an asymmetric effect on model outcomes that favor OA. To confirm the asymmetric impact of DRM, we conducted a test with symmetrical strategies while varying the degree of resource asymmetry between OA and CA. The data presented below were produced with both OA and CA by using the *focus evenly* resource strategy with the sanction strategy set to 0. Test parameters are summarized in Table 6. For comparison, the test was run with and without the DRM active. Results of the model symmetry test on average are presented in Figures 8 and 9.

TABLE 6 Symmetry test conditions

OA and CA Strategy		Resources	
Resource	Sanction	OA	CA
Focus evenly	0	100	0
Focus evenly	0	100	100
Focus evenly	0	0	100

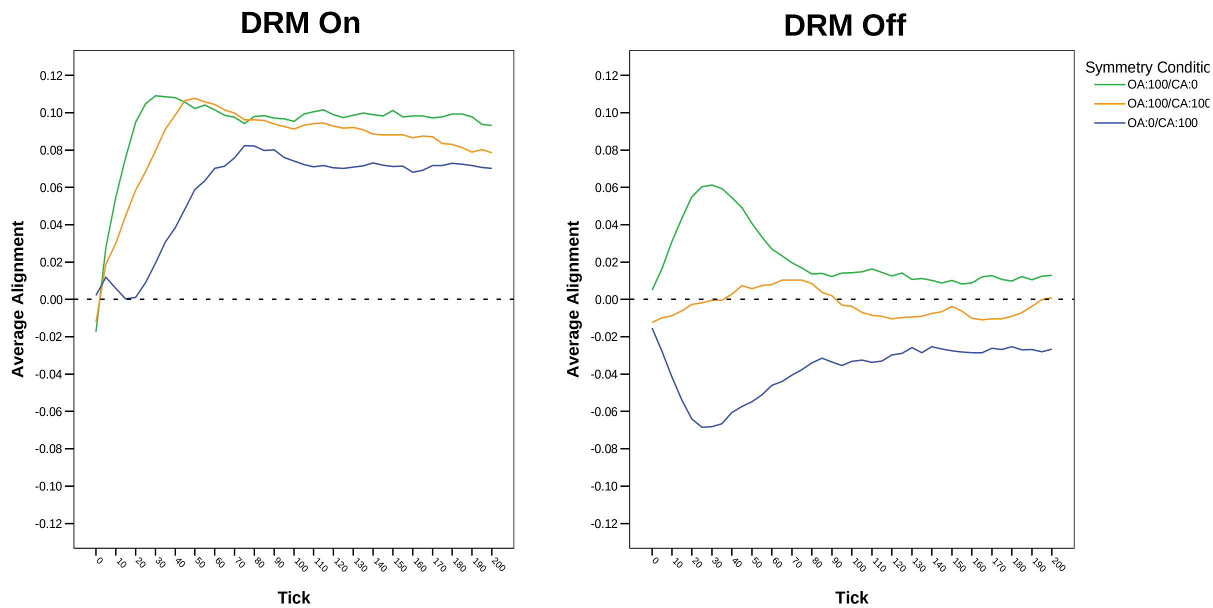

FIGURE 8 Effect of resource asymmetry on alignment over time

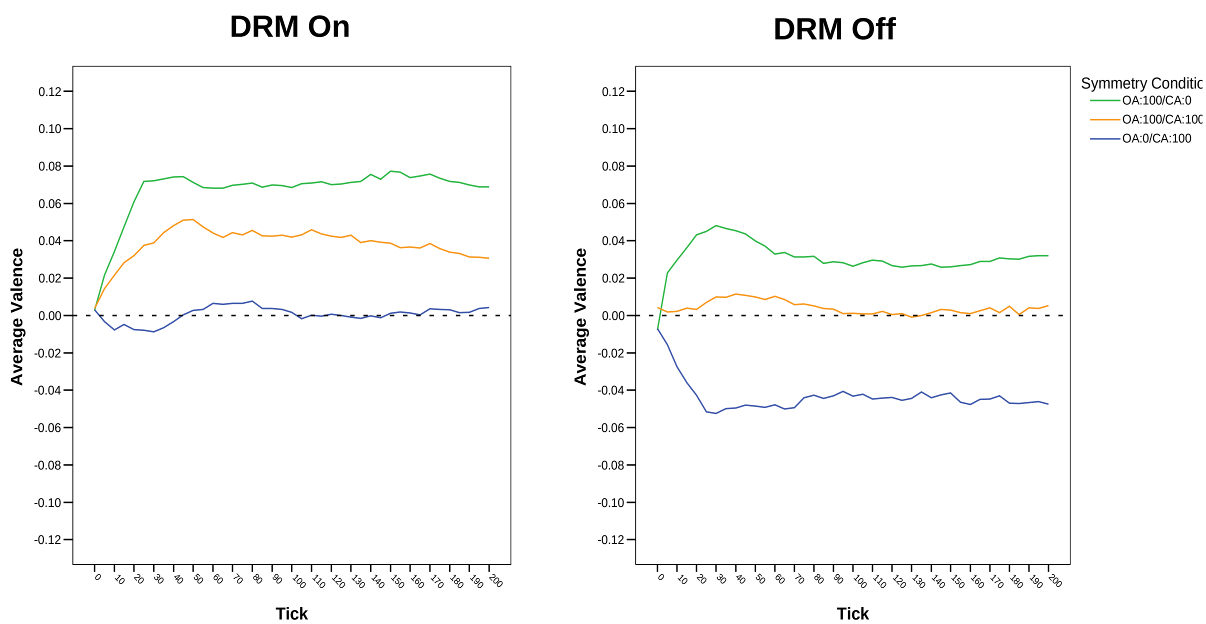

FIGURE 9 Effect of resource asymmetry on valence over time

Resource asymmetry has visible effects whether or not DRM is activated, but the positive OA bias of DRM is clearly evident; even when the OA has *no resources* (blue line), average alignment climbs to 0.8 and valence remains near zero. When DRM is deactivated, the model shows the expected symmetry in both alignment and valence.

Discussion of Mechanism Tests

The ability to model the relationship between (1) privately held emotions and beliefs (valence, or hearts and minds) and (2) publicly expressed support (alignment) is a key component of the Occupation Dynamics Model. DRM regulates how disparities between alignment and valence are either resolved or acted upon in the model. Both the main and interaction effects tests reported above indicate that DRM is introducing a positive bias into the model. The source of the unintended DRM bias can be traced to the fact that DRM shifts valence when alignment is greater than zero and shifts alignment when alignment is less than zero. Figure 10 shows the axes and directions of valence and alignment shift due to the DRM (see Figure 4 for reference).

The two shaded triangles with the red arrows represent those zones in alignment-valence space where dissonance adjusts negatively (i.e., in favor of CA). Seventy-five percent of the alignment-valence space generates positive movement in alignment or valence in favor of the OA. One additional unanticipated effect of dividing the alignment-valence space at alignment = 0 is on agents located in the lower right-hand quadrant where alignment is between 0 and –1 and valence is between 0 and 1. Here, as expected, agents shift alignment in a positive direction in response to dissonance. However, once these agents reach the horizontal axis of alignment = 0, their movement shifts negatively along the axis toward valence = 0. The convergence to zero valence in the bottom left quadrant appears to explain why DRM's independent effect on valence is slightly less strong than its effect on alignment. Of course, this assumes that agents are unaffected by other mechanisms, thereby allowing DRM to run its course. Nonetheless, the effect remains strong enough to suggest omitting the mechanism from the model until a solution is identified, implemented, and tested.

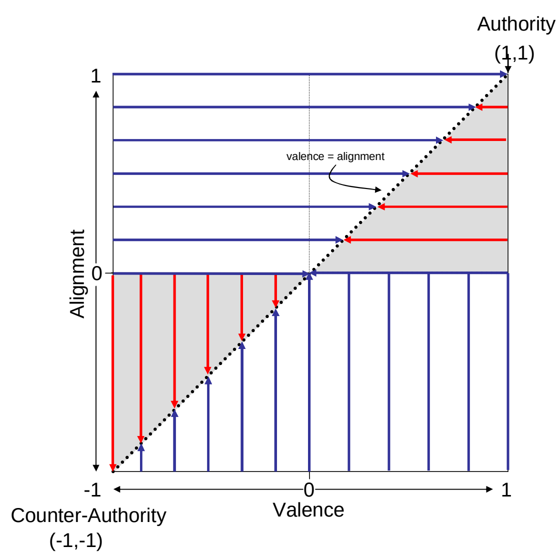

FIGURE 10 Identifying source of DRM bias

EFFECTS OF NETWORK AND POPULATION SIZE

In this section, we evaluate the effect on the model of varying the size of the population and the number of neighbors. We use the same strategy and three resource conditions as those used in the model symmetry test described in the section called Discussion of Mechanism Tests.

Network Density Effects over Time

Neighbors play an important role in the model as agents adjust their valences to reflect how their friends and enemies were treated by the authorities and adjust their alignment to be more like their friends and less like their enemies. It is not clear, however, what effect that increasing the *size* of agents' reference groups will have on model outcomes, since it is the distribution of friends and enemies in one's reference group rather than the size per se that determines agent actions. We conducted four tests varying the minimum and maximum number of neighbors. Neighbor settings and the corresponding observed average number of neighbors are presented in Table 7.

TABLE 7 Average number of neighbors by network setting

Network Neighbor Setting	Observed			
	Min.	Max.	Mean	Std. Error
0	0.00	0.00	0.00	0.00
1 to 5	4.4	5.52	4.94	0.00
5 to 10	12.00	14.00	13.00	0.00
10 to 15	21.00	23.00	22.00	0.00

For each network setting, we test the effects of network size, with both OA and CA initialized with identical strategies (*focus evenly*, sanction strategy = 0) while levels of resource asymmetry are varied. DRM is deactivated. Each configuration is run 50 times for a total of 200 ticks. Because neighbors can, in principle, diverge quite significantly in alignment prior to the activation of the imitation mechanism, having more neighbors increases the distribution of available referents across the alignment spectrum, essentially watering down the effect. Having zero neighbors is like having one referent — the self — which takes IM out of play while magnifying the effect of VM. VM's influence is increased because the average valence toward neighbors, which approaches zero as the number of neighbors increases, climbs to 0.8 when no neighbors are assigned (the value of valence toward self). Figure 11 presents the effects that varying network size have on average alignment and valence over time across the three resource asymmetry conditions.

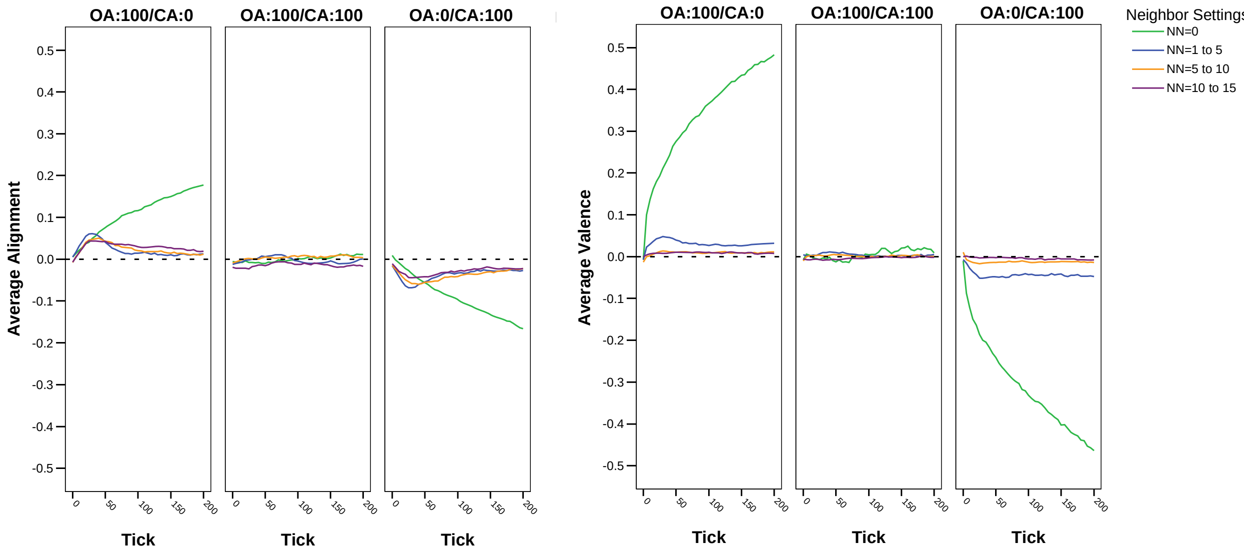

FIGURE 11 Effect of network size on average valence and alignment over time, no DRM

In line with basic expectations of model symmetry, outcomes of the network test favor the authority with greater resources and are indeterminate when both authorities are equal. Also in line with expectations, the magnified effect of VM on average valence when agents have zero neighbors is clearly evident. The effect of zero neighbors on alignment is more difficult to explain. Since having zero neighbors effectively deactivates IM, the effect on alignment can be caused only by CBM. In the same way that increasing network size dampens the effect of VM on valence by lowering the average valence toward neighbors from 0.8 with no neighbors to 0 as the number of neighbors increases, IM damps the effects of CBM, which otherwise has a strong pull on alignment in the direction of resource superiority. The effect of increasing the number of NNs is therefore similar for both alignment and valence, except that the effect is generated by different pathways: With each increase in network size, the effects of VM are decreased, while the constraining effect of IM on CBM is increased.

Effects of Population Size over Time

To test the effect of population on the outcomes in the baseline model, we ran four tests with increasingly larger populations (50, 100, 200, and 400 agents). Population is expected to have an effect on the value of resources; as population size increases, the relative effectiveness of resources used to punish and reward individuals — that is, the number of individuals that can be punished and rewarded — decreases. At the same time, the amount of resources available via public support increases. Each of the four population sizes was tested in relation to the same three resource conditions used in the tests of network size above and with DRM deactivated. The only difference is that for the population tests, the CA receives, in addition to its fixed endowment, public support of five resource points for every individual having an alignment below –0.5. Results are presented in Figure 12.

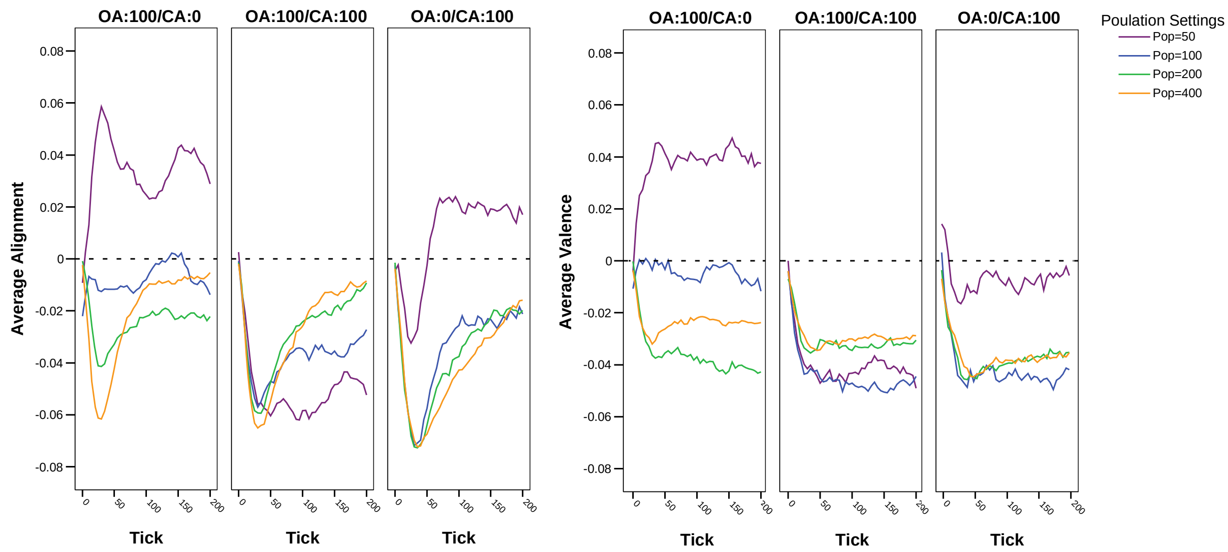

FIGURE 12 Effect of population size on average alignment and valence over time, no DRM

As expected, population size affects model outcomes; however, while the graphs that make up Figure 12 show variation, it is important to note that the absolute range is very small, since all values remain close to zero. Still, the effects are largely as expected: increasing the size of the population confers a resource advantage to CA through public support. The effect of population size on public support for CA is presented in Figure 13.

Already at a population of 100, the CA generates enough resources from public support to top the OA's fixed endowment, even when the CA receives no fixed endowment. Only when the population is 50 and the CA receives no fixed endowment (OA:100/CA:0) does the OA have the advantage. The effect of superior resources is not uniformly beneficial for the CA: When the OA has no resources, a small population size and CA resource superiority generate alignment in favor of the OA. The exact cause of this positive effect on alignment is difficult to pinpoint, but the effect is very small in any case.

EXPLORING ODM: BASELINE SIMULATION FINDINGS

To explore the full range of strategic interactions possible in the model, we ran the fully specified model (minus DRM) by using a baseline assumption of a more powerful OA competing with a weaker CA receiving public support from individuals holding strong private sympathies against the occupation.

For each of the 81 strategy combinations, data on output variables were collected from 50 runs at 200 ticks recording initial conditions and then at every fifth tick, for a total of 2,050 observations per strategy combination (50 runs × 21 recorded ticks) and 166,050 observations per model configuration (81 strategy combinations × 2,050). The number of runs was determined by examining mean variance in output variables at different run configurations. At 50 runs, the mean standard error across all strategy conditions for average

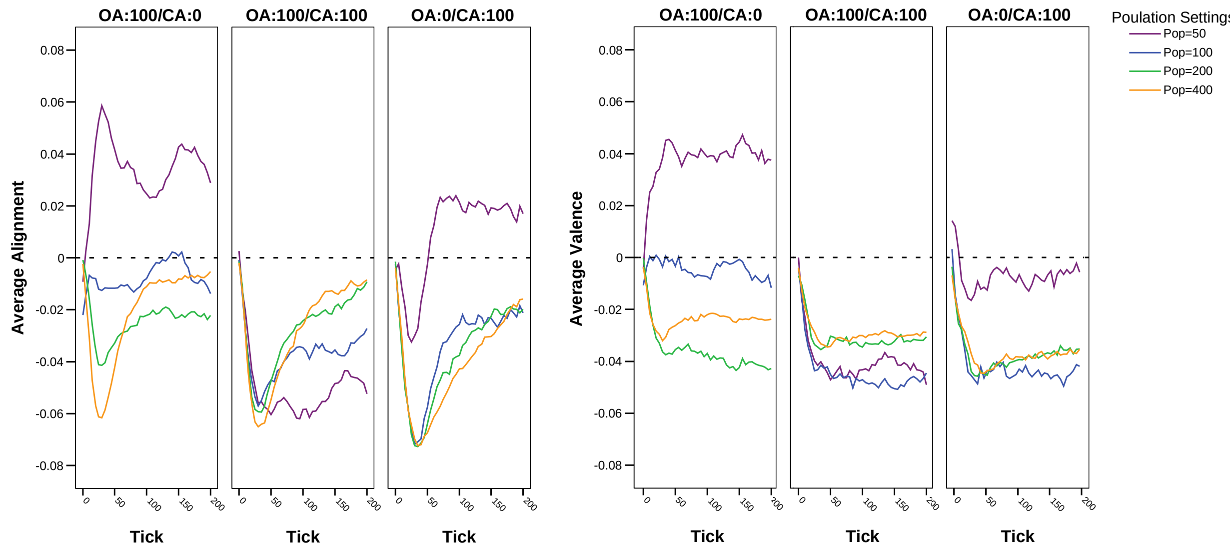

FIGURE 13 Effect of population size on CA resources from public support alone, no DRM

valence and average alignment at tick 200 was 0.01, indicating a high level of precision in estimating mean outcomes. Data were collected in tab-delimited text files and stored in a Microsoft Access 2003 database for preliminary analysis and export to the statistical software package SPSS Base 13.

In this and all subsequent model configurations, alignment and valence toward authorities were uniformly distributed between –1 and 1 at model initialization to control for distributional effects. The full parameter specification of the baseline model is presented in Table 8.

Because the model is being run without DRM, findings are necessarily provisional, with a more detailed model testing and evaluation planned in the future when all mechanisms can be included. Nonetheless, even without DRM, the model produces results that are suggestive of real-world insurgency dynamics. Baseline findings are presented below for outcomes at tick 200 in the baseline model.

As described under the heading "Occupation and Counter Authorities" in the second section, both the OA and CA have nine strategy combinations composed of three resource strategy options and three sanction strategy options with which to administer punishments and rewards to individuals in the occupied public in order to generate support; these generate a total of 81 strategy combinations (9 × 9). Table 9 presents outcomes in valence and alignment by resource strategy alone (i.e., it ignores for the moment the further specification of sanction strategy). Instead of using average alignment and valence, we use the percent of agents with valence or alignment greater than or equal to zero, which represents greater detail on the distribution of alignment and valence than does a simple mean.

When the baseline model is run with DRM deactivated, mean outcomes for both alignment and valence across all strategy combinations tend toward 50%, with only small amounts of variation introduced by the interaction between resource asymmetry (the OA starts

TABLE 8 Occupation dynamics simulation parameter settings

Parameter		Setting
Authorities	OA strategy	All (9)
	CA strategy	All (9)
	OA resources	150
	CA resources	50
	Public support	CA only, 5 pts for v_{ia} of ≤ –0.5
Agents	Population	100
	Valence toward authority (v_{ia})	Uniform dist., [–1,1]
	Alignment toward authority (a_{ia})	Uniform dist., [–1,1]
	Resources	Normal dist., mean 50, st. dev. 25
	Network neighbors (NNs)	1 to 5
	Valence toward neighbors (v_{ij})	Uniform dist., [–1,1]
Mechanisms	VM	On
	CBM	On, AG = 5%
	IM	On, IF = 50%
	DRM	Off
Model	Runs	50
	Ticks	200

TABLE 9 Percent of agents with alignment and valence of ≥0 at tick 200, baseline model

			CA Resource Strategy			
			Neutrals	Friends/ Enemies	Even	Mean
OA Resource Strategy	Neutrals	Alignment	49.76	47.83	47.78	48.46
		Valence	49.98	50.29	49.72	50.00
	Friends/ enemies	Alignment	52.10	49.51	50.22	50.61
		Valence	50.24	49.77	49.06	49.69
	Even	Alignment	51.98	49.55	50.43	50.65
		Valence	50.15	50.77	50.07	50.33
	Mean	Alignment	51.28	48.96	49.48	
		Valence	50.12	50.28	49.62	

with three times the resources of the CA) and resource strategy (i.e., how the available resources are to be concentrated in the occupied public). Looking at mean outcomes presented in Table 9, OA achieves its best average result if it selects the *focus evenly* resource strategy, with slightly more than 50% of agents having alignments and valences greater than or equal to zero. The same holds true for CA: Irrespective of what resource strategy OA selects, CA achieves its best average outcome if it selects the *focus evenly* resource strategy. When sanction strategy is added to resource strategy, authority interactions and outcomes can be assessed at a higher degree of granularity. Table 10 presents outcomes in percent of agents with alignment and valence greater than or equal to zero for all 81 strategic interactions. Given the right mix of resource and sanction strategy, the CA can significantly influence results in its favor. For example, if the OA chooses *focus evenly* with the sanction point set to alignment = 0 (OA's dominant strategy), CA can still finish ahead if, by anticipating OA's strategy, it selects *focus on neutrals* with a sanction point set to 0.

When the same outcomes are presented from the perspective of the OA in a scatter plot of valence on the horizontal (x) axis and alignment on the vertical (y) axis (Figure 14), variation in outcome due to strategic interaction are readily visible. In the scatter plot, resource strategies are marked by color (green = focus on neutrals, blue = focus on friends/enemies, and red = focus evenly), and sanction strategies are marked by shape (triangle = 0.5, circle = 0, and star = –0.5). Markers in the top half of the scatter plot depict outcomes where the percentage of agents that both feel and express some level of support for the OA is greater than zero. What is remarkable about the distribution of outcomes is that even with an initial resource asymmetry between the OA and CA, the CA can apply strategies that neutralize this advantage. This finding is suggestive of real-world occupation/insurgency dynamics, in which insurgencies adjust their strategies to find the Achilles' heel of their much more powerful adversaries.

CONCLUSIONS AND FUTURE DIRECTIONS

This paper represents our first thorough verification of the Occupation Dynamics Model, and it exposes the inner workings of the model under a variety of parameter conditions. We began by describing the parts of the model: the parameters and mechanisms that translate

TABLE 10 Percent of agents with alignment and valence of ≥0 at tick 200, 81 strategy combinations, baseline model

					CA Resource Strategy									
					Neutrals			*Friends/Enemies*			*Even*			
					CA Sanction Strategy									
					0.5	*0*	*-0.5*	*0.5*	*0*	*-0.5*	*0.5*	*0*	*-0.5*	***Mean***
OA Resource Strategy	*Neutrals*	**OA Sanction Strategy**	*0.5*	Alignment	52.22	43.92	50.54	42.64	46.48	47.28	45.6	44.26	47.74	**46.74**
				Valence	39.08	44.9	49.98	41.96	44.78	50.92	42.08	44.94	48.82	**45.27**
			0	Alignment	59.42	48.9	54.08	54.62	53.06	52.42	56.4	52.56	52.48	**53.77**
				Valence	47.14	49.22	54.28	46.24	52.08	53.84	45.7	51.36	53.36	**50.36**
			-0.5	Alignment	49.28	39.2	50.32	47.32	44.72	41.94	46.98	39.96	**44.02**	**44.86**
				Valence	50.66	52.66	61.92	50.02	52.74	60.06	48.4	51.38	**61.42**	**54.36**
	Friends/ Enemies		*0.5*	Alignment	**58.8**	46.8	54.58	46.92	52.4	50.66	54.66	45.62	50.62	**51.23**
				Valence	**41.04**	45.96	51.5	37.62	47.6	50.8	39.06	43.72	49.04	**45.15**
			0	Alignment	57.36	46.04	52.26	50.08	47.62	48.94	51.82	45.4	52.16	**50.19**
				Valence	49.3	47.8	54.02	45.62	48.98	53.88	45.5	46.9	54.68	**49.63**
			-0.5	Alignment	52.36	45.28	55.44	48.96	48.36	51.62	50.5	45.88	55.32	**50.41**
				Valence	50.54	53.54	58.44	49.6	52.76	61.1	50	54.1	58.54	**54.29**
	Even		*0.5*	Alignment	56.88	46.96	51.48	44.2	48.42	48.36	49.18	44.3	50.2	**48.89**
				Valence	40.66	46.68	50.9	41.18	45.34	50.46	37.12	45.56	50.9	**45.42**
			0	Alignment	61.46	45.54	57.5	56.62	54.26	54.9	58.82	50.6	55.82	**55.06**
				Valence	48.4	47.34	56.04	46.64	52.78	55.24	45.4	50.4	55.7	**50.88**
			-0.5	Alignment	49.98	42.96	55.06	48.08	47.32	43.76	49.52	44	51.4	**48.01**
				Valence	49.34	53.46	58.56	50.48	53.92	60.92	50.3	53.02	62.26	**54.70**
			Mean	**Alignment**	**55.31**	**45.07**	**53.47**	**48.83**	**49.18**	**48.88**	**51.50**	**45.84**	**51.08**	
				Valence	**46.24**	**49.06**	**55.07**	**45.48**	**50.11**	**55.25**	**44.84**	**49.04**	**54.97**	

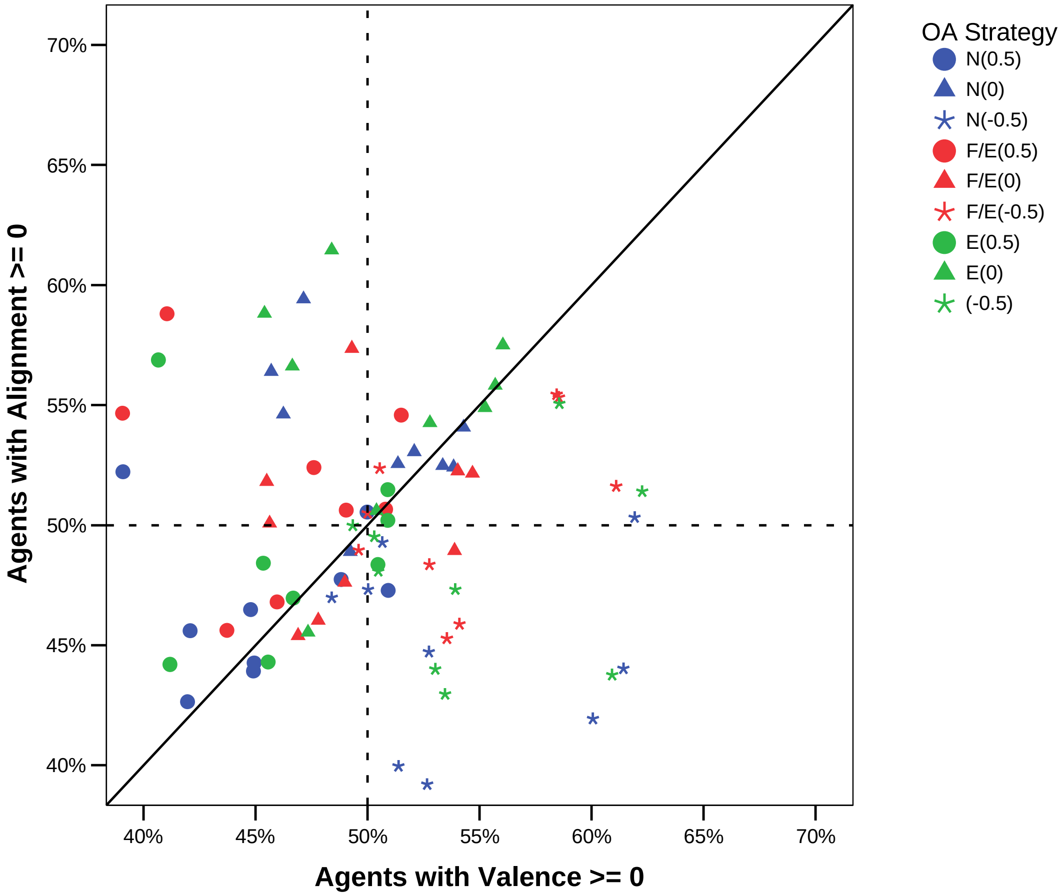

FIGURE 14 Scatter plot of percent of agents with alignment and valence of ≥0 at tick 200, all 81 strategy combinations, baseline model

authority strategies into agent outcomes. We then presented a series of tests designed to evaluate mechanisms and their impact on the model. We found that an unintentional asymmetry built into one of the mechanisms, DRM, introduced bias into the model. When DRM was deactivated, the model behaved in line with expectations. We also tested the effects of network and population size, and we presented a baseline model configured to represent a notional occupation setting. The baseline model test revealed that strategic interaction represents a key element in determining authority outcomes; even a materially weaker authority could gain the advantage in specific instances of strategic interactions. This finding has implications for the study of insurgent strategy selection and will likely play a key role in future iterations of the model in which authorities dynamically adjust their strategies in order to maximize their desired outcomes.

However, before we can pursue the implications of the baseline model with any confidence, more investigation into the mechanisms is needed to ensure they accurately reflect our expectations. A thorough understanding of the mechanisms — especially DRM — is particularly important (as stated above) for the mechanisms to translate authority strategies into outcomes in the occupied public. As expected, the testing exercise uncovered a number of puzzles with respect to how the various mechanisms work and interact, raising questions of internal validity that will need to be pursued before the complexity of the model continues to be expanded. Thus, we stress the detailed investigation of internal validity prior to addressing external validity. Even if the model produces interpretable and reasonable results, if the mechanisms that generate those results do not function as expected, then the explanatory purchase of the model is diminished.

We plan to pursue two lines of further testing. First, we need more finely grained measures to capture changes in the distribution of agents over time. In particular, current measures miss the degree to which agents are clustered around boundaries — most importantly, with respect to zero alignment or valence. What appear to be great victories (e.g., 80% of agents with an alignment of ≥ 0 may, upon closer inspection, be revealed to be only marginal victories, with the vast majority of those 80% hovering just above zero. Second, we plan to continue testing and adjusting mechanisms in isolation, fine-tuning them to meet our expectations, and then to proceed to test their interact effects.

Finally, while we continue to investigate and improve the functioning of the present four mechanisms, we plan to continue developing ODM as well. There are additional innovations that can provide the focus for current and future development of ODM. These innovations can be divided into two types: complex and interpretive. In complex innovation, components are introduced to extend the model so it includes additional characteristics that affect the solidarity of contending groups. In the near term, these include dynamic networks, the emergence of multiple authorities (without the mirroring assumption described in footnote 2), and the explicit representation of occupier publics as the source of sanctioning resources and the target of insurgent strategies.

In interpretive innovation, the agents assess the meaning of communications and actions, and larger historical events emerge from the interaction among agents (Sallach 2003a,b). Initially, this innovative approach will be applied to strategy selection of relevant authorities. Later it will be extended to occupied and occupier publics, as they assess policies and alternatives within informal settings.

ACKNOWLEDGMENT

This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Cederman, L.-E. (1997). Emergent Actors in World Politics: How States and Nations Develop and Dissolve. Princeton, NJ: Princeton University Press.

Cederman, L.-E. (2001). “Agent-based Modeling in Political Science.” The Political Methodologist 10(1):16–22.

Cederman, L.-E. (2004). Articulating the Geo-cultural Logic of Nationalist Insurgency v. 2. Unpublished. Center for Comparative and International Studies, Swiss Federal Institute of Technology.

Edelstein, D.M. (2004). "Occupational Hazards: Why Military Occupations Succeed or Fail." International Security 29(1):49–91.

Epstein, J.M. (2002). "Modeling Civil Violence: An Agent-based Computational Approach." PNAS 99(Suppl. 3):727–250.

Festinger, L. (1962). A Theory of Cognitive Dissonance. Stanford, CA: Stanford University Press.

Grossman, H.I. (1991). "A General Equilibrium Model of Insurrection." The American Economic Review 81(4):912–921.

Kuran, T. (1991). "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." World Politics 44:7–48.

Kuran, T. (1995). Private Truths, Public Lies: The Social Consequences of Preference Falsification. Cambridge, MA: Harvard University Press.

Kuran, T. (1998). "Social Mechanisms of Dissonance Reduction." In P. Hedström and R. Swedberg, Social Mechanisms: An Analytical Approach to Social Theory. Cambridge, UK: Cambridge University Press, pp. 147–171.

Lustick, I.S., et al. (2004). "Secession in Multicultural States: Does Sharing Power Prevent or Encourage It?" American Political Science Review 11(3):367–394.

Olson, M. (1993). "Dictatorship, Democracy and Development." American Political Science Review 87:567–570.

Pepinsky, T.B. (2005). "From Agents to Outcomes: Simulation in International Relations." European Journal of International Relations 11(3):367–394.

Sallach, D.L. (2003a). "Interpretive Agents: Identifying Principles, Designing Mechanisms." In Proceedings of the Agent 2003 Conference on Challenges in Social Simulation, presented by Argonne National Laboratory and The University of Chicago, Oct. 2–4.

Sallach, D.L. (2003b). "Social Theory and Agent Architectures: Perspective Issues in Rapid-discovery Social Science." Social Science Computer Review 21:179–195.

Skaperdas, S. (2001). An Economic Approach to Analyzing Civil Wars. Workshop on Civil Wars and Post Conflict Transitions, University of California, Irvine.

Tákacs, K. (2001). "Structural Embeddedness and Intergroup Conflict." Journal of Conflict Resolution 45(6):734–769.

Thomson, N. (2001). J: The Natural Language for Analytic Computing. Baldock, UK: Research Studies Press.

Tilly, C. (1985). "War Making and State Making as Organized Crime." P.B. Evans, D. Rueschemeyer, and T. Skocpol, eds. Bringing the State Back in Cambridge, UK: Cambridge University Press, pp. 169–191.

SIMULATIONS OF EGALITARIAN SOCIETIES WITH COMPARISON TO OBSERVATIONS

S.M. YOUNGER,* Los Alamos National Laboratory, Los Alamos, NM

ABSTRACT

Multi-agent simulation was used to study normative behavior in model egalitarian societies (i.e., those without centralized leadership). We simulated populations of 100 agents in finite landscapes, such as one might find on isolated islands. Agents moved in search of food, produced offspring, and ultimately died of hunger or old age. They remembered and shared action-generated reputations of other agents, and these reputations influenced future interactions. The aggregate of agent reputations, called mutual obligation, monitored sharing-generated social cohesion. Various methods of sharing, the effect of tolerance to theft, and the effect of homicide and revenge were simulated. We found that social cohesion was maximized for indiscriminant sharing rather than sharing designed to optimize individual fitness. When reputation was a factor in mate selection and when some tolerance of past transgressions was allowed, populations were stable only for very low or very high values of tolerance. In between, there was a high probability of population collapse. Societies optimized their probability of success by excluding a major segment of the population from homicide and revenge. These results are compared to observations of a number of egalitarian cultures around the world.

Keywords: Egalitarian society, multi-agent simulation, reciprocity, violence

INTRODUCTION

Egalitarian societies offer interesting test cases for social simulation in that they are typically small, exist in relative isolation, lack complex political structures, and demonstrate a variety of cultural patterns. Typical egalitarian societies number in the few hundreds of persons, well within the reach of many simulation techniques. The isolation of desert bands or of island peoples makes boundary conditions more straightforward than when several cultures closely interact. Social behavior in egalitarian cultures is dominated by the individual agent, inviting the systematic study of various rules of behavior or other agent models. Finally, egalitarian societies around the world offer substantial cultural diversity so as to constitute a rich basis of comparison for simulations.

Significant ethnographic data exist for egalitarian societies in a variety of environments, from resource-poor deserts to resource-rich tropical islands. Most important, several anthropologists have undertaken to collect data that permit alternate social models to be compared on an objective footing. These comparisons present an excellent opportunity to test agent models against real-world data.

---

* *Corresponding author address*: Stephen M. Younger, Los Alamos National Laboratory, Mail Stop B210, Los Alamos, NM 87545; e-mail: syounger@hawaii.edu.

One of the complications of such comparisons is that the entities involved are sometimes qualitatively different in nature. For example, in studying the sharing of food, it is straightforward to measure the caloric value of the food but it is more difficult to measure the "social value" ascribed to the sharing. However, by examining the predictions of various sociological models and comparing them to observation, qualitative assessments can sometimes be made to support or reject hypotheses. Simulation is especially helpful in this regard in that it enables systematic examinations of the effects of different behavioral models with a comparison to what is found in real-world societies. In this paper we examine three topics — sharing, tolerance, and violence — and compare the results of simulations to observations of egalitarian societies.

METHODOLOGY USED IN THE SIMULATONS

A detailed description of the simulation method used here can be found in Younger (2003, 2005a,b); the last reference contains pseudocode of the major algorithms. We modeled a population of 100 agents on a 20×20 grid containing five sources of food. The simulation proceeded through a series of time steps in which agents decided their individual course of action on the basis of their hunger and their relationships to other agents.

The food sources were replenished at a rate of 20 food units per time step so that an average population of 100 agents could be sustained. Food was enduring, so food units not used in one time step remained for use in the future. Agents moved around the landscape in search of food, and when they found a food source, they remembered its location and the amount of food present. Agents could sense food and other agents to a distance of five squares in each direction, a sensory range that prevented them from seeing the entire environment at one time. When an agent was at a food source, it consumed food so that its hunger was reduced to zero and collected up to 100 food units to carry for later consumption.

An agent required one unit of food per time step and died of starvation if its need for food exceeded 200 points. The maximum age to which an agent could live was 4,000 time units. All simulations reported here were run for 40,000 time steps or 10 agent lifetimes, and the results in the tables and figures are averages of 20 such runs.

Agents were divided into two normative categories: sharing and stealing. Sharing agents shared whatever food they carried with all collocated agents; stealing agents who were not carrying food stole food carried by another collocated agent. (A more complex algorithm including theft was used in some scenarios and is described below.) An interaction matrix, imx(j,k), tallied agent interactions. When agent k shared with agent j, the amount shared was added to imx(j,k). When agent k stole from agent j, the amount stolen was subtracted from imx(j,k). The interaction matrix element thus represented a form of normative reputation of agents. When two agents met, they shared normative reputations of all other agents by averaging the interaction matrix elements connecting them to those other agents. The sum of all of the interaction matrix elements connecting agents in the population was termed the mutual obligation and represented the sharing-generated social cohesion of the model society.

Female agents chose a mate upon reaching the reproductive age of 1,000 time units. They chose the unmatched male with whom they had the highest interaction matrix element. Mating was monogamous and for life. At each time step, there was a probability of conception, set to 0.004. Offspring were born immediately, with no gestation period, and had no knowledge of other agents or of the landscape. The normative character of the mother (sharing or stealing) was inherited by the new agent.

More information on the model and the effect of various choices of parameters can be found in Younger (2005a,b).

RESULTS OF SIMULATIONS COMPARED TO OBSERVATIONS OF EGALITARIAN CULTURES

Sharing in Small Societies

Sharing was a ubiquitous phenomenon in egalitarian societies. In societies where the acquisition of food was sporadic and occurred in large parcels, such as those that hunted game too large to be consumed by an individual or family unit, sharing helped ensure that everyone was fed regardless of who felled the prey. In many other situations, notably in the case of resource-rich tropical islands, there was no need to share, yet sharing occurred all the same. Sharing served to create a network of mutual obligation within the society that was an important component of social cohesion. When every person was in some way indebted to every other person, there was an enhanced sense of belonging and an expectation that one would be cared for in a time of need. This attitude was exemplified among the Semai of Malaysia (Robarchek and Robarchek 1992), who regarded belonging to the group as an essential element of life in an uncertain world.

There are a number of means by which people can choose to share — from indiscriminant sharing that is independent of the sharer's relationship to the recipient to focused sharing done in expectation of comparable return. Taken to its extreme, the later form of sharing approaches trade. To examine various types of sharing, we simulated a society of 100 agents that either shared or did not share. (There was no theft in this scenario.) Non-sharers represented "free riders" who derived benefit from others without the cost of contributing any food of their own. The initial population was evenly divided between sharers and non-sharers. Four models were examined:

- Indiscriminant sharing wherein an agent shared with whoever was at the same location, regardless of relationship or past history;
- Sharing only with the head of a household;
- Sharing only within the sharer's family unit (mother, father, spouse, offspring, sibling); and
- Sharing only with other sharing agents.

The results of the simulations are given in Table 1. We found the mutual obligation, which in our model represented sharing-generated social cohesion, was maximized for indiscriminant sharing. In each of the other cases, less sharing occurred, so that the network of mutual obligations generated by the receipt of gifts from others was reduced.

TABLE 1 Mutual obligation for various models of sharing within a gathering society

Model for Sharing	Mutual Obligation	Standard Deviation
Indiscriminant Sharing	330	58
Sharing Only with Head of Household	160	21
Sharing Only within Family	190	21
Sharing Only with Other Sharing Agents	210	20

Bliege et al. (1997) did a quantitative study of sharing of hard-to-obtain turtle meat on Mer Island, located in the Torres Straight off the northern coast of Australia. In that resource-rich environment in which an individual could easily satisfy his needs and in which sharing was not required, they found that hard-to-obtain turtle meat was shared with whomever happened to be nearby, regardless of kin or social relationship. In fact, the probability of sharing was inversely proportional to the distance of the sharer to the potential recipient. There was no attempt to direct meat to those who might provide some future advantage, such as the families of prospective marriage partners, and there was no consideration of whether the recipient had ever shared with the giver.

Kaplan and Hill (1985) observed a similar pattern among the Ache of Paraguay. They found that sharing did not follow an inverse relationship with kinship. They did not find that non-sharers received less of a share than sharers. The simulations thus support the observations that sharing in egalitarian cultures played an important role in building solidarity within a population.

Tolerance

All societies possess a set of behavioral norms that govern the actions of individuals and, in many cases, groups. A key question in evolutionary social dynamics is whether there is a preferred set of normative guidelines that improve the survival probability of a population. Boehm (1999) notes that the normative systems of egalitarian people the world over are remarkably similar and, in particular, that they all seem to display a remarkable intolerance to non-normative behavior. For example, it is common in such cultures that transgressions are immediately responded to by the aggrieved party, sometimes by ridicule and sometimes by violence. Most often, such sanctions are conducted at the individual level, between the two people involved, rather than at the group level. Why is this, and why don’t such people display more tolerance toward non-normative behavior?

We investigated this issue by simulating a population of 100 agents, the initial group being equally divided between those who shared and those who stole. Further, we made the selection of mates dependent upon the reputation of the agents. When it came time for a female to mate, she chose the male with whom she had the highest interaction matrix element, which in our model represented the reputation of the agent. If the prospective mate had a reputation below a certain tolerance level, treated as a variable in the simulations, then that agent would be rejected as a mate. Thus one would expect agents who shared frequently to have a high reputation and hence have a high probability of being chosen as a mate. Conversely, an agent who stole would have a lower probability of being chosen. Note that sharing and theft had opposite near-term and long-term consequences. Sharing detracted from short-term survival in that food was given away, but there was a long-term advantage in finding a mate. Stealing increased the short-term survival probability by allowing an agent to take food from another, in essence providing another source of food, but there was a long-term disadvantage in finding a mate. Figure 1 shows the survival probability of the total population vs. the tolerance level.

The population survived when tolerance was either very low or very high. In between, there was a significant probability of population collapse. For low tolerance, agents with a reputation for theft were effectively excluded from the mating pool and were thus unable to pass along their “theft gene” to the next population. Over several generations the population evolved to include only sharing agents. (Recall that sharing and stealing behavior was inherited from the mother.) Conversely, when tolerance was very high, there was no long-term advantage to sharing, and the short-term advantage of theft prevailed. In between, we found that the subpopulation of sharing agents disappeared as a result of being preyed upon by thieves and that once those sharing agents were gone, the stealing agents could not find mates among themselves. This effect is illustrated in Figure 2, which shows the fractions of sharing and stealing agents vs. tolerance.

While our model is simple compared to human egalitarian societies, it demonstrates that tolerance to transgressions can have negative effects when reputation is important in mate selection. It is interesting that all known egalitarian societies practice strict intolerance to individual transgressions, in accord with the results of the simulations.

Homicide and Revenge

Homicide and revenge were significant contributors to adult deaths in many egalitarian societies. It was not uncommon for homicide and warfare to account for several tens of percent

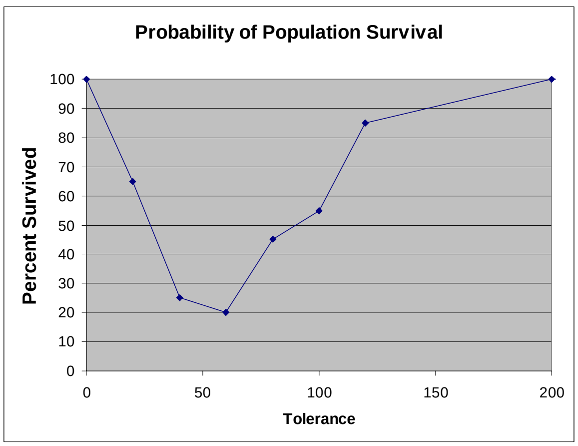

FIGURE 1 Probability for survival of the population until the end of the run

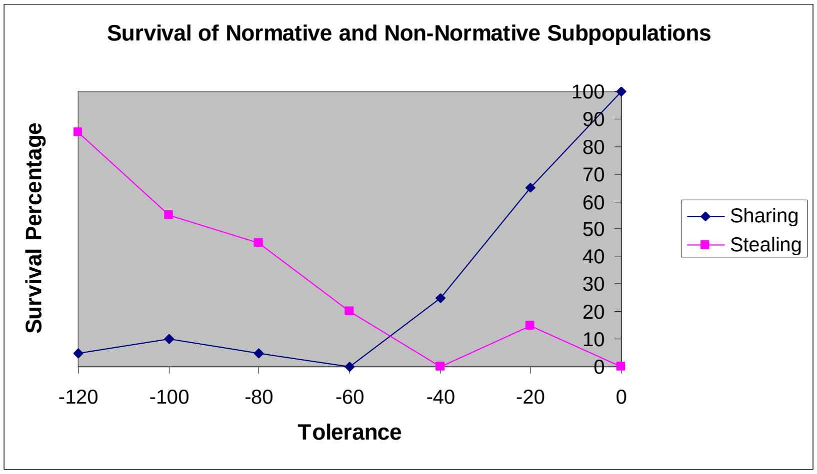

FIGURE 2 Fraction of sharing and stealing agents as a function of tolerance

of all adult deaths (Keeley 1996). For example, among the Gebusi of New Guinea, Knauft (1987) found that about one third of all adults died as a result of violence. Among the Waorani of the Amazon, the homicide rate was over 60% (Yost 1981). Patterns of violence varied widely among indigenous peoples. Otterbein (2000) found that only a fraction of the societies in his cross-cultural study killed females captured in raids. On the other hand, the Gebusi killed men, women, and children with equal frequency. Boehm (1999) conjectures that many societies proscribe violence within the social group, but Kelly (1987) finds that violence with the residential community is common. Merely stating that these differences are "cultural" ignores the question of whether there is some underlying systemic reason for them.

We simulated two types of violence in a population of 100 agents: homicide committed during the act of theft, and violence committed in revenge for a previous transgression. The agents were divided into two equal social groups. In this simulation, we employed a version of "situational ethics," wherein an agent would share if its hunger relative to the maximum allowed before starvation was less than an altruism parameter A. An agent stole if its hunger was greater than its altruism parameter and more than the quantity (1 – G), where G was an aggression parameter. Both A and G were in the interval zero to one, so that agents with high A were likely to share and agents with low A and high G were likely to steal. The success of a theft depended on G and another parameter, F, which described the fighting ability of the agent. If the attacker had higher G and F, then theft occurred without fighting. If the defender had higher G and F, then no theft occurred. If the attacker was more aggressive (higher G) but had less fighting ability (lower F), then it died in the attack. If the attacker was less aggressive but had greater fighting ability, then the attacker killed the defender and took the defender's food. In this scenario, we did not make reputation a factor in mate selection.

Revenge occurred when an agent encountered another agent against whom it held a negative reputation. Here the agent with the higher fighting ability won the conflict. Whenever a killing occurred, whether during theft or by revenge, an amount equal to an agent lifetime was deducted from the interaction matrix element of all members of the victim's village who were collocated with the killing. This could result in a cycle of revenge, wherein one killing would be in revenge for a previous one, with the original cause of the dispute long forgotten.

The results are given in Table 2 and demonstrate that without excluding some major portion of the population from homicide and revenge, there is a significant probability of population collapse. It mattered less what portion was excluded as long as there were enough members in that portion to limit the total amount of violence.

Not all transgressions are serious enough to result in blood revenge. We studied tolerance before revenge and found that even small amounts of tolerance — less than what would be required to forgive a single theft — were sufficient to greatly reduce the rate of violent deaths. We also studied the effect of higher population density on the murder rate and found that, while violence did increase, its negative effects were overtaken by the positive effects of more frequent interactions between agents.

Ecological factors sometimes result in increased non-normative behavior within a population. The Ik of Uganda are an example of a population for whom the norms of sharing and group solidarity broke down when the traditional hunting grounds of the tribe were deemed off limits. In this case, family members kept food to themselves and stole from others; the spirit of cooperation almost completely disappeared (Turnbull 1972). We simulated this effect by reducing the amount of food that replenished the food centers in our environment and found that the result was a significantly increased rate of killing committed in the act of theft. A comparison of the results of simulations to ethnographic observations is given in Table 3. These and other results of simulations of violence and revenge in egalitarian societies are discussed in more detail in Younger (2005b).

TABLE 2 Results of excluding different segments of the population from violence and revenge[a]

Subpopulation Excluded from Violence	None - -	- - Fe	- G -	Fa - -	- G Fe	Fa - Fe	Fa G -	Fa G Fe
Survival rate (%)	35	40	30	10	60	35	55	90
Deaths due to old age (%)	70	75	68	71	71	71	71	74
Deaths due to hunger (%)	4	13	4	3	16	17	18	19
Deaths due to violence (%)	12	3	11	8	3	3	4	2
Deaths due to revenge (%)	14	9	17	19	10	9	7	5
Total mutual obligation	380	330	400	420	340	310	430	370

[a] Fa means that violence and revenge were forbidden within the family, G within the group, and Fe among females. The last column represents a situation where violence and revenge were permitted only against males of the other social group. Each entry represents an average over 20 runs, where only those runs that had a nonzero population at the end of the run were included in the average.

TABLE 3 Comparison of simulations of homicide and revenge with ethnographic observations

Simulation Result	Ethnographic Observations	Comments
Violence and revenge contributed substantially to mortality and reduced the overall survival rate of the population.	Violent deaths accounted for tens of percent of the total recorded deaths among the Copper Eskimos, Gibusi, Waorani, and other indigenous peoples.	Violence is a population control mechanism in some egalitarian societies.
Excluding significant segments of the population from violence and revenge improved the survival rate of the total population.	Kapauku excluded females from violence. There is little violence among females in Kunimaipa society. Some primate and human societies proscribe violence within immediate social group.	Many societies discourage violence among significant parts of the population.
Tolerance before revenge increased the survival rate of the total population.	Peaceful societies (e.g., Semai) have high levels of tolerance. Violent societies (e.g., Yanomomo) have low levels of tolerance.	Tolerance reduces the rate of revenge killing.
Increasing the population density increased the survival rate of the total population, even though revenge killings increased.	Keeley (1996) found that population density and the rate of violence were not correlated.	A higher survival rate in simulations is a result of more mating opportunities. Simulations omit control mechanisms that limit violence in real societies.
The rate of violence increased when food scarcity was introduced.	Scarcity reduced sharing within a group and, in extreme circumstances, increased antagonism and theft within the group. The Ik of Uganda are a particular example of theft increasing in times of scarcity.	Scarcity of food increased the rate of violence, consistent with ethnographic observations.

DISCUSSION

Simulation provides a useful methodology for testing various assumptions about relating normative behavior to small societies. In particular, rule-based simulations allow hypotheses to be tested in a systematic manner and the results compared to real societies. If the simulation agrees with nature, then there is support for the hypothesis. If there is substantial disagreement, then one must look at the underlying assumptions in the model to find the cause. While our simulations are very simple compared to even the “simplest” human culture, and while the detailed modeling of human behavior must cope with the fundamentally stochastic nature of social interactions, they may still provide a framework to help improve our understanding of how individuals and societies behave. In this sense, simulations are analogous to cross-cultural studies of real societies in that the conclusions are general rather than specific to one society. However, it is also possible to model a single society in detail, including realistic birth rates, food sources, and behavior patterns. Such a simulation of a Pacific society is in progress and will be described in a later report.

REFERENCES

Beliege Bird, R.L., and D.W. Bird, 1997, "Delayed Reciprocity and Tolerated Theft: The Behavioral Ecology of Food-Sharing Strategies," Current Anthropology 38, 49–78.

Boehm, C., 1999, Hierarchy in the Forest: The Evolution of Egalitarian Behavior, Cambridge, MA: Harvard University Press.

Kaplan, H., and K. Hill, 1985, "Food Sharing among Ache Foragers: Tests of Explanatory Hypotheses," Current Anthropology 26, 223–246.

Keeley, L.H., 1996, War before Civilization: The Myth of the Peaceful Savage, Oxford: Oxford University Press.

Kelly, R.L., 1995, The Foraging Spectrum: Diversity in Hunter-Gatherer Lifeways, Washington, DC: Smithsonian Institution Press.

Knauft, B.M., 1991, "Violence and Sociality in Simple Human Societies: Homicide among the Gebusi of New Guinea," Current Anthropology 28, 457–500.

Otterbein, K.F., 2000, "Killing of Captured Enemies: A Cross Cultural Study," American Anthropologist 41, 439–443.

Robarchek, C.A., and C.J. Robarchek, 1992, "Cultures of War and Peace: A Comparative Study of Waorani and Semai," in J. Silverberg and J.P. Gray (Eds.), Aggression and Peacefulness in Humans and Other Primates, New York, NY: Oxford University Press.

Turnbull, C., 1971, The Mountain People, New York, NY: Touchstone.

Yost, J.A., 1981, "Twenty Years of Contact: The Mechanism of Change in Huao (Auca) Culture," in N.A. Whitten (Ed.), Cultural Transformations and Ethnicity in Modern Ecuador, Urbana, IL: University of Illinois Press.

Younger, S.M., 2003, "Discrete Agent Simulation of the Effect of Simple Social Structures on the Benefits of Resource Sharing," Journal of Artificial Societies and Social Simulation; available at http//www.soc.surrey.ac.uk/JASSS/6/3.

Younger, S.M., 2005a, "Reciprocity, Sanctions, and the Development of Mutual Obligation in Egalitarian Societies," Journal of Artificial Societies and Social Simulation; available at http//www.soc.surrey.ac.uk/JASSS/8/2.

Younger, S.M., 2005b, "Violence and Revenge in Egalitarian Societies," Journal of Artificial Societies and Social Simulation; available at http//www.soc.surrey.ac.uk/JASSS/8/4.

DISCUSSION

Computational Social Theory

(Social and Conflict Dynamics, Saturday, October 15, 2005, 9:45–11:45 a.m.)

Chair and Discussant: *Irving Birkner, The University of Chicago*

Emergent Agents and the Simulation of Political Unrest: Application to Palestinian Political Coalitions

Irving Birkner: This session starts off with Lawrence Kuznar from Indiana Purdue University in Fort Wayne.

Larry Kuznar: Thank you very much. As always, I'm very glad to be in this community; I'm relatively new to the mod-sim community. Today, I'm going to talk about emergent agents, which seems to be the flavor of the conference, in the way that we all grapple with creating truly emergent agents in our various models. It's applied to the Palestinian case of the emergence of political coalitions, although today's presentation is very much on the methodology of how you would go about doing this. At the end, I'll have a few rough comparisons to make the Palestinian case, but my political science colleague, Jaime Toole, is digging into some databases to see if we can more closely tie the emerging coalitions to actual Palestinian coalitions. That will give further validation in the future, or refutation, as the case may be.

[Presentation]

Birkner: We'll take some questions for Professor Kuznar in just a second. I have a quick and probably very simple question. I saw some correlation at the end between what you predicted in Islamic Jihad and Hamas and PLO. Does your model account for differences between affiliations to Islamist groups, or with the data, do you see different kinds of people affiliating themselves with Islamist groups or nationalist groups? Can you say that for certain, or are you lumping them all into these groups together?

Kuznar: Excellent question. Right now we're just lumping them together. It's an issue of risk-proneness: how radical is the group in its politics? Of course, you could be radical in a nationalist way, radical around a religious ideology. However, that's what my colleague, Jaime Toole, is digging into databases to see if we can go back to the data outputs from the model and find correlations of the sort that you're suggesting, to see if we can have even closer correlations than the very rough statements I've made here today.

Lars-Erik Cederman: This talk was very interesting. I want to ask you about the sigmoid functional form. We seem to be talking a lot about functions today.

Charles Macal: I do have an equation for that, if you're interested.

Cederman: It seems that you are focusing on undulations around some kind of overall trend. I mean, you're looking at CDFs, cumulative density functions. Looking at rank is the flip side of that. The economists who have looked at the income distribution have been debating whether it's all about some kind of Pareto distribution or a log-normal distribution. I'm surprised that you are focusing on the undulations or, as I say, the variations around that main trend. Can you comment briefly on what happened to the Pareto distribution and these other usual suspects?

Kuznar: I'll do two things. I'll comment on that, but we haven't systematically looked at what density functions best match these empirical distributions. They could very well be Pareto. There's a lot of log-normal new trends and whatnot. However, the key here is that it's not the overall trend that matters; it is the oscillations. That's where you get the interesting behavior, and that's what our empirical studies are showing. It's those oscillations, those upswings that seem to really matter. I love big trends and whatnot, but in this case, the oscillations seem to be more important, so that's why I developed this particular functional form. I'll take another question, but I'd be happy to talk about the functional form, too.

Carl Johnston: Carl Johnston from George Mason. On that very narrow topic, I've found that sigmoidal-type utility functions work very well in healthcare in terms of risk-seeking and risk-avoiding behavior, and insurance-buying behavior. I think there's something there, and I was surprised — I mean, you're evidently doing original work here; you're not relying on some other person, but I think somebody ought to be formalizing this type of sigmoidal shape. I hope it's you.

Kuznar: I'll look forward to talking with you afterward. Yes, we actually came up with an exposigmoid function that we have a methodology, an algorithm for fitting to data. Basically, it's *e* raised to a function, which is in part a linear function, but also a trigonometric polynomial. We use different Fourier transform techniques to estimate that.

Robert Reynolds: I have a question about your coordination game. Basically, the way you give it, it's a one-shot thing. In terms of adaptive systems, though, you're doing it over a sequence of time steps, and therefore you get into what's called the iterative game. Do you do the iterative game here?

Kuznar: Yes, it's very much iterative, but there's no memory.

Reynolds: Okay. That's what I'm getting into because once you get the memory, it's something else. That's important because in terms of your coalitions, I see how your game coalitions can be produced, but in an iterative game over time, they also can be taken apart, or the players can withdraw. I don't see any of your coalitions breaking down such that the producement somehow loses favor down in the base and collapses. If you went to the iterative game where there is a memory and bring in a third operation, you have basically cooperate, defect, and then, rather than just defect, deconstruct this alliance, or let's break this alliance right now. I think alliance-breaking is an important thing; otherwise, you get these coalitions that remain a bit too long.

Kuznar: I'd like to comment on that. We don't have that kind of static development with the coalitions. My math colleagues thought it was going to happen. I said, "Trust me. It won't." And it didn't. Actually, they're very dynamic because everything revolves around a Nash optimum choice of mixed strategy, but then the probabilities with which they'll make those choices change. Plus, the agents themselves evolve through time. Their preferences evolve through time, so you're going to have agents who are very radical at one time, but if through some sequence, they try and get into one of those concave sections of a curve, they lose that radicalism, and that drives a lot of the dynamism. This means that there was a lot of dynamism where we didn't have black holes of coalitions that they all fell into. What you're suggesting would be another thing to model — individuals trying to rebel against the coalition itself, which would be a fine thing to do, but it's another complication that you'd have to go to.

Kostas Alexandridis: Kostas Alexandridis, from Purdue. I don't know how important it is for the research you're doing, but there is an apparent discontinuity on realism in many of the high-income professions. Because of the poor educational system in Palestine, they are educated in Europe; they bring back to the culture a European perspective, as opposed to the lower income and lower professions, which are grounded in the reality of the area.

Kuznar: Yes. Let me respond in this way. Robert Axelrod argues for the KISS approach: Keep It Simple, Stupid. You know, let's have simple models. I recently wrote an article for *Social Science Computer Review* on hi-fi models (to be published). There is this tension, isn't there, always as to how realistic our models should models. How simple should our models be? This was a KISS approach. I've done hi-fi models as well. Let's say you actually wanted to use this for policy. In that case, you would probably want to go to a more high-fidelity-type model, in which case you would start programming in those sorts of specifics. If your concern is wealth coming from expatriates, that's one thing to program, or, as you're suggesting, new ideas, new mental models, in which case you might want to go to an interpretive agent or something like that to see how contact with the outside might change a worldview, and therefore the options an agent would want to consider.

But I would say, "Okay, give me some tasks for which we need that high fidelity, and then let's try to make a more realistic model." I don't know if that makes sense. Based on your talk, you are engaged in more high-fidelity-type models and decision-making among farmers and whatnot.

Birkner: Okay. I think we'll stop it there and begin our next presentation.

Kuznar: Thank you.

Multiple Attitude Dynamics in Large Populations

Wander Jager: Good morning to all of you. My name is Wander Jager, University of Groningen, the Netherlands. In Groningen, we have a nice group of people working with social simulations, and our main emphasis is how to translate existing behavioral theory into agent rules. I'm trained as a social psychologist. We have a lot of theory, of course. The trick is how to simplify the abundancy of social theories in a set of simple rules for agents and combine that with empirical data. Currently, we have projects underway on stock market dynamics, self-organization and themes, inflation diffusion, product markets dynamics, and crowd and riot control. There's a huge variety of projects we are doing.

Today I do not want to talk about those projects, although I would be glad to send you some information. If you are interested, send me e-mail, and I can present you with some material on that. Today I want to focus on a project together with Fred Amblard of the University of Toulouse on what we are doing on multiple attitude dynamics in large populations.

[Presentation]

Unidentified Speaker: It looks like things are gravitating toward your edges.

Jager: In a way it does. We have on one opinion dimension and some more polarization than on the other dimension. That's why a single run, of course, is not representative. What we did is conduct a correlation analysis on a number of runs, and we saw and the correlation indicates it ranges from –1 to +1, indicating how strong position A is related to B. What we find typically is that several runs show that the attitudes are uncorrelated, but we also see a number of runs where the correlation is in the range of almost 0.4. You know, a correlation of 0.4 is generally in social sciences accepted as "Whoa, we really got an effect. That means something." Here, though, it's purely coincidental. We thought that was very interesting. More interesting, of course, would be to find out what happens if we have central processing of A and peripheral processing of B. Here we have the same situation, but on B is the peripheral processing…. Who has an idea what will happen? Can you guess?

Unidentified Speaker: Did they come to correlate?

Jager: Right. And we see what happens. This is very simple, of course. We see in this case on the top right in the figure that a line is skewed down, indicating that we've got almost a pure effect negative correlation.

[Presentation Concludes]

Birkner: Thank you. It was very interesting. I would ask, "What extent does emotion play in the process that you're modeling?" I think of a couple things in terms of Converse's work on mass belief system where he shows that the public becomes more and more fragmented in how coherently they can articulate a mass belief system. It also brings to mind Grady and Snyderman's work about affectivity heuristic, where they end up saying, which is, of course, a form of emotion, that we all sound very opinionated; we all sound like we know the details on quite complex policy issues when we're not paying very much attention. And so they talk about a kind of triangulation based on affectivity. I'm wondering the extent to which you think it's necessary to explicitly model emotion as part of this process and whether some kind of field of actors, symbols, significant objects, beliefs, and so forth that is held together by varying levels of emotional intensity might be a useful way of approaching it.

Jager: Well, I think that emotion plays a very important role in these kinds of processes. The question is, "Is it necessary to include emotion explicitly in these kinds of models?" I think if you look at existing research on these processes, it may contribute to formalizing the model.

For example, I was talking about the *U* and *T* and how easy you contrast versus assimilate another's position. It's well known from field studies that in times of crisis the difference between *U* and *T* becomes smaller, so there's a much faster transition from assimilating to compressing. In previous work, we formalized this in a model. Typically, we observed a strong polarization effect, where first of all the population tended toward assimilative effects, heterogeneous opinions. If we simulated some kind of crisis resulting in this effect, we then saw a strong polarization. I think in a very simplistic way this would express how emotion or an incident would affect these kinds of dynamics.

Robert Reynolds: Bob Reynolds, Wayne State University. A very interesting talk. I have a comment. Often there is more than one political party. If you took another meta agent who is in fact an adversary to Jacques Chirac and placed him into the scenario, would you get a dampening or amplification effect? In other words, would you get this bipolar result sooner or later or not at all? Have you looked at that?

Jager: No, we haven't looked at that. First of all, I'm afraid it would become a second book.

Reynolds: A second book, okay.

Jager: It would mean a lot of experiments, of course, but, indeed, I think this may provide a testing bed for experimenting with different meta actors.

Reynolds: Right.

Jager: In addition, I think, this is a very simple approach in that we have only local actors versus meta actors, although there's a lot in between. You can also distinguish between relative power of agents, the susceptibility for change, and things like that. I think that's a natural thing to develop in these kinds of simulations.

Reynolds: That gives a lot of opportunity for symbiosis between the different meta agents.

Jager: Yes, absolutely. It's a very good idea indeed.

László Gulyás: I have a quick comment that might lead to your third book. You mentioned in the last slide that you are planning to conduct experiments on different kinds of meta works, but all of those seem to be static ones, and there's this interesting work by Cohen, Axelrod, and Riolo on iterative games on different kinds of networks. What it shows is that if you start varying the network during the run, interesting effects happen, like convergence disappears and such. I wonder whether that would be somewhat similar here if you start adding noise, for example, or changing the neighborhood structure.

Jager: That's a very interesting point indeed because you touch on the issue of the emergence of subcultures. It might be very well the case that people group together on an issue that they find very important. Then they are confronted with diversity on other issues, but because of peripheral processing, they might tend toward growing together on that dimension as well. Although it's not related to the first dimension, they tend to cluster also on other dimensions. So you may have a group of people being quite similar on various dimensions that basically are unrelated.

For example, I have a group of friends, and we talk about some things that are very important, but I'm also affected by their choice of cars, their preferences for beer, or whatever. All these things might grow together in some kind of clustering of joint behavior on various fields. I think that would be very interesting using this kind of work to study the emergence of cultures.

Birkner: Are there any other questions? Okay. Thank you.

Contesting Hearts and Minds: A Baseline Model of Occupation Dynamics in Military Occupations

Birkner: Our next presentation is by Keven Ruby from the Department of Political Science at The University of Chicago.

Keven Ruby: I'm a graduate student in the Department of Political Science at The University of Chicago, and I've been working on this project with colleagues at Argonne National Lab who are probably somewhat more well-known in this community than I am: Veena Mellarkod, David Sallach, and Chick Macal.

Basically, we've been interested in the idea of solidarity dynamics in conflict situations; our particular frame of reference has been occupation dynamics, although it's been pointed out to us that many of these dynamics would play out in other political situations as well. One of the reasons why we wanted to model occupation dynamics is that oftentimes you have a situation, in particular where insurgencies develop, where you have different competitors for local authority, or authority among, so to speak, the occupied public. Often we don't know what the actual effect of strategies and policies is based on the observed support of the public. That's because we are simultaneously operating an external process of support, while also operating an internal process of valance, as we call it, toward these authorities. These may be motivated by different factors. In addressing the issue in this way, we build on the work of, certainly in political science, on Timur Kuran. Also in the case of agent-based modeling, where there's been a growing interest in the question of insurgency, ethnic violence, and the like, we're building on the work of folks such as Lars-Erik Cederman, Joshua Epstein, and Ian Lustig.

[Presentation]

Birkner: We'll take some questions for Keven now.

Ana Carrie: I know you've got 45,000 patterns already, but here's an idea. You have resource asymmetry in favor of the occupier. Have you thought about knowledge asymmetry in favor of the people themselves? Have you thought about the fact that the resistance might have better information because they're embedded rather than being separate?

Ruby: Yes. In fact, we've had many discussions in implementing various additions to the model, and that's definitely one of them. Actually, it makes it very interesting because it effectively changes the nature of the asymmetry. Basically, one of the aspects of the dynamic strategies, when we implement that, is the idea that the authorities have a much poorer vision of what the alignments are, for example, of agents, while the counter-authority knows exactly who it is who is supporting and will be more targeted and therefore more effectively sway outcomes.

John Sullivan: As a follow-up to the last question, it would also seem to me that, based on one of the assumptions that you made, which is 'the enemy of my enemy is my friend,' that

that's rather tentative, and that there may be reluctance to share information between various cells. My assumption is you have various cells operating more or less independent of one another here. Is that correct?

Ruby: That's actually a good question. It would be useful to have a metric for just how interconnected all of the networks are. But the way the network generation mechanism works is that it leads to, I believe, no isolated sets of friends and enemies, so that there's always an overlap in the social network, which is the way that the mechanism produces the effects. Of course, 'the enemy of your enemy is your friend' is a heuristic simplification that we've used in this baseline version of the model. But, in fact it would be interesting to hear any kind of refinements or things that you might have to suggest.

Kuznar: I have a question about your valuing of rewards and punishments, and I've got a justification of sorts for it. Did you examine the differential effects of rewards versus punishments? It seems that they're both on the same scale, which goes back and forth. The reason that I bring this up is, first of all, for the same kinds of rewards and punishments we would expect prospect theory effects, like loss aversion, to occur. The other issue is that there are some punishments that are going to be qualitatively in another thing. It's one thing to say, "Hey, I'll give you a job." It's another thing to say, "I'm going to torture your family." I was wondering if you looked at different effects of reward versus punishment.

Ruby: Not yet. That is also, along with the information asymmetries, one of the things that we're working at incorporating. I mentioned at the beginning of the presentation that we're using rewards and punishments as a bit of heuristic, including things like long-term rewards, such as building schools and infrastructure. There are also short-term rewards and even direct cash payments. All quite reasonably have different effects, and exactly as you say, with punishments as well. So that's something for the next iteration of the model when we add layers of complexity.

Birkner: Are there any other questions? Thank you then.

Simulations of Egalitarian Societies with Comparison to Observations

Stephen Younger: Good morning, everyone. I'd like to talk with you about some simulations I've been doing over the last couple of years on egalitarian societies, which are a little different than what we've talked about so far.

[Presentation]

Lee Hoffer: Can you tell us how you define egalitarian societies in the societies you're looking at as data for your simulation? This is fascinating work, by the way. I really find it interesting.

Younger: Lacking any central leadership, all the agents make their own decisions.

Hoffer: No chiefs or big men or anything?

Younger: No chiefs, no big men, no. Just all by themselves.

Hoffer: Have you looked at how — one of the things in the exchange, at least among dealers and customers, is that when dealers share drugs with users, there's this transmission of status; status comes back to the dealer. I know in anthropology that this is popular, especially with Salens, that there is a status increase for those individuals who give freely to their underlings, who can't, of course, repay that gift. I was wondering if you had looked at that. It would be interesting to work into your simulation.

Younger: Yes, I have looked at that, and I call it prestige, but I've looked at it in a Trobriand Island, Malanovski way of selecting a leader. Typically, it's a complicated story with regard to leadership, but I'm finding that it doesn't make a lot of difference.

Armando Geller: Armando Geller, the Swiss Military Academy. I liked your paper. I have one question and one comment. First, how did you get your agent rules?

Younger: I tried to create a set of rules that mimicked generic behavior in egalitarian societies. I view the simulations almost as equivalent to a cross-cultural study. So I extracted from a cross-cultural reading of egalitarian societies those rules that seemed to fit; no more than that, I'm afraid.

Geller: That's what I expected, actually. Have you ever thought about trying to simulate, even if these are small group societies, groups like in the Kosovo or the Pashtuns in Afghanistan, where they have very fixed codings of their living together, like the Pashtun Wally or the Konun in the Kosovo?

Younger: Yes. Theoretically, that is the goal of my research, or at least some people think it is. The reason that I started with egalitarian societies is there's a fair amount of data, and I wanted to start very simply. Because Pashtun Wally, for example, is a set of normative rules, you could think of using this method to describe that. At some point, I want to do that, but I'm building my way up to that.

Kostas Alexandridis: This is real exciting work, and I was wondering, in corroborating evidence from island biogeography, especially in the Pacific with the Polynesian expansion and things like mainland/island relationship in metapopulation dynamics, shifting those scales you think will help understanding more about those populations?

Younger: I think so, and I should say that I'm also with the Department of Anthropology at University of Hawaii. There is a lot of information that can be used, and there is a richness in the Pacific, so there are some islands — my favorite isolated island is Pukapuka in the Cook Islands, which is all by itself, as is Rapi nui, Easter Islands. There are others in the Marquesas where there are a number of islands within a quick canoe trip of one another, and there are tremendous contrasts between them. So, yes, there is a rich fund of information there.

Gabriel Istrate: Gabriel Istrate, Los Alamos. Just a quick comment to the previous question. Recently, I heard a talk by another Los Alamos colleague, Ed MacKerrow, on simulation of Pashtun societies.

Younger: I should have mentioned that. Ed's done a lot of work on that….

Istrate: There was a one-day workshop in Santa Fe, part of LACSI 2005. LACSI is Los Alamos Computer Science Institute.

Kuznar: Wonderful paper. One thing wasn't clear to me. Was there any link between the resources an agent captured and their ability to share. Was that part of the model at all?

Younger: What do you mean by 'the resources captured'?

Kuznar: How did the agents get resources so that they could share them?

Younger: They went to the food center. I should have mentioned that. They went to the food center, and they ate, and then they picked up some stuff to carry around. When they got hungry, they would eat it, but if they weren't particularly hungry, depending on their altruism, they would share it. If they were really hungry and they had low altruism, they'd steal it.

Kuznar: Okay, so there was some contingency on agents' ability to share and whether or not they … .

Younger: Yes.

Kuznar: Did you think of looking at the distribution of, if you will, wealth stores in food and their ability to gain reputation? That would be really key because, as you know, that's an underlying current in all of this. The reason why hunters in particular compete to become good hunters is so that they can, in the long run, use those reputations to gain more accesses to …

Younger: Yes. Now, one thing that I did that I didn't talk about is that I modeled a situation where the food supply went steadily down, like in a famine. I found that the level of violence went up, but it didn't go up as much as one might think because at a certain point, the agents just didn't have time, and nobody was carrying anything to steal. It was hand-to-mouth just finding food for yourself. There have been some studies of famine in small societies that indicate that, but it's much more complicated with humans.

Kuznar: Yes, I liked that issue about the warfare and the low resources. In political science, I believe, there's a resource mobilization theory that says, "It doesn't matter what your motives are. What do you have to conduct warfare?" I think that relates. I like that touch.

Younger: Yes. I did a number of different models of warfare based on wanting power, revenge, or greed, and actually it doesn't make any difference how you start the war. Once it starts, it takes its own course.

Craig Stephan: I found this a very fascinating talk. I had one question as to your model. Can your model be applied to nonhuman societies? I'm thinking particularly of chimpanzee societies, where you might have the same sort of interactions of sharing and warfare and the like.

Younger: Yes, there are some similarities. However, and Joanna [Bryson] knows this much better than I do, there are hierarchies in nonhuman primate societies that I haven't represented here. There are some similarities, like raiding and stealing of females and other sorts of things. So, yes, you could apply it. I haven't gotten around to that yet.

Birkner: All right. I think that seems like a good place to stop. We're going to break for lunch now.

Charles Macal: I'd like to thank the speakers in the last session, a very fascinating set of talks. And I'd like to thank Irving Birkner from The University of Chicago Center for International Studies for being the session chair.

Patterns and Actions

POWER-LAW-LIKE DISTRIBUTIONS: A PRACTICAL SURVEY

A. CARRIE,* Department of Economics,
University of Dublin, Trinity College, Dublin, Ireland

ABSTRACT

Unrelated datasets from biology, economics, computer science, and many other disciplines follow power law distributions, characterized by a straight line in the log-log rank-frequency plot. This universality, along with the tempting prospect of a common underlying generative process, has attracted significant research interest. Upon closer inspection, many of these datasets show slight or pronounced curvature. In light of this, several alternative distributions have been proposed in the literature. The lure of the power law, however, is extremely strong, and these alternatives are rarely fitted. This paper reviews these alternative distributions and fits them to a standardized collection of power law datasets. The practicalities of fitting these distributions are discussed. The hope is that presenting these distributions in a user-friendly and systematic format and testing them against some canonical datasets will facilitate their use within the power law literature.

Keywords: Power law distribution, DGX, lognormal distribution, maximum likelihood estimation, rank-frequency plot

INTRODUCTION

Many datasets that are described as following a power law (i.e., having a linear probability distribution in log-log coordinates) do, in fact, show some inconvenient curvature. There are several alternative distributions that allow for such curvature in the literature; however, it is not straightforward for researchers to utilize these because they do not share consistent notation or statistical methodology. Ideally, we should have access to a readily available toolkit of skew distributions that are straightforward to use and interpret and that can be easily compared with each other. In this paper, I undertake an informal survey of some of these skew distributions that I hope may contribute to the development of such a toolkit. This work is at an early stage, so references should be made to other sources before using any result included here. I first define the terminology and notation used throughout this paper. I then describe the statistical distributions, their functional forms, and how their parameters may be estimated from sample data. I work through the Beowulf dataset in detail as an example. This is followed by summary results for several datasets and finally by concluding remarks.

* *Corresponding author address*: Ana Carrie, Department of Economics, Trinity College, Dublin 2, Ireland; email: carriea@tcd.ie.

NOTATION AND TERMINOLOGY

Observations $x_1,\ldots,x_n$ are sorted from largest to smallest so that the subscript i corresponds to the rank of the observation, the largest observation being ranked first and the smallest n'th. The value of an observation is known as a frequency. The origin of this unfortunate convention is that in many cases, particularly in the early terminology-formation years, the values being studied were frequencies of occurrence. For instance, our first sample dataset will be one studied by Zipf: the frequency of words appearing in the text of Beowulf. Our vector of observations is also converted into the form $f_1,\ldots,f_m$ (frequency) and $c_1,\ldots,c_m$ (count), where c_i represents the count of observations equal to f_i, such that the sample size $n = \sum_{i=1}^{m} c_i$. That is, if our original vector of observations is (30, 5, 3, 2, 2, 1, 1, 1), then we will have a frequency vector (30, 5, 3, 2, 1), with each value appearing once, and a count vector (1, 1, 1, 2, 3). Summing the values in count (1 + 1 + 1 + 2 + 3 = 7) tells the total sample size. The use of "count" and "frequency" in this very specific manner is problematical, since these words are interchangeable in everyday speech. If you are familiar with using the function count ( ) in a summary calculation, this should help keep you oriented correctly. In general, it is advisable to read any use of this terminology very carefully to make sure of the author's intended meaning.[1]

THE DISTRIBUTIONS

Power Law

The probability density function (PDF) of the power law distribution is given by:

$$f(x) \propto x^{-(k+1)}, \tag{1}$$

where the parameter k determines the steepness of the slope. The probability mass function in the upper tail (PMUF) is

$$\overline{F}(x) \propto x^{-k}. \tag{2}$$

Taking logarithms of $y = x^{-k}$ gives $\log y = -k \log x$, which illustrates the trademark power law linear relationship in log-log coordinates (Adamic 2005). In principle, one could determine the distribution parameter k either by fitting a straight line to PDF data as given by Equation 1 or to PMUF data as given by Equation 2. In practice, the PDF does not yield reliable results, so the PMUF is used.

The power law can also be fit by using rank-frequency data. In a log-log rank-frequency plot, the parameter k is derived from the slope $-b$ by $k = 1/b$. In the traditional Zipf distribution, both b and k are equal to 1 (Adamic 2005). Note that there is no intercept term in the PMUF regression, since the line fitted must pass through (1,1), which is (0,0) in the log-log plot. For rank-frequency data, we fit a straight line with an arbitrary intercept, the fitted intercept giving us the scale of the object with rank 1.

---

[1] I strongly considered referring to frequency data by another name. However, the term "rank-frequency plot" is widely used, so it was preferable to explain and use the terminology.

Unfortunately, all of the above apply only to continuous data. Our data are discrete; in fact, in the Beowulf example and many others, the data take only integer values. Following the example of Bi et al. (2001), we can derive the discrete probability function (point-mass function) as follows:

$$p(x) = \frac{x^{-(k+1)}}{\sum_{x=1}^{\infty} x^{-(k+1)}} \tag{3}$$

or

$$p(x) = \frac{x^{-(k+1)}}{\varsigma(k+1)} \,, \tag{4}$$

where $\zeta(n) = \sum_{i=1}^{\infty} \frac{1}{i^n}$ is the Riemann zeta function (Weisstein 2005). We will use maximum likelihood to determine the distribution parameter *k*. When independent, identically distributed data are assumed, the likelihood function is

$$L(k) = \prod_{i=1}^{n} P(x_i) = \prod_{i=1}^{n} \frac{x_i^{-(k+1)}}{\zeta(k+1)} \ . \tag{5}$$

The log likelihood function has a simpler form:

$$l(k) = -(k+1) \sum_{i=1}^{n} \log x_i - n \log\left[\zeta(k+1)\right] \,, \tag{6}$$

or in terms of count-frequency data:

$$l(k) = -(k+1) \sum_{i=1}^{m} c_i \log f_i - \left(\sum_{i=1}^{m} c_i \right) \log\left[\zeta(k+1)\right] \ . \tag{7}$$

Lognormal/DGX

The most well-known alternative distribution is the lognormal. While the power law has a straight line in log-log coordinates, the lognormal is parabolic. The PDF is

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}x} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2} \right] \ . \tag{8}$$

The discretized form of the lognormal distribution, known as the discrete Gaussian exponential (DGX) (Bi et al. 2001), has this PDF:

$$p(x) = \frac{A(\mu,\sigma)}{x} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2} \right], \tag{9}$$

where *A* is a normalization constant given by:

$$A(\mu,\sigma)=\left\{\sum_{j=1}^{\infty}\frac{1}{j}\exp\left[-\frac{(\ln j-\mu)^2}{2\sigma^2}\right]\right\}^{-1}. \tag{10}$$

When independent, identically distributed data are assumed, the log likelihood function is

$$l(\mu,\sigma)=n\ln A(\mu,\sigma)-\sum_{i=1}^{n}\left[\ln x_i+\frac{(\ln x_i-\mu)^2}{2\sigma^2}\right]. \tag{11}$$

By substituting count and frequency data, the log likelihood expression becomes

$$l(\mu,\sigma)=\left(\sum_{i=1}^{m}c_i\right)\ln A(\mu,\sigma)-\sum_{i=1}^{m}c_i\left[\ln f_i+\frac{(\ln f_i-\mu)^2}{2\sigma^2}\right]. \tag{12}$$

The DGX reduces to the power law as $\mu\to-\infty$.

Stretched Exponential

The stretched exponential distribution is a generalization of the exponential distribution. The PDF is defined as

$$f(x)=c\left(\frac{x^{c-1}}{x_0^c}\right)\exp\left[-\left(\frac{x}{x_0}\right)^c\right], \tag{13}$$

with the cumulative distribution function (CDF) being

$$F(x)=P(X\le x)=\exp\left[-\left(\frac{x}{x_0}\right)^c\right] \tag{14}$$

for $c\le 1$. When $c = 1$, this reduces to the exponential distribution (Laherrère and Sornette 1998).

The stretched exponential produces a straight line when the natural logarithm of the rank is plotted against observed values raised to the power *c*:

$$x_i^c=-a\ln i+b. \tag{15}$$

The three parameters of the distribution are *a*, *b,* and *c*, with $x_0=a^{\frac{1}{c}}$.

The authors provide no algorithm for fitting the stretched exponential. Thus far, the simplest method I have found is one of brute force. Allow c to take each of the values in (0.001, 0.002, …, 0.999, 1.000), or the required search precision, and proceed to fit the linear model specified in Equation 15 to the vector of observations $x_1^c,\ldots,x_n^c$. Choose the value of c that corresponds to the highest regression R^2, and a and b are then obtained from the corresponding linear model.

Parabolic Fractal

The parabolic fractal is another second-order polynomial extension of the linear power law, but while the lognormal is a parabola in log-log frequency-count, the parabolic fractal is a parabola in log-log rank-frequency:

$$\log x_i = \log x_1 - a \log i - b(\log i)^2 . \tag{16}$$

When $b = 0$, this reduces to the power law. Since a concave parabola has a maximum value, the theoretical maximum observation (regardless of sample size) can be calculated as follows:

$$x_{max} = x_1 e^{\left(\frac{a^2}{4b}\right)} . \tag{17}$$

The parabolic fractal can be fit by using linear regression on log i and (log i)^2.

A future task is to develop discretized versions of both the stretched exponential and the parabolic fractal so that they can be directly compared with the discrete power law and DGX.

Other Distributions

This is not an exhaustive list, and new distributions are being developed all the time, such as the double Pareto, which has two straight-line segments connected at a transition point (Mitzenmacher 2003) rather than a single straight line as in the standard Pareto/power law.

SAMPLE DATASET

Beowulf, one of the earliest surviving poems in English, was a source text for Zipf's study of the frequency with which words appear in the written language (Zipf 1965). The text of Beowulf was obtained from Project Gutenberg, and a word count list (concordance) was prepared, the start of which is shown in Table 1.

Table 2 has word frequencies in the first column and the number/count of words that appear with said frequency in the second column. There are 1,611 words that appear only once in the text, and the most common word (THE) appears 1,587 times. Rank is also given for the highest-ranked observations:

TABLE 1 Concordance of Beowulf

Frequency	Word
1	ABANDONED
1	ABEL
2	ABIDE
1	ABJECT
3	ABLE
4	ABODE
6	ABOUT
2	ABOVE
1	ABROAD
2	ACCURSED

TABLE 2 Frequency, count, and rank data for Beowulf

Frequency	Count	Rank	Word
1	1,611		
2	548		
3	293		
4	180		
5	115		
6	93		
7	61		
8	49		
…	…		
163	1	8	HIM
222	1	7	FOR
229	1	6	WAS
276	1	5	WITH
321	1	4	THAT
408	1	3	HIS
636	1	2	AND
1587	1	1	THE

A natural first step in our exploration is to graph the count-frequency data in both linear and logarithmic scale (Figure 1). The log-log graph on the right suggests a linear relationship, albeit with rather messy data for high-frequency words. Figure 2 shows that this plot is not suitable for curve fitting. Although the values obtained for *k* will not be correct for our discrete data, by way of illustration, Figure 3 shows the CDF and rank-frequency plots.

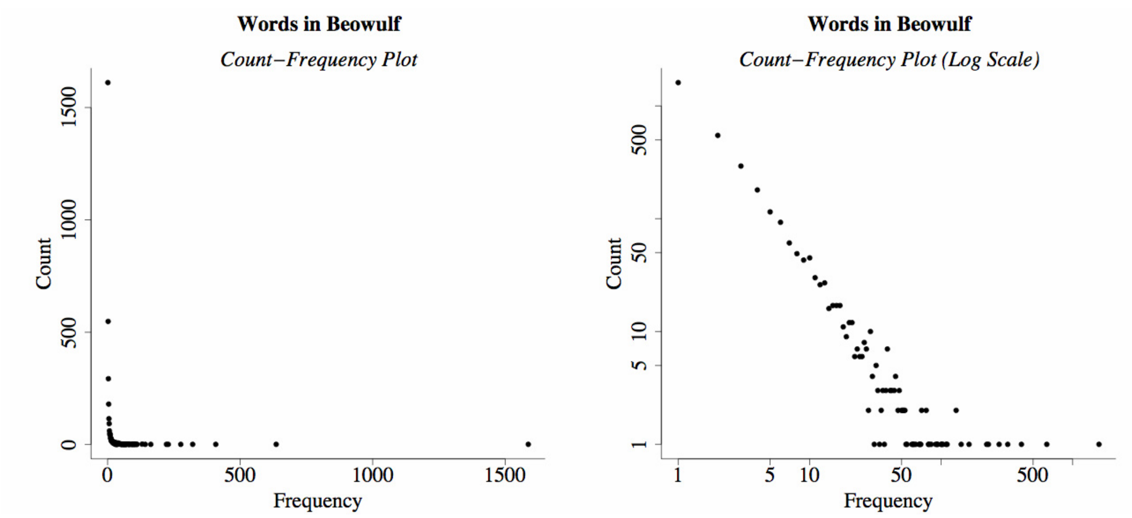

FIGURE 1 Count and frequency data in linear and logarithmic scale

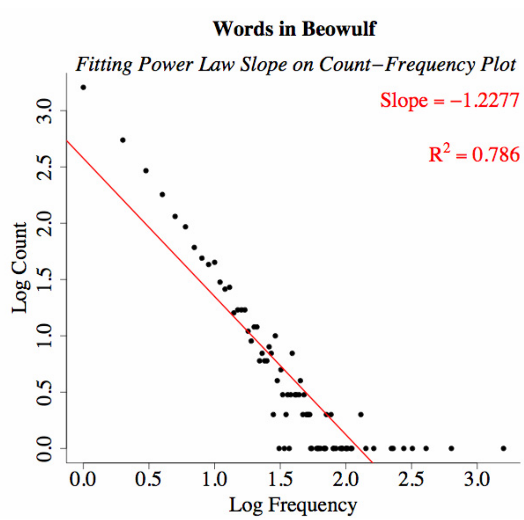

FIGURE 2 Fitting power law slope on count-frequency plot

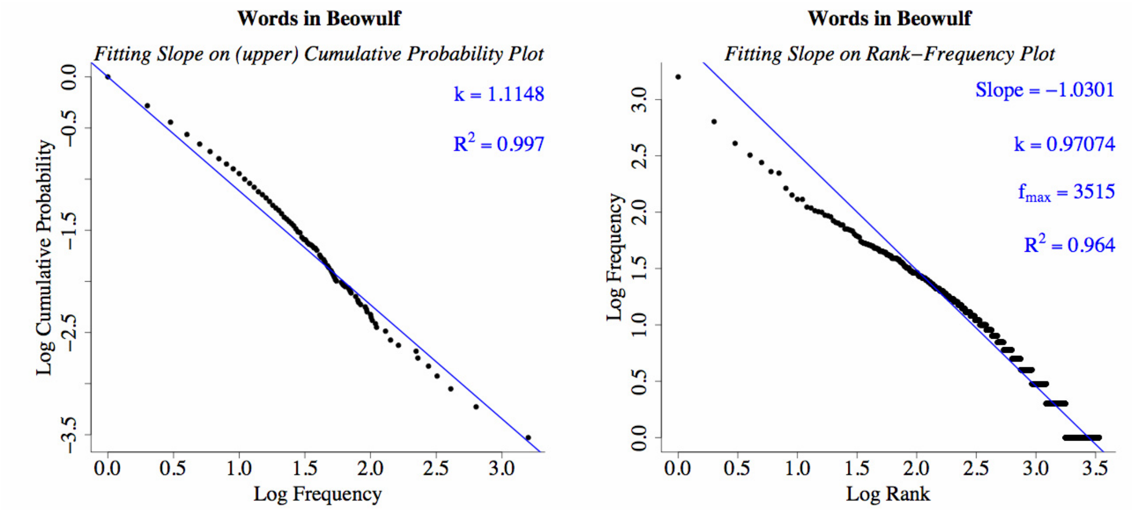

FIGURE 3 CDF and rank-frequency plots

We can see some evidence of curvature in the CDF plot, and even more in the rank-frequency plot. We will now calculate the power law distribution parameter k and the DGX distribution parameters μ and σ by maximizing the respective log likelihood functions. To compare these two distributions, we can define an error statistic (denoted *ERR*), which is a straightforward extension of the mean squared error (MSE):

$$ERR = \sum_{i=1}^{m} \frac{[nP(f_i) - c_i]^2}{m}. \tag{18}$$

We see that the DGX has a much lower error statistic; it is 12, compared with 358 for the power law. Since the DGX is, in effect, a generalization of the power law, this is to be expected.

Moving on to the stretched exponential and parabolic fractal (Figure 4), we will be able to compare them with each other, but not, for now, with the discrete power law and DGX (Figure 5). We can define an error statistic for rank-frequency data by

$$ERR = \sum_{i=1}^{n} \frac{[F(i) - x_i]^2}{n}, \tag{19}$$

where $F(i)$ is the predicted frequency for the observation of rank i.

For this dataset, the stretched exponential and parabolic fractal give similarly shaped fitted curves and similar error statistics. We see that both curves miss the handful of highest-ranked observations by a considerable amount. This may be a feature of the rank-frequency plot, which has n data points and thus places more emphasis on common, small events; in the PDF plot, these small events are aggregated together.

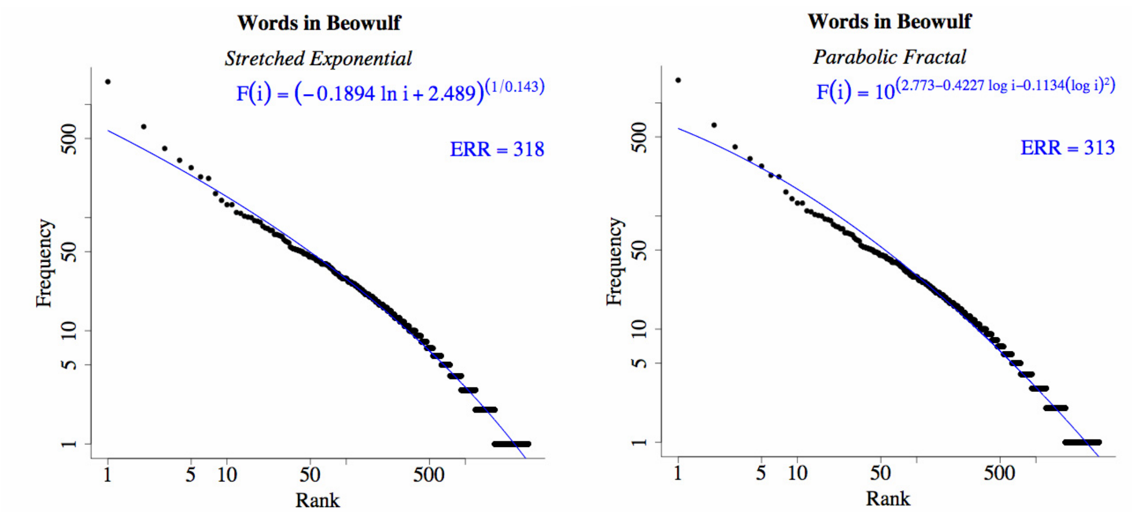

FIGURE 4 Stretched exponential and parabolic fractal

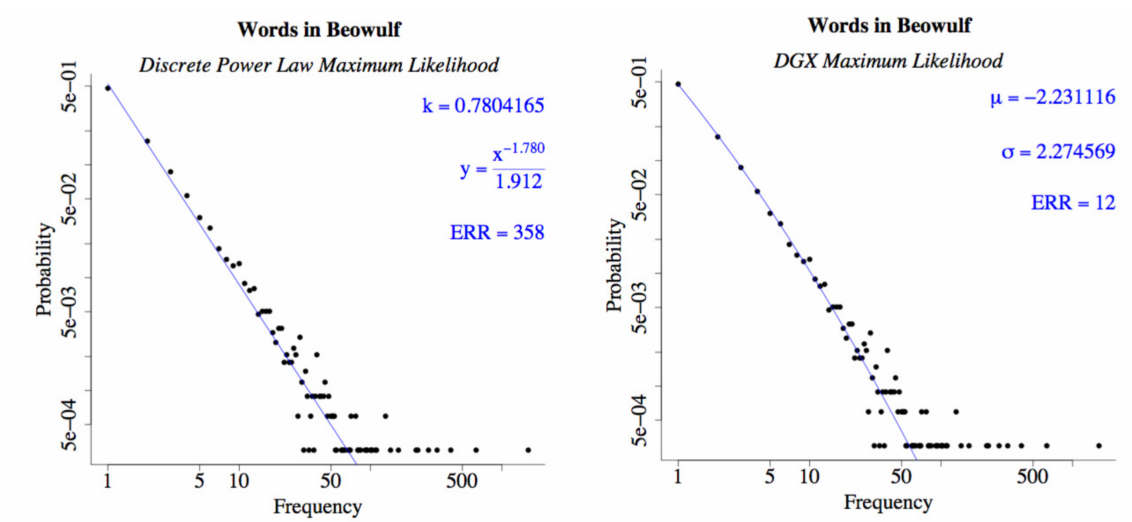

FIGURE 5 Discrete power law and DGX

OTHER DATASETS

Genera and Species of Snake

The number of species in a genus, for a family of plants or animals, has a skew distribution. I present two datasets here. One is from Yule (1925), which was quoted from an earlier work by Willis, which collated the data from the *Catalogue of the Snakes in the British Museum* by G.A. Boulenger, published in 1893 (Figure 6). The other is an updated version with 2005 data (Uetz and Heidelberg 2005) (Figure 7). There are 293 genera and 1,475 species in the 1893 dataset, and 463 genera and 3,002 species in the 2005 dataset.

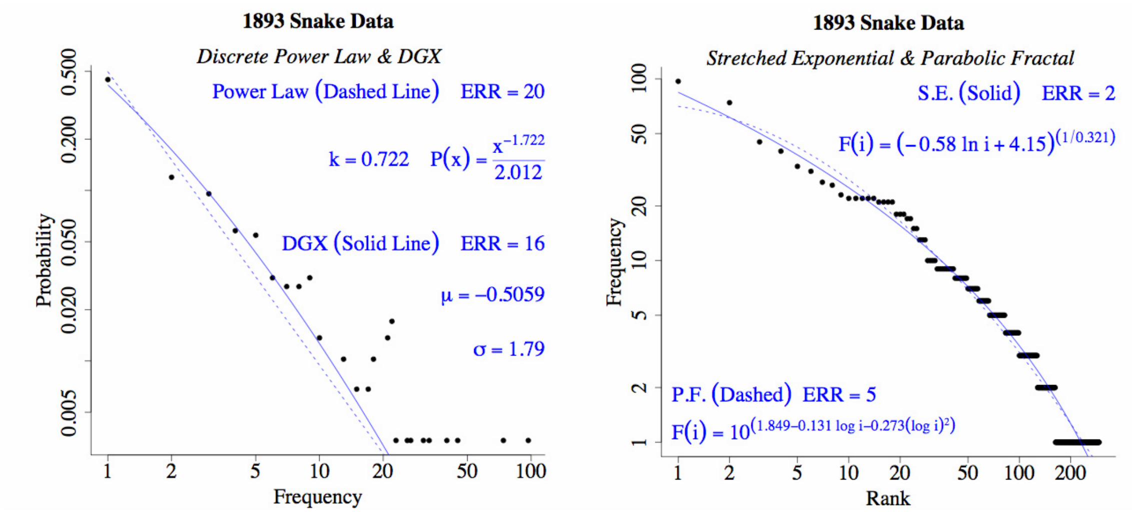

FIGURE 6 Year 1893 snake data

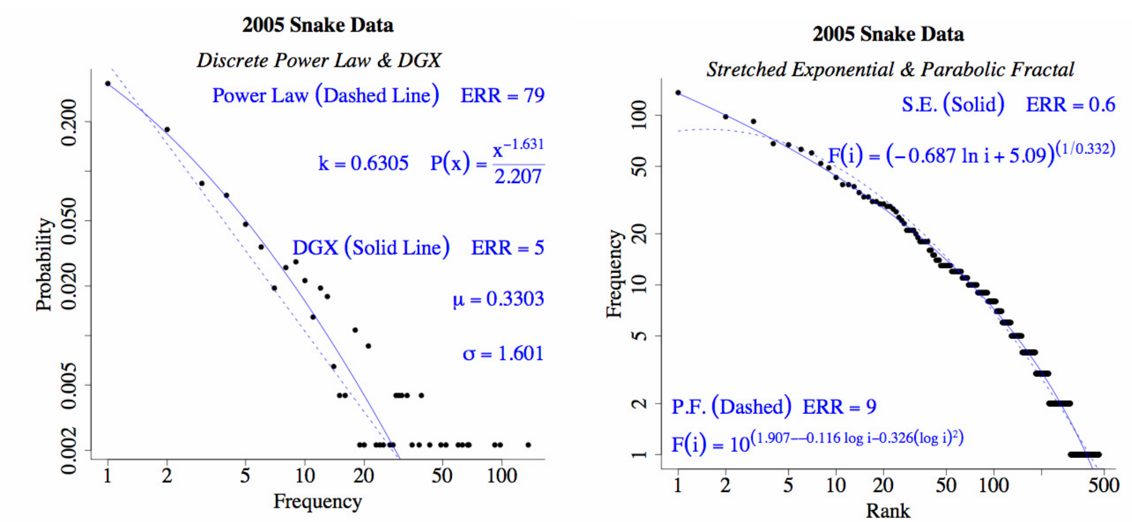

FIGURE 7 Year 2005 snake data

For both 2005 and 1893 data, we see that the DGX has a lower error statistic, as we expect. The stretched exponential is also a better fit in both cases, and in the 2005 data, it seems to match even the largest events. The parabolic fractal, in addition to having a poor fit, also has a positive coefficient for log(i) in the 2005 data, which violates its specification.

U.S. Cities

Here we look at the distributions of population in U.S. cities with more than 100,000 people. The DGX again outperforms the power law. The parabolic fractal in this instance has a slightly lower error statistic than the stretched exponential, but this is probably not a significant difference (Figure 8).

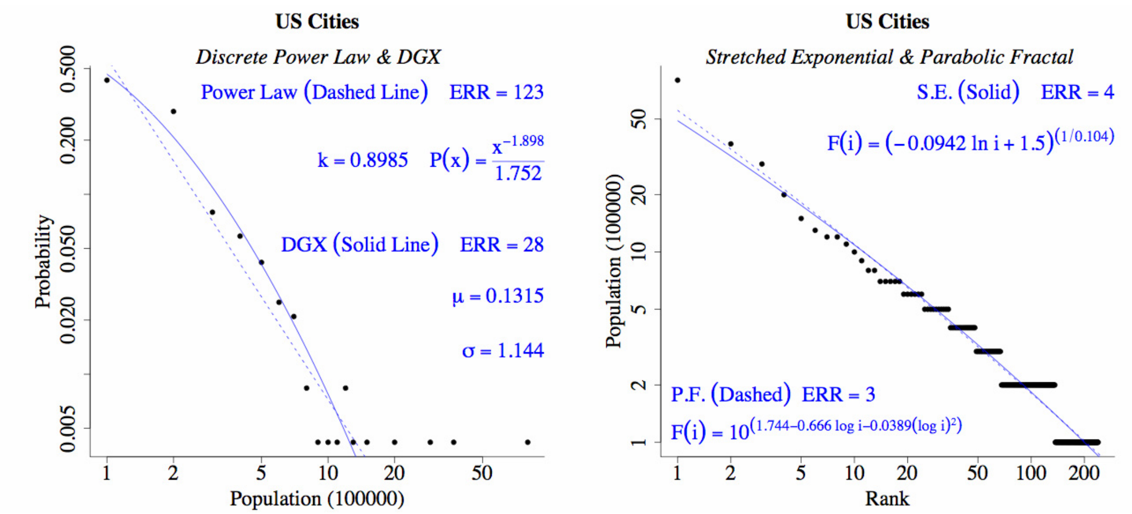

FIGURE 8 U.S. city sizes

Isle of Man Surnames

The DGX has an unfortunate quirk that means that for certain data, the normalization constant takes a very long time to converge. This means that the computations can be prohibitively time-consuming. Hence, for data pertaining to the distribution surnames of families living in the Isle of Man, we show only the stretched exponential and parabolic fractal (Figure 9). Again, the parabolic fractal has a negative coefficient for one of its terms, which is not valid according to its definition. Visually, the stretched exponential seems to fit both extremes, but it misses the curvature in the middle of this dataset.

DISCUSSION AND CONCLUSIONS

With so much academic interest in power laws, much more research is needed on the probability distributions that describe skew data, including the development of standardized criteria for discriminating between alternative distributions. Many of the conventional tools are based on a distribution having a finite mean or following a Gaussian error distribution; thus, they are not helpful for dealing with skew data. Because this research interest is interdisciplinary, consistent terminology and notation are all the more crucial. Distributions tend to be invented in response to a particular research problem and so have “baggage” from the academic or industrial realm in which they arose.

The DGX and discrete power law were fit by using maximum likelihood, and their distribution functions explicitly are acknowledged the discrete nature of the data. The stretched exponential and parabolic fractal were fit by using linear regression, implicitly assuming continuous data, and they were fit in the rank-frequency plot rather than a frequency-count or frequency-PDF plot. It is not clear at this point whether these two approaches will turn out to be complementary, each highlighting different and useful aspects of the data, or whether one will emerge to be “correct.”

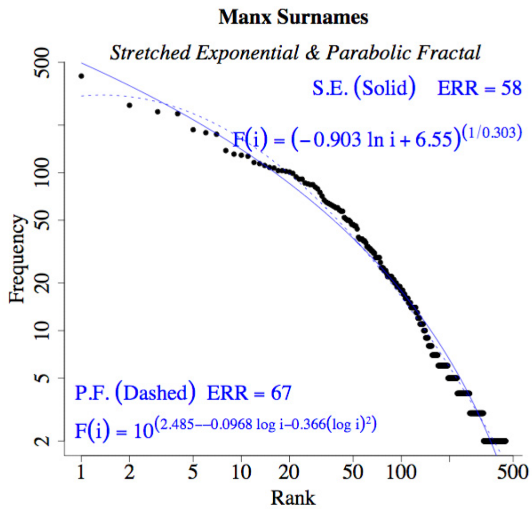

FIGURE 9 Distribution of surnames in the Isle of Man

There may or may not be a “best” alternative distribution that is a better fit for all or nearly all datasets. For the datasets considered here, the DGX was a better fit than the discrete power law, which was expected, since the DGX is a generalization of the power law. The parabolic fractal proved problematic, since it should be strictly decreasing, but for several datasets, the fit produced by linear regression led to negative values for the “a” coefficient. The stretched exponential did not have this difficulty, and it had a better or comparable error statistic to the parabolic fractal.

I am looking forward to continuing this work and incorporating additional distributions and datasets. There are plenty of practical and theoretical challenges involved in working with skew distributions, and the development of a statistical methodology will be a vital component of research in the years to come.

REFERENCES

Adamic, L.A., 2005, *Zipf, Power-laws, and Pareto — A Ranking Tutorial;* available at http://www.hpl.hp.com/research/idl/papers/ranking/ranking.html. Accessed May 26, 2005.

Bi, Z., et al., 2001, *The “DGX” Distribution for Mining Massive, Skewed Data*, San Francisco, CA.

Laherrère, J., and D. Sornette, 1998, “Stretched Exponential Distributions in Nature and Economy: ‘Fat Tails” with Characteristic Scales,” *The European Physical Journal B* 2, 525–539.

Mitzenmacher, M., 2003, “A Brief History of Generative Models for Power Law and Lognormal Distributions,” *Internet Mathematics* 1, 226–251.

Uetz, P., and EMBL Heidelberg, 2005, *The EMBL Reptile Database;* available at http://www.reptile-database.org/. Accessed Sept. 19, 2005.

Weisstein, E.W., et al., 2005, "Riemann Zeta Function," in *Mathworld — A Wolfram Web Resource;* available at http://mathworld.wolfram.com/RiemannZetaFunction.html.

Yule, G.U., 1925, "A Mathematical Theory of Evolution, Based on the Conclusions of Dr. J.C. Willis, F.R.S.," *Philosophical Transactions of the Royal Society of London, Series B* 213, 21–87.

Zipf, G.K., 1965, *Human Behaviour and the Principle of Least Effort*, New York: Hafner Publishing Co.

TOLERANCE AND SEXUAL ATTRACTION IN DESPOTIC SOCIETIES: A REPLICATION AND ANALYSIS OF HEMELRIJK (2002)

H. LEHMANN,* J.J. WANG, and J.J. BRYSON, University of Bath, Bath, United Kingdom

ABSTRACT

Most primate societies are characterized by hierarchical dominance structures. Males are usually dominant over females, but in periods of sexual attraction (during the females' period of tumescence) male "tolerance" toward females rises. Hemelrijk (2002) shows in a model that this tolerance is created as a side effect of the rise of female dominance during periods of sexual attraction. This rise, in turn, is the consequence of the more frequent approaches of males toward females during these periods. In Hemelrijk's model, the males gain no benefit from tolerating females, and they only do so at high aggression levels as a kind of "respectful timidity," because some of the females have become dominant over them.

This paper replicates and examines the results of Hemelrijk's study. We have found that some of Hemelrijk's results are highly reliant on aspects of the model that are not well supported by the current primate literature. We analyze the mechanisms underlying her results and suggest data that should be sought from observation logs of real primate colonies that would support or overturn the model.

Keywords: Agent-based modeling, hierarchy, primates, sexual attraction, social system, tolerance

INTRODUCTION

In this paper, we examine the best-established AI model of primate social systems, Hemelrijk's DomWorld (Hemelrijk 1999a,b, 2000, 2002). Hemelrijk models a large amount of primate behavior by using an incredibly simple model of social interactions based on spatial locations. In this paper, we replicate DomWorld, which allows us to examine the mechanisms underlying the system. We pay particular attention to the results from Hemelrijk (2002), the explanation of the increase of male tolerance experienced by females when they are sexually receptive (in tumescence). This particular experiment, situated in a wider model of differences between species in classifications of primate social structures, gives us a great deal of insight into the validity of Hemelrijk's approach.

We begin this paper by describing the primate social data to be explained and then by reviewing Hemelrijk's contributions. We then present our replication and our initial insights into the working of the DomWorld mechanisms. Finally, we discuss the validity of the model and propose specific data to look for that will either support or undermine the DomWorld model.

* *Corresponding author address*: Hagen Lehmann, Department of Computer Science, University of Bath, Bath, BA2 7AY, United Kingdom; email: h.lehmann@cs.bath.ac.uk.

BACKGROUND

Most primate species are highly social. They live in structured societies that can be characterized as having more or less steep dominance hierarchies. A steep hierarchy is one in which individuals would never consider violating rank. For example, a lower-ranked individual would not take any food in the presence of a higher-ranked individual. In a more shallow hierarchy, dominant animals show greater tolerance of subordinate behavior, and considerations of rank play less of a role in ordinary action selection. The difference between these social structures have been most studied in macaque societies (see Thierry et al. 2004 for a recent review). Societies characterized by steep hierarchies are often referred to colloquially as *despotic*, while those with the less rigid dominance structures are called *egalitarian*. When a dominant animal allows subordinate animals to take advantage of resources in its presence, the dominant animal is said to be expressing *tolerance*.

Tolerance is considered one of the most basic forms of conflict resolution (de Waal and Luttrell 1989). It might be difficult to see tolerance as an action to be selected, since it seems more like a form of inaction. However, if an agent is very inclined to preserve resources (including its own social rank), then expressing tolerance can require considerable inhibition of strong inclinations. In some species, for example, this is achieved by the apparently deliberate averting of gaze or even moving away from a resource in order to avoid witnessing a desired event, such as allowing a juvenile throwing a tantrum to feed. This shift in visual attention is necessary if witnessing such an event would automatically trigger an emotional/species-typical response that would, in turn, prevent the completion of the feeding.

The structure of a primate society is also correlated with a number of other characteristics (de Waal and Luttrell 1989; Thierry 2000; Hemelrijk 2002). Societies that are more despotic also tend to have more violent or aggressive interactions. On the other hand, there tend to be fewer conflicts in these societies than in egalitarian societies. In egalitarian societies, there are more frequent conflict interactions, but many of these involve no injury or violent dispute. For example, they may involve only hissing or snatching.

In most primate hierarchies, males are usually dominant over females because of their greater size, strength, and aggression. However, during the female sexually attractive period of tumescence, chimpanzee males, for instance, allow females priority in food access (Yerkes 1940). This has been explained as probably being a cognitive strategy — an exchange for copulation — which is adaptive in that it also therefore produces offspring (Goodall 1986; de Waal and Luttrell 1989; Stanford 1996).

Hemelrijk and her colleagues have proposed a cognitively minimalist explanation of this change in behavior. Hemelrijk claims that there is no statistical evidence for such exchanges for food (Hemelrijk et al. 1992), and neither is there any increase in related offspring (Hemelrijk et al. 1999). Hemelrijk (2002) demonstrates a model where such a change in dominance occurs in despotic societies even without any benefit for the males, but as a simple consequence of the higher frequency of dominance interactions between the sexes brought on by the males' attraction to the females.

Hemelrijk claims that in her models, under the condition of high aggression intensities, males show tolerance toward females. Her evidence of tolerance is that, in her model, in times of sexual attraction, females may achieve ranks higher than males, while in other times, they do not.

Females are modeled as initially 50% weaker than males and persistently 20% less aggressive, which explains why such outcomes are improbable in general. However, once an animal achieves a higher rank, its power is assumed (in these models) to also increase.

Hemelrijk explains her findings as a side effect of the higher frequency at which males approach females. Normally, animals tend to avoid invading each other's personal space and triggering a conflict unless they are of a higher rank than the animal they are approaching. However, in times of sexual attraction, Hemelrijk's males ignore rank in approaching females. Further, in Hemelrijk's model, the outcome of a dominance interaction is highly influenced by the extent to which it was unexpected. Thus if a very-low-ranking female happens to win a competition (there is always a small chance of success, with the probability being inversely proportional to the discrepancy in rank), then she will suddenly achieve a much higher rank.

Consequently, the opportunity for a low-ranking female to win an interaction will rise as more males approach her. Thus, she could become more dominant then some of the males, who will nonetheless continue approaching her, consequently likely increasing her rank as they fail in their subsequent dominance disputes. Therefore this "tolerance" is more a "respectful timidity" toward higher-ranking females. The males will approach but not attack simply because the female has a higher rank.

Thus a behavior typically described as complex or even cognitive could, according to Hemelrijk's model, arise without any corresponding cognition. This change could be introduced to the species through a single exogenous factor, such as the availability of food resources, if this leads to an increase in aggression. This higher aggression then leads to a more despotic society in which in the periods of sexual attraction, the dominance of the females rises, as shown in the model and explained above.

Many researchers have expressed skepticism about Hemelrijk's work because of her anti-cognitivist stance. People who work closely with apes feel that it is "obvious" that the animals have some cognitive capacity, or at least that when humans express very similar behavior, they subsequently report having been in a cognitive state.

Because we were curious about Hemelrijk's model and wished to understand it better, and because no version of DomWorld is freely available on line, we replicated Hemelrijk's work. In so doing, we were able to examine the assumptions behind the model and find out what aspects of the model were critical to its success in replicating primate behavior.

METHODS

Hemelrijk's model consists of a small troop of chimpanzees living near each other and occasionally having aggressive interactions, which result in shifts in dominance rank. After the model has run for a while, quantitative descriptions of the agents' relationships are taken, such as the steepness of the dominance ranking hierarchy or the average centrality of an agent within its troop. These measurements are then compared to measurements made of real chimpanzees in natural situations to judge the quality of the model as a hypothesis of their behavior.

The Model World

Our simulation was based on the model described by Hemelrijk (2002). She wrote her version in Object-Pascal and Borland Pascal 7.0. We used NetLogo 2.1, because it, being a purpose-built modeling tool, provides a relatively easy, high-level language for quickly constructing models and visualizing results. The world in which the agents interact is wrapped around on all sides and therefore resembles the geometrical structure of a torus. This is to avoid border effects and enable the agents to move in every direction. As described by Hemelrijk, this space is of a size 200×200 units. It is a continuous space — agents have real-valued locations and can move in any of 360 directions. When an experiment starts, the agents set initially at random locations within a 30×30 parcel of this space. Each agent has a forward vision angle of 120 degrees (that is, it "sees" or attends to agents that are 60 degrees to either side of its direction of forward motion) and a maximum perception range (*MaxView*) of 50 units. Consequently, at the beginning of the simulation, each agent will need to do no more than turn around to see all the other agents in the simulation. The visual limits restrict the amount of things that the agent is likely to attend to at any particular time.

Agent motion and social interaction is determined by a number of additional threshold parameters:

- A near-perception range, *NearView*, of 24 units. Agents feel comfortable as long as they see some other agent within this range. If they do not, but they do see an agent (that is, one is within MaxView), then they will go toward that agent.

- A personal space parameter, *PerSpace*, of 2 units. Agents within this range of each other will have a dominance interaction.

- A search angle of 90 degrees. Agents rotate this amount if they can see no one within their MaxView.

- A waiting period. After an agent moves around or engages in a dominance interaction, it is assigned a random waiting time before it performs its next action. The waiting period simulates foraging or resting in the wild — constant dominance interactions are not only unnatural but also make the troop so chaotic that spatial measurements of troop coherence and rank have no meaning. The waiting period is abbreviated when the agent observed a dominance interaction within its NearView. This is in accordance with observations in real animals, since in primate groups, nearby fights are likely to trigger active behavior in individuals (Galef 1988).

In our experience, the model does not appear overly sensitive to most of the parameter values, although at the same time, none of them can be eliminated and still maintain the action-selection model. However, the mode *is* particularly sensitive to the organization of the waiting period. This is because many dominance interactions would not happen if the relatively subordinate animals were able to avoid the relatively dominant one, but because only one animal tends to be moving at a time, the dominant one can invade the personal space of the subordinate.

In the simulations dealing with the impact of female tumescence on their dominance ranking, there is one additional parameter, *attraction*, which is either *on*, indicating that all the females are tumescent, or *off*, indicating that none of them are.

The Interaction Structure

The interactions in the model are classified into two groups. One class consists of grouping interactions, the other consists of dominance interactions. These two classes resemble the two forces that in nature, on one hand, drive groups apart, and on the other hand, hold them together in order to stabilize them (c.f. Reynolds 1987).

For the grouping interactions, Hemelrijk gives a set of four rules:

1. An agent that observes another agent within its personal space may perform a dominance interaction, depending on its own rank and the rank of the other agent. For such an interaction, first the nearest potential opponent is chosen. After an interaction, the winning agent moves one unit toward its opponent, while the loser turns around 180 degrees, plus or minus an angle drawn randomly from 45 degrees, then moves two units away.

2. If the agent does not detect anyone in its personal space but can see other agents within its NearView, then — in trials without attraction — it moves one unit forward on its present course. In the attraction condition, if VirtualMale can see VirtualFemale, they will change their direction toward the nearest visible VirtualFemales and then move one unit forward.

3. If the agent detects no other agents within NearView, but there are agents within its MaxView range, then it changes direction toward the nearest one and moves one unit toward it.

4. If there are no other agents within MaxView, the agent turns in a search angle of 90 degrees at random to the right or left.

The dynamics of the simulation are such that, for any agent, there will always be at least one agent still in MaxView in some direction. Occasionally the troop splits, but the agents always reunite shortly. Given the rate of motion of the troop, the maximum duration of the waiting period, and the large difference between MaxView and NearView, no single individual can become “lost” from the troop.

In nature, dominance interactions between primates are characterized by the competition for resources, such as food or potential mates. In order to gain stable access to such resources, the different individuals within a group try to establish a rank in hierarchy that is as high as possible. This is achieved by constant interaction, which Hemelrijk calls in her paper a “long-term ‘power’ struggle.” In the model, there are no resources specified, and the only trigger for interactions is spatial distance. The agents start “fighting” when another agent is within their personal distance and the rank of the other is lower or equal to their own rank. The agent “estimates” its chances to win, and if its chances seem good, then it engages in the competition (see following text).

Since the dominance values *within* each sex are equal at the beginning of a simulation, the outcome of every single interaction influences the chances of winning the next one. Such a system is self-reinforcing and has been shown empirically in many animal species (Hemelrijk 2000).

The formula for determining the outcome of a dominance interaction was modeled after Hogeweg (1988) and Hemelrijk (1999b). Each agent has a certain dominance value, which is readjusted after every "fight" the agent gets involved in. We called this value *Dom* according to Hemelrijk's notation. This variable is correlated both to the agent's rank and its ability to win an interaction. If one agent finds another agent in its PerSpace, it compares its own Dom-value with the Dom-value of the other. If its own value is higher or equal to the other, it "estimates" that it has good chances to win and will therefore interact. The outcome of the interaction is calculated with the following formula (from Hemelrijk 2002, page 734):

$$w_i = \begin{cases} 1 & \dfrac{Dom_i}{Dom_i + Dom_j} > Random(0{,}1) \\ 0 & else \end{cases}, \tag{1}$$

where Random(0,1) produces a random real value between 0 and 1.

In this calculation, w_i is the value that determines whether agent *i* has lost or won. Here 1 means victory and 0 means defeat. The relative dominance value is compared with a randomly drawn number between 0 and 1. If it is greater than the drawn number, the agent wins. This means that higher an agent's rank is relative to its opponent's, the more likely the agent is to win.

After a dominance interaction, the dominance values of both agents are adjusted according to the outcome by using roughly the same information:

$$\begin{aligned} Dom_i &= Dom_i + \left(w_i - \frac{Dom_i}{Dom_i + Dom_j} \right) * StepDom\,, \\ Dom_j &= Dom_j + \left(w_i - \frac{Dom_i}{Dom_i + Dom_j} \right) * StepDom\,. \end{aligned} \tag{2}$$

The only exception to the above equations is that the lowest possible Dom-value is set to 0.01 in order to keep the Dom-values positive.

Hemelrijk calls this system for determining dominance values a *damped positive feedback system*, since, in the case of winning, the dominance value of the higher-ranking agent goes up only slightly, but if the lower-ranked agent wins, its dominance value undergoes a great change. This is intended to reflect the fact that it is very unlikely for a low-ranking individual to win an interaction with a high-ranking one. Thus ranking is not changed much by an expected outcome, but it changes greatly for an unexpected one.

The amount of rank shift is also affected by another value: *StepDom*. This value Hemelrijk uses to represent the intensity of the "aggression" (or violence) of the interaction, which she hypothesizes also correlates to the impact the interaction has on ranking. She uses a high StepDom value to represent the level of aggression in despotic species, and a low StepDom value to represent the level in egalitarian ones. Values for StepDom can vary from 0 to 1 but are held constant within any give simulation, since they are considered to be determined by species. Although Hemelrijk calls this value “aggression,” note that it has no direct impact on the probability or outcome of an interaction (see Equation 1). Rather, its impact is only indirect through its long-term impact on the dominance values, which do determine both whether and how well an agent fights.

Another important element for correlating Hemelrijk’s models to the real world is understanding her *coefficient of variation of dominance values*. This coefficient indicates the average variation between dominance ranks of the individuals in the troop. Hemelrijk interprets this coefficient as an indication of how despotic or egalitarian a society is. Her hypothesis is essentially that there isn’t a qualitative difference in how monkeys in an egalitarian society treat their superiors versus how those in a despotic one do, but rather that every agent will show an equal amount of respect for a troop mate with twice its dominance value. Thus Hemelrijk represents a despotic society as one with an unambiguously “steep” dominance hierarchy (with a great difference in rank between individuals) and represents an egalitarian one as having relatively ambiguous rankings.

Experimental Setup

For our attempted replications, we used the parameter settings Hemelrijk uses in several studies (Hemelrijk 1999a, 2000). We used eight agents in a troop, four of each sex ($N = 8$). As explained earlier, each agent had an personal space of 2 (*PerSpace* = 2), a vision angle of 120 degrees, a maximum perception range of 50 units (*MaxView* = 50), and near-perception range of 24 units (*NearView* = 24). The search angle was 90 degrees, the fleeing distance was 2 units (*fleeD* = 2), the fleeing angle was 45 degrees at a random direction away from the opponent, and the chasing distance was 1 unit (*chased* = 1) in the direction of the opponent.

To resemble the difference in physical strength between males and females, both sexes started out with different winning or loosing tendencies; that is, the DomValues of females were half that of males (*virtual females* = 8, *virtual males* = 16). Also, females have only 80% of the aggression intensity (*StepDom*) of males. The experiment was conducted under four different conditions. We used two level of aggression to correlate with the two types of social interactions witnessed in different primate species. In the high level, the StepDom value of males was 1 and that of females 0.8. In the low aggression level, the StepDom value of males was 0.1 and that of females 0.08. These two aggression conditions were each run under two conditions of *sexual attraction* (either turned on or off) 10 times each, resulting in a total number of 40 runs. Each run was 42,800 time units long.

RESULTS

Our results match Hemelrijk’s results to the extent that we used the same analysis, which we largely did in order to test the replication. The first figure shows a comparison between the number of interactions performed by virtual females during the different conditions. In the graph, the total number of aggressive interactions initiated by virtual females is compared for all four different conditions used in the experiment.

In Figure 1, we can see that the number of virtual female dominance interactions increases significantly under conditions with sexual attraction in both intensities of aggression (Mann-Whitney, N = 10, U = 0, $p < 0.001$, two-tailed, Mann-Whitney, N = 10, U = 0, $p < 0.001$, two-tailed). That means females are involved in considerably more interactions when they are attractive. The aggression level amplifies the result, even though this effect for the aggression is rather weak (Mann-Whitney U-Test, N = 10, U = 24, p < 0.049, two-tailed).

Figure 2 shows the dominance of virtual females as the sum of the number of males ranked below each female at different times in different conditions. We can see that, as reported in Hemelrijk, female dominance under conditions with a high aggression level increases over time but stays constant under conditions with a low aggression level.

Figure 3 is the classic Hemelrijk result. It shows the distribution of the coefficient of variation of dominance values for both sexes (see discussion in previous section). If aggression is high, there will be a steeper hierarchy (i.e., the difference between rank values will be larger). This is true both within and between sexes. Attraction amplifies this result, despite the fact that some females may outrank some males under this condition.

Figure 4 shows the change of dominance values for both sexes under conditions with high and with low levels of aggression. With high aggression, a constant change in the dominance structure is noticeable as greater and greater differentiation/steepness in the hierarchy. With low aggression, there is only a very small change in the dominance values. This creates a very stable hierarchy where the females never gain a higher positions in the group.

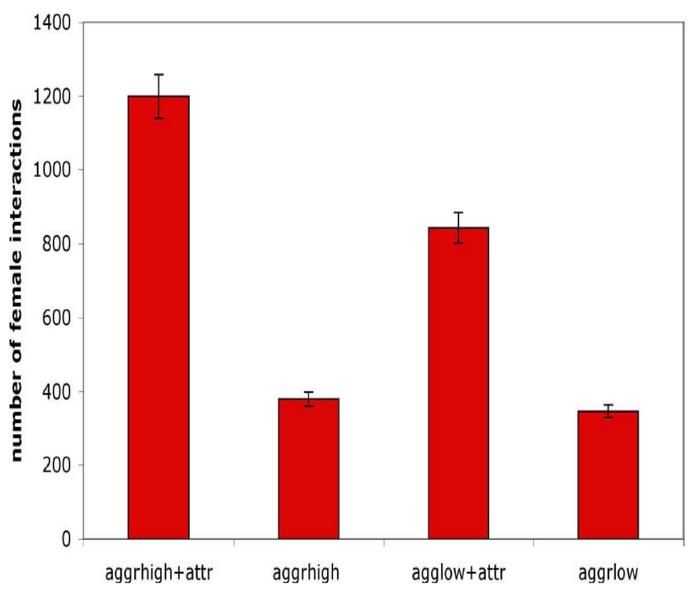

FIGURE 1 Total number of female interactions under different conditions (*aggrhigh+attr* = high aggression + attraction; *aggrhigh* = aggression high + no attraction; *agglow+low* = aggression low + attraction; *aggrlow* = aggression low + no attraction.)

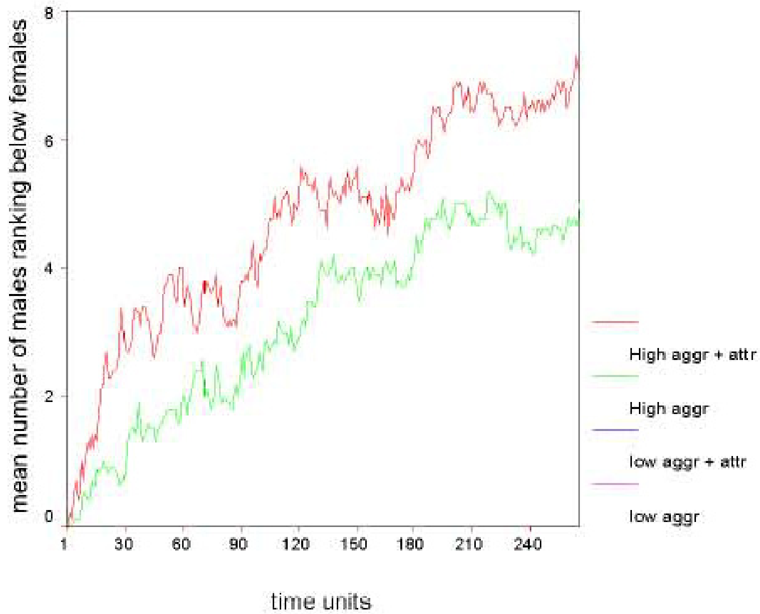

FIGURE 2 Dominance of virtual females as the sum of the number of males ranked below each female at different times under different conditions

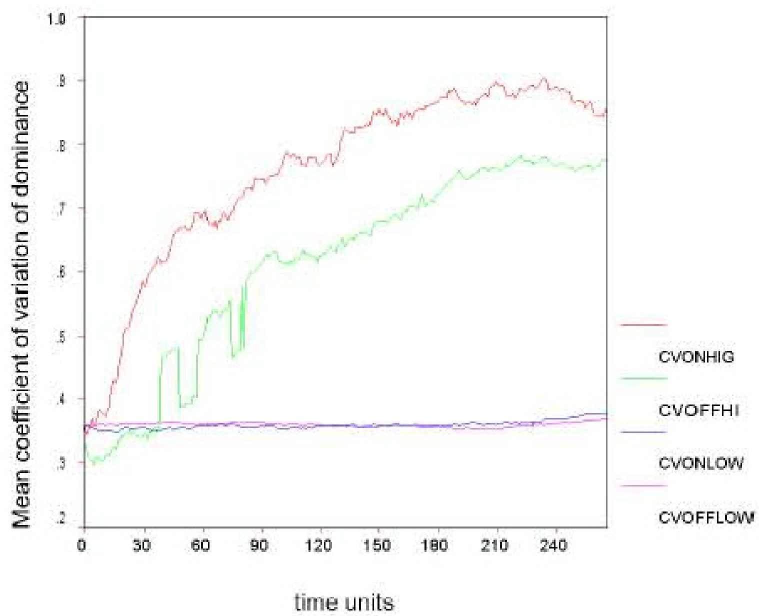

FIGURE 3 Distribution of the coefficient of variation of dominance values under different conditions for both sexes

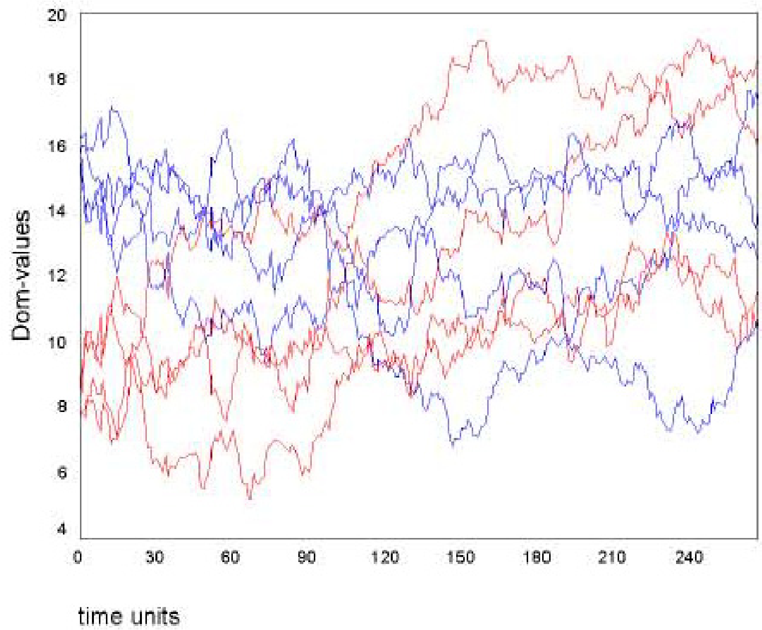

(a) High level of aggression

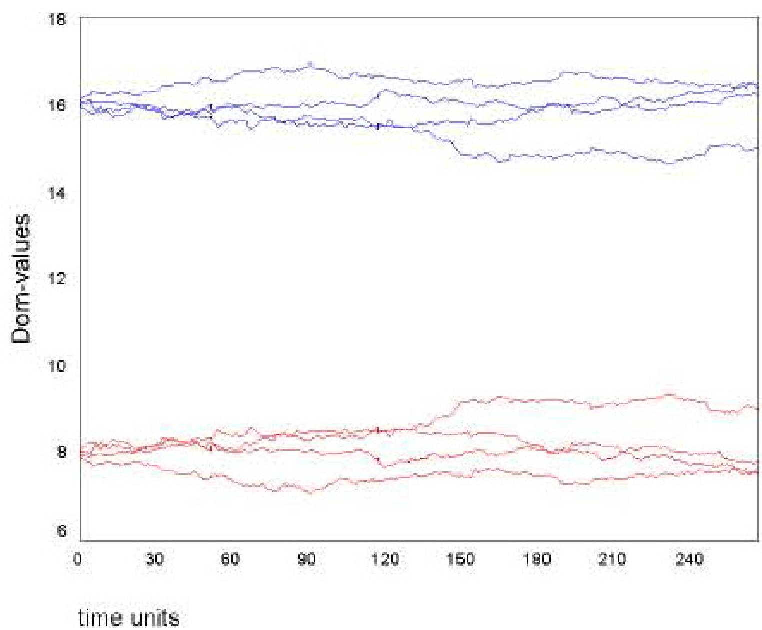

(b) Low level of aggression

FIGURE 4 Distribution of dominance values at a high and low level of aggression (Under both conditions, the males start off initially higher than the females.)

The conclusion of these results is that only in groups with a high level of aggression are females able to gain higher positions in the social hierarchy. Attraction amplifies this effect but plays a secondary role.

DISCUSSION

Our results show the same structure as the results in the original study (Hemelrijk 2002, Figure 3A on page 739 and Figures 4A, B, and C on page 741) and can therefore be seen as a replication. In general, the diversity of different dominance values between individuals increases if there is a high aggression level existing within the population. Under conditions with low aggression levels, this effect does not appear, even though the results in this model show that the increase in interactions between virtual females and virtual males depends not on the increased level of aggression but on the existence of female attraction. In this first result, we can see that the level of aggression has no (or, at best, only very little) influence on the number of interactions between the individuals, yet under both conditions, with female attraction, the increase in interactions is significant.

The most interesting effect is the change in dominance values toward more dominant females and, as a possible consequence, a change in group structure. This connection between higher interaction frequency and the dominance value change Hemelrijk claims in her article (see page 742) could be a simple explanation for the observed natural phenomenon of male tolerance toward females in their period of sexual attractiveness. Given our understanding of Hemelrijk's model derived from our replication, we will now examine these claims more closely.

One of the strengths of agent-based modeling (ABM) is its ability to demonstrate whether theories of the origin of behavior can be explained by a given model of how an agent selects its actions. In particular, as with the rest of science, there is an emphasis in ABM on looking for the simplest possible explanation that fits the data. We look for the origins of complex behavioral patterns on a social level as emergent from simple behavior in the individual.

We need to realize, however, that this is not only a case of following the principle of parsimony for reasons of the philosophy of science, it may also be a case of looking for our keys under the light of the street lamp rather than over in the dark where we lost them. Complex individual behavior is difficult to program, takes a long time to execute in simulation, and then is difficult to analyze. So we may have a strong bias toward looking for overly simple solutions. Thus, while on one hand, we need to be open-minded and sure to understand correlations where we find them, on the other hand, we cannot allow our biases to blind us to a situation where data may not fit the predictions of our model. Guarding against this bias is just as important as guarding against its opposite — the overly cognitive explanations.

The Hemelrijk model we have replicated seems to be a good analogue system for macaque behavior. Her DomWorld model shows that apparently complex behaviors in primate societies (like "male tolerance" or "female assertiveness") can be created in computer-generated primate societies with only a few simple assumptions about individual behaviors. The effect of female dominance appears, for example, under conditions with high aggression and is consolidated by a high level of attractiveness in the females. Hemelrijk notes the difference between this and the classical explanations for this phenomenon, which propose exchanges involving food for sexual opportunities (Goodall 1986). Hemelrijk's model does not include any food or sex yet still leads to analogous results.

Now that we have a working model, we can try to understand exactly where and how these phenomena "emerge." We can now analyze what the critical factors of the model are and look for biological correlates that would either prove or disprove the model.

The effect of the model is based on two major assumptions:

1. The self-reinforcing effect of domination and

2. The fact that females attract males in their time of tumescence, but that males are not attractive to females.

The first assumption relates to the fact that the dominance value *DOM* of an individual *i* (operationalized as the ability to win a fight) increases with a victory and decreases with a defeat. Although this self-reinforcement is a well-known phenomenon that has been studied extensively in laboratory animals such as mice, we are somewhat skeptical of the exact extent to which this model depends on these factors. In Hemelrijk's model, the strength of the effect is determined by the dominance ranking of the opponent, the level of aggression (i.e., the step-value assigned to this species), and chance. The result of a fight is calculated with Equation 1, repeated here:

$$w_i = \begin{bmatrix} 1 \dfrac{Dom_i}{Dom_i + Dom_j} > Random(0,1) \\ 0 \quad else \end{bmatrix}. \tag{3}$$

Again as a reminder, the dominance level after a fight is calculated with Equation 2:

$$Dom_i = Dom_i + \left(w_i - \frac{Dom_i}{Dom_i + Dom_j} \right) * StepDom\,,$$

$$Dom_j = Dom_j + \left(w_i - \frac{Dom_i}{Dom_i + Dom_j} \right) * StepDom\,.$$

As we emphasized earlier, Hemelrijk has defined the factor StepDom to mean aggression. An individual therefore increases its ability to win a fight (its dominance) most if it wins against an individual with a preferably much higher dominance level and if the aggression level in the group is high.

Aggression is therefore the crucial value that decides within the system how far an individual can go up or fall down in the hierarchy as the result of a single fight. This is largely the basis of the reinforcement effect of domination. But to what extent does this effect exist in nature? Hemelrijk's text only mentions observations of bumblebees and other computational models as examples (page 743 f). Thinking about it in a intuitive way, it might be plausible that self-confidence about winning a fight increases if one wins against someone much stronger. Further, we know that even in adult mammals, growth hormones can be triggered by success in social competitions. Nevertheless, in a real fight, the body's size and strength are at least as important as the psychological status of the individual.

To test the validity of Hemelrijk's model, we need to use the documented history of dominance hierarchies in real animals. We need to look carefully at the relatively rare events in which a lower-ranked animal bested a higher-ranking animal and see what the impact was on the troop's dominance structure before and after. We should look in particular for the following factors:

- If one agent defeats another that vastly outranks it in a dominance interaction, do the two agents immediately change ranks within the troop? In other words, is an unexpected outcome from a fight likely to have a very significant effect? If this is true, it would validate the use of relative dominance values in Equation 2.

- In comparing across species, does it take fewer interactions to advance rank in a despotic species? If this is true, then it would justify the use of StepDom in Equation 2.

- Within species, if a fight is more violent (e.g., if blood is drawn compared to mild beating, or if there is mild beating compared to a nonphysical interaction), does it have more impact on dominance hierarchy? If this is so, then it makes sense to refer to StepDom as "aggression," and it would further validate its use in Equation 2.

- Are females more likely to engage in fights when they are tumescent? If not, then this model cannot account for their increased dominance.

- Do females only become dominant during their tumescence in despotic species? Given that in Hemelrijk's model, the prime indication of increased dominance for females is the males' increased tolerance of them, discriminating an increase of rank in an egalitarian species may be difficult, since these species are by definition tolerant toward all group members. But it is a prediction of the model.

- Is it true that when an animal in an egalitarian species is *clearly* outranked by another animal, those two animals' interactions will be similar to two more nearly ranked animals in a less egalitarian species? Or *is* there a qualitative difference in how different species behave with respect to dominance hierarchies? The answer to this question will serve to validate whether the steepness of the dominance hierarchy is a good representation of despotism/egalitarianism.

Of course, this is complicated by the fact that establishing a dominance hierarchy is never easy. It's not clear that every animal will agree on the current hierarchy; indeed, some animals will behave differently with respect to others depending on what other animals are present (Harcourt 1992). However, many groups work diligently to attempt to establish these sorts of records, so we can hope to test these predictions.

We need to also look critically at the second basic assumption: the idea that female primates attract male primates when they are in their fertile days. This is obviously true, but sexual attraction is bidirectional and therefore influences the grouping behavior of females as well. Of course, it is possible that the male attraction is strong enough to overwhelm the data, or even that just putting high male attraction is a good approximation for mutual attraction. However, the question remains as to whether the mechanism exploited by the model — increased conflict leading to a higher probability of an occasional lucky win by the female that immediately catapults her high into the dominance hierarchy — is at all plausible.

CONCLUSIONS

We have presented a replication of Hemelrijk (2002) and an analysis of how her model works. We have also presented a critical list of suggestions for testing the validity of the mechanism. We suspect that the rules for determining dominance from the outcome of dominance battles are not sufficiently realistic and cannot fully explain the change in female dominance rank on their own. If we are right, then this model may need additional factors to explain this phenomenon, possibly including a cognitive state sufficient for the traditional theories of reciprocation.

ACKNOWLEDGMENTS

Several MSc and undergraduate students over the years have explored using NetLogo in the AmonI group for this task. We owe particular gratitude to Wang (2003). All software used in this paper is available upon request or from our Web site. This paper originally appeared in the proceedings, *Modeling Natural Action Selection*, edited by J.J. Bryson, T.J. Prescott, and A. Seth, AISB Press, 2005.

REFERENCES

de Waal, F.B.M., and L. Luttrell (1989). Toward a comparative socioecology of the genus *macaca*: Different dominance styles in rhesus and stumptailed macaques. *American Journal of Primatology*, 19:83–109.

Galef, B.G J. (1988). Imitation in animals: History, definition, and interpretation of data from the psychological laboratory. In T. Zentall and B.G.J. Galef, editors, *Social Learning: Psychological and Biological Perspectives*, pp. 3–25. Erlbaum, Hillsdale, NY.

Goodall, J. (1986). *The Chimpanzees of Gombe: Patterns of Behavior*. Harvard University Press.

Harcourt, A.H. (1992). Coalitions and alliances: Are primates more complex than non-primates? In A.H. Harcourt and F.B.M. de Waal, editors, *Coalitions and Alliances in Humans and Other Animals*, Chapter 16, pp. 445–472. Oxford.

Hemelrijk, C.K. (1999a). Effects of cohesiveness on intersexual dominance relationships and spatial structure among group-living virtual entities. In D. Floreano, J.D. Nicoud, and F. Mondada, editors, *Proceedings of the Fifth European Conference on Artificial Life (ECAL99)*. Springer.

Hemelrijk, C.K. (1999b). An individual-oriented model on the emergence of despotic and egalitarian societies. In *Proceedings of the Royal Society London B: Biological Sciences*, 266:361–369.

Hemelrijk, C.K. (2000). Towards the integration of social dominance and spatial structure. *Animal Behavior*, 59(5):1035–1048.

Hemelrijk, C.K. (2002). Self-organization and natural selection in the evolution of complex despotic societies. *Biological Bulletin*, 202(3):283–288.

Hemelrijk, C.K., C.M. Meier, and R.D. Martin (1999). "Friendship" for fitness in chimpanzees? *Animal Behavior*, 58:1223–1229.

Hemelrijk, C.K., G.J. van Laere, and J.A.R.A.M. van Hooff (1992). Sexual exchange relationships in captive chimpanzees? *Behavioral Ecology and Sociobiology*, 30:269–275.

Hogeweg, P. (1988). MIRROR beyond MIRROR: Puddles of LIFE. In *Artificial Life: Santa Fe Institute Studies in the Sciences of Complexity*, pp. 297–316. Addison-Wesley, Reading, MA.

Reynolds, C.W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34.

Stanford, C.B. (1996). The hunting ecology of wild chimpanzees: Implications of the evolutionary ecology of pliocene hominids. *American Anthropologist*, 98:96–113.

Thierry, B. (2000). Covariation of conflict management patterns across macaque species. In F. Aureli and F.B.M. de Waal, editors, *Natural Conflict Resolution*, Chapter 6, pp. 106–128. University of California Press.

Thierry, B., M. Singh, and W. Kaumanns, editors (2004). *Macaque Societies*. Cambridge University Press.

Wang, J.J. (2003). Sexual attraction and inter-sexual dominance among virtual agents — Replication of Hemelrijk's DomWorld model with Netlogo. Master's thesis, University of Bath, Department of Computer Science.

Yerkes, R.M. (1940). Social behavior of chimpanzees: Dominance between mates in relation to sexual status. *Journal of Compartative Psychology*, 30:147–186.

INTERPRETIVE AGENTS: A HEATBUG REFERENCE SIMULATION

D.L. SALLACH,* Argonne National Laboratory, Argonne, IL,
and The University of Chicago, Chicago, IL
V.S. MELLARKOD, Texas Tech University, Lubbock, TX,
and Argonne National Laboratory, Argonne, IL

ABSTRACT

Interpretive heatbugs (IHB) is a reference application designed as the first implementation of the interpretive agents (IA) research program. As described in previous papers, the IA initiative identifies three mechanisms (prototype reasoning, situation definition, and orientation accounting) to create an interpretive architecture that supports agents whose actions are oriented by meaning.

The reference application uses the familiar heatbugs environment, where bugs require a temperate zone and flee from settings that are uncomfortably hot or cold. Because each bug emits a small amount of heat, congregations of bugs initially create needed warmth, but then overcrowding creates excessive heat, with the feedback creating unstable churning patterns.

To this dynamic of temperature fluctuation driving bug movements, IHB adds the capability of bugs to assist or undercut each other and ethnic and religious identities that mediate the decisions to help or hinder. IHB intends to provide an exemplar for a wide variety of cultural interactions, including, not only genocide and ethnic cleansing but also diversifying markets and constructive interdependencies.

OVERVIEW

Agent simulation is a methodology that has been an important innovation in the social sciences. Through the device of distributed, endogenously motivated software agents, a range of social processes have been simulated in interesting ways. The variety of ways in which microinteractions produce familiar large-scale effects has been a source of insight and held promise for future research.

If the insight arising from simple models has been the fruit of embryonic agent simulation, it also constitutes a ceiling. The first generation of agent simulation has been largely characterized by simplicity in rules, agents, relationships, motivational structures, and resulting processes. The challenge for modelers is to retain the clarity of simple models while, at the same time, extending them in order to more fully capture the fluidity of social relationships and processes.

Many artificial intelligence research programs have sought to computationally effectively emulate natural intelligence in its fullness. The interpretive agent (IA) research program is based

* *Corresponding author address*: David L. Sallach, Decision and Information Sciences Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439; e-mail: sallach@anl.gov.

on the premise that it is possible to capture the nonlinearity of interactive dynamics arising from the processes of social construction and interpretation without first establishing strong forms of artificial intelligence (Sallach 2000, 2003). More specifically, interpretive agents are context-sensitive and meaning-oriented, registering the flow events and shifting their goals, intentions, and actions accordingly (Mellarkod and Sallach 2005a; Sallach and Mellarkod 2004, 2005).

The IA architecture is based on three interleaved mechanisms: prototype reasoning, situation definition, and orientation accounting. Prototypes provide a calibrated implementation of bounded rationality. The process by which participants define situations generates a contextual framework and provides a focus of activity. Participants and relevant nonparticipants maintain and, to some extent, share orientations toward a significant group, object, symbol, etc. Accordingly, as communications and actions are generated, they must take such constraints into account and adjust accordingly.

Ultimately, all interpretive mechanisms in the IA research program are based on geometrical models of meaning (Gardenfors 1999; Widdows 2004). Taken broadly, the latter attributes meanings on the basis of their proximity to a reference point in a conceptual or semantic space. Reference points are defined relative to exemplary and/or idealized concepts, where both serve as reference points. Prototype reasoning, as implemented in this reference application reported here, utilizes reference points based on empirical exemplars.[1]

A REFERENCE APPLICATION

The "heatbug" application has played a significant role in the history and emergence of agent simulation. It provides a simple, controlled example that has served to illustrate basic design and techniques. The goal of the present paper will motivate an analogous reference application for interpretive agent simulation.

A reference application is one that defines and implements generic mechanisms for a specific type of problem domain. These implemented mechanisms provide an exemplar for developers who want to construct an application of that type. Specifically, developers can take the generic mechanisms of the reference application and customize them to address the requirements of the detailed structures and dynamics of a particular scientific, institutional, or business area.

The application "interpretive heatbugs" (IHB) is conceived as a reference application for how interpretive agents can be designed and developed in a relatively simple setting; it provides a way to illustrate the concepts and mechanisms and acts as a modest example of how they may be implemented. The described reference application is implemented in the J mathematical programming language (Thomson 2001; Peele 2005).

The present discussion summarizes how the IA mechanisms are designed and the assumptions on which the application is based. The focus is on what aspects of meaning orientation in social settings are used to motivate the design and the way that this reference application is intended to serve as a bridge to social simulation models and applications. More

---

1 The mechanisms for both types of reference points are actually similar, attributing meaning on the basis of proximity.

detailed design considerations, and the implementation of the mechanisms and their interaction, are discussed in a parallel paper (Mellarkod and Sallach 2005).

Interpretive heatbugs inhabit a setting comparable to the original heatbug setting. Agents seek a location with a comfortable temperature within a “heatspace,” experiencing discomfort from excesses of heat and cold. Each bug emits a small amount of heat, which, in aggregate, contributes to temperature differences and evolution. interpretative heatbugs are assigned regions of comfort, discomfort, and extreme discomfort that are unique to them. In the baseline implementation reported here, they assume other bugs share equivalent zone sensitivities. Interpretive heatbugs have a basic energy metabolism that is more rapidly depleted in areas of higher discomfort. If their energy drops below zero, they can only undertake movement on alternate ticks. In later versions, we may introduce bug demographics; in this case, bugs without energy resources will die.

In addition to movement, which is governed by discomfort levels, bugs have three additional capabilities and proclivities: asking for energy, giving energy, and shoving their way into more comfortable locations. In the baseline model, the three actions are defined by exogenous rules, except that many of the governing rules fire relative to conceptual prototypes of how *#similar, #nice,* or *#tough* pertinent bugs are,[2] where the latter are defined in Table 1.

TABLE 1 Pertinent Bugs in IHB Rules

Activity	Pertinent Bugs
Asking for energy	Bugs that are potential donors
Giving energy	Bugs requesting energy
Shoving	Bugs attempting to enter the same cell

The *#similar*, *#nice*, and *#tough* prototypes[3] are defined by the idiosyncratic training and experience of the individual bug; therefore, each has a unique interpretation of the concept. However, all such attributions are mediated through two geocultural forms of social structure: ethnicity and religion. Ethnicity becomes a locus of identity and in/out group dynamics. Religious identification serves as a source of tighter or looser value commitment. Religion also has more finely grained distinctions that are known to members and to those who know them well but not to casual outsiders. All such complexities, including behavior that can be observed in particular circumstances, define distinctions that can be used by the reference point reasoning of the bug.

2 By convention, words that stand for conceptual prototypes are marked by a pound (#) sign. The hash mark is meant to suggest the radial structure of the prototype it represents. For both the designers and the bugs, a conceptual prototype is a *region* within the space of experience and thus semantic space as well. With the multiple dimensions and semantic variety implied by its radial structure, the region cannot be fully or adequately conveyed by a single label.

3 The particular conceptual prototypes used in the baseline design are designed to illustrate the operation of conceptual and computational mechanisms, not to express an articulated social or psychological theory.

IHB is a reference application and is thus meant to provide an example for richer social models. Additional kinds of structure (e.g., age, gender, status, wealth) can be added and taken into account by prototype inference. Available actions can become more calibrated, more closely aligned with social and historical issues, and, ultimately, be given prototypical form. In this way, IHB is designed to provide a catalyst for a new generation of social simulation models.

Modern history is replete with examples of genocide, ethnic cleansing, human rights violations, and movements for civil and/or minority rights in which ethnicity and religion have served as interpretive queues for mass behavior. IHB is designed to facilitate the construction of models that locate such historical events within interpretive social processes.

ABSTRACT SOCIAL STRUCTURE

Social structure is a historically pervasive formation that envelopes and shapes all social action. Its effects are subtle and complex, varying in form and effect, yet there are relatively simple commonalties that underlie its diverse manifestations. A synthetic model of social structure will undertake to integrate simplicity and complexity within a formal model. The complexity of the model will allow expression of the texture of empirical social life; its simplicity will facilitate inference about structural dynamics.

Social structure can be modeled at multiple levels. It can be defined abstractly, so that *diverse historical structures* (e.g., serfdom, patrimonialism, slavery) *can be compared* on pertinent criteria. An advantage of developing an abstract definition of social structure is that social structures *before and after a transformation can be explored*. Second, the social structure of any given historical conjuncture can be *framed in historical context* or *articulated in greater depth*, as illustrated by numerous stratification examples (e.g., Frazier 1957; Warner 1960; Dumont 1970; Zeitlin and Ratcliff 1988).

Development of an abstract concept of social structure can contribute to social modeling. Specifically, to the extent that historically unique structures can be understood as variations on an abstract concept and articulated accordingly, the ability to apply a common model to highly diverse social phenomena enhances our ability to model them comparatively.[4] One purpose of the IA research program is to investigate how microinteraction can produce large-scale spatially distributed structures.

Regarding stratification patterns, perhaps no contemporary theorist has articulated a model as synthetic across levels as has Randall Collins. Within an emphasis on theoretical coherence and cumulation, Collins’s contributions can be seen to lie in three primary areas: (1) the location of stratification processes within a broader social context, including the significance of interaction rituals (1987); (2) the recognition and articulation of the role of emotion in stratification dynamics (1981,1990); and (3) a recurrent focus on the integration of macro and micro processes (1981, 1988, 2000).

---

4 Ultimately, abstract models can be used to *generate* distinct patterns, including forms of social structure that may not have historically existed but, given a theoretical concept of how social structure is composed, are possible.

Macro Context

Whether the conception is abstract or concrete, social structure is located within an enveloping context. Collins (1988, pp. 395–397) describes three dimensions of this macro milieu as space, time, and number. These three dimensions may be considered as an informational context within which abstract social structure can be defined. While space and time may reasonably be indicated as points or extents, number is inherently more complex. Collins (1988, p. 394) describes it as the "number of people or situations involved," but, as a dimension of social structure, number must be more fully described.[5]

Space is the dimension along which geographical dispersion occurs. Dispersion may take, inter alia, the form of genetic inheritance, migration, contagion, diffusion, or imitation. The spatial dimension includes race and ethnicity and also multiple layers of cultural forms. Regardless of the means of dispersion, civilization, nation, language, religion, and various cultural traditions, rituals, and practices all spread geographically. They may be regarded as layers of social differentiation, branching through space-time (cf., Cavalli-Sforza 2000). The integration of geocultural spatial layers is an essential step toward representing the complexity of social structure in coherent ways.

The IHB reference application implements a simple form of geocultural social structure in which a notional interaction among heatbugs manifesting diverse ethnic and religious patterns is modeled. Empirically, the relationship between these two geocultural structures can be fairly complex, as suggested by Figure 1, which uses census data to show the interaction between ethnicity and religion in the United Kingdom.

Geocultural layers can converge as well as diverge. Marriage and progeny can unite two ethnic groups. Children can be taught to be fluent in a second language (Laitin 1994). An innovative religious movement may borrow from and emulate a competing religious tradition. Accordingly, geocultural evolution can be best represented as a network with the potential for both divergence and convergence (rather than as a hierarchical structure).

Although not currently represented in the IHB reference application, for the sake of completeness, the other two dimensions of abstract social structure are briefly summarized as well. Time is the dimension in which social activity occurs. While the content of human activity is ceaselessly creative (Pareto 1980; Joas 1996), forms of activity recur as well. Such recurrence is recognized and functionally codified in the division of labor (Durkheim 1933; Luhmann 1982; Turner 1995; Mark 1998). One of the most fundamental divisions creates institutions that become semiautonomous from the larger community: the state in prehistory, religion in antiquity, and the economy in modernity. Each emergent institution is further functionally differentiated in historically and culturally specific ways, forming a complex network. In most cultures, there are also age- and gender-based aspects of the division of labor, yielding what might be called a biofunctional form of differentiation.

The third contextual dimension of social structure is based on the accumulation of resources. The types of resources accumulated in history have been highly diverse. Classically,

---

[5] Informally, for the purpose of the present discussion, "number" is applied to the resources that partially constitute stratification processes.

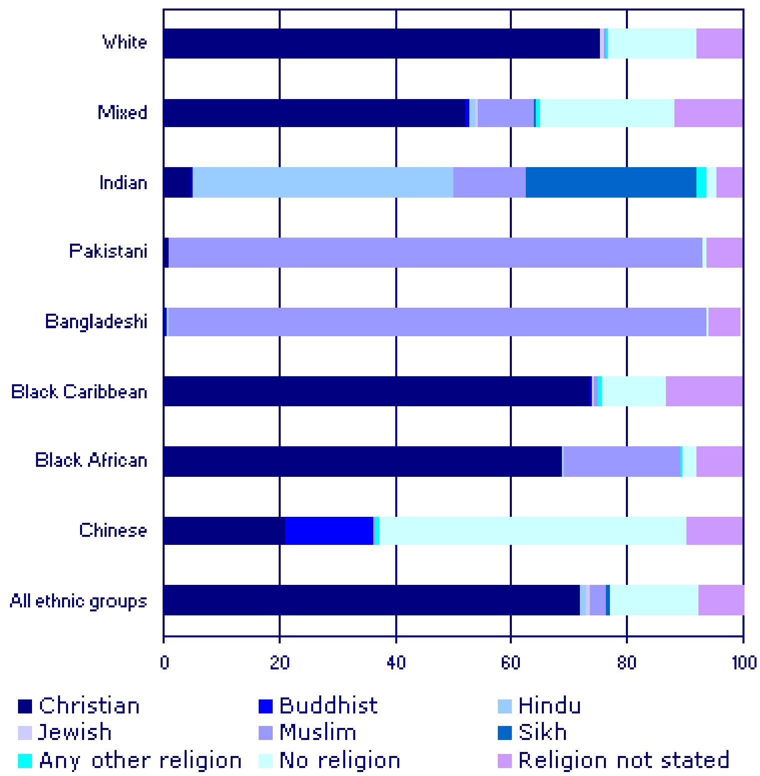

FIGURE 1 Ethnicity and religion in the United Kingdom

Weber (1968) distinguishes the accumulation of economic, social, and political resources in the form of class, status, and party. Such a high-level classification, however, does not begin to suggest the range of resource types that have been accumulated, including social deference, cattle, land, slaves, sexual access, political office and influence, precious metals, symbolic certificates of business ownership, and electronic currencies (Collins 1976, 1987).[6]

Accumulation is inherently hierarchical. In specific historical settings, each resource has formed a dimension of the situated stratification system in which parallel systems of accumulation intertwine in structured ways. It is evident that in modern society, there are numerous accumulation hierarchies, manifesting subtle and dynamic interactions that together form a complex system of stratification (Blau 1977; Zelizer 1994). Complexities, however, are present in stratification systems of simpler societies as well. Investigations reveal that such historically specific complexities should not be underestimated (Lenski 1966; Dumont 1970; Dirks 2001).

Viewed in the broadest comparative context, dispersion in space, functionality in time, and accumulation of resources together produce a diffuse coordinate system that is vast and complex but never actually encountered in any historical setting (Figure 2). On the

6 A stratification system also shapes the cognitive framework through which the world is comprehended (cf., Sallach 1974; Smith 1987; Sewell 1992).

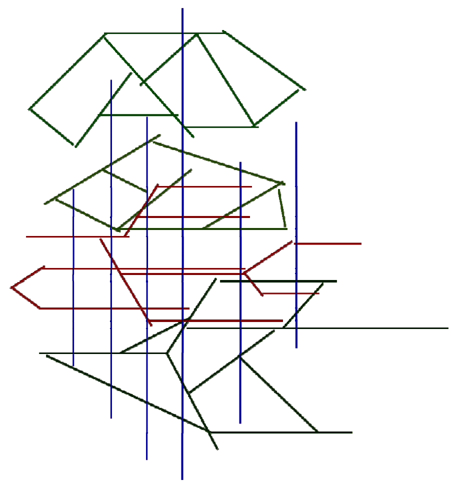

FIGURE 2 Three abstract dimensions of social structure

contrary, a coordinate system defined by macrocontext is a product of scientific theory. Defined in abstraction, *it provides a means of comparing and generalizing across cultures and time periods*.

Abstract social structure defines an analytical context in which historical social structures arise, evolve, and sometimes disintegrate. Situated structures, on the other hand, define a locus in which microinteraction draws upon and shapes the macro patterns of structure (Smith 1987; Scheff 1990; Sewell 1992; Collins 2000).

BUG ETHNICITY AND RELIGION

The reference application considered here implements two forms of geocultural structure: ethnicity and religion. The purpose is to provide forms of social structure the attribution of which shapes interactions and also serves as a generic model for applications that require richer forms of social interaction.

The number and relative size of multiple ethnicities are exogenously defined, and, subject to those constraints, bugs are randomly assigned an ethnic identity. Each ethnicity also has a defined value on a "clusivity" dimension, which varies from +1 (highest inclusivity) to –1 (highest exclusivity). The clusivity value determines the center of a range from which individual clusivity is randomly assigned.

Similarly, the number and relative size of multiple religions are also exogenously defined and randomly assigned. However, religions differ from ethnicities in two ways. First, rather than clusivity, religious identity is mediated by multiple (currently two) forms of religiosity, which determines the extent of various religious influences on the individual. Each religion is randomly assigned a value on a nice/tough dimension,[7] and religiosity determines how closely the (religious) group intensity determines individual aggressiveness and generosity values.

A second difference concerns the fact that religions can have a subreligion as well. Subreligions have distinct (and controlling) nice/tough values. However, members of other religions cannot perceive subreligion distinctiveness until or unless they have been neighbors with members of that subreligion for a specified period of time. This perceptual limitation tends to blur the visible relationship between religious identity and action; thus, it contributes to the diverse concepts of religions that compete in the larger population.

Bug action rules (that govern asking, giving, and shoving) are also mediated by ethnic and religious identities and perceptions. Rules are not the only way that prototype concepts might be translated into actions, but, in the reference application, they illustrate how casual observation, and the prototypes formed thereby, can contribute to the calibration of agent concepts and responses (Figure 3).

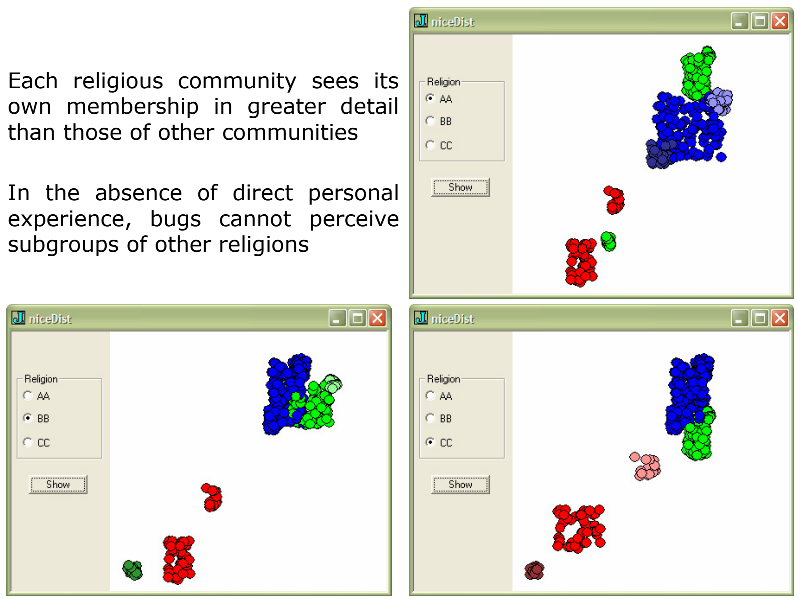

FIGURE 3 Fractal visibility

PROTOTYPE REASONING

Since antiquity, it has been assumed that the concepts employed by the human mind can be described in Aristotelian form; that is, that objects are organized by *genus* (class) coupled with *differentia* (distinguishing characteristics sufficient to produce an unambiguous definition). A bird, for example, is sometimes defined as a biped (*genus*) with feathers (*differentia*).

---

7 These two values are viewed differently, of course, depending on the perspective of the bug. Nice bugs tend to view tough bugs as mean, while tough bugs are inclined to view nice bugs as weak.

In the final quarter of the twentieth century, however, cognitive science research called the Aristotelian model into question. Seminal and well-replicated studies reveal that human conceptual structures are organized in terms of family resemblances (Rosch 1978; Heit 1997). A prototype or exemplar defines a reference point, relative to which other examples are classified in terms of their similarity, along radial dimensions of difference. Stated differently, prototypes may be regarded as a focal object or event that serves as the reference point for objects or events that are more or less similar.

An entire concept referenced by its prototype is the subject of proximity-based reasoning. Rosch (1983) refers to the overall cognitive process as reference point reasoning. The latter incorporates the typicality of any given instance relative to the radial structure of the concept as a whole. In a given situation, various prototypes may be comparatively assessed with regard to which is most appropriate for understanding the entity and/or situation at hand. When considered in conjunction with the Miller (1956) constant, an early formulation of bounded rationality constraints, the latter defines a (slightly variable) constraint that controls the number of prototypes considered in such comparisons.

In the reference application reported on here, interpretive heatbugs note the behavior of their neighbors and, on the basis of their observations, including the neighbors' ethnic and religious markers, construct *#nice, #similar,* and derivative prototypes that can then be used in applying their distinctive rules to new situations.

GENERATING SOCIAL COMPLEXITY

As Simon (1996, page 53) hypothesized, "The apparent complexity of [human] behavior over time is largely the reflection of the complexity of the environment in which we find ourselves." This insight is both a challenge and potential source of reassurance about the potential for the emergence of a truly scientific sociology. To the extent that modest but social skills enable adaptation (as well as mere coping) in strikingly diverse but always complex environments, our modeling task may be more tractable than we sometimes fear.

There are numerous sources of social complexity, and they interact in ceaseless flux. In expansive contexts, rationality bounded in space and time assures that generalizations are idiosyncratic, shared in only limited ways, and constantly evolving. Frameworks enabling coordination must be socially *achieved* by using relaxed expectations and/or carefully selective imputations of commonality. An emphasis on coordinated action as an *accomplishment* results in a healthy reorientation of social analysis toward shared *capabilities,* and they are what must (remain to) be modeled. The present initiative seeks to forge a reachable exemplar along this path.

Among our preliminary insights along the present path is that extensive and generative initialization is essential for producing the needed complex social setting and its derivatively subtle agents. It is also clear that continued progress in visualization tools is vital to illuminate regions that are important but not directly observable in empirical settings.[8] Chief among such areas is the amorphous field of intentionality. A very fluid attractor system (cf., Juarrero 2000), indeed, will be required to adequately express the interactive dynamics of intentionality.

---

8 Bell (2004) calls them "beables."

CONCLUSION

IHB is a reference application that serves as an initial implementation of the IA research program. The initiative has developed a computational implementation of three social mechanisms (prototype reasoning, situation definition, and orientation accounting) that emulate agents whose actions are *oriented by meaning*.

While this reference application uses the familiar heatbugs environment, where bugs prefer a temperate zone while escaping from settings that are uncomfortably hot or cold, such exogenous constraints serve primarily to establish the dynamic context for interpretive interaction. IHB adds the capability of bugs to help or hinder each other, mediated by modestly complex ethnic and religious identities that shape their situated responses. Prospectively, IHB provides an exemplar for diverse cultural interactions, including not only genocide and ethnic cleansing but also diversifying markets and civilizational interdependencies.

ACKNOWLEDGMENT

This work is supported by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Bell, J.S., 2004, *Speakable and Unspeakable in Quantum Mechanics*. New York: Cambridge University Press.

Blau, P.M., 1977, *Inequality and Heterogeneity: A Primitive Theory of Social Structure.* New York, NY: Free Press.

Cavalli-Sforza, L.L., 2000, *The Great Human Diasporas: The History of Diversity and Evolution.* Boston, MA: Addison-Wesley.

Collins, R., 1981, "Three faces of cruelty: Toward a comparative sociology of violence," pp. 133–158 in *Sociology since Midcentury: Essays in Theory Accumulation*. New York, NY: Academic Press.

Collins, R., 1987, "Interaction ritual chains, power and property: The micro-macro connection as an empirically based theoretical problem," pp. 193–206 in J.C. Alexander, B. Giesen, R. Munch, and N.J. Smelser, eds., *The Micro-Macro Link.* Berkeley, CA: University of California Press.

Collins, R., 1988, *Theoretical Sociology.* New York, NY: Harcourt.

Collins, R., 1990, "Stratification, emotional energy and the transient emotions," pp. 27–57 in T.D. Kemper, ed., *Research Agendas in the Sociology of Emotions.* Albany, NY: SUNY Press.

Collins, R., 2000, "Situational stratification: A micro-macro theory of inequality," *Sociological Theory* 18:17–43.

Dirks, N.B., 2001, *Castes of Mind: Colonialism and the Making of Modern India*. Princeton, NJ: Princeton University Press.

Durkheim, E., 1933, *The Division of Labor in Society*. New York, NY: Free Press.

Dumont, L., 1970, *Homo Hierarchicus: The Caste System and Its Implications.* Chicago, IL: University of Chicago Press.

Frazier, E.F., 1957, *Black Bourgeoisie: The Rise of a New Middle Class in the United States.* New York, NY: Collier Books.

Gardenfors, P., 2000, *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.

Heit, E., 1997, "Knowledge and concept learning," in K. Lamberts and D. Shanks, eds., *Knowledge Concepts and Categories*. Cambridge, MA: MIT Press.

Joas, H., 1996, *The Creativity of Action.* Chicago, IL: University of Chicago Press.

Jurgens, M., 2002, *Index Structures for Data Warehouses.* New York, NY: Springer.

Laitin, D.D., 1994, "The Tower of Babel as a coordination game: Political linguistics in Ghana," *American Political Science Review* 88 (September):622–634.

Lenski, G., 1966, *Power and Privilege: A Theory of Social Stratification.* New York, NY: McGraw Hill.

Luhmann, N., 1982, *The Differentiation of Society.* New York, NY: Columbia University Press.

Mark, N., 1998, "Beyond individual differences: Social differentiation from first principles," *American Sociological Review* 63 (June):309–330.

Mellarkod, V.S., and D.L. Sallach, 2005a, "Meaning-oriented social interactions: Conceptual prototypes, inferences and blends," presented at Lake Arrowhead Conference on Human Complex Systems.

Mellarkod, V.S., and D.L. Sallach, 2005b, "Interpretive heatbugs: Design and implementation," in C.M. Macal, M.J. North, and D. Sallach (eds.), *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms,* ANL/DIS-06-1, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 13–15.

Miller, G.A., 1956, "The magical number seven, plus or minus two: Some limits on our capacity for processing information," *Psychological Review* 63:81–97.

Pareto, V., 1980, *Compendium of General Sociology*. Minneapolis: University of Minnesota Press.

Peele, H., 2005, *Mathematical Computing in J: Introduction.* Hertfordshire, U.K.: Research Studies Press.

Rosch, E., 1978, “Principles of categorization,” in E. Rosch and B.B. Lloyd, eds., *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum.

Rosch, E., 1983, “Prototype classification and logical classification,” in E.K. Scholnick, ed., *New Trends in Conceptual Representation: Challenges to Piaget’s Theory?* Hillsdale, NJ: Lawrence Erlbaum.

Sallach, D.L., 1974, “Class domination and ideological hegemony,” *Sociological Quarterly* 15 (winter):38–50.

Sallach, D.L., 2000, “Classical social processes: Attractor and computational models,” *Journal of Mathematical Sociology* 24 (winter):245–272.

Sallach, D.L., 2003, “Interpretive agents: Identifying principles, designing mechanisms,” in *Proceedings of the Agent 2003 Conference on Challenges in Social Simulation*, presented by Argonne National Laboratory and The University of Chicago, Oct. 2–4.

Sallach, D.L., and V.S. Mellarkod, 2004, “Prototype inference and social interaction,” in C.M. Macal, D. Sallach, and M.J. North (eds.), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Sallach, D.L., and V.S. Mellarkod, 2005, “Accounting for agent orientations,” presented at Lake Arrowhead Conference on Human Complex Systems.

Scheff, T.J., 1990, *Microsociology: Discourse, Emotion and Social Structure.* Chicago, IL: University of Chicago Press.

Sewell, W.H., 1992, “A theory of structure: Duality, agency and transformation,” *American Journal of Sociology* 98 (July):1–29.

Smith, D.E., 1987, *The Everyday World as Problematic: A Feminist Sociology*. Boston, MA: Northeastern University Press.

Simon, H.A., 1996, *The Sciences of the Artificial.* Cambridge, MA: MIT Press.

Thomson, N., 2001, *J: The Natural Language for Analytic Computing.* Baldock, U.K.: Research Studies Press.

Turner, J., 1995, *Macrodynamics: Toward a Theory on the Organization of Human Populations.* New Brunswick, NJ: Rutgers University Press.

Warner, W.L., 1960, *Social Class in America: The Evaluation of Status.* New York, NY: Harper & Row.

Weber, M., 1968, *Economy and Society: An Outline of Interpretive Sociology*. New York, NY: Bedminster Press.

Widdows, D., 2004, *Geometry and Meaning.* Stanford, CA: CSLI Publications.

Zeitlin, M., and R.E. Ratcliff, 1988, *Landlords & Capitalists: The Dominant Class of Chile.* Princeton, NJ: *Princeton* University Press.

Zelizer, V.A., 1994, *The Social Meaning of Money: Pin Money, Paychecks, Poor Relief, and Other Currencies*. New York, NY: Basic Books.

INTERPRETIVE HEATBUGS: DESIGN AND IMPLEMENTATION

V.S. MELLARKOD,* Texas Tech University, Lubbock, TX,
and Argonne National Laboratory, Argonne, IL
D.L. SALLACH, Argonne National Laboratory, Argonne, IL,
and The University of Chicago, Chicago, IL

ABSTRACT

Agent-based modeling has been proven to be an effective strategy for expressing social behavior. In order to extend the veracity of social models and capture related phenomena, we introduce the concept of interpretive mechanisms in social agents. The central concept is to allow an agent to view others and the environment through his own interpretation of events and situations. There are three main mechanisms that we introduce [Sallach 2003] to capture interpretive behavior: prototype inference, social accounting, and situation definition. Agents use the mechanisms in their response cycles to infer situated meaning.

Our current work involves building a small interpretive application that can be used as a basic exemplar for the use of interpretive mechanisms in social modeling. Similar to the role played by heatbugs as a basic example for agent-based modeling; we introduce interpretive heatbugs (IHBs), which can be used as a first example for interpretive social agent modeling. This paper discusses the design and implementation strategies of interpretive heatbugs. The implementation is done in J programming language, which we are investigating as a potentially effective language for the exploration of interpretive agents.

Keywords: Agent-based models, interpretive agents, prototype inference, social accounting, situation definition, array programming

INTRODUCTION

Since antiquity, it has been assumed that the concepts employed by the human mind can be described in Aristotelian form (i.e., objects are organized by genus [class] coupled with differentia [distinguishing characteristics sufficient to produce an unambiguous definition]). A bird, for example, is sometimes defined as a biped (genus) with feathers (differentia).

In the final quarter of the twentieth century, however, cognitive science research called the Aristotelian model into question. Seminal and well-replicated studies reveal that human conceptual structures are organized in terms of family resemblances (Rosch 1978; Heit 1997). A prototype or exemplar defines a reference point, relative to which other examples are classified in terms of their similarity, along radial dimensions of difference. Stated differently, prototypes may be regarded as a focal object or event that serves as the reference point for objects or events that are more or less similar.

* *Corresponding author address*: Veena S. Mellarkod, Department of Computer Science, Texas Tech University, 2500 Broadway, Lubbock, TX 79409; email: mellarko@cs.ttu.edu.

An entire concept, referenced by its prototype, is the subject of proximity-based reasoning. Rosch (1983) refers to the overall cognitive process as reference point reasoning. The latter incorporates the typicality of any given instance relative to the radial structure of the concept as a whole. In a given situation, various prototypes may be comparatively assessed as to which one is the most appropriate for understanding the entity and/or situation at hand. When considered in conjunction with the Miller (1956) constant (an early formulation of bounded rationality constraints), the latter defines a (slightly variable) constraint that controls the number of prototypes considered in such comparisons.

The interpretive heat bugs (IHBs) application is a reference example for introducing interpretive mechanisms in agent-based modeling. This paper deals with the design and implementation of this reference application. Three mechanisms — prototype reasoning, situation definition, and orientation accounting — are the basis for introducing interpretive behavior in agents. IHBs note the behavior of their neighbors, and, on the basis of their observations (including the neighbors' ethnic and religious markers), they construct #nice, #similar, and derivative prototypes that can then be used in applying their distinctive rules to new situations. The following text briefly introduces an IHBs example; this is followed by a discussion of the design methodology.

INTERPRETIVE HEATBUGS

The IHBs concept consists of a world as an environment and bugs (agents) who live in this world. The world is a place where heat gets diffused across the surroundings according to standard heat diffusion laws (similar to heatbugs). The bugs are more complex than the regular heatbugs application; they have a religion and ethnicity. They also have religiosity and clusivity factors that portray the depth of their religious and ethnic beliefs, respectively. They have a level of aggressiveness and generosity that depend on their religious and ethnic beliefs. The agents output heat to the surroundings at constant periods of time. They also prefer a temperature zone at which they are comfortable. There are other temperature zones where they feel mildly uncomfortable (hot or cold), and at the rest of the temperatures, they are extremely uncomfortable, to the point of distress. The agents prefer being in their comfort zones; so they find the best possible neighborhood around them and try and move to that place. When confronted by another agent trying to get the same place, the agent analyzes the situation and determines whether it wants to shove others or not. The shove rules, which depend on the agent's religion and ethnicity, help him to decide. The shove rules make use of the situation that the agent currently perceives and depend on several prototypes, such as #nice and #similar. The strength of the shove depends on the agent's aggressiveness. The agent with maximum shove strength wins the place, and others get pushed backed to their original places. The agents also possess resources that they lose or gain depending on their zonal situation.

DESIGN OF IHB

The design of IHBs integrates the regular heatbug design and the response cycle of an agent. In particular, the parts that emit heat and diffuse heat are the same in both the applications. Apart from the regular initialization of parameters, IHB initialization also deals with the initialization of the prototypes the agents possess. There are two prototypes being considered for this application: #nice and #similar. Other prototypes, like #mean and #tough, can be incorporated while extending this application example to a real domain. The flowchart in Figure 1 describes the operational sequence of the application. The dotted line encompasses the step function, which is repeatedly executed to run the simulation.

At initialization, agents are randomly assigned a religion, subreligion, and ethnicity according to the distributions expressed. The initialization graphical user interface (GUI) allows a user to mention the parameters necessary for the simulation. The arbitrary religions and subreligions used are shown here:

Religion	Subreligions
AA	AA, BA, CA
BB	AB, BB, CB
CC	AC, BC, CC

The agents are also divided into different ethnicities — X, Y, and Z — where X implies exclusive, Z implies inclusive, and Y denotes no preference. The agents have resources in the nature of health that they could use.

The simulation proceeds as follows. At every step, agents emit heat, and the heat diffuses in the environment according to the standard heat diffusion laws. An agent finds a desirable position in the Moore neighborhood and would like to move there. When he takes a half step toward the desired position, he can see other agents (if any) who also want to move to the same position. The agent sees this as a new situation and would like to decide whether he still wants to pursue moving to the desirable position. There are several factors that define this situation: the agent's current position (i.e., whether he is comfortable, uncomfortable, or extremely uncomfortable); the competing agents that are involved; the agent's knowledge about these

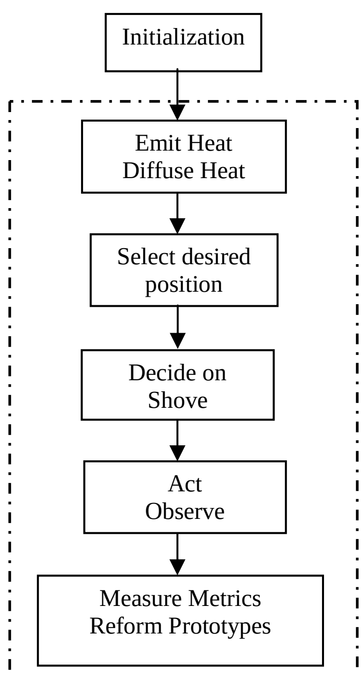

FIGURE 1 IHB flowchart

agents; and the agent's religion, subreligion, and ethnicity, which shape the agent's shove rules, which, in turn, govern the decision of whether to pursue the goal in the picture. The shove rules used are shown in Table 1. These rules use prototypes of the agents and their knowledge about the other agents that they acquire over time. After the agent analyses the situation and decides his particular response, he acts accordingly. Finally, agents observe others' actions, record them, and create/update their prototypes.

Initialization

Major parts of initialization deal with creating the initial prototypes with which the agents start. This is seen as parental knowledge contributed to the child about other ethnicities, religions, subreligions, etc. Each agent has two prototypes that need to be initialized. The #nice prototype consists of three dimensions: nz, ns, and nn. The #similar prototype consists of six dimensions: religion, subreligion, ethnicity, and sz, ss, and sn metrics. The dimensions or attributes associated with each prototype and their definitions are defined in section classification metrics. The definitions of some of these metrics use prototype clusters, which can, in turn, be generated only when the metric values are calculated. This recursive definition makes the initialization of prototypes nontrivial.

A sample set of agents is used to run the initialization. These agents undergo the same steps of simulation, with little variations, in order to build prototypes. Initialization is divided into three parts denoted as parts a, b, and c. Part a is very much like regular heatbugs with a little variation (i.e., an agent knows another agent's religion, subreligion, and ethnicity when he is in the agent's Moore neighborhood). This part is executed for an arbitrary number of times when the agents move around sequentially (as in heatbugs) and become familiar with their neighbors. Part b deals with the execution of shoving actions when an agent's requirements in shove rules get satisfied and his aggressiveness is greater than a particular level. These actions are observed by the other agents around. The metrics calculated are nz and sz, which do not require the clusters for calculation. Part b is executed for an arbitrary number of times, and the agents then have a collection of observed acts. These acts are used to find an initial similarity prototype that also uses the religion, subreligion, and ethnic dimensions. So after Part b is executed, similar prototype clusters are formed with dimensions: Religion, subreligion, ethnicity, sz metrics, and nice prototypes consist of only one dimension, which is nz. In Part c, the simulation is executed as described in the flowchart, and the other metrics that require prototype clusters use the partial clusters formed in Part b as their basis and build upon them. At the end of the initialization, the environment is again cooled, the bugs are again given arbitrary positions, and the prototype clusters that have been built are used for the agents as the initial seed.

Shove Rules

The shove rules are organized by bug ethnicity and religion. They are designed to not illustrate any essential qualities but instead to generate the diverse patterns of action, as well as uncertainty, regarding the relationship between the social structure and patterns of action. These patterns are intended to correspond to the complexities of naturally occurring cultures.

TABLE 1 IHB shove rules[a]

Prototype Shove Rules	XX		YY		ZZ	
	Rule 1	Rule 2	Rule 1	Rule 2	Rule 1	Rule 2
AAA Requisite Except	Distress None	Zone improve *#Nice #simEth*	Zone improve *#Nice* bug		Distress None	Zone improve CLUS >> CRV (& min *#simRel* & in pDistress)
AAB Requisite Except	Distress *#Nice #simEth*		Distress CLUS < CRV & *#nice #simEth* CLUS > CRV & *#nice* bug		Distress more *#nice* bugs	
AAC Requisite Except	Zone or large improve more *#nice #simEth*		Distress None	Large improve CLUS > CRV & *#nice* bug CLUS < CRV & *#nice #simEth*	Zone improve None	Large temp improve *#simRel* & in pDistress
BBA Requisite Except	Distress Very *#simBoth*		Distress *#nice* bug		Distress & all *#mean*	
BBB Requisite Except	Distress All *#simEth*		Distress CLUS >0 & *#nice* bug CLUS <0 & *#nice #simEth*		Distress CLUS >>0 & very *#simRel*	
BBC Requisite Except	Zone improve None	Temp improve More very *#simBoth*	Temp improve CLUS< 0 & (more *#simEth* & in pWorse zone)		Temp improv None	

TABLE 1 (Cont.)

Prototype Shove Rules	XX		YY		ZZ	
	Rule 1	Rule 2	Rule 1	Rule 2	Rule 1	Rule 2
CCA						
Requisite	Zone improve	Temp improve	Zone improve	Temp improve	Zone improve	Temp improve
Except	None	More very *#simBoth*	None	CLUS < 0 & all *#simEth*	None	Very *#simRel*
CCB						
Requisite	Zone improve		Distress	Zone improve	Zone improve	
Except	More *#simEth*		None	CLUS < CRV & (*#simEth* & in pDistress) CLUS > CRV & (very *#simRel* & in pDistress)	*#nice* very *#simRel* & in pDistress	
CCC						
Requisite	Zone or large improve		Zone or large improve		Zone or large improve	
Except	More *#simEth*		CLUS < CRV & more *#simEth* CLUS > CRV more very *#simRel*		CLUS >> CRV & very *#simRel*	

[a] pDistress and pWorse stand for "projected" distress and worse, respectively, meaning it is "distress" and "worse" from the decision-making bug's perspective. CRV is the cluster reference value used to subdivide the bug category. In baseline IHB, CRV = 0. Italics indicate prototypes.

IHB rules have two parts: prerequisites and exceptions. Prerequisites concern heat conditions and heat tolerance. They are exogenously and stochastically defined by ethnic and religious groups, but they are also distinctly determined for each bug. Exceptions address how the bugs recognize other bugs, particularly with regard to what categories of bugs are (or are not) taken into account in the decision to shove (or not shove). Exceptions rely heavily on #similar and #nice prototypes and their derivatives (e.g., #mean, #simEth, #simRel).

Prototypes are calculated by using hierarchical agglomerative cluster analysis using a Euclidean distance metric. Bugs classify a baseline of cases during initialization (analogous to parental instruction) and continue to observe and classify the actions of other bugs during the simulation run. These clusters form the basis of the bug's prototypes and are used to uniquely activate the bug's shove rules.

Classification Metrics

A bug's prototype clusters represent the bug's notion/idea/concept of other bugs. This is calculated by observing the actions of other bugs. A bug can observe the actions of others in two ways: as an observer who actively looks at a desired cell and the situation in that cell or as an actor who acts (by shoving, etc.) for a particular cell and observes others' actions related to that cell. Let us refer to a specific bug (either an observer or an actor) as a "self" that is actively classifying the actions of other bugs of interest and creating a world of its own thoughts, interpretations, and conclusions about the others. This bug's prototype clusters are these concepts. The two prototypes being dealt with here are #similar and #nice. The others of interest are named as neighbors.[1]

The self observes the actions of the neighbors and calculates some metrics as a means of interpreting their actions. The action of a neighbor is toward other neighbors (which may or may not include self). While calculating the metrics, self looks at a neighbor's action with respect to himself and others affected by the action. (Thus, there may be multiple actors in a situation.) Each acting neighbor is analyzed separately, and metrics for each one are calculated. During a given assessment, the neighbor being observed is in the role of an actor, while others in the same observation can be seen as reactors. While the action of a neighbor is being assessed by self, others are reactors. The resulting metrics can be sophisticated and/or widely divergent. We use two types of simple metrics to define the actions: #nice and #similar. The next few paragraphs define the current metrics.

Nice Metrics

The nice clusters are defined by using a set of three metrics: nz, ns, nn. The following discussion considers each of them in detail. We start with the nz metric. It captures an actor's niceness as shown toward his neighbors in extreme temperatures. Sympathetic actions of an actor are given more recognition in extreme heat zones than in comfortable zones. More precisely, the nz metric says, "Not shoving another in extreme situations is nicer than not shoving another in more comfortable situations." The degree of niceness is measured by using an nz metric table. The action of an actor with respect to others' situations is given a niceness measure between –1 and 1 by using Table 2.

This is the nz metric related to the not-shove action. The value in the first column and bottom row of 0.9 can be read as follows: "The actor who did not shove another in order to move from a distress situation to a comfortable situation is 0.9 nice according to the nz metric." A similar table representing metrics for the shove action is shown in Table 3. As mentioned earlier, these metrics can be as complex as necessary. Here, for the sake of clarity, we chose to keep it simple. It should be noted that the operative definitions are perspectival. For example, the self

1 One point to note is that these "neighbors" do not need to be self's Moore neighbors.

TABLE 2 Not-shove niceness metric

Not Shove From ↓ To →	Comfortable	Uncomfortable	Distress
Comfortable	0.2	0.15	0.1
Uncomfortable	0.6	0.5	0.4
Distress	0.9	0.6	0.2

TABLE 3 Shove niceness metric

Shove From ↓ To →	Comfortable	Uncomfortable	Distress
Comfortable	-0.8	-0.85	-0.9
Uncomfortable	-0.4	-0.5	-0.6
Distress	-0.1	-0.4	-0.8

sees an actor to be in a distress situation and wants to move to a comfortable situation. The actor, on the other hand, may have a different definition of distress and thus might be comfortable in his present position and hence not want to shove.

The ns metric deals with how actors behave toward similar bugs. These statements apply: “Not shoving a dissimilar bug is nicer than not shoving a similar bug,” and “Shoving a similar bug is regarded as being less nice.” The metric is approximated in Figure 2.

Finally, the nn metric deals with how actors behave toward nice bugs. “Not shoving a bug that is less nice is nicer than not shoving a bug that is more nice,” since not shoving a nicer bug is expected. This metric is shown in Figure 3.

Similarity Metrics

Apart from religion, subreligion, and ethnicity factors, the following three metrics are used to define the similarity clusters: sz, ss, and sn. The sz metric represents the similarity in dealing with the situation. The self looks at the action of the actor and its situation (zone and reactors) and decides whether it (self) would act in the same way. The similarity/dissimilarity in action is represented by using a metric table called the difference metric. The ss metric represents the current similarity of the actor to self as seen by self. The metric is calculated as SCDsa, which is the similarity cluster distance between self and the actor. The sn metric represents the similarity of the actor with respect to self in niceness level. It is calculated as NCDsa, which is the nice cluster distance between self and actor.

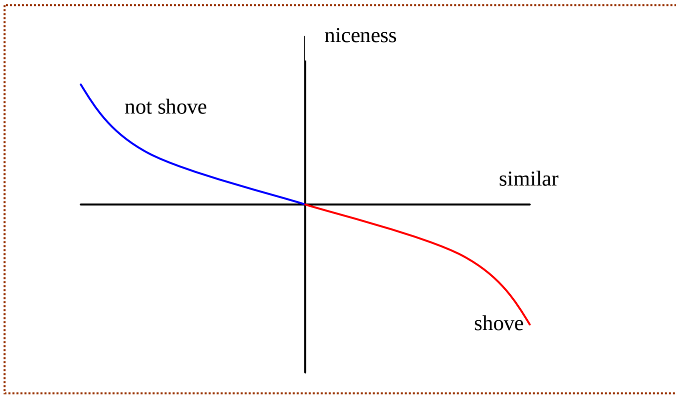

Average $[(SC_{dan} – SC_{davg}) * I * \alpha_s]/ SC_{davg}$.
For shove: I = 0 if $(SC_{dan} – SC_{davg}) > 0$ and I = 1 if otherwise.
For not shove: I = 1 if $(SC_{dan} – SC_{davg}) > 0$ and I = 0 if otherwise. SC_{dan} is the similar cluster distance between the actor and reactor, SC_{davg} is the cluster average distance of the similar clusters of self, I is an indicator function defined as above, and α_s is the salience given by self to this metric and is set to 1 as the default.

FIGURE 2 Notional interaction of two prototypes

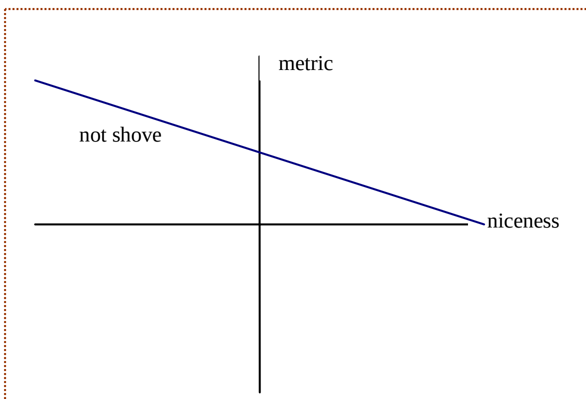

The metric is calculated as Average (NCDnv), where NCDnv is the nice cluster distance between the acceptor and self's view of very nice cluster.

FIGURE 3 Relative niceness metric

PROTOTYPE CLUSTERS

The actions executed by the neighbors are recorded by proximate observers in two tables: events and event-members. These tables are then used to calculate the metrics defined above. There are five values recorded: from-zone, to-zone, actor, action, and event-num. The from-zone and to-zone are the actor's current position and desired position (for which he shoves or not) as seen by the observer/self. The action is shove or not-shove. The event-num keeps track of all the actors in a particular event or situation that the self observes. This is used while calculating the metrics: to recognize the actors and reactors. The metrics are then used to form the prototype clusters.

A random selection within the Miller magic number range (7± 2) is used to form the clusters. The cluster algorithm used is hierarchical agglomerative clustering. The clusters thus formed are used in shove rules. For example, the shove rule 2 for AAAA and X bug says, "If the move gives a zone improvement, then shove, unless there is a #nice and #similar bug with the same ethnicity." Thus, if self finds itself in a situation where there are bugs competing for the desirable position, it will shove unless it recognizes one of the competing bugs as a nice and similar bug with ethnicity X. A bug is considered nice if it is in the very #nice cluster or the next-closest cluster to it, as attributed by the self. A bug is considered similar if the self and the bug are members of the same or adjacent #similar clusters. A fast join, which is the intersection of the two clusters, is performed. If a competing bug is a member of this similar-nice joining, then the self does not shove and remains in its place.

IHB IMPLEMENTATION

IHB implementation was done in an array programming language called J (Thomson 2001; Peele 2005). J is a mathematical language containing high-level primitives useful for building complex programs in fewer lines of code. This implementation was also a test for experimenting with using J as a language for agent-based modeling and simulation.

The initial user interface allows the user to change some of the parameters for the simulation. The heatbug interface showing the heatbugs and the environment looks like a standard interface of heatbugs. It is shown in Figure 4a. The colors of the bugs depend on their religions and subreligions. The greens are AAs, with three shades of green for three subreligions. The yellows are BBs, and the blues are CCs. Lighter and darker shades are given for each color to show three subreligions in each group. The shapes define the ethnicity to which the bugs belong. The Xs, who tend to be exclusive, are rectangles; the Zs, who tend to be inclusive, are circles; and the Ys, who are neither, are oval.

The resources of the bugs increase when they are comfortable and decrease when they are in distress. These resources combine with the bugs' aggressiveness factor to give strength to their shoves. The bug with the highest strength in shoving wins the position, and the resources get depleted by the amount used. If there is a tie, a random contender wins the position. If all the agents decide not to shove, then none of them move into the position.

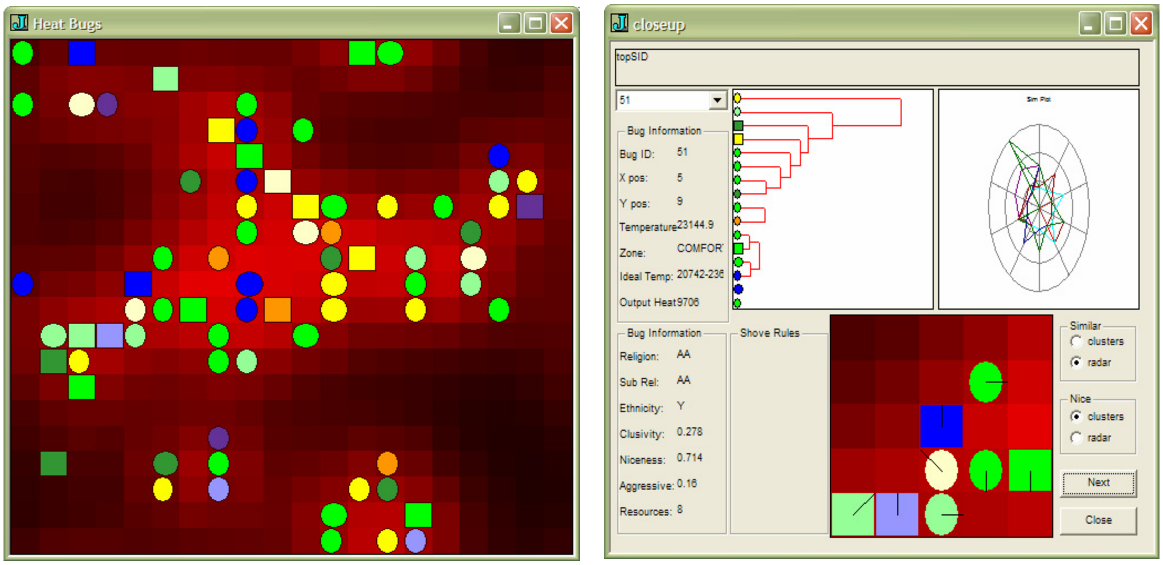

FIGURE 4 Interpretive heatbugs (a) and CloseUp interfaces (b)

Interaction Walkthrough

The interpretive agent's interaction with the environment is fairly complex. Because it is difficult to understand the mechanisms in operation, it is useful to get a close view of interesting interactions and their effects. A CloseUp interaction walkthrough visualization was designed to work toward this goal. A snapshot of the interface is shown in Figure 4b. The CloseUp can be opened by double-clicking any bug on the main window. The bugs in the extended Moore neighborhood of range two are shown in the lower window. Each of the bugs in this CloseUp can be analyzed by looking at his internal properties and his nice and similar clusters, or the radar plots of their dimensions can be displayed. The shove rules are shown, and a user can walk through the four steps involved in the whole process by clicking the next button. In the first step, the bugs emit heat, and the heat diffuses. In the second step, the desired neighboring Moore cell is found. The third step presents a new situation to the bugs, wherein they can see who the competitors who are aspiring for the same position are, and they decide whether to use force to get the position. The fourth step implements the actions, and the bugs observe the situation from their viewpoints and re-categorize their clusters accordingly. Using the CloseUp interface, a user can analyze the microinteractions among the agents. Each agent in the CloseUp can be selected from the drop-down menu, and all the bug's attributes, including its current prototype definitions, are displayed. We believe that such interfaces will, in general, help in understanding details at the microlevel and thus help in analyzing emergent behaviors.

CONCLUSIONS

Interpretive agents in agent-based modeling represent a new area of research that we believe has the potential to decrease the grain and increase the veracity of social models, thereby increasing the potential for representing nonlinear and other phenomena of interest. We introduce a reference application as a resource for interpretive agent research and describe the design and implementation of this application.

ACKNOWLEDGMENT

This work is sponsored by the U.S. Department of Energy, Office of Science, under contract W-31-109-Eng-38.

REFERENCES

Blau, P.M. 1977. *Inequality and Heterogeneity: A Primitive Theory of Social Structure.* New York, NY: Free Press.

Collins, R. 1981. "Three faces of cruelty: Toward a comparative sociology of violence," pp. 133–158 in *Sociology Since Midcentury: Essays in Theory Accumulation.* New York, NY: Academic Press.

Collins, R. 1987. "Interaction ritual chains, power and property: The micro-macro connection as an empirically based theoretical problem," pp. 193–206 in J.C. Alexander, B. Giesen, R. Munch, and N.J. Smelser, eds., *The Micro-Macro Link.* Berkeley, CA: University of California Press.

Collins, R. 1990. "Stratification, emotional energy and the transient emotions," pp. 27–57 in T.D. Kemper, ed., *Research Agendas in the Sociology of Emotions*. Albany, NY: SUNY Press.

Collins, R. 2000. "Situational stratification: A micro-macro theory of inequality," *Sociological Theory* 18:17–43.

Heit, E. 1997. "Knowledge and concept learning," in K. Lamberts and D. Shanks, eds., *Knowledge Concepts and Categories*. Cambridge, MA: MIT Press.

Joas, H. 1996. *The Creativity of Action.* Chicago, IL: The University of Chicago Press.

Jurgens, M. 2002. *Index Structures for Data Warehouses.* New York. NY: Springer.

Mellarkod, V.S., and D.L. Sallach. 2005. "Meaning-oriented social interactions: Conceptual prototypes, inferences and blends," presented at the Lake Arrowhead Conference on Human Complex Systems.

Miller, G.A. 1956. "The magical number seven, plus or minus two: Some limits on our capacity for processing information," *Psychological Review 63*:81–97.

Peele, H. 2005. *Mathematical Computing in J: Introduction.* Hertfordshire, UK: Research Studies Press.

Rosch, E. 1978. "Principles of categorization," in E. Rosch and B.B. Lloyd, eds., *Cognition and Categorization.* Hillsdale, NJ: Lawrence Erlbaum.

Rosch, E. 1983. "Prototype classification and logical classification," in E.K. Scholnick, ed., *New Trends in Conceptual Representation: Challenges to Piaget's Theory?* Hillsdale, NJ: Lawrence Erlbaum.

Sallach, D.L. 2000. "Classical social processes: Attractor and computational models," *Journal of Mathematical Sociology 24*:245–272 (winter).

Sallach, D.L. 2003. "Interpretive agents: Identifying principles, designing mechanisms," in *Proceedings of the Agent 2003 Conference on Challenges in Social Simulation*, presented by Argonne National Laboratory and The University of Chicago, Oct. 2–4.

Sallach, D.L., and V.S. Mellarkod. 2004. "Prototype inference and social interaction," in C.M. Macal, D. Sallach, and M.J. North (eds.), *Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence*, ANL/DIS-05-6, co-sponsored by The University of Chicago and Argonne National Laboratory, Oct. 7–9.

Sallach, D.L., and V.S. Mellarkod. 2005. "Accounting for agent orientations," presented at the Lake Arrowhead Conference on Human Complex Systems.

Sallach, D.L., and V.S. Mellarkod. 2005. "Interpretive agents: A heatbug reference simulation," in C.M. Macal, D. Sallach, and M.J. North (eds.), *Proceedings of the Agent 2005 Conference on Generative Social Processes, Models, and Mechanisms*, ANL/DIS-06-1, co-sponsored by Argonne National Laboratory and The University of Chicago, Oct. 13-15.

Scheff, T.J. 1990. *Microsociology: Discourse, Emotion and Social Structure.* Chicago, IL: The University of Chicago Press.

Thomson, N. 2001. *J: The Natural Language for Analytic Computing.* Baldock, UK: Research Studies Press.

DISCUSSION

Computational Social Theory

(Patterns and Actions, Saturday, October 15, 2005, 1:00–3:00 p.m.)

Chair and Discussant: *William Griffin, Arizona State University*

Charles Macal: Before we begin this first session of the afternoon, I have couple of slides to talk about the proceedings and the papers. First, the proceedings for all Agent conferences are available on the conference website, and they can be referenced. They all have ISBN numbers. Along with accessibility, we've tried to make them as real as possible. Also, we record the question and answer sessions and make them available as transcripts in printed form. This makes it possible for us to go back to 1999, for example, and see what someone said about validation or to review other comments. In some cases, of course, we are still saying much the same today, six years later. In other cases, you can see that the field as a whole has made progress.

Some people use these proceedings as part of their courses, as supplemental material, either because of the papers or the discussion that followed. The discussion reproduced for the Agent conferences is very unique, and I think it's a valuable adjunct to the proceedings for the conference — the record of the conference — that we would very much like to maintain. Now I'm going to turn the session over to Bill Griffin.

Power-law-like Distributions: A Practical Survey

William Griffin: The first speaker is Ana Carrie from Trinity College in Dublin, and she's going to talk about power-law distribution. She's going to offer some alternatives and some interpretations. For those of us who do a lot of this work, power-law distributions always come up. Do our data fit? I mean natural-occurring data. How do we interpret the frequency of events? Do they fit a log-law plot is basically the big question?

Ana Carrie: Hello. My name is Ana Carrie. I'm a Ph.D. student in the Department of Economics at Trinity College in Dublin, Ireland. Some of you might recognize my opening slide. It's the library of the training school for Jedi knights in the most recent *Star Wars* movie; it also bears a striking resemblance to the long room in the old library in Trinity College.

This is a summary of my talk today, and those of you at Nexus might recognize this. This was a panel in my poster. A lot of us have come across power-law distributions in our work. Today, I'm looking at some of the statistical issues that arise in dealing with power-law distributions, and there is a difficulty because a power law is basically a straight line log-log-plot. That's a very simple concept. It's very easy to get seduced by that straight line — just fit a straight line and not look at it any further. I'm not particularly qualified to do this, and I'm not the only person who has done this. It just came up in the course of my work that I wanted to look at some of the alternative distributions. Today, I want to talk about some of the practical issues that arise when you're working with this type of data.

[Presentation]

Griffin: Thank you. Questions?

Macal: I'd like to make a comment. I can see where there's a real need and a utility for this kind of thing with regard to agent modeling in the following sense. In some of our models, and one Keven Ruby mentioned this morning, agents in the occupational dynamics model have friendship, well, basically friendship networks, that we generate based on the existing algorithms in the literature, whether it's the Jin-Girvan-Newman algorithm or other types of algorithms. But those algorithms are approximations to generating social networks, which have an apparent structure that looks similar to real social networks. The distributions that you're coming up with, however, could be applied to our simulated or algorithmically created social networks and determine what best characterizes what we've created via the algorithm. This compilation of power-law-type distributions could be quite valuable in the broader agent community.

Carrie: Yes, I think we're still thinking in terms of normally distributed errors, you know, when for a lot of us, the normal distribution is in no way relevant. Yes, I think it would be useful. Power laws are a buzzword, which is good to a point, but I think now we have to get past the buzzword stage and start making sure we're actually doing good statistics.

I also think that if we're going to be generating or creating models, then, yes, decide which distribution our parameters should come from and have a good choice and know what the different generative models are underlying. We've got preferential attachment and several others for power-laws. We have some interesting relations between the generative models for the log-normal, and certainly if we're going to be creating, why not populate our agents with random draws from some of these other distributions as well as uniform or normal or something. That's a good point.

Kostas Alexandridis: Well, leaning toward more nonparametric estimation properties, I understand what you're saying, especially in terms of ordinal-type data, but when we started looking at nominal databases, and especially in relation to agent-based modeling where the model simulates some kind of real-world process, there has to be some kind of physical explanation. Do you have any comments on what you think those distributions might simulate and how they can be interpreted in the real world?

Carrie: One thing that's happened in the literature is that there are several physical models that generate power-law distributions. One thing people are implicitly saying is that they have these data that look like a power law, and therefore, it's preferential attachment or it's one of these underlying physical models. I think what's going to be important is that before we go down the road of commenting on the physical process, we need to be more careful with our statistics. Generally there are several different physical processes that could lead to a certain distribution, and they're very interesting. I think the physical models are very interesting. We want to know when the same physical model is producing a power law in two completely different areas — one in biology and one in astrophysics. If it's the same physical process that's doing that, it's very interesting.

So generally, sometimes there is more than one physical process that's going to lead to an identical distribution. I do think that being able to say a bit more quantitatively what our distribution is, is going to help us with making those interpretations where they might be appropriate.

John Sullivan: Yes. I'm a little confused on your motive, though. Are you actually trying to understand by looking at the underlying data or details of underlying physical processes? Or is this merely representational, so somebody can use it and at some point call it up? If it is the latter, then why not use some kind of spline fit or something like that? It appears in the table and it's convenient and easy to use.

Carrie: I think the motivation for what you're going to do with this is that everybody is going to have their own…. Well, I think some people are going to want to look at the process, while other people are just going to, like you say, fit a line to it. I think that where I'm probably coming from in my own personal motivation is just wanting to be very, very precise. I think too many people are just saying that it's a power law, so let's fit a straight line and talk about the slope. Are they making further conclusions or not? Maybe they're not. But a lot of people are using the term very loosely. So my very, very specific physical motivation was that there are all these different distributions, and people are talking about — well, the data's not quite linear. Then what? So, at a very practical level, it's a case of giving people a list of things to do for the 'then-what' situation. So let me fit one of these alternatives.

I think it's up to the researchers to decide whether they'd like to look at the physical interpretation or not. Like you say, if they don't, then, sure, just fit a line to it. I think, though, if we can standardize our terminology and if we're talking about a power law, we really mean a straight-line fit with a certain confidence, or we mean a parabolic line fit with a certain confidence, or an exponential curve. I suppose this is just a thing to kick-start people to be a bit more precise in what they mean, and then they can decide at that point. Personally, I would like to get into the physical process aspect of it because I think that's interesting. But I think this has to come first.

Griffin: Thank you. Very good.

Tolerance and Sexual Attraction in Despotic Societies: A Replication and Analysis of Hemelrijk (2002)

Griffin: Next we have Joanna Bryson who will talk about tolerance and sexual attraction in despotic societies.

Joanna Bryson: Thank you. The models I'm going to present here were originally written by JingJing Wang and then they were refined and finished by Hyde and Hagen Lehmann. Hagen will be the first author in the paper that will be in the proceedings. My talk today is from the extended journal version, where I'm going to be the first author.

[Presentation]

Griffin: Very good. Thank you. Are there any questions?

Ana Carrie: Did you actually answer your question about why the agent-based models weren't accepted into the literature?

Bryson: Yes, but maybe I didn't go into that enough. Basically, when you talked to people, they didn't believe that the monkeys are that noncognitive, but, said that she has a model and she has equations, and we don't know what to do. So it was a combination of not matching their instincts from what they've seen and finding the way the data were presented and accessible so they couldn't say in a scientific way what was wrong. I think that's why I'm trying to help them because they should have been able to go in.

You'll see in the paper that you have the six questions but not the answers. When I presented that paper to Bernard Thierry, he said that he now knew the answer. Bernard has also closely collaborated with Charlotte Hemelrijk, so the point I'm trying to make is that we're just trying to move on. This is what science is about — that we all keep trying to find errors in our own work and errors in other people's work so that we keep getting closer to the truth. I'm in no way trying to undermine anybody. But, hopefully, people will start referencing this stuff because they'll understand it better.

Venkatesh Mysore: Venkatesh Mysore from New York University. I have a very general question. What other models are available? What avenues of modeling other than agent-based modeling are available for primate sociology? Why should we choose agent-based modeling?

Bryson: Well, okay. In general, the biologists are at a very, very descriptive level, so in a way, they don't do much modeling. Does this or does this not describe a species? They keep looking for simple mathematical rules about it or characterizations. Jessica Flack, who used to work with Frans De Waal for her Ph.D., did a bit of work with an economist, where they came up with criteria for trying to find the discriminations between these different things. But it's very, very qualitative right now.

Like I said, I think this area is ripe for agent-based modeling. One of the interesting things, compared to doing human work, is that I think it's useful to understand the basics of primate motivation, and we are primates.

People can go out and watch primates all day and see every single close interaction they have. They count everything. It's amazing what you get graduate students to do just so they can be out in the jungle because they love the monkeys. It's like a soap opera. Every day they come back, and it's weird when you see primatologists get back together. The first thing they talk about is like, what, who, which monkey has had a baby and which monkey started hitting the other monkey. It's totally a soap opera. Anyway, you can collect data you could never collect on humans. Basically, if a human was there, the other humans would start behaving differently, but we get really, really good data on social interaction out of nonhuman primates.

Griffin: Joanna, I have a question for you. In your role as a modeler, what's the correspondence between what you bring to the primatologist and what — is this a rule, or I can give you a rule? You give me the behavior, and I'll give you a rule. What's the correspondence there?

Bryson: Well, honestly, I hope that being a modeler is just one thing in my repertoire. I hope that I *am* a scientist. I think to do a good job going out there and being able to cope. So when we went out and worked with Thierry, and, again, Thierry had been aware for a very long time about how this — in fact, he was one of the few people who did reference their work and who's been aware of her work for a long time. But he still hadn't understood it as well before he read the paper. Then we were able to sit down, and I thought that we're going to have to generate and do tests. We'll generate some models, and he'll tell us if they're right or not. But he actually was able to come down and sit down and think computationally, and the three of us — Hagen, myself, and Bernard —were able to sit around and come up with the next model.

So it is both testing — I mean, the ripe area is that so many people are drawn by scenario diagrams and never test them to see what behavior they would actually produce. So as people that are running the experiments for scientists, we can then be …. I think once you've built that, it's a good way to start. Once you do that, though, you're going to start coming up with another hypothesis. If you can do that collaboratively with the domain experts, then obviously you're participating in the process.

Larry Kuznar: Larry Kuznar from Indiana-Purdue, Fort Wayne. I would like to applaud your whole approach because, and building off Bill's comments, for those of us that do field research and are kind of domain experts, the first thing I have to say to modelers is that they should please give us our due. However, as we work through these models, and I've had the same experience: we begin to discover things, just as you pointed out, and then we're making some assumptions that might not be empirically warranted.

Bryson: Exactly.

Kuznar: That becomes a blueprint for the next field project. Lee, as an anthropologist, you're probably familiar with *Notes and Queries*, which, of course, has been out of print for a very long time. It was a blueprint of what an anthropologist should gather when going to the field. I really think that modeling process can become a blueprint, especially for problem-oriented research.

Bryson: Yes. I thought when we presented these six questions that they were going to have to hire a grad student to go out into the field and find out this stuff. It was interesting, though, because of the fact that the three things that we don't know, we're not going to know, right? It was such a low-frequency event that we couldn't gather enough data. So we could write something, but we'd have to write something for the next century. Something like we're going to observe 10 troops and see how … you know, we wouldn't get the numbers to have good stats. It's just impossible. And then the other three already had the data. He just pulled out as a domain expert. He just pulled out papers, and most of them were published papers because we have to cite those, but he also had tons of unpublished data he was just pulling out of these directories, and he knew what was going on. He could either motivate someone to publish the data that they'd been letting sit on the shelf, or quite often that stuff was out in the literature and people just didn't see the correlation. But, as you say, we're finding those holes then getting people to see that there is another way to apply that information. It just helps you to understand the problem better.

Carrie: Those changes in rank don't happen very often. Is that in itself a criticism of the model?

Bryson: Yes, absolutely. Yes. No, it wasn't one that I'd set out in the initial six points, and I tried to make that clear from that slide, but I didn't do a good job. The reason that the questions didn't make sense was because of something that was even more basic than we had even thought to ask. That is, does it actually happen? No, it doesn't happen.

Griffin: One more question.

Mysore: We are having a similar problem. We're trying to model disease by talking to doctors. Looking at their perception of a disease and how it evolves at times is very different from how a modeler time codes the disease. I also get the same feeling that, while we are primarily modelers, we are helping to crystallize some thoughts that enhance how we relate to the advancement of science.

Bryson: I totally agree, and I have one other thing about that. In the UK, a lot of people in AI were trying to build things to help biologists, and they were completely failing. They couldn't understand why the biologists weren't looking at their brilliant robots and whatever. I think, being a modeler, that you're an ambassador, and it's your job to go out and learn the language of the locals. They aren't going to come to you. But to make that work, you do have to go out and find out how to communicate with them. One of the things I've done is grab Ph.D. students out of that society, teach them, and make them the translator. But yes. Thank you for your comment.

Interpretive Agents: A Heatbug Reference Simulation

Interpretive Heatbugs: Design and Implementation

Griffin: Next up is David Sallach who will talk about new heatbugs, a new kind of heatbug.

David Sallach: Two years ago, I had a conceptual paper about how to build interpretive agents that would be oriented by meaning, but without it being something like a full AI application. Now we are moving to the point of having a working model of this, and that's what this paper and the next paper will report on today. Basically, I'm going to try to motivate it and give you a sense of the kind of mechanisms and how they operate and where they exist in overall design space. The next paper will focus more on implementation.

In the first place, this is designed to be a reference application, so that the mechanisms are defined as a reference that can be then plugged into a variety of different types of models. I'll briefly review the concepts behind interpretative agents, talk about the strategy of this implementation, and then get into the specifics of the process.

[Presentation]

Griffin: Very nice. Excellent. I had the opportunity to read David's paper before the conference, and so there's a lot he left out. Obviously, you will get a chance to read that. Are there any questions for David before we go on to Veena? Is your talk following up exactly on that? Do you want to hold up questions then?

Sallach: We could.

Griffin: That's what I was thinking.

Sallach: It might be better because the second talk might answer some of the questions.

Griffin: Yes, so Veena's going to do the implementation or the technical part of that. We'll hold questions off until she has finished.

Veena Mellarkod: Okay. This presentation is a follow-up on David's talk, but it is more about the design and implementation of the heatbugs that you were hearing about. We are interested in introducing interpretive mechanisms into agent-based models. Let me first give a very small recap for everybody.

Interpretive capabilities allow an agent to view others and its environment through its own interpretation of certain situations. For instance, if two agents are performing or executing the same action, then another heatbug can look at these two actions and interpret them both differently. That is one of our ideas, one of the things that we want to capture. We have designed interpretive heatbugs by introducing these interpretive capabilities using the regular heatbug design. We think that it's a basic exemplar for interpretive mechanisms in social modeling.

[Presentation]

Robert Reynolds: Bob Reynolds, Wayne State University. Very interesting. Could I ask this question to both of you? It seems like your agents have a shared ontology. But, in fact, each ontology is a set of concepts with metrics to assess them. Based on this, they build up, I would say, somatic networks where they characterize objects relative to their observations on these metrics. And so each individual develops a somatic network or maybe a frame-based system describing their own information. So they're building up their interpretations of the generic ontology.

One of the interesting things I was thinking about is if — but right now they bump into each other, but they don't exchange their interpretations. It would be interesting when agents bump into each other to look at exchanging some aspect of their somatic network to see whether or not you can get some generalized cultural phenomena emerging from that interaction. And the whole idea, too, is that you may have some contradictions in perspective that might have to be dealt with in that prospect too.

Sallach: No, absolutely. According to what we're thinking, they'll never hand each other a prototype. I mean, we're not going to have any omniscient agent, but they will have to have cooperative action, and there's going to have to be a certain similarity. That may mean that each of them may have to relax some of their prototype concepts to be able to work together, but they may not completely know what they've relaxed. I mean, they agree to relax it, but they don't know what it is that they've relaxed or how strongly they hold to that kind of thing. So the basic — certainly, one of the directions we want to head in, is that cooperative action is an achievement, and not just an automatic ….

Unidentified Speaker: *De facto*….

Sallach: Yes, exactly.

Bryson: I have a quick question. I share your dislike for tagging, and I like the direction you're going with this. I may have missed this in the first talk, but are you comparing this against a fixed data set? You mentioned a lot of things that were similar, but is there something you're trying to replicate, or is it mostly just an exploration of the space?

Sallach: Yes, at this point, we're not comparing it to a fixed data set. At this point, this is an attempt to say that we should pick a simple application and show how these mechanisms would work. Focusing on a more empirical framework would be the next step.

John Sullivan: Yes, it's clear that this particular model has behaviors that are becoming rather complicated, but interesting. But coupling what this model is capable of doing with some of the things that Joanna was talking about earlier, particularly in regard to modeling primate behavior, is there anything interesting that could be learned from the reverse problem? That is, could you have someone watch a display like this and infer the behavior from that, especially those who want to make a living or do research in a field where they're looking at animals and what have you, to try to glean insights or understanding of their behavior?

Sallach: It's an interesting idea. I hadn't thought about that.

Mysore: I just want to ask you to reason what is, in a way that's easy for me to understand, in terms of interpretive reasons. One way of thinking about agents and interaction is that they have knowledge and belief states; they have a model of the world that they live in, and they continually update interactions in this model. Would it be possible to compare with that kind of notion as to having a belief state or a knowledge state of the environment, and what are the communication channels that you effectively provide, and how does this knowledge and belief state change? Similarly, another notion is that an agent is a bunch of parameters that govern the agent's behavior. You could have meta-rules that govern the evolution of these parameters with interactions. Would it be possible to summarize what you've done in terms of this?

Sallach: In the first place, I think that it's closer to the first example, your first case, right now. I would say that one of the things that we're trying to do is to get away from discreet, well-bounded categories and get into a framework where there is a more or less continuous space. That's not to say that there are no discreet elements within it, but you've got a variegated space where a number of the dimensions are continuous and where the decision concerning which parts of the space are relevant to what is endogenous, and what is taken into account. The agent orientation is also unique. I mean, it's based on the individuated experience of the bug. And in that combination you're getting what psychologically you might call an attribution — you're attributing which categories are relevant for what purposes. What is unique is that the bugs are not going to be making attributions in the same way, but it's closer to the internal conceptual model right now.

Mysore: I have one more unrelated remark to make about the advantage of the heatbugs. First of all, it's great that you're developing more reference applications because it really quickens the process of somebody who's just a computer scientist who's trying to get in like me, take a bunch of examples to make the quota, and get what I'm doing. So I really appreciate your effort. This is just a comment about the choice of colors for the agents.

So the heatbugs model is appealing because you can see and understand what is going on, but the moment you introduce too many colors and shapes, you cannot get any quick idea of the global behavior of what's happening. Maybe you can think about something to do with the color scheme there. Thanks.

Sallach: Okay. Well, we're between three and nine colors, right? We've really got three broad colors for the three religions, but then because we've got subreligions, it ends up being somewhat of a spectrum.

Griffin: Dave, I've got a question. Following up on Bob's earlier question, do you see this moving in … well, you started out simple, and it gets complex. I know that you don't plan your agents to hand off information or to omnipotently acquire information. Do you see yourself in the future moving where they'll hand off small bits of information based on similarities or some other criteria?

Sallach: Yes, we do, and in fact we're interested in utilizing the Juarrero attractor model of intentionality, where you have successive states of instantiation of intent. I mean, in early examples, one important issue is how it maps to action. We've got a simple mapping to action now. All you can do is shove, ask, and give, right? But we want to move toward prototype concepts for action as well, so that they can be graduated. You can shove vigorously, or you can shove guiltily. So, there are degrees and nuances, etc. And so, mapping it to a real environment, facing a more empirical kind of issue, you would want to have appropriate calibrations.

I think that, even though it's a different methodology, there still is present in this some of the action selection kinds of concepts, but it's trying to make it graduated, rather than discreet. Yes, I definitely think that there will be communication among the agents.

Kuznar: Back to the notion of sharing information, though, there needs to be a mechanism and a reason why they would, as opposed to just letting them bump into one another and hand off information because you're back to the tagging, which I agree is not a good model for cultural transmission.

Sallach: Yes. That's why I mentioned about the role of similarity— there will be rules about why this must occur, and how much effect that related rule may have ….

Macal: Okay. I'd like to thank Bill Griffin in particular for running the session. Thank you, Bill, and the authors and speakers as well. Dave's going to mention one more thing, and then let's take a five-minute break.

Sallach: I just want to mention some upcoming conferences. the NAACSOS conference will be held in Notre Dame in the last week of June. I also want to call your attention to the fact that next summer in August is the First World Congress on Social Simulation. It is a joint activity among the ESSA, NAACSOS, and PAAA, so it actually *is* a world congress. I believe that the due date for papers is in February but, in any case, we'll leave this up on the board. You can see the website. The truth is that if you search on "wcss2006 kyoto," you'll find the details. But I did want to mention it so that everyone here is aware of it.

Griffin: Thank you.

Evolutionary Models

CAN MANY "LITTLES" MAKE A "MUCH"? ONE APPROACH FOR TRANSFORMING UNDERSPECIFIED THEORY INTO AGENCY-ORIENTED RULES AND BEHAVIORS

D. SAUNDERS-NEWTON,* L. ROSZMAN, and A. FRANK,
BAE Systems AIT, Arlington, VA
R. AXTELL, The Brookings Institution, Washington, DC

ABSTRACT

While models of social agents and complexity are powerful tools for understanding societal phenomenon, appropriate and credible observation and interpretation of model output requires the lens of theory. Unfortunately, many compelling theories of social structure and dynamics are not specified in a fashion that allows for their easy instantiation as agent-based models. This paper describes some of the challenges faced by one team of researchers as they attempted to exploit the insights inherent in cultural evolution theory by converting an agent-based model of social formation, fragility, and dissolution.

Keywords: Social complexity, cultural evolution, agent-based modeling, societal fragility

INTRODUCTION

Models of social agents and complexity are powerful tools for thinking about societal phenomenon. By eschewing modeling norms such as extreme reductionism and aggressive parsimony, social scientists and inquirers can consider social structure and dynamics in a fundamentally different fashion that links micro and macro levels of observation. However, this methodological approach must confront a number of nontrivial challenges. Of particular interest to the authors is the transformation of grand theories on the evolution of social complexity (e.g., the frameworks of cultural evolution and collapse as expressed by Kent Flannery or Joseph Tainter) into an agent-based model (ABM). These grand theories are typically underspecified in terms of agency. Thus, it is important to consider a scheme for transforming theory pertaining to aggregate (global-level) features of society and culture into an agency-based (local-level) formalism.

The grand theory underlying this research effort focuses on the notion that states — particularly nation-states — emerge from increasing levels of socio-political/socio-physical complexity. This conceptual frame emphasizes the belief that a transition within and across the hierarchy of a society either toward increased complexity (e.g., the transition from "chiefdom" to a nation-state) or away from a complex form such as the nation-state is an anticipated and likely transformation. While scholars can question whether the aforementioned evolutionary imperative is a "truth," theories of evolving social complexity provide a compelling description of social

* *Corresponding author address*: Desmond Saunders-Newton, BAE Systems Advanced Information Technologies, 3811 North Fairfax Drive, Suite 500, Arlington, VA 22203; e-mail: desmond.saunders-newton@baesystems.com.

transformations with historical and contemporary significance. While the aforementioned framework of cultural evolution is compelling because it is grounded in general systems theory, it is difficult to instantiate as an agent-based formalism. What was necessary was a translation of the high-level dynamics of such cultural/civilization evolutionary theories into the language of micro-motivated agents (i.e., theories of individual choice, obligation, and opinion formation).

This paper describes the approach pursued by the authors to instantiate a grand theory of social/cultural transformation in order to better understand nation-state structure and dynamics. To that end, the paper presents a general overview of the motivation for modeling state formation, fragility, and dissolution (failure). This overview is followed by a description of the specific model formalism forwarded in this modeling effort, as well as the underlying ABM instantiation. The authors conclude this paper with general lessons for translating global-level theory into agency-level rules, routines, and dynamics.

OVERVIEW: THE NECESSITY OF MODELING THE STATE

The particular problem motivating the creation of this particular model was that of political instability and state failure. While such terms have multiple meanings, the general consensus among applied researchers concerned with supporting defense planners and strategic analysts clustered around the collapse of legitimacy on the part of the central government and the emergence of armed regional- or national-level rivals to centralized authority.[1] While the result of state failure is the loss of territorial sovereignty and the monopolization of the means of violence, such outcomes can be the result of multiple complex processes that include economic, environmental, cultural, geopolitical, technological, demographic, and other forces. This particular effort started as a Defense Advanced Research Projects Agency (DARPA)-funded exploratory effort called PCAS,[2] and was extended with internal research and development funds.

DARPA's venture into state failure as a research area marks an important turning point in defense and security policy. Traditionally, the U.S. Department of Defense (DoD) has focused its resources on deterring and defeating strong states — nuclear and conventional military forces that directly threatened U.S. citizens, territory, and allies. While the threat posed by nuclear and conventional military forces remains, other nontraditional threats have increased in quantity and quality. Irregular warfare in the form of insurgent movements, international terrorism, and the proliferation of weapons of mass destruction (WMD) has become a growing threat (Snow 2004). Moreover, most of these threats do not emanate from the traditional sources of ideologically opposed strong states; rather, they arise from discontented segments of societies residing in states that are too weak to control their populations. Thus, as states weaken and fail, their populations may challenge the legitimacy of central government and engage in hostile attacks against targets within and without the country.

Developing technology to support the analysis of state failure and political stability is new neither to the military nor the security community in general. In the mid-1960s, the

---

1 Other challenges to legitimacy included how precursor events, such as pandemics or natural disasters, could result in a loss of sovereignty or governance capacity.

2 Pre-conflict anticipation and management "seedling/sapling" effort; AFRL Contract (FA-8650-05-C-7243).

U.S. Army conducted a major social science research program on insurgency that covered a range of methods, including ethnographic field research, statistics, and computational modeling. Although this program was quickly cancelled for policy and methodological reasons, its initiation demonstrated a long-standing belief in the value of social science research to the defense community that had developed out of the use of behavioral and communications sciences during World War II (Bray 1962; de Sola Pool 1963; Horowitz 1974; Knorr 1964; Deitchman 1976). Likewise, other organizations within the defense and international security community have also sought to develop technologies to support the assessment of political stability and prevent state failure. These efforts ran the gamut from the design and development of dynamic and statistical models, to systems for indications and warning, to virtual collaborative environments for coalition building and planning preventive actions; examples include the Carnegie Commission on Preventing Deadly Conflict, the Political Instability Task Force, the Conflict Early Warning Systems, and the Pre-Conflict Management Tools Program (Carnegie Commission 2005; Alker et al. 2001; Frank 2005). Indeed, political assessment and the attempt to identify and intervene in weak and failing states before a crisis occurs were priorities for the DoD even before 9/11 demonstrated the linkages between state failure and international terrorism (Rumsfeld 2001).

STATE FORMATION AS SOCIAL COMPLEXIFICATION

As noted earlier, this particular research was concerned with identifying instances of state fragility, failure, or dissolution. In support of this state fragility/failure modeling effort, the authors developed an ABM that exploited the theories pertaining to the evolution and collapse of social complexity. This model was named SOET (i.e., Societies, Organizations, Elites, and Territories) on the basis of the notion that states are societies composed of elites who manage information and formal organizations over specified territorial bounds.

The authors developed a model of state failure based on a bottom-up process of state formation and societal fragility — processes of increasing and decreasing levels of social complexity. This theoretical process was instantiated as an ABM, a methodology particularly well suited for exploring the dynamics of decentralized, distributed systems, such as the formation or dissolution of states based on individual and institutional decision making. This approach adopted a theoretical frame of increasing and decreasing social complexity grounded in the anthropological literature on state formation and societal collapse. This theoretical framework was selected for several reasons — most important of which was its ability to provide insights into the dynamics of state failure and provide an institutional model of state health.

It is worth noting that the dynamics of state failure remain relatively unexplored formally; indeed, the dynamics of political systems remain underrepresented in within social science research when compared to comparative and cross-sectional perspectives (Pierson 2004). While there exists a general consensus that states fail and governments collapse as a result of a process that unfolds over time, little consensus exists on how fast these processes occur, the sequencing of events within the processes, or the kinds of warning that decision makers may be able to acquire in order to organize successful interventions. For example, numerous studies of revolutions, social movements, and societal collapse note the importance of political, intellectual, and ideological elites; however, the times at which they enter the process and the particular effects that they have on the process of state failure vary (Arendt 1965; Brinton 1965; Tainter 1988; Skocpol 1997; Goldstone 2003; Tilly 2004).

Alternatively, a large body of work is focused on the empirical indicators of state failure, and massive collections of statistical data have focused on the search for correlations or variables that differentiate stable from unstable states. These studies include the Correlates of War research project, the Political Instability Task Force, the Carnegie Commission on Preventing Deadly Violence, and others. These efforts are largely statistical and, while differing in their details, generally classify states in similar broad ranges of stability based on available empirical data and political, economic, and social indexes.[3]

Even though statistical and empirical work on state failure has largely converged on a core set of indicators of state failure, several unexplained dynamical factors have yet to be appropriately addressed. Moreover, definitional and theoretical clarity has yet to be achieved, causing researchers to question the fidelity of their findings; states may have multiple paths to failure, and these different paths have yet to be adequately explored or modeled (King and Zeng 2001).

Given that most of the models in the area of state failure have emphasized the search for statistical patterns, our research team decided to focus on the dynamics of state formation and failure in order gain better insights into the processes by which populations come together under stable political authority and which lead to a collapse of that authority.

Theoretical Foundation: Cultural Evolution

This research effort adopted a definition of state failure that was grounded in anthropological theories of the state and social complexity.[4] This disciplinary foundation was significant, because it regards the state as a functional construct possessing particular institutional properties. Thus, the BAE Systems' model sought to differentiate between societies that are organized into states and those that are organized into other social and political structures such as tribes and chiefdoms.

The primary theoretical frame of the BAE Systems' model was based on Kent Flannery's anthropological research on the evolution of social complexity. Flannery's model describes how changes in social complexity, as measured through evolving social institutions, occur over time and transition from bands to tribes to chiefdoms to states (Flannery 1972, pages 399–426). Flannery's theory was selected for two reasons. The first reason was the deficiencies of output models of state failure that dominate the study of state failure within political science. The second reason was that Flannery's model has been difficult to instantiate formally on the basis of

---

3 For a sampling of statistical investigations into state failure and societal fragility, and associated indexes see *The Correlates of War*, available at http://www.correlatesofwar.org, accessed on September 12, 2005; *The Carnegie Commission on Preventing Deadly Conflict*, available at http://www.wilsoncenter.org/subsites/ccpdc/index.htm, accessed on September 11, 2005; *Political Instability Task Force*, available at http://www.cidcm.umd.edu/inscr/stfail, accessed on September 11, 2005; *Polity IV Project*, available at http://www.cidcm.umd.edu/inscr/polity, accessed on September 12, 2005; and *Freedom House*, available at http://www.freedomhouse.org, accessed on September 12, 2005.

4 It is important to note that the term "social complexity" is not used to imply whether a society is sophisticated, nor to speak to the qualities of the individuals living within it. The term discusses the particular institutional organization and practices of a society as they work to achieve collective ends such as food production, distribution of wealth, economic management, the enforcement of social norms and political laws, etc.

mainstream quantitative methods; thus, instantiating it computationally would constitute a significant methodological development, demonstrating the unique abilities of computational social science methods.

Most models of state failure focus on whether the state delivers a given set of services to its citizens (i.e., they emphasize its outputs). Indeed, one leading scholar in the area of state failure notes, “It is according to their performances — according to the levels of their effective delivery of the most crucial political goods — that strong states may be distinguished from weak ones, and weak states from failed or collapsed ones” (Rotberg 2004). Likewise, another leading scholar argues, “Why do states collapse? Because they can no longer perform the functions required for them to pass as states” (Zartman 1995, page 5). While this approach is effective in making normative assessments, it ultimately suggests that states fail as a result of their internal weakness or their policy decisions. Thus, even strong states, such as Nazi Germany or Stalinist Russia, which did not support free markets or protect individual rights, would be classified as weak or failed because they adopted policies that ran counter to the preferences of liberal democracies.[5]

Output-based models of state failure experience several logical difficulties. As a result, these models produce a confusing array of inconsistent and incompatible results. For example, output models differentiate between weak, failed, and collapsed states, and note that failed states are weak while weak states may not have failed and may not fail in the future (Rotberg 2004, pages 1–25). Likewise, states that experience genocide or politicide are regarded as having failed.[6] Yet, the organizational complexity and discipline required to commit these acts are high. Thus, strong states that pursue policies of genocide or politicide are regarded as having failed, despite the fact that they are organizationally and ideologically sophisticated enough to mobilize efficient, yet horrific, campaigns against targeted groups (Kavka 1986; Smith 2005a). The conclusion is that the state has not failed because it is weak, but because it pursues political ends that run counter to the normative standards of contemporary liberal democracies.

Given the deficiencies of output models of state failure, the authors sought to create a model that provided an empirical and operationalizable definition of the state and is therefore less subject to claims and interpretations of state stability based on normative or contextualized nuance. As a result, the research team turned to the notion of social complexity more commonly represented in anthropology. Studies of social complexity emphasize the internal structure of societies rather than their outputs. Societies are categorized broadly into bands, tribes, chiefdoms, and states — where each category has particular organizational properties. According to Flannery, the features of each level of social complexity are as follows:

---

[5] It is important to note that while Nazi Germany eventually collapsed, it took the combined military and economic effort of the world’s great powers to defeat it. Likewise, Stalinist Russia eventually succumbed to the internal weaknesses identified in 1947, but this process took four decades to complete and bifurcated the world militarily, economically, and politically in the form of a Cold War. In both cases, authoritarian states displayed significant organizational and technological capabilities that discredit any notion that these were weak states based on their ideology or policy choices. For an assessment of the internal problems of Stalinist Russia, see George F. Kennan, writing as X (Kennan 1947).

[6] State Failure Task Force, *Phase III Findings*, available at http://www.cidcm.umd.edu/inscr/stfail/SFTF%20Phase%20III%20Report%20Final.pdf, accessed on September 15, 2005.

- *Bands.* These are simple egalitarian societies that are segmented along lines of kinship and marriage. Leadership within bands is limited and ephemeral, and the division of labor is based on age and sex. Ceremonies, whether religious or political, are ad hoc, and occur only when sufficient time and people are available. Bands are most commonly found among hunters and gatherers and are regarded as the dominant form of social order prior to 10000 BCE.

- *Tribes.* These are relatively large egalitarian societies whose membership extends beyond lines of kinship. Leadership in tribes is weak and largely based on personality and individual loyalty. Ceremonies are conducted on a schedule, occurring regularly on a "calendric" basis. Tribes maintain a weak sense of property rights, as land and property are owned within familial structures. The first tribes are believed to have emerged in 7000 BCE.

- *Chiefdoms.* Chiefdoms are larger than tribes and display inegalitarian distributions of, and access to, resources. Social status is hereditary, and land and property transfer from one generation to another within the family. Social stratification allows for the emergence of an elite class that manages official social rituals, and the position of the chief is institutionalized — it exists regardless of the individual who occupies it. However, while the chief occupies a settled office, the administration is filled by people who are personally loyal to the chief. The first chiefdom is believed to have emerged in 5500 BCE.

- *States.* States are highly stratified societies with institutionalized bureaucracies, and landownership and property rights. States possess strong centralized governments, and the bureaucracy is occupied by a professional class that is divorced from bonds of kinship. States maintain a near monopoly over the means of violence, and elites have advantageous access to resources. A small percentage of the population of states is involved in the production of food, while others perform specialized crafts and services (Flannery 1972, pages 401–404).

Although anthropologists have contested the precise meaning of these terms, noting that societies categorized in one way have often displayed features of higher or lower stages of complexity, the community has nevertheless accepted the general contours of a scale of social complexity based on the internal organization of the society (Tainter 1988, pages 28–31; Blanton et al. 1993, pages 10–19).

The implications for distinguishing between states and other social organizations are important. Leading theories of warfare emphasize the manipulation of adversary social and physical networks and the isolation and removal of enemy leadership (Frank 2004). Understanding the level of social complexity within a society enables the analysis of leadership structures and the underlying social structures; these structures will respond differently based on the isolation or removal of their leadership, and the leadership will respond differently to economic, military, and environmental crises. Indeed, the difficulties encountered by the U.S. military in Iraq reveal the complexities of manipulating societal structures through the use of force and the removal of the leadership. By assuming that Iraq's internal organization was that of a state, military planners concluded that its governing and economic institutions could continue to operate despite the removal of individuals loyal to Saddam Hussein and the Bath Party (Bodansky 2004; Mann 2004). However, by viewing Saddam Hussein's Iraq as a chiefdom lacking an institutionalized, professional bureaucracy, and managed on the grounds of political and personal loyalty, the expected effects of the leadership's isolation and removal become quite different.[7] Indeed, the individual chief, or what other anthropologists have referred to as the "Big Man," is so dominant that his removal, and the removal of those aligned with him, creates a power vacuum and causes the society to collapse into smaller, less complex social units, such as bands, tribes, and smaller chiefdoms (Tainter 1988, pages 25, 38).

On the basis of anthropological models of social complexity, BAE Systems developed a model that examines the process by which societies transition into and out of states. Thus, Flannery's theory produces a bi-directional process in which societies increase and decrease in social complexity. State formation is the process by which societies develop institutionalized, hierarchal organizations of political, economic, and military management, while state failure is the loss of these attributes.

Underspecification and Theoretical Adaptation

Given the nature of Flannery's theory, and the limitations of traditional modeling formalisms, formally testing Flannery's evolutionary ideas has been difficult even though it has served as a leading theory of societal evolution for more than three decades (Owen 2005). To instantiate Flannery's ideas on changes in social complexity into a computational model capable of exploring social dynamics and producing outcomes with enough specificity as to be testable and useful within the context of the DARPA effort, Flannery's theoretical scheme was expanded. This expansion was necessary to formalize behavioral properties of individuals and organizations into algorithms. Once in algorithmic form, these behaviors were used to populate an ABM discussed in detail later.

From a formal perspective, model underspecification occurs when the model identifies more variables than it possesses rules for — whether those rules are mathematical equations or behavioral/procedural algorithms. Flannery's theory provides a description of the dynamic that produces increases in social complexity, but it does not explicitly delineate behavioral rules for the organizations and individuals that compose the society. Therefore, while it is known that individuals and organizations interact in ways that dynamically alter the social structures in which they reside, the specific causal path between individual action and societal outcome is not specified by Flannery. To address this gap in model specification, BAE Systems conducted a focused literature search for meso-level theoretical models that were both compatible with Flannery, in that the direction of causation within the models was identical to Flannery's, and addressed levels of interaction and analysis lower than those described by Flannery. In addition, Flannery's theory was bolstered by using literature that focused on the collapse of social complexity.

Literature emphasizing the meso level of analysis focused on issues of social composability and the institutionalization of societal functions into formal and informal

---

7 Experts on Iraqi political and military organization have noted that Saddam Hussein's Iraq had become increasingly managed and organized based on bonds of kinship and personal loyalties after operation Desert Storm in 1991 and the subsequent assassination attempts on Saddam Hussein (Baran 1998).

organizations, and the role of elites in motivating and directing collective action. One source of particular value was the work of Michael Mann, who examined the rise of the state and the elites that manage it through the formation and control of four different networks: informational, economic, military, and political (Mann 1986). Mann's work specified the relations between the types of power that reside within society, and that are effectively used by elites to achieve political objectives.

In addition to Mann, other work on opinion formation was used to examine the relationships between elites and to model the convergence or divergence of interests based on institutional and personal affiliations. By endowing elites with multiple identities, their strategies and actions occur within a social context, and the activation and deactivation of resident identities is both a determinant of behavior and an outcome of increases and decreases in social complexity (Lustick et al. 2004).

Finally, because Flannery's theory is primarily directed at a society's accumulation of social complexity, theories of societal collapse were also used to further refine and specify the model. The primary texts used to examine the issue of collapse were Joseph Tainter's *The Collapse of Complex Societies* (Tainter 1988), which specifically deals with sudden or short-term losses of social complexity, and Jared Diamond's *Collapse: How Societies Choose to Fail or Succeed* (Diamond 2005), which largely deals with environmental change and the effects of environmental destruction on polities. Both of these texts argue that social fragility can develop rapidly, and that the loss of social complexity can occur suddenly, perhaps within a single generation.

These additional texts are important because they link the process by which states form to the paths by which they fail. For example, Tainter argues that collapse can be regarded as the reversal of the process of state formation (Tainter 1988, page 38). Likewise, Diamond argues that collapse occurs as a result of failures in collective decision-making — in particular, the failure to anticipate problems, the lack of awareness that problems have arrived, and conflicts of interest within the group's membership or between elites and society (Diamond 2005, pages 419–440). Linking state formation to government institutional design and behavior and decision-making patterns and priorities — the attributes and capabilities with which states confront threats to their cohesion — has been a long-standing tradition within social science research and remains a valuable and fruitful research area (Machiavelli 1981; Ayoob 1995; Smith 2005b).

Given the advantages of Flannery's theory, BAE Systems determined that it was far better to bolster it with supporting meso-level models, rather than find an alternative theory, because of the fruitfulness and novelty of Flannery's ideas. The shortcomings of Flannery's original work (i.e., theoretical underspecification) were addressed by selecting theories from anthropology and other social sciences to create a more complete view of the behavior of the system at lower levels of analysis while remaining true to Flannery's macro-level emphasis on social complexity.

To create these various levels of hierarchy, i.e. macro and meso, it was necessary to use the multi-agent modeling methodology for social structure to emerge and reflect the dynamics of increased complexity. Thus, starting from the micro level (i.e., the level of agents and their ability to exploit environmental resources), the effort allowed for observing the formation of social structure (meso-level) in the form of societal networks (ideology, economic, military/coercive, and political) that collectively illustrated complex relationships analysts could associate with an organizational form such as a "state." An illustration of this relationship is shown in Figure 1.

Of even greater interest to the authors was illustrating conditions that suggested societal fragility. Involved in this case was the emergence of networks of elites capable of reducing the "authority" of the government, as measured by nongovernmental cliques with access to environmental resources. Consistent with Flannery's theory, one could interpret these cliques as being equivalent to emerging chiefdoms that challenge or reduce the "power" of formal governance networks. It is assumed that at some critical yet undetermined threshold, the governmental networks become weak enough that the nongovernmental cliques could substantially reduce perceived legitimacy. The reduction could be in the form of governmental collapse or the emergence of the cliques as de facto, and necessary, societal institutions.

In addition to providing an elegant means of instantiating the respecification of Flannery's theory, the ABM methodology was selected for instantiating SOET, for several additional reasons. First and foremost, ABMs are particularly attractive for modeling emergent properties and situations where activities at one level of analysis produce behaviors and structures at higher levels of analysis that cannot be predicted based on the average properties of the lower level components (Axelrod 1997; Resnick 2000). Thus, in the case of SOET, the ABM is based on modeling the arrangement of self-interested elites into organizations that induce transitions into and out of differing governing structures. Through this process of micro-level interactions, society can be viewed as an emergent property that results from the disaggregated, uncoordinated actions of elites; and the state is considered to be structures in which powerful individuals work through existing institutions and network structures remain stable as individuals move through the power structure.

The second advantage of an ABM is that it allows for the instantiation of numerous social theories of behavior, many of which have not been formally represented or tested due to the

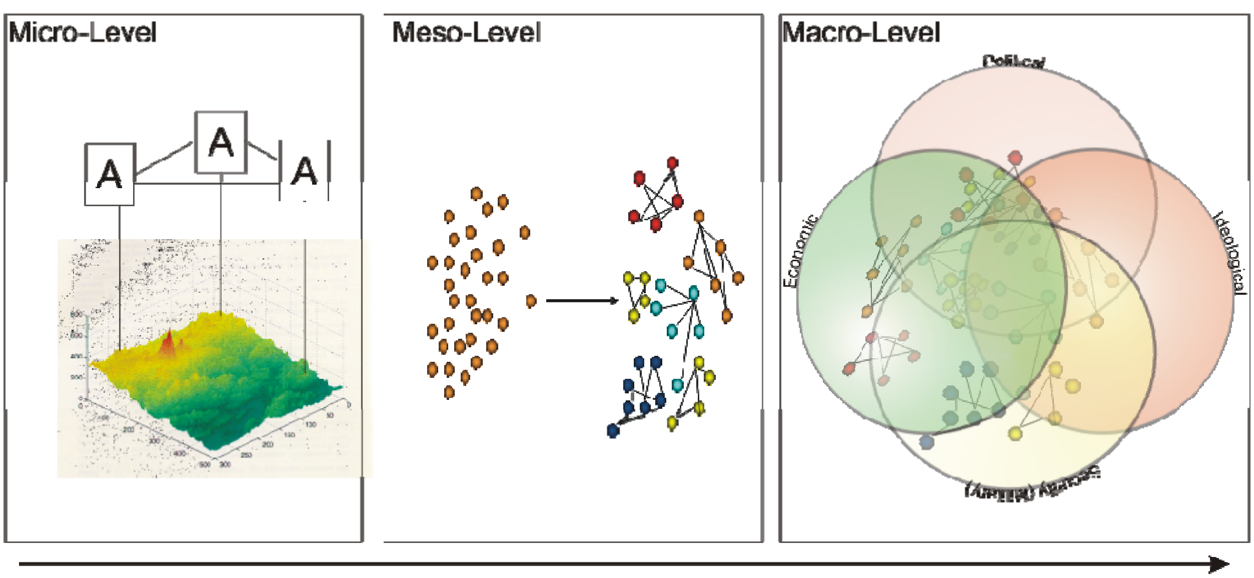

FIGURE 1 Illustration of multi-level modeling formalism

limitations of older modeling formalisms. Computational methods in general and ABMs in particular allow for the algorithmic representation of social behaviors that through recursion and mutation generate social phenomena (Axtell 2000; Epstein, forthcoming). Thus, an ABM serves as a methodological innovation that allows for the development, instantiation, and testing of a social body of theory that is grounded in social interactions, rather than inferencing based on statistical regularities, extrapolations of time-series data, or mathematical simplifications that assume closed-form solutions to social problems.

By instantiating SOET as an ABM a new form of knowledge can be created. Because ABMs provide insights into the dynamics of systems based on individual-level acts of agency, as opposed to stocks and flows of large aggregates, model users can gain a better understanding of when a system's behavior is likely to change as a result of individual and collective decision-making. This is important because statistical models generally extrapolate based on known data, implicitly asserting that the future will resemble the past and that the causal mechanisms within the society are stable. Therefore, while statistical methods can interpret empirical data to show a society's present condition, they cannot predict that society's trajectory should the underlying structure change as a result of changing dynamics — an ABM provides an insight into these dynamics.

A third advantage of the ABM formalism is its ability to support scenario planning, hypothesis testing, and other forms of exploratory analysis and credible model exploitation (Bankes 1993; Lempert et al. 2003; Saunders-Newton 2006). Although the PCAS "sapling" emphasized the modeling of a small number of nations, the longer term intent of the DARPA effort was concerned with developing an ability to explore scenarios and examine potential futures that result from alternative policy choices. ABM provides an attractive, dynamic simulation environment where alternative data sets, behavioral rules, agent or system attributes, etc., can be explored. Linking these simulation "hooks" to policy levers allows for the systematic exploration of outcome spaces based on potential policy options. The ABM formalism allows analysts and policy-makers to reach beyond today's information and explore alternative futures, identifying paths to desirable, stable structures and the indicators of troublesome outcomes.

Brief Technical Description of SOET

The actual realization of the SOET formalism required the development of a software environment inclusive of agents (individual and organizational), their interactions between each other and various resource landscapes. To that end, the BAE Systems team represented "nation-state" dynamics in terms of each agent's ability to transform resources into goods that would enable it to pursue various goals. The action of creating products comes at a cost to the agent, and often require negotiating with other agents. The transactions between agents themselves and the environment are represented — and managed — by the connections. The underlying dynamics of SOET are further illustrated in Figure 2.

The initial model exploited the simulation engine of a Java-based ABM development environment called Repast.[8] The Java agents are unique to the BAE Systems approach. Future

---

[8] Repast is the acronym for the Recursive Porous Agent Simulation Toolkit. It is an open-source software environment that can be downloaded at http://repast.sourceforge.net.

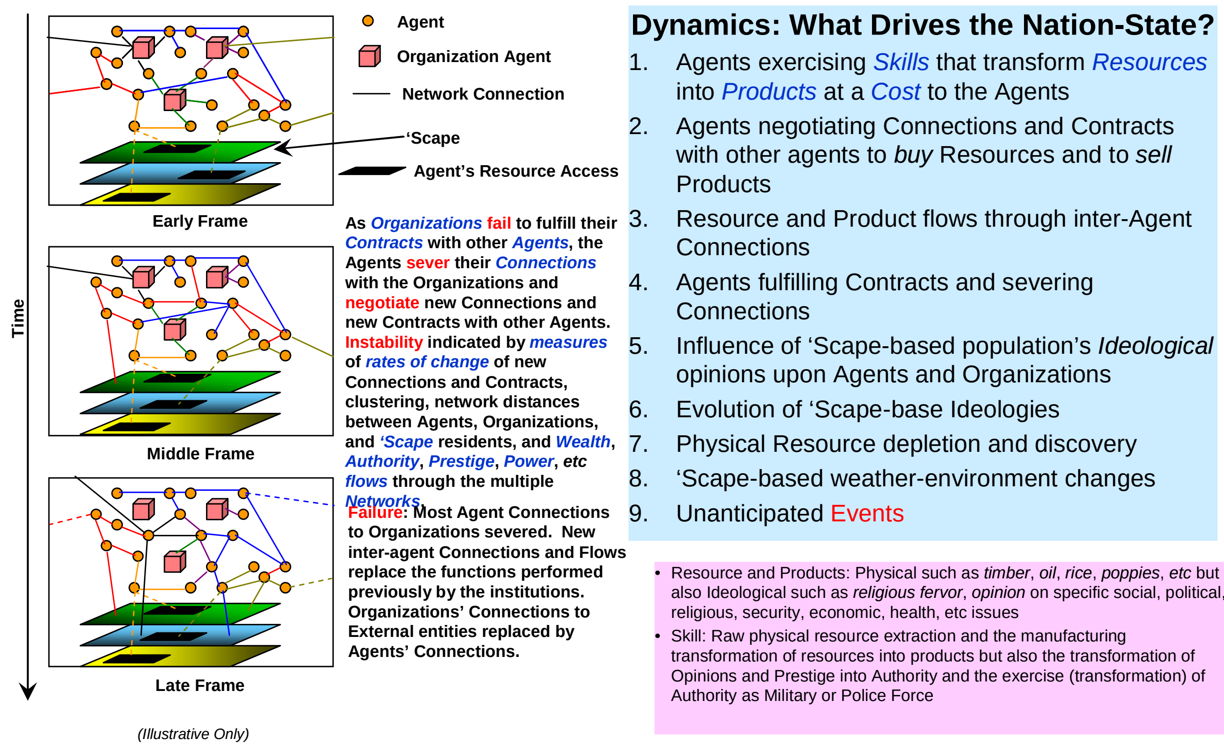

FIGURE 2 Illustration of dynamics associated with nation-state

instantiations of this model will likely make use of a custom-crafted ABM development environment that will allow for more efficient development efforts inclusive of an ability to easily edit nation-state or societal configurations, as well as multi-processor instantiations.

CONCLUSIONS

Compared to traditional methodological approaches of modeling and interpreting societal fragility and state failure, the SOET approach was extremely novel. By exploiting the power of multi-agent simulation and revisiting and reconsidering some of the major theories of cultural evolution, the authors found a means of examining the dynamics of state formation as it relates to fragility and possible state dissolution. This evolutionary approach differed greatly from the statistical approaches that tend to characterize this area of study. It is the opinion of the authors that it is difficult to understand why a state fails without actively considering how it came to be fragile. Statistical models do not typically allow for the easy representation and consideration of such concepts.

The challenge in this case is that theories for interpreting the evolution of societies over time are not well specified for instantiation as agent-based models. Flannery's theory, as an example, reflects systems-level thinking, and as such suggests the use of a system dynamics approach. However, the "stock, flow, feedback" metaphor of this methodological approach is not an appropriate means for considering both state formation and evolution. Moreover, Flannery's theory does not speak to the underlying dynamics that give rise to choices by the agents who comprise the social forms he delineates, i.e. bands, tribes, chiefdoms, and states. Thus, even if it were possible to represent these notions in the most likely methodological approach, the instantiation of agency is not aided by Flannery's work.

Thus, efforts to convert many of the intuitively appealing theories of social outcomes and dynamics will benefit greatly from a more systematic approach for representing the agency that is implicit in the theoretical frameworks posited by many social thinkers. This will aid researchers, analysts, and decision-makers greatly in exploiting the rich corpus of social thought that has been created over recorded time.

REFERENCES

Alker, H.R., T.R. Gurr, and K. Rupesinghe, eds., 2001, *Journeys through Conflict: Narratives and Lessons* (Lanham, MD: Rowman & Littlefield).

Arendt, H., 1965, *On Revolution* (New York, NY: Penguin Books).

Axelrod, R., 1997, *The Complexity of Cooperation* (Princeton, NJ: Princeton University Press).

Axtell, R., 2000, *Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences*, Working Paper No. 17 (Nov.), Center for Social and Economic Dynamics; available at http://www.brookings.edu/dybdocroot/es/dynamics/papers/agents/agents.pdf (accessed Sept. 22, 2005).

Ayoob, M., 1995, *The Third World Security Predicament: State Making, Regional Conflict, and the International System* (Boulder, CO: Rienner).

Bankes, S., 1993, *Exploratory Modeling for Policy Analysis* (Santa Monica, CA: RAND).

Baram, A., 1998, *Building Toward Crisis: Saddam Husayn's Strategy for Survival* (Washington, DC: Washington Institute for Near East Policy).

Blanton, R.E., S.A. Kowalewski, G.M. Feinman, and L.M. Finsten, 1993, *Ancient Mesoamerica* (New York, NY: Cambridge University Press).

Bodansky, Y., 2004, *The Secret History of the Iraq War* (New York, NY: Regan Books), p. 47.

Bray, C.W., 1962, "Towards a Technology of Human Behavior for Defense Use," *The American Psychologist*, Vol. 17, pp. 527–541.

Brinton, C., 1965, *The Anatomy of Revolution* (New York, NY: Vintage).

Conflict Early Warning Systems; available at http://www.usc.edu/dept/LAS/ir/cis/cews (accessed Sept. 11, 2005).

de Sola Pool, I., ed., 1963, *Social Science Research and National Security* (Washington, DC: Smithsonian Institution).

Deitchman, S.J., 1976, *The Best Laid Schemes* (Cambridge, MA: MIT Press).

Diamond, J., 2005, *Collapse: How Societies Choose to Fail or Succeed* (New York, NY: Viking).

Epstein, J.M., forthcoming, *Generative Social Science: Studies in Agent-Based Computational Modeling* (Princeton, NJ: Princeton University Press).

Flannery, K.V., 1972, "The Cultural Evolution of Civilizations," *Annual Review of Ecology and Systematics*, Vol. 3, pp. 399–426.

Frank, A.B., 2004, "Get Real: Transformation and Targeting," *Defence Studies*, Vol. 4, No. 1 (spring), pp. 64–86.

Frank, A.B., 2005, "Pre-conflict Management Tools: Winning the Peace," *Defense & Technology Paper*, No. 11 (Feb.).

Goldstone, J.A., ed., 2003, *Revolutions: Theoretical, Comparative, and Historical Studies* (Belmont, CA: Thompson).

Horowitz, I.L., ed., 1974, *The Rise and Fall of Project Camelot* (Cambridge, MA: MIT Press).

Kavka, G.S., 1986, *Hobbesian Moral and Political Theory* (Princeton, NJ: Princeton University Press), pp. 159–160.

Kennan, G.F., writing as X, 1947, "The Sources of Soviet Conduct," *Foreign Affairs*, Vol. 25 (July), pp. 566–582.

King, G., and L. Zeng, 2001, "Improving Forecasts of State Failure," *World Politics*, Vol. 53 (July), p. 625.

Knorr, K., 1964, *Foreign Intelligence and the Social Sciences* (Princeton, NJ: Center for International Studies, Princeton University).

Lempert, R.J., S.W. Popper, and S.C. Bankes, 2003, *Shaping the Next One Hundred Years: New Methods for Long-term Quantitative Policy Analysis* (Santa Monica, CA: RAND).

Lustick, I.S., D. Miodownik, and R.J. Eidelson, 2004, "Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It?" *American Political Science Review*, Vol. 98, No. 2 (May), pp. 209–229.

Machiavelli, N., 1981, *The Prince* (New York, NY: Penguin Classics), George Bull trans., pp. 33–44.

Mann, J., 2004, *Rise of the Vulcans* (New York, NY: Viking Press), p. 360.

Mann, M., 1986, *The Sources of Social Power, Volume I: A History of Social Power from the Beginning to A.D. 1760* (New York, NY: Cambridge University Press).

Owen, B., *More Theories of State Formation*; available at http://bruceowen.com/emciv/34104s15.htm (accessed Sept. 7, 2005).

Pierson, P., 2004, *Politics in Time: History, Institutions, and Social Analysis* (Princeton, NJ: Princeton University Press).

Political Instability Task Force; available at http://www.cidcm.umd.edu/inscr/stfail (accessed Sept. 11, 2005).

Resnick, M., 2000, *Turtles, Termites, and Traffic Jams* (Cambridge, MA: MIT Press).

Rotberg, R.I., 2004, "The Failure and Collapse of Nation-States: Breakdown, Prevention, and Repair," in R.I. Rotberg, ed., in *When States Fail: Causes and Consequences* (Princeton, NJ: Princeton University Press), p. 4.

Rumsfeld, D., 2001, "Prepared Testimony of U.S. Secretary of Defense, Donald H. Rumsfeld, Senate Armed Services Committee," June 21; available at http://armed-services.senate.gov/statemnt/2001/010621rumsfeld.pdf (accessed Sept. 11, 2005).

Saunders-Newton, D., 2006, "When Worlds Collide: Reflections on the Credible Uses of Agent-based Models in International and Global Studies," in N.E. Harrison, ed., *Complexity in World Politics: Concepts and Methods of a New Paradigm* (New York: SUNY Press, summer), pp. 229–250.

Skocpol, T., 1997, *Social Revolutions in the Modern World* (New York, NY: Cambridge University Press).

Smith, B., 2005a, discussion, Department of Political Science, University of Florida (spring).

Smith, B., 2005b, "Life of the Party: Regime Breakdown and Persistence under Single-party Rule," working paper, *World Politics*, Vol. 57 (April), pp. 421–451; available at http://www.clas.ufl.edu/users/bbsmith/LoP.pdf (accessed Sept. 21, 2005).

Snow, D.M., 2004, *National Security for a New Era* (New York, NY: Pearson Longman).

Tainter, J.A., 1988, *The Collapse of Complex Societies* (New York, NY: Cambridge University Press).

The Carnegie Commission on Preventing Deadly Conflict; available at http://www.wilsoncenter.org/subsites/ccpdc/index.htm (accessed Sept. 11, 2005).

Tilly, C., 2004, *Social Movements, 1798–2004* (Boulder, CO: Paradigm).

Zartman, W.I., 1995, "Introduction: Posing the Problem of States Collapse," in W.I. Zartman, ed., *Collapsed States: The Disintegration and Restoration of Legitimate Authority* (Boulder, CO: Rienner).

KNOWLEDGE SWARMS: GENERATING EMERGENT SOCIAL STRUCTURE IN DYNAMIC ENVIRONMENTS

R.G. REYNOLDS,* B. PENG, and X. CHE,
Department of Computer Science, Wayne State University, Detroit, MI

ABSTRACT

Our previous work on real-valued function optimization problems had shown that cultural learning emerged as the result of meta-level interaction or swarming of knowledge sources (i.e., "knowledge swarms") in the belief space. These meta-level knowledge swarms induced the swarming of individuals in the population space (i.e., "cultural swarms"). The interaction of these knowledge swarms also produced emergent phases of problem solving at the population level that reflected an algorithmic process and resulted in the emergence of individual roles within the population: explorers and exploiters. Roles similar to this have been observed in animal populations and labeled "producers" and "scroungers," respectively (Barnard and Sibly 1981). Here we investigate the impact of environmental dynamics on the spatial and temporal aspects of role emergence. Specifically we generate a repeated shift in the resource landscapes at different intervals and note that this adds new distinctions within the previous role structure. That is, environmental complexity induces an increase in the complexity of social roles within a given system through the knowledge swarming process.

Keywords: Cultural algorithms, role emergence, cultural swarms, social intelligence, marginal value theorem

INTRODUCTION

Recently, a number of socially motivated algorithms have been used to solve optimization problems. Some of the example algorithms are the particle swarm algorithm (PSO) (Kennedy and Eberhart 1995), ant colony algorithm (ACO) (Dorigo et al. 1996), and cultural algorithm (CA) (Reynolds 1978, 1994). These three algorithms all use a population-based model as the backbone of the algorithm and solve problems by sharing information via social interaction among agents in the population.

Figure 1 expresses each of these approaches in terms of both a space and a time continuum over which the social interactions take place. Notice that both the ant and particle swarm approaches can be found near the lower left end of this continuum, with the social interaction between individuals taking place within limited temporal and spatial dimensions. For example, in particle swarm, agents can exchange their direction of movement and velocity locally with other agents. In the ant algorithm, agents locally exchange information in terms of the density and gradient of a "pheromone" substance that marks their trail. The pheromone

* *Corresponding author address*: Robert G. Reynolds, Department of Computer Science, Wayne State University, 409 State Hall, Detroit, MI 48202; e-mail: reynolds@cs.wayne.edu.

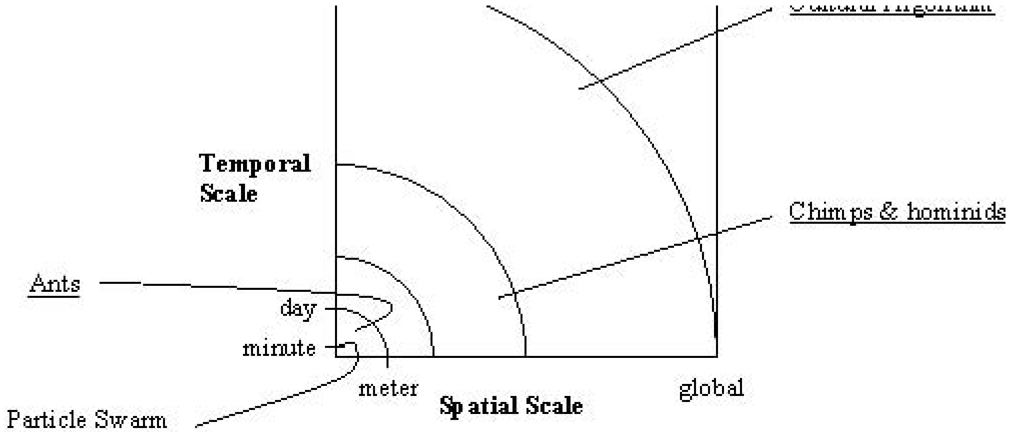

FIGURE 1 Scale of social interaction (The emergent properties depend on the scale at which the interaction takes place.)

chemical is deposited by an ant as it moves along a trail. The frequency of use of a trail is indicated by the amount of pheromone that is deposited relative to its degradation in the environment over time.

CAs, on the other hand, allow agents to interact in many different ways by using various forms of symbolic information reflective of complex cultural systems. The basic CA allows individuals to communicate via a shared belief space. The shared space stores five basic types of information that can be shared cognitively or symbolically.

Cultural Algorithms

A CA is a class of computational models derived from observing the cultural evolution process in nature (Reynolds 1978, 1994). A CA has three major components: a population space, a belief space, and a protocol that describes how knowledge is exchanged between the first two components. The population space can support any population-based computational model, such as genetic algorithms and evolutionary programming. The basic framework is shown in Figure 2.

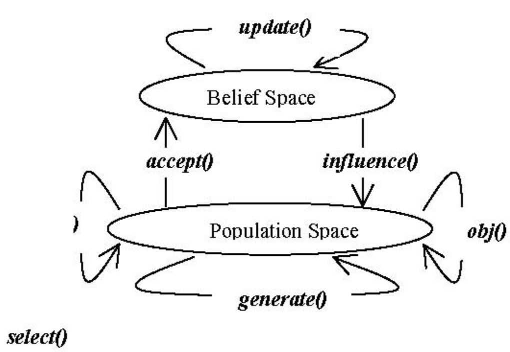

FIGURE 2 Framework of a cultural algorithm

A CA is a dual inheritance system that characterizes evolution in human culture at both the macro-evolutionary level, which occurs within the belief space, and the micro-evolutionary level, which occurs in the population space. Knowledge produced in the population space at the micro-evolutionary level is selectively accepted or passed to the belief space and used to adjust the knowledge structures there. This knowledge can then be used to influence the changes made by the population in the next generation.

What makes a CA different from the PSO and ACO approaches is the fact that a CA uses five basic knowledge models in the problem-solving process rather than just one or two locally transmitted values. There is ample evidence from the field of cognitive science that each of these knowledge models is supported by various animal species (Wynne 2001; Clayton et al. 2000), and it is assumed that human social systems support each of these models as well. The knowledge sources include normative knowledge (ranges of acceptable behaviors), situational knowledge (exemplars or memories of successful and unsuccessful solutions, etc.), domain knowledge (knowledge of domain objects, their relationships, and interactions), history knowledge (temporal patterns of behavior), and topographical knowledge (spatial patterns of behavior). This set of categories is viewed as being complete for a given domain in the sense that all available knowledge can be expressed in terms of a combination of one of these classifications.

Problem Statement

The CA has been studied with benchmark problems (Chung and Reynolds 1998) as well as applied successfully in a number of diverse application areas, such as modeling the evolution of agriculture (Reynolds 1986), concept learning (Reynolds 1994), real-valued function optimization (Jin 1999; Reynolds and Saleem 2005), re-engineering of semantic networks (Rychlyckyj 2003), and agent-based modeling of price incentive systems (Ostrowski and Reynolds 2002), among others.

While successful, the relative complexity of the knowledge sources and their interaction made it difficult to determine why CAs worked so well. Alternatively stated, under what conditions will such systems successfully solve a given problem, and what social structures will emerge along the way? The emergence of these structures in both the population and belief space can be viewed as signs of a successful problem-solving process.

In this paper, we attempt to develop answers to these questions. To do this, we begin by examining how CAs solve resource optimization problems within an experimental environment. In our investigation here, we employ a simulated cones world environment developed initially by Morrison and De Jong (1999) and extended here. Within this world, resources are viewed as being distributed in piles (cones) on the landscape (Sugarscape style; see Epstein and Axtell 1996).

In our paper, we use five different knowledge sources to direct the agents. Each knowledge source is a model for an agent's behavior. Since the belief space consists of five different knowledge sources or models, the question at each time-step is how to assign agents to the various models. The key here is that each knowledge source has an expression in two-dimensional (2D) space in terms of a bounding box characterized by a midpoint and standard deviation in the x and y directions. If we view each box as analogous to a resource patch in the environment, every knowledge source model can be viewed as a predator searching for prey in a given patch.

Since our cones world problems can then be described as foraging problems within a search space, we use a framework within the CA for the selection of a given knowledge source by an agent inspired by theoretical results from studies of foraging theory in population biology. Specifically, agents select different knowledge sources on the basis of what we characterize as "the marginal value of information." The inspiration for this comes from the classic work by Charnov (1976) concerning the "marginal value theorem." In certain situations, agents using the marginal value theorem were able to optimize their long-term resource intake within an environment. Simply stated, the marginal value theorem says that an agent stays within a location (patch) on the landscape until the current resource gain is less than the average expected value. It then moves to another patch.

Agents are then attracted to different knowledge sources on the basis of how successful individuals using each model are. In a previous paper (Reynolds and Peng 2005), we showed that this approach produced two classes of individuals — explorers and exploiters — depending on the particular knowledge models that they tended to use when the cones were configured in a static environment. As it turns out, these distinctions have been observed in naturally occurring animal populations as well. Barnard and Sibly (1981) identify "producers" who are engaged in finding resource patches and "scroungers" who exploit the found resources. Thus, our model was able to show how these roles might have emerged as a result of the knowledge swarming process within their shared belief space.

However, a question remains. What impact does the environment have on the emergence of these roles? In other words, do similar roles emerge in dynamically changing environments? In order to investigate this, we make the cone configurations in our model dynamic but predictable. A resource come is placed in each of four quadrants, and the cones are interchanged in counterclockwise fashion at regular intervals. We will show that each of the five knowledge sources or models adjusts it patch size and dynamics in a rather complementary fashion to exploit these dynamics. As a result, we see the emergence of subgroups of individuals within both the exploiter and explorer classes based on the models that they select to control their movement. In other words, the addition of environmental variability offered the agents more opportunities in terms of their knowledge models than were offered in the static case. This resulted in the production of a more complex social structure.

The next (second) main section describes the cones world environment. The third section describes the CA system configuration and how the marginal value theorem is employed here to adjust the patches for each knowledge source. The fourth section describes the simulation environment and experimental dynamics. The next section presents our results and describes some emergent properties of the social system, and the last section gives our conclusions.

THE CONES WORLD ENVIRONMENT

Our test problems are generated by a multi-modal problem generator DF1 (Morrison 1999), in which a "field of cones" of different heights and different slopes are randomly scattered across the landscape. The landscape is given by:

$$f\left(< x_1, x_2, ..., x_n >\right) = \max\left(H_j - R_j * \sqrt{\sum_{i=1}^{n} \left(x_i - C_{j,i}\right)^2} \right)$$

where $< x_1,\ldots,x_k >$ represent points in the landscape, n specifies the number of cones in the environment, and k is the number of dimensions. Each cone is independently specified by its location $< C_{j,1},\ldots,C_{j,n} >$, its height H_j, and its slope R_j. The cones are then "blended" together by using the max function to form the search surface.

The main reason that we picked this generator is that by changing its parameters, it can generate test functions over a wide range of surface complexity and problem dynamics. This enables us to evaluate our model in a more flexible and systematic way. An example 2D landscape is shown in Figure 3.

KNOWLEDGE SOURCES AND THE MARGINAL VALUE THEOREM

In this section, we briefly discuss the five knowledge sources used in the belief space and then motivate their integration in the optimization and search process by using the marginal value theorem. Each of the five knowledge sources or models has been observed to be cognized and used in various nonhuman species as a basis for encoding their social knowledge (Wynne 2001; Clayton et al. 2000).

The Five Knowledge Sources

For each knowledge type, we elaborate on its definition, the data structure, and how it is updated. Throughout the description, we use the symbol n for the number of parameters of the optimization problem. It is often referred to as dimensions of an optimization problem.

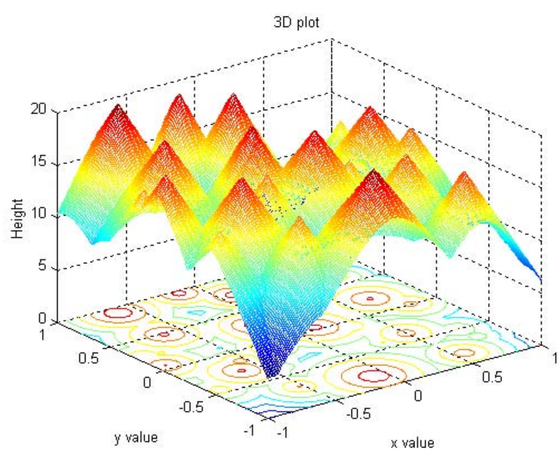

FIGURE 3 DF1 with n = 50, H ∈ (1, 10), R ∈ (8, 20), and in 2D space [(–1.0, 1.0), (–1.0, 1.0)]

Situational Knowledge

The situational knowledge source was first proposed by Chung (1997) for real-valued function optimization problem-solving in static environments. Situational knowledge contains a set of exemplars taken from the population. The data structure of the situational knowledge is represented as a list of exemplar individuals, as shown in Figure 4.

Each exemplar contains a value for each parameter and the fitness value for this exemplar. The situational knowledge will be updated either by adding the population's best individual to the situational knowledge if it outperforms the current best or reinitializing it when environmental change is detected. Situational knowledge represents exemplars or examples for other individuals to follow. These are case studies or events that are the basis for others behavior (Wynne 2001).

Normative Knowledge

Normative knowledge was also introduced by Chung (1997). It is represented as a set of intervals, and each is viewed to be a promising range for good or socially acceptable solutions for a parameter. The normative knowledge data structure for *n* variables is given as follows in Figure 5.

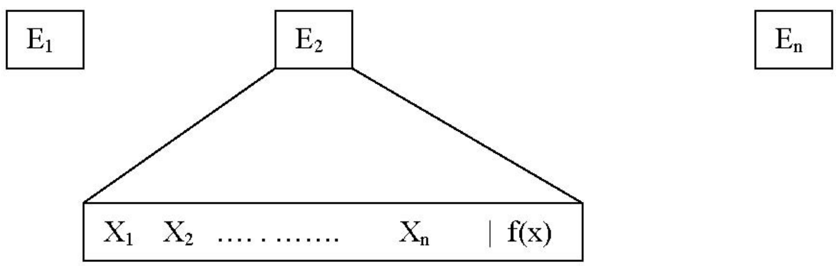

FIGURE 4 Structure of situational knowledge

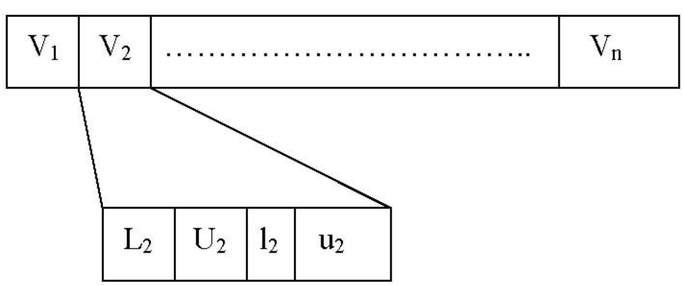

FIGURE 5 Structure of normative knowledge

For each variable V_i, the data structure contains the upper and the lower bounds, l_i and u_i, and the performance values for individuals in the upper and lower bounds, L_i and U_i. Normative knowledge is updated by shifting the variable ranges and updating the corresponding performance values to reflect changes in the environment.

Topographical Knowledge

Topographical knowledge, also called regional schemata (Jin 1999), is represented in terms of a multidimensional grid or array with cells in the grid described as $c_1,\ldots,c_i,\ldots,c_n$, where c_i is the cell size for the *i*th dimension. There is strong evidence for the ability of difference species to process 2D data displays. The data structure representation is an array of size *n*, where *n* is the number of cells in the mesh. Each cell in the data structure contains a lower and an upper bound for the *n* variables $[(l,u)1,\ldots,(l,u)n]$, indicating the ranges associated with the best solutions found in that cell so far, and a pointer to its children, as shown in Figure 6.

The topographical knowledge structure is initialized by the random placement of agents within cells in the grid and by creating a list of best cells. The update occurs when a cell is divided into subcells when an accepted individual's fitness value is better than the best solution in that cell, or if the fitness value of the cell's best solution has increased after a change event is detected. Topographical knowledge provides a spatial or array framework in which environmental patterns can be identified and adjusted for.

Domain Knowledge

Domain knowledge was introduced into the CA (Reynolds and Saleem 2005) in order to solve dynamic optimization problems. Domain knowledge was designed to reason about local dynamics. especially in terms of the prediction of gradients of incline or decline. Its data structure is shown in Figure 7.

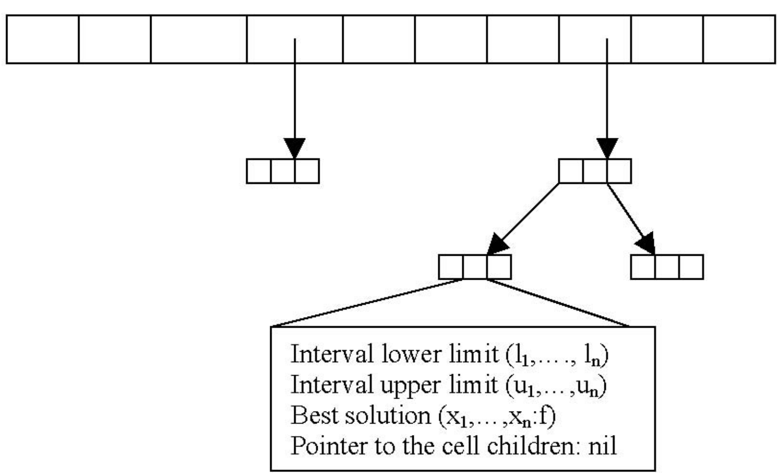

FIGURE 6 Structure of topographical knowledge

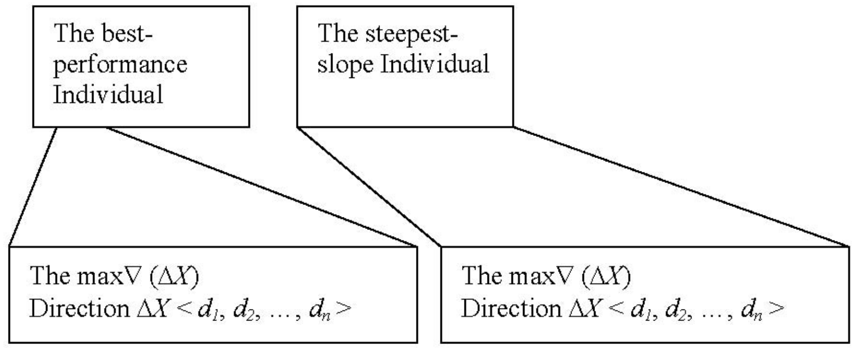

FIGURE 7 Structure of domain knowledge

Here domain knowledge consists of the domain ranges for all parameters and the best examples from the population space, similar to the situational knowledge representation above. Here domain knowledge is used to predict trends in the resource landscape both statically and dynamically. For example, given the cones world, if an upward slope or gradient is detected, then like an ant following a pheromone trial, one predicts a source increase. Likewise, domain knowledge can be used to predict the locations of resources in a dynamic environment. For example, if the amount of resources at a point is under the influence of a single cone, and if the slope at that point changes, then so has the point. One can predict the amount of shift necessary to place the slope at that point. This will allow agents using this model to make predictions about the future locations of a cone in the dynamic case.

History Knowledge

History knowledge was developed (Reynolds and Saleem 2005) in order to reason about global dynamics and to facilitate backtracking or retracing steps. It contains information about sequences of environmental changes in terms of shifts in the distance and direction of the known resource cones in the search space. Its cognitive origin comes from episodic memory (both in humans and animals), which is a type of memory based on personal experience. It stores information about events and temporal-spatial relations among those events (Clayton et al. 2000). While domain knowledge is focused on the interpretation of a resource shift locally in terms of geometrical or gradient considerations, history knowledge provides a more global perspective of the change. It computes the average change in parameter values within a region, the window size, and predicts the direction of the shift in the optimum from the previous position. The knowledge data structure representation is shown in Figure 8.

Here w represents the number of change events stored and $(ds_1,\ldots,ds_n)$ and $(dr_1,\ldots,dr_n)$ represent the average environmental changes in distance and direction, respectively, for each one of the n parameters. e_1 through e_w are change events. The history knowledge is updated after every change event by updating the history list and the moving averages for each parameter.

History knowledge is implemented as a list of up to m temporal events/points on the search path $\{P_1, P_2,\ldots, P_m\}$. m is the size limit of the history list, and each $P_j = < p_j, 1, p_j, 2, \ldots, p_j, n >$ represents a significant point on the search path.

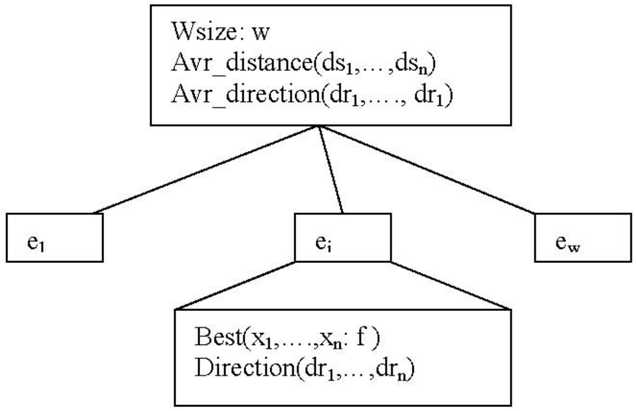

FIGURE 8 Structure of history knowledge

Communication Protocol

The communication protocol of a CA system is composed of two functions. The acceptance function determines which individuals are used to impact the belief space, and the influence function determines how the belief space influences the population space in generating a new solution.

Acceptance Function

The acceptance function determines which individuals and their behaviors can impact the belief space knowledge. It is often specified as a percentage of the number of current individuals, ranging between 1% and 100% of the population size, and based on selected parameters such as performance. For example, we can select the best performers (e.g., top 10%), worst performers (e.g., bottom 10%), or any combinations.

Influence Function: Using the Marginal Value Theorem

The choice of influence function has a great impact on the problem-solving process. Some influence functions are more successful than others, as measured by the success of the agents that each has influenced in the past. Early influence functions randomly applied the five knowledge sources to individuals in the population in order to guide their problem-solving process.

A good search approach should optimize the rate at which the available resources are processed by the foraging agents as they search for the optimum food search. While the distribution is continuous, it was observed that at each time-step that the individuals generated by each knowledge source using a normal distribution could be described by a “bounding box” or patch with a given central tendency and standard deviation. For example, in Figure 9, notice the shifting of the patch for situational knowledge from one location on the landscape to another. In fact, the original patch orientation is rotated and then translated toward the optimal point “+” over time.

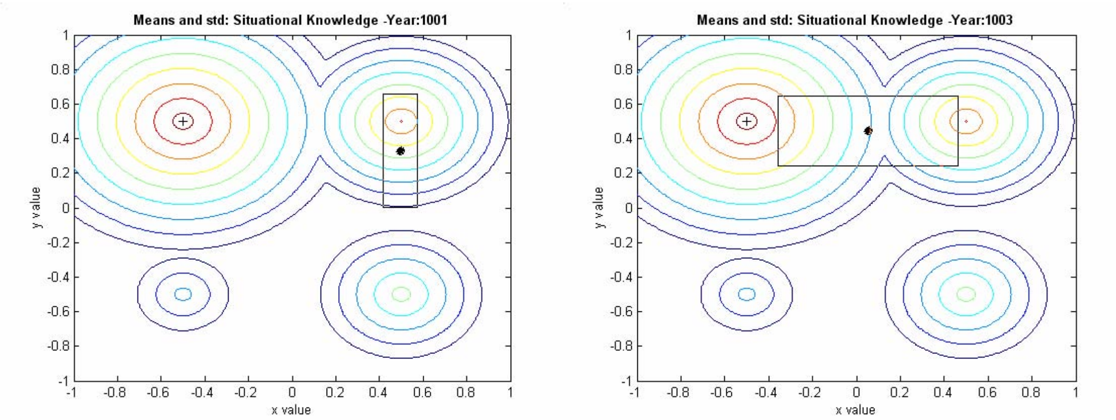

FIGURE 9 Situational means and standard deviation at year 1001 and 1003

In foraging theory, it had been shown that the use of the marginal value theorem is able, under certain conditions, to optimize the long-term average rate of energy intake within a patch-base environment (Charnov 1976). The principle behind the marginal value theorem is that residence time in a patch by a forager affects the expected energy gain. The marginal value principle states that the forager should reside in the patch "until the intake rate in a patch drops to the average rate for the habitat . . . it is the 'moving-on threshold' intake rate that is important" (Stephens and Krebs 1986, page 31). The forager, when doing so, will maximize the average long-term energy intake of the individual. One of the key assumptions is that the gain function associated with a patch is initially increasing but eventually negatively accelerated. Other assumptions are shown in Figure 10 taken from Stephens and Krebs (1986).

ASSUMPTIONS

Decision

The set of residence times for each patch type, t_i for patch type i.

Feasible choices: For all patch types $0 \le t_i < \infty$.

Currency

Maximization of long-term average rate of energy intake.

Constraint

C.1 Searching for and hunting within patches are mutually exclusive activities.

C.2 Encounter with patches is sequential and is a Poisson process.

C.3 Encounter rates when searching are independent of the residence times chosen.

C.4 Net expected energy gain in a patch is related to residence time by a well-defined gain function [$g_i(t_i)$] with the following characteristics:

(i) Change in energy gain is zero when zero time is spent in a patch.

(ii) The function is initially increasing and eventually negatively accelerated.

C.5 Complete information is assumed. The forager knows the model's parameters and recognizes patch types, and it does not acquire and use information about patches while foraging in them.

FIGURE 10 Summary of the patch model (based on Stephens and Krebs 1986)

Figure 11 describes the calculation for a single patch. This figure is taken from Stephens and Krebs (1986, p. 30). There are two quantities plotted on the abscissa: travel time or placement effort and patch residence time. Each of the knowledge sources in the influence function is viewed as a predator. Travel time increases from the origin (vertical line) to the left, and patch residence time increases from the origin to the right. The gain function shape exhibits an initial increase and then escalating decrease. The optimal residence time can be found by constructing a line tangent to the gain function that begins at the point 1/λ on the travel time axis. The slope of this line is the long-term average rate of energy intake, because 1/λ is the average time required to travel between patches. When the travel time is long (1/λ2), then the rate-maximizing residence time ($\hat{t}_2$) is long. When the travel time is short (1/λ1), then the rate-maximizing residence time ($\hat{t}_1$) is shorter. Here travel time is a constant amount that represents a model time-step.

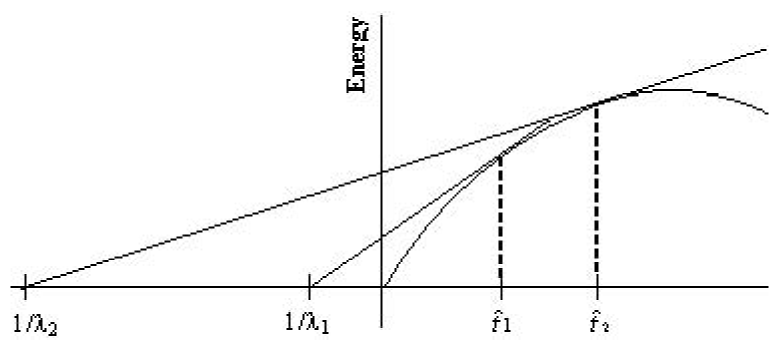

FIGURE 11 Marginal value theorem in the one-patch-type case

In this paper, we use an influence function based on the marginal value theorem as discussed in Peng (2005). The marginal value theorem is implemented here in terms of a roulette wheel function. The size of a knowledge source area under the wheel is a function of its ability to produce above-average gains. Each of the five knowledge sources, predators, is initially given 20% of the wheel area with which to generate its patch.

The likelihood of using one of the knowledge sources as a model for agent movement is based on the size of the area under the wheel, and the area for a knowledge source (predator) is adjusted on the basis of the performance of those agents it influences. At every time-step, each of the agents in the population is influenced by one of the knowledge sources on the basis of a spin of the wheel. The agent then moves to a position within the patch or bounding box associated with the selected knowledge source. The gain produced by the agent there is then recorded for the predator there.

The performance of a knowledge source can then be generated by computing the average fitness value of all individuals generated by each knowledge source. The average fitness value of individuals generated from using a specific knowledge source (predator) is:

$$avr_i = \frac{\sum_{j=1}^{k} f_j(x)}{k},$$

where k is the number of individuals generated via the knowledge source and is the fitness value of individual j.

Now each influence operator is assigned an area on the roulette wheel relative to its average performance, computed above, over the average performance for all of the influence functions:

$$p_i = \frac{avr_i}{\sum_{j=1}^{n} avr_j},$$

where p_i is a percentage on the roulette wheel assigned to influence operator i; and n is the number of influence operators used in the system.

When the value for a patch falls below the average, the area under the wheel will approach 0, and few individuals will be placed in that patch. However, its patch dimensions can be affected by the other active patches and new patch dimensions produced. If the patch shift is successful, the gain for the knowledge source will increase, and its share of the wheel will be become larger. At the same time, other knowledge sources will be experiencing a decrease in gain, and their areas will shrink.

Thus, with a gain function that increases initially and then decreases exponentially, we should get a phased pattern of knowledge use, where as some patches begin to fail, others are getting more individuals and increasing; however, with too much exploitation, they begin to fail and the cycle repeats itself. In the next section, we provide an example of how the bounding boxes for each of the knowledge sources (predators) shift during the course of the problem-solving process.

Experiment Settings

We set up a dynamic problem on the basis of the cones world problem described in Section 2. As a starting point, we construct a baseline landscape with four cones of different heights and slopes, where each is placed in one of four quadrants of the Cartesian plane (Figure 12). The cones can overlap on the basis of their slopes. Then the cones are shifted periodically 90° in the counterclockwise direction, so that every four 90° shifts form a complete rotation (Figure 12). We use this pattern to examine how the different knowledge sources react to this patterned movement and to observe the roles that emerge. Here, the cones are moved every 200 generations for a total of eight shifts or two cycles around after 1,600 generations.

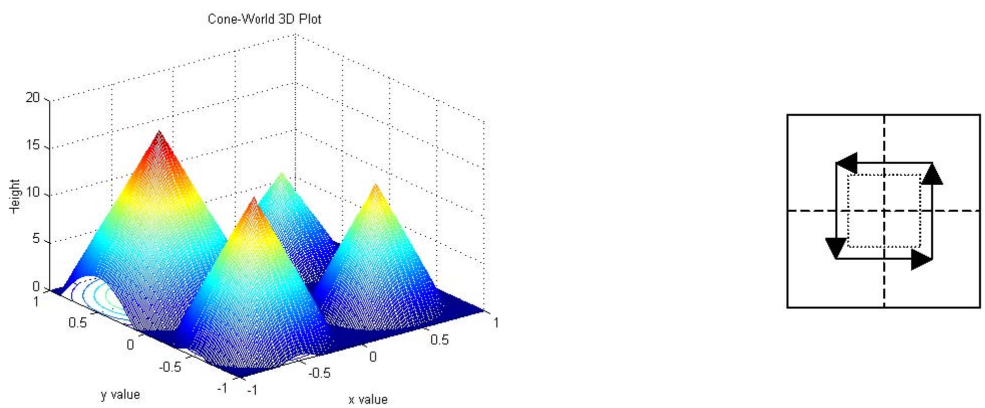

FIGURE 12 Dynamic cones world and the moving pattern

In all of our experiments, the parameters for CA were set as follows: the population size was set to 200, and the maximum number of generations between shifts was set to 200.

RESULTS AND DISCUSSION

We discuss the behavior of each knowledge source in terms of the placement of the individuals they influence in the search. The term "patch" is used here to indicate the area in which a given knowledge model is likely to place individuals. A patch is identified statistically by its center or the average location of those individuals that use it as a model this time-step, and its size is bounded by the standard deviation of the locations of these individuals. Figures 15 through 19 give the changing location of the center of the patch for each knowledge source over the 1,600 time-steps over the 2D grid in response to the movement of the cones. An arc connects the patch center at one time-step to the patch center at the next.

Here agents are given a number of knowledge models that they can use. Each time, the agents look at their previous performance and select the best model for them. From this perspective, we can see the strategy associated with each of our five knowledge models here. For example, if we look at Figure 13, we see that center for the topographic patch moves around the center of the region, within a fairly constrained radius. The strategy behind this knowledge model is to basically place individuals at the whole region so that overall intake of resources by all individuals for this strategy will be relatively constant. From this standpoint, it is an exploratory knowledge strategy that is sampling the current total environment.

Looking next at situational knowledge, Figure 14, we see a completely different scenario. This knowledge model is tracking the optimum. Notice that there are two lines connecting each of the dense regions. This corresponds to the fact that the maximum valued cone travels around the circuit in two trips, and that situational knowledge jumps quickly from its old location to its new one and then exploits the new one intensely, given the tight radius at each location. There is variability in that the cones, when moved to the new quadrant, are placed randomly within it (not always in the same spot). Situational knowledge is an example of an exploitive knowledge model. In our traveling band analogy, it corresponds to the avid fans that follow their favorite group from venue to venue.

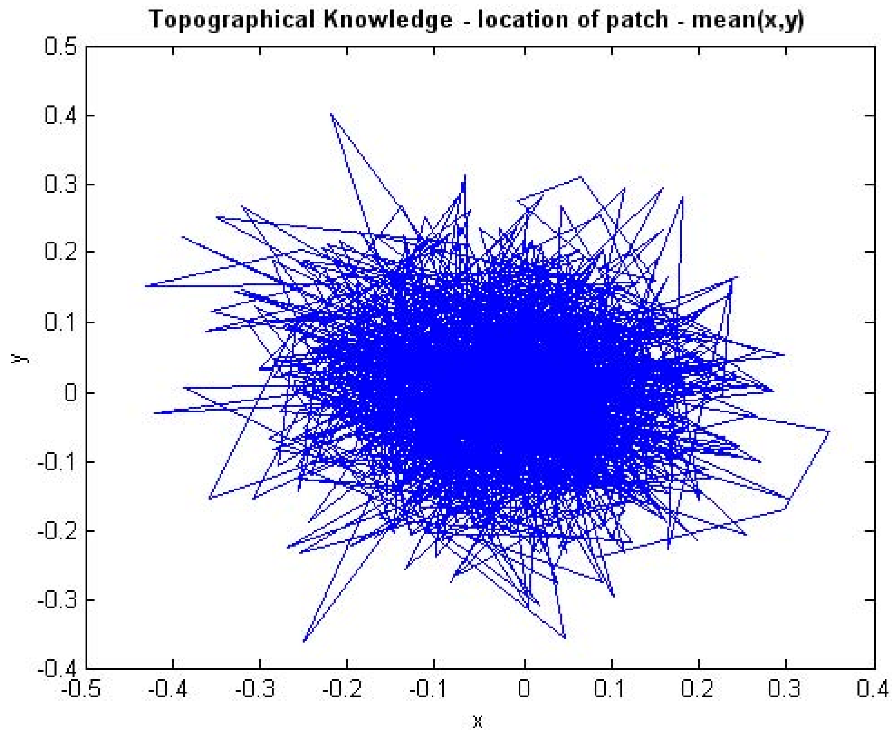

FIGURE 13 Change in the location of the center of the bounding box for topographic knowledge over 1,600 time-steps

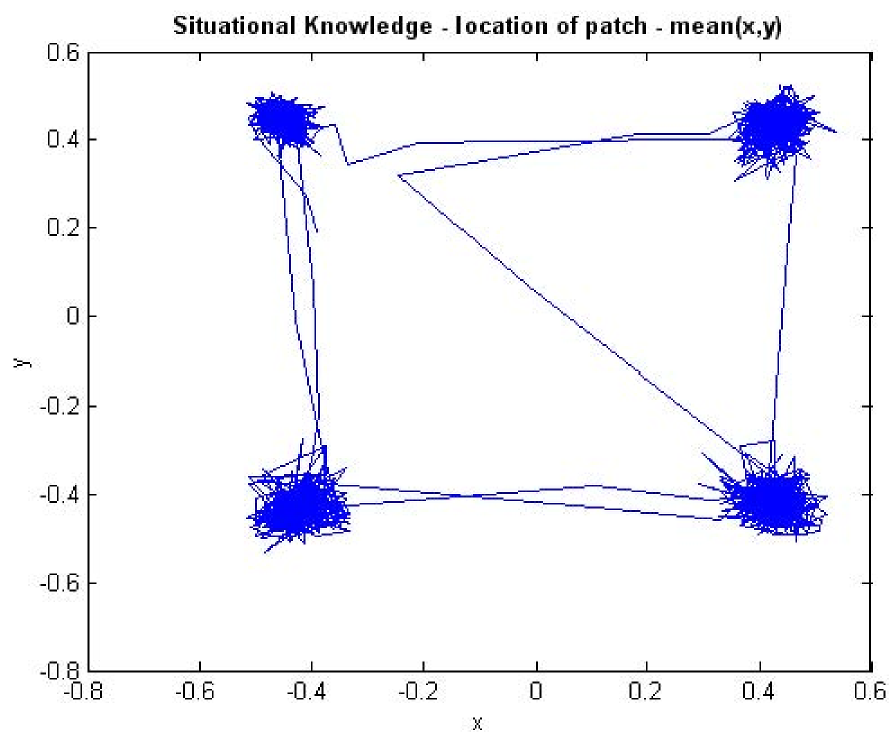

FIGURE 14 Change in the location of the patch center for situational knowledge over 1,600 time-steps

The movement of the domain knowledge patch is given in Figure 15. Domain knowledge was classified as an exploitative knowledge source in previous work, like situational knowledge. However, note that while it locates individuals on each of the four quadrants following the best one, there is an interesting emergent behavior here. There is a diagonal trace between the third and first quadrant. What this means is that this knowledge source is allowing those individuals to follow a shortcut from quadrant three back to quadrant one from where the best peak started. One reason for this is that given the popularity of the exploitation approach, there are many individuals attracted to the best peak by the time it gets to the third quadrant. The domain knowledge model uses the gradients to predict where the best cone will be going, and the individuals take a shortcut to get there ahead of the group.

This strategy is akin to those individuals who, when the number of exploiters becomes dense, are able to use their knowledge to project future venues and strive to be the first to arrive at the new opportunity. While this is an exploitive knowledge model, it has a relatively small population size, because if everyone were to take this choice, it would become disadvantageous. This works well when only a few are able to do it.

In Figure 16, we see the trajectory for normative knowledge. This knowledge was viewed as a type of exploratory knowledge source. Here we see that unlike topographic knowledge, it

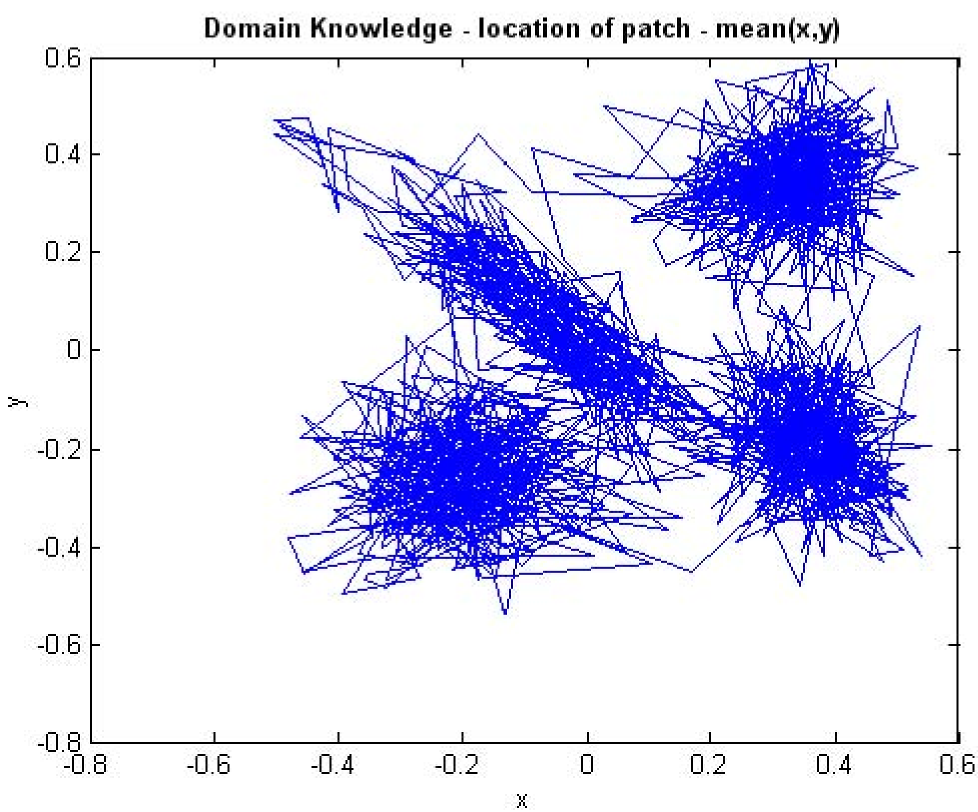

FIGURE 15 Change in location of the patch center for domain knowledge over 1,600 time-steps

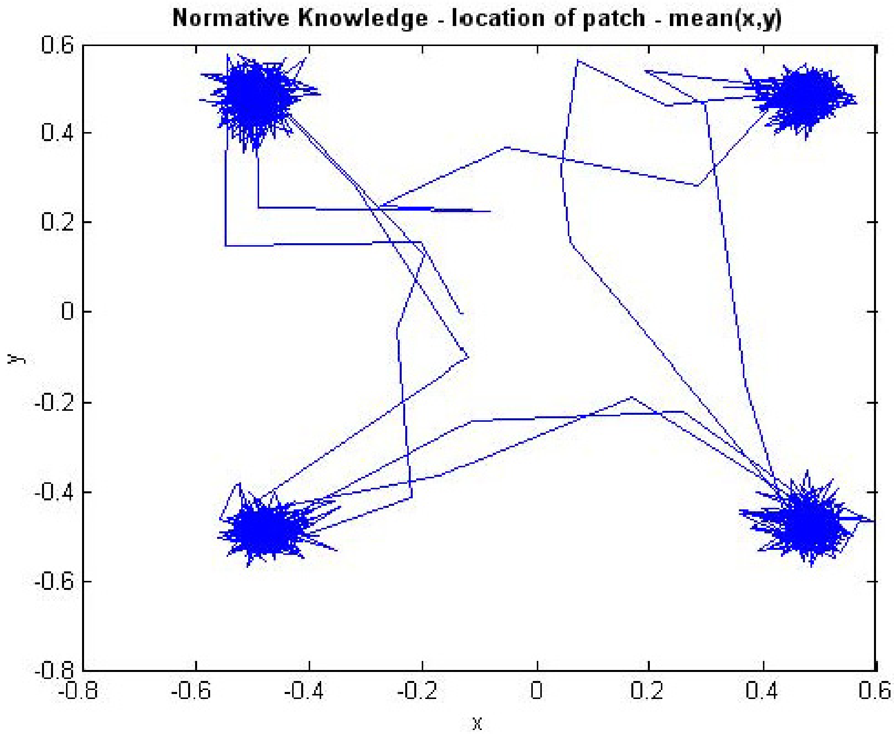

FIGURE 16 Change in location of the patch center for normative knowledge over 1,600 time-steps

moves from quadrant to quadrant following the optimal cone. Notice that again there are two connecting arcs between each quadrant, signifying the two cycles in our experiment. But now the arcs have a more meandering nature than that of the situational knowledge source. Whereas topographic knowledge explores the entire region, normative knowledge explores a subregion and is able to produce bounding boxes that connect one peak area to another. The key difference between normative knowledge and situational knowledge is that individuals using normative knowledge find and exploit the peak first, with situational knowledge directing individuals to follow. So while they both track the agents who follow their models to the best peak, they do it at different rates.

Finally, history knowledge is tracked in Figure 17. This knowledge model is focused on learning the pattern sequence. One can see that the radius of the patches is more spread out than that of situational and normative knowledge. Those using the history approach are able to encode trends and learn from them.

Other distinctions can be made between the knowledge models on the basis of parameters such as patch size, patch capacity, patch performance, and patch location. Looking first at the patch size in Figure 18 shows that topographic knowledge has the largest patch size, with normative and situational knowledge having the smallest sizes. All three patch sizes are relatively stable during the 1,600-year period. Both history and domain knowledge patch sizes experience an increase in size at the onset of a new cycle. Once the trajectory of the new cycle is determined, history knowledge returns to a stable patch size, similar to situational and normative knowledge. Domain knowledge, however, continues to decrease in patch size cycle, exploiting the cumulative changes in gradients.

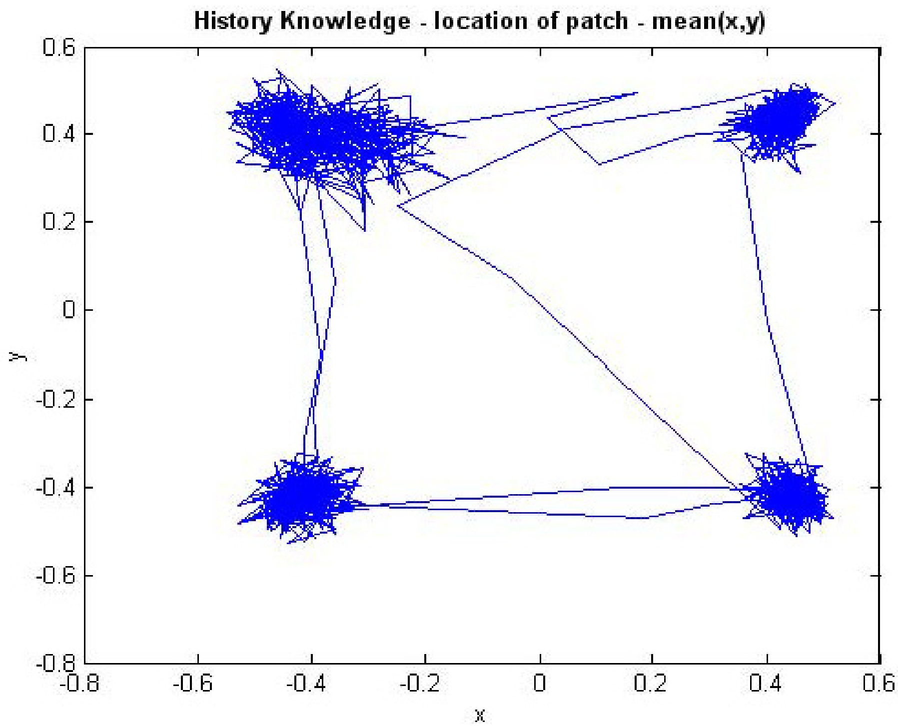

FIGURE 17 Change in location of the patch center for history knowledge over 1,600 time-steps

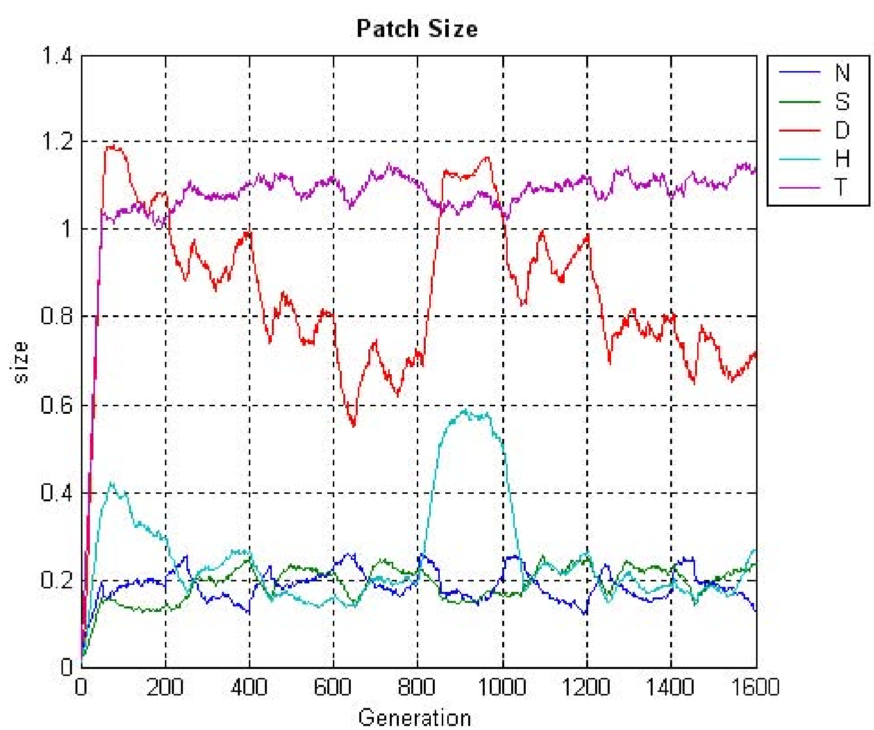

FIGURE 18 Patch size for all five knowledge sources over 1,600 time-steps

Patch capacity (Figure 19) corresponds to the number of individuals occupying each of the patches. Since the population is 200, and since not all individuals are found exactly within a patch, the total is somewhat less than 200. What is interesting here is at the onset of a cycle, history and domain knowledge recruit the most individuals, since they have information that can be used to predict the pattern of change. Topographic knowledge recruits a fairly constant number of individuals. However, situational and normative knowledge recruit more as the cycle continues, inheriting individuals from perhaps the domain and history models.

In Figure 20, patch performance is observed. What is interesting here is that each knowledge model exhibits a gain function that increases and then begins to decrease exponentially, but the adjustments take place at different frequencies. For example, topographic knowledge exhibits this shift at the onset of each cycle and is stable in between. The cycle for domain knowledge is longer (around 800 years) and shifted from the origin. Normative, situational, and history knowledge exhibit a higher frequency of change, around every 200 cycles. But the shifts are complementary in the sense that history and situational knowledge are going down, while normative is shifted so it is going up at the same time. This reflects that wave pattern of occupation mentioned earlier.

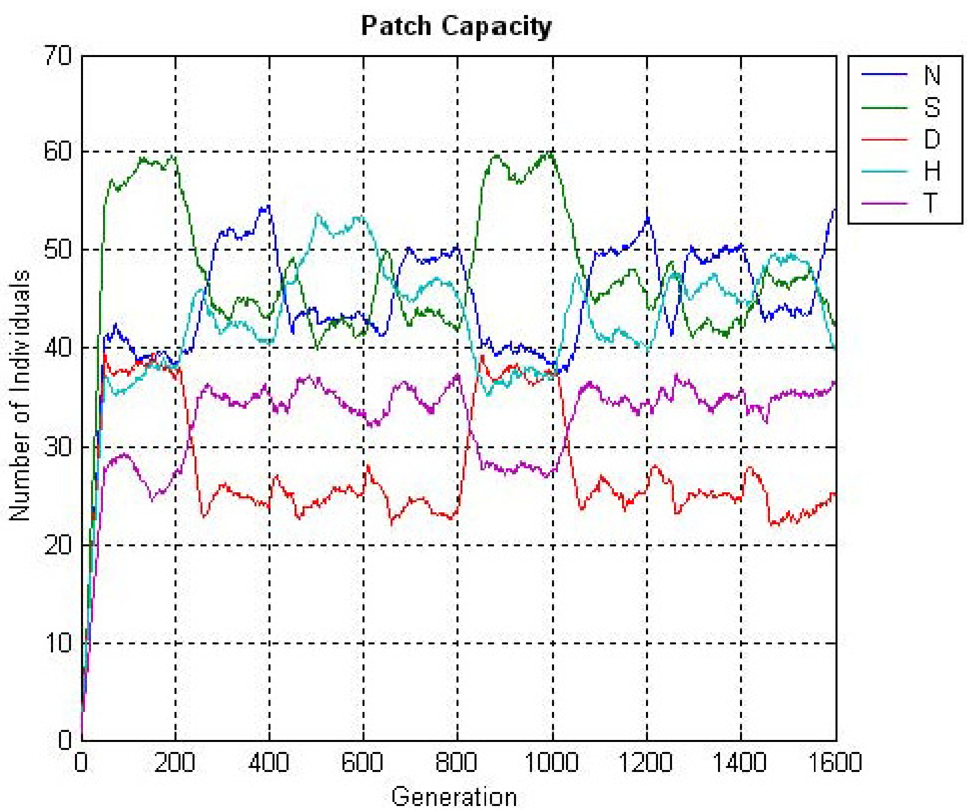

FIGURE 19 Patch capacity for all five knowledge sources over 1,600 time-steps

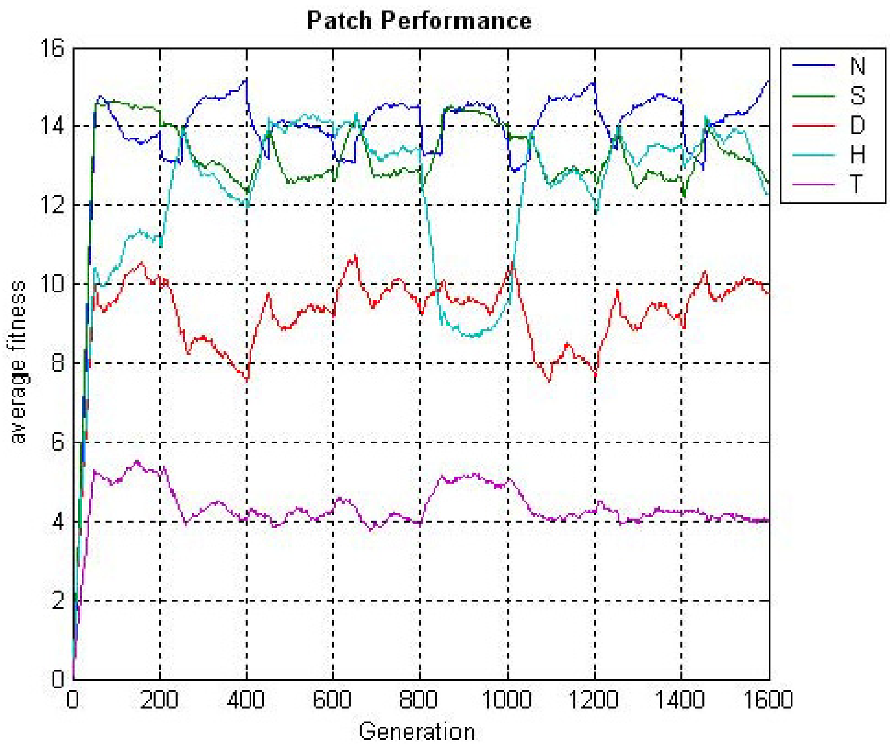

FIGURE 20 Patch performance for all five knowledge sources over 1,600 time-steps

In Figures 21 and 22, we observe the location of the patch centers in terms of the x and y coordinates, respectively. Notice that topographic knowledge is located consistently at the center of the region. However, we noticed earlier that its patch size increases and decreases over time, reflecting the need to produce more exploration. Situational and normative knowledge are quickly relocated into the center of the new quadrant at each phase. However, once they are in a quadrant, their patch size then changes, typically getting smaller. Domain knowledge, on the other hand, exhibits a hedging affect, which means its patch center tends to hedge back toward quadrant one as the cycle continues.

CONCLUSIONS

In this paper, we have investigated the impact of environmental dynamics on roles available to individuals. Here, the addition of a cyclical dynamic component to the model allowed each knowledge model to exploit a different aspect of the environmental dynamics. In previous reports that used a static configuration, certain knowledge sources, such as history and domain knowledge, had too little information to apply their expertise; therefore, those individuals that were controlled by those knowledge models exhibited behavior that was similar to other exploiter knowledge sources.

Here the additional information provided by the environmental dynamics allowed both history and domain knowledge to generate a pattern of patch movements, which was able to

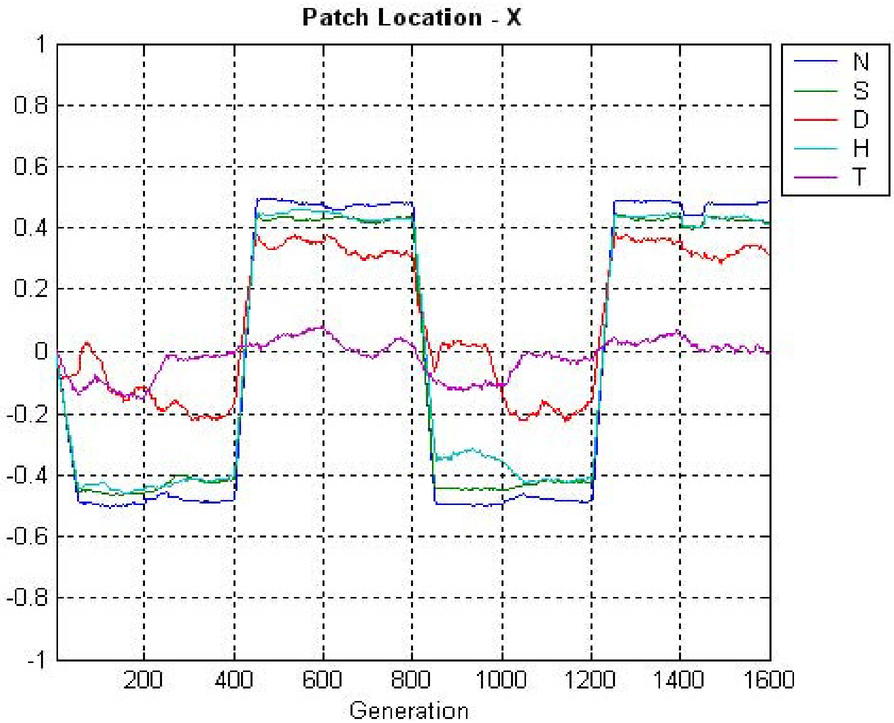

FIGURE 21 The x axis location of the patch center for all five knowledge sources over 1,600 time-steps

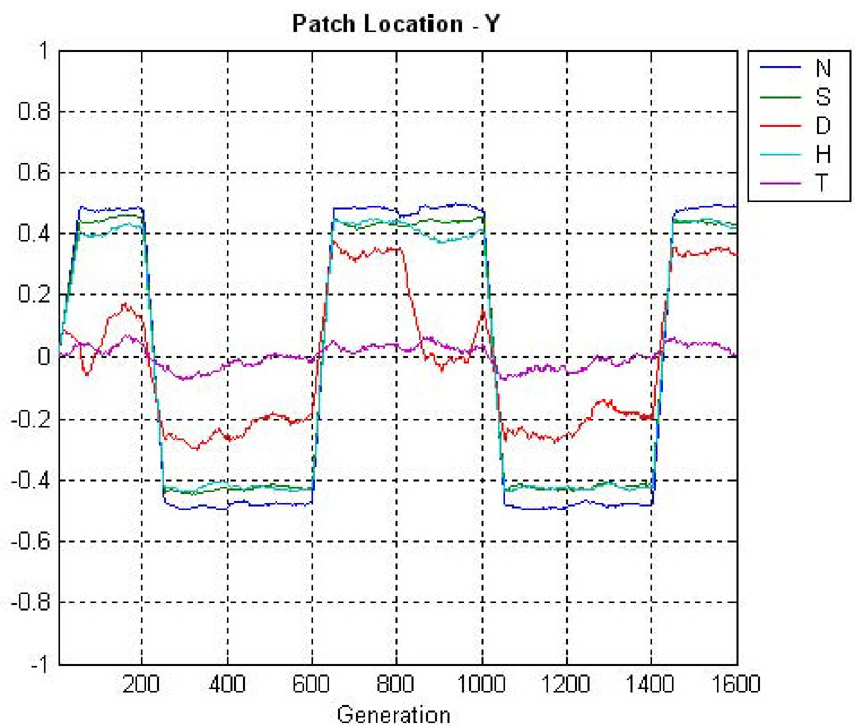

FIGURE 22 The y axis location of the patch center for all five knowledge sources over 1,600 time-steps

guarantee them some success with a small group of individuals. Thus, three distinct exploiter strategies emerged, when there was only one before. History and domain knowledge were able to predict aspects of the dynamics and use them in different ways. Both were able to arrive at the new optimal quadrant before situational knowledge, and both would move out as the agents driven by the situational model moved in. However, while history knowledge moved to the next patch, domain knowledge demonstrated the ability to hedge its path and move ahead to patches farther down the route. This behavior became clear as the number of individuals following a patch increased after year 600 and 1200.

Likewise, the cyclical nature of the environmental dynamics caused the two exploratory knowledge sources to also differentiate their behaviors. While topographic knowledge focused on a central location, normative knowledge mined the related regions but adjusted its patch location as more exploiters were attracted there. In future work, we anticipate that by adding in different dynamics, we will affect the mix of strategies that emerge.

REFERENCES

Barnard, C.J., and R.M. Sibly (1981). "Producers and Scroungers: A General Model and Its Application to Captive Flocks of House Sparrows." *Animal Behavior* 29:543–550.

Beckers, R., J.L. Deneubourg, and S. Goss (1992). "Trails and U-turns in the Selection of the Shortest Path by the Ant Lasius Niger." *Journal of Theoretical Biology* 159:397–415.

Charnov, E.L. (1976). "Optimal Foraging: the Marginal Value Theorem." *Theoretical Population Biology* 9:129–136.

Chung, C. (1997). *Knowledge-based Approaches to Self-adaptation in Cultural Algorithms,* Ph.D. Thesis, Wayne State University.

Chung, C., and G.R. Reynolds (1998). "CAEP: An Evolution-based Tool for Real-valued Function Optimization Using Cultural Algorithms." *International Journal on Artificial Intelligence Tools* 7(3):239–291.

Clayton, N.S., D.P. Griffiths, and A. Dickinson (2000). "Declarative and Episodic-like Memory in Animals: Personal Musings of a Scrub Jay," in *The Evolution of Cognition.* Edited by C. Heyes and L. Huber, MIT Press, Cambridge, MA.

Colorni, A., M. Dorigo, F. Maffioli, V. Maniezzo, G. Righini, and M. Trubian (1996). "Heuristics from Nature for Hard Combinatorial Optimization Problems." *International Transactions in Operational Research* 3(1):1–21.

Dorigo M., V. Maniezzo, and A. Colorni (1996). "Ant System: Optimization by a Colony of Cooperating Agents." *IEEE Transactions on Systems, Man, and Cybernetics* 26(1):29–41.

Epstein, J., and R. Axtell (1996). *Growing Artificial Societies.* MIT Press/Brookings Institute, Cambridge, MA.

Goss, S., S. Aron, J.L. Deneubourg, and J.M. Pasteels (1989). “Self-organized Shortcuts in the Argentine Ant.” *Naturwissenschaften* 76:579–581.

Holland, J.H. (1998). *Emergence.* Addison-Wesley Press, Reading, MA, pp. 1–10.

Hölldobler, B., and E.O. Wilson (1990). *The Ants.* Springer-Verlag, Berlin, Germany.

Hu, X., C.R. Eberhart, and Y. Shi (2003). “Engineering Optimization with Particle Swarm,” in *Proceedings of 2003 IEEE Swarm Intelligence Symposium,* Indiana, IEEE Press, pp. 53–57.

Iacoban, R., R.G. Reynolds, and J. Brewster (2003). “Cultural Swarms: Modeling the Impact of Culture on Social Interaction and Problem Solving, ” in *Proceedings of 2003 IEEE Swarm Intelligence Symposium,* Indiana, IEEE Press, pp. 205–211.

Jin, X., and R.G. Reynolds (1999). “Using Knowledge-based Evolutionary Computation to Solve Nonlinear Constraint Optimization Problems: A Cultural Algorithm Approach,” in *Proceedings of the 1999 Congress on Evolutionary Computation,* Piscataway, NJ, IEEE Service Center, pp. 1672–1678.

Kennedy, J., and R.C. Eberhart (1995). “Particle Swarm Optimization,” in *Proceedings of the IEEE International Conference on Neural Networks,* Perth, Australia, IEEE Service Center, pp. 12–13.

Kennedy, J. (1999). “Small Worlds and Mega-Minds: Effects of Neighborhood Topology on Particle Swarm Performance,” in *Proceedings of the 1999 IEEE Congress on Evolutionary Computation*, Piscataway, NJ, IEEE Service Center, pp. 22–31.

Kennedy, J., R.C. Eberhart, and Y. Shi (2001). *Swarm Intelligence.* Morgan Kaufmann Publishers, San Francisco, CA.

Maniezzo, V. (2000). *Ant Colony Optimization: An Overview.* Available at http://www3.csr.unibo.it/~maniezzo/didattica/Vienna/ACOintro.pdf. Accessed March 12, 2005.

Morrison, R., and K. De Jong (1999). “A Test Problem Generator for Non-stationary Environments,” in *Proceedings of the 1999 Congress on Evolutionary Computation,* Piscataway, NJ, IEEE Service Center, pp. 2047–2053.

Ostrowski, D.A., T. Tassier, M. Everson, and R.G. Reynolds (2002). “Using Cultural Algorithm to Evolve Strategies in Agent-based Models,” in *Proceedings of the 2002 IEEE World Congress on Computational Intelligence,* Wakiki, HI, pp. 2–7.

Peng, B. (2005). *Knowledge and Population Swarms in Cultural Algorithms for Dynamic Environments,* Ph.D. Thesis, Wayne State University.

Reynolds, R.G. (1978). “On Modeling the Evolution of Hunter-Gatherer Decision-making Systems,” *Geographical Analysis* 10(1):31–46.

Reynolds, R.G. (1986). “An Adaptive Computer Model of Plan Collection and Early Agriculture in the Eastern Valley of Oaxaca,” in *Guila Naquitz: Archaic Foraging and Early Agriculture in Oaxaca, Mexico*, K.V. Flattery, Editor, Academic Press, pp. 439–500.

Reynolds, R.G. (1994). “An Introduction to Cultural Algorithms,” in *Proceedings of the Third Annual Conference on Evolutionary Programming,* World Scientific Publishing, pp. 131–139.

Reynolds, R.G., and B. Peng (2005). “Knowledge Learning and Social Swarms in Cultural Systems.” *Journal of Mathematical Sociology* 29:1–18.

Reynolds, R.G., and S.M. Saleem (2005). “The Impact of Environmental Dynamics on Cultural Emergence,” in *Perspectives on Adaptations in Natural and Artificial Systems.* Oxford University Press, pp. 253–280.

Rychlyckyj, N., D. Ostrowski, G. Schleis, and R.G. Reynolds (2003). “Using Cultural Algorithms in Industry,” in *Proceedings of the 2003 IEEE Swarm Intelligence Symposium,* IEEE Press, pp. 187–192.

Saleem, S. (2001). *Knowledge-based Solution to Dynamic Optimization Problems Using Cultural Algorithms,* Ph.D. Thesis, Wayne State University.

Stephens, D.W., and J.R. Krebs (1986). *Foraging Theory.* Princeton University Press, Princeton, NJ.

Wynne, C.D. (2001). *Animal Cognition — The Mental Lives of Animals.* Palgrave Macmillan, Great Britain.

ADVERSARIAL ANALYSIS OF EVOLUTIONARY MODELS AND MULTI-AGENT SYSTEMS
(Toward theoretical foundations for generative social science)

G. ISTRATE,* Los Alamos National Laboratory, Los Alamos, NM

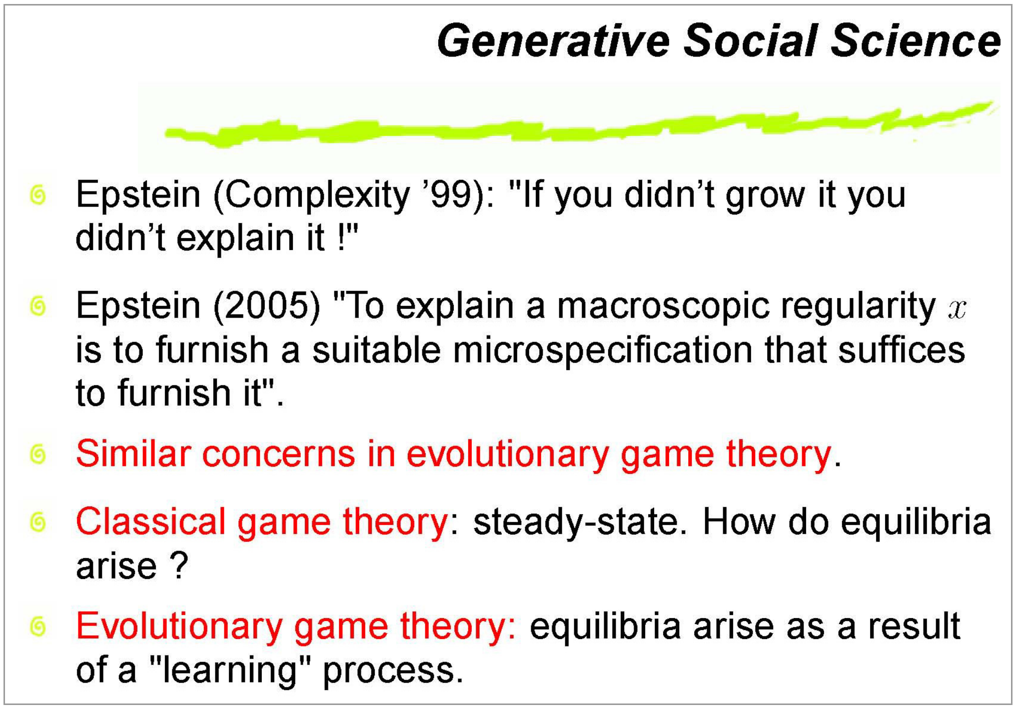

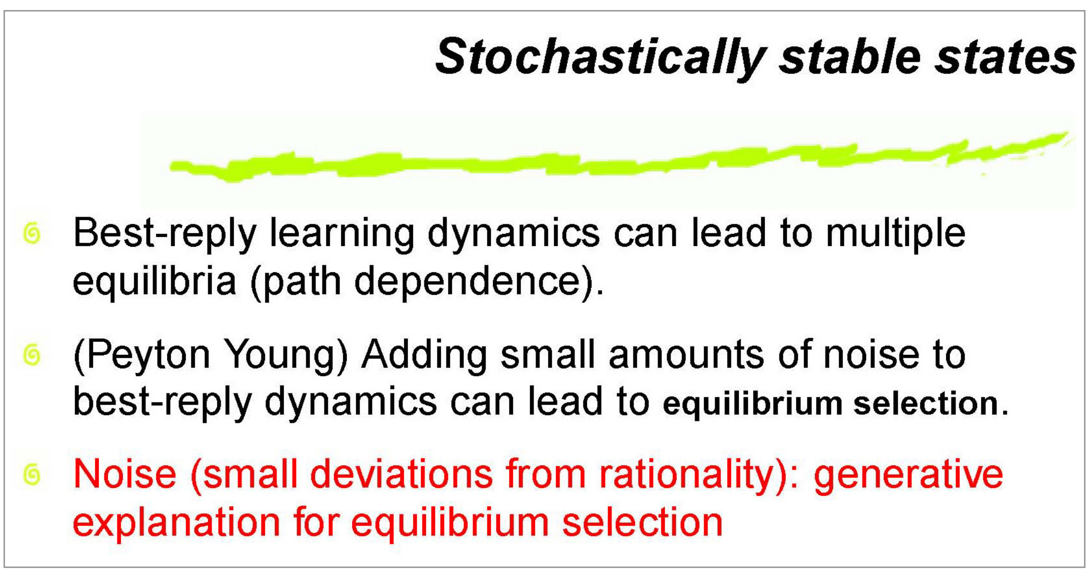

* *Corresponding author address*: Gabriel Istrate, CCS-5, Basic and Applied Simulation Science, Los Alamos National Laboratory, P.O. Box 1663, Mail Stop M997, Los Alamos, NM 87545; e-mail: istrate@lanl.gov.

Games/simulations as dynamical systems

- Multiagent simulations: interacting, nondeterministic dynamical systems.
- Robustness concerns: specification of interaction network, scheduling, dynamics.
- Most models assume some form of random scheduling. Not really plausible. Scheduling can make a difference (Huberman and Glance). Theory ?
- Approach: increase robustness of the models by considering adversarial scheduling.

Approach (II)

- Start with base case result with random scheduling.
- Isolate properties of random scheduling.
- Gradually eliminate some of these properties ...
- ... Until base case result no longer true.
- Identify feature of scheduler responsible for the failure.
- Eliminate this property (result holds again), etc.
- In the course of this process: *add more realistic features, more robust restatement of results.*

Setup

- Population games (Blume): agents at the vertices of a graph. Each agent has a *state*.
- When agent scheduled, play a game against some of its neighbors. Changes state as a result of game playing.
- Scheduler: specifies what agent can get scheduled at what time.

Example I: emergence of institutions

Individual Strategy and Social Structure:

An Evolutionary Theory of Institutions

by H. Peyton Young

- Examines how "economic and social institutions coordinate people's behaviors in various spheres of interaction"

Example I: emergence of institutions

strategies	A	B
A	a,a	c,d
B	d,c	b,b

- Strategy A is a *strict risk-dominant equilibrium*.
- That is $a - d > b - c > 0$.
- Selection of risk-dominant equilibria: Harsanyi and Selten.

Specifying dynamics

- When scheduled *agents play using the same strategy against each of their neighbors*.
- $\nu_i(z, \overline{x}_{-i})$, the payoff of the i'th agent, should he play strategy z with its neighbors is

$$\nu_i(z, \overline{x}_{-i}) = \sum_{(i,j) \in E} w_{i,j} a_{z,x_j}.$$

- If agent i is the one to update, $\overline{x}$ is the joint profile of agents' strategies, and $z \in \{\mathbf{A}, \mathbf{B}\}$ is the candidate new state,

$$p^\beta(x_i \to z | \overline{x}) \sim e^{\beta \cdot \nu_i(z, \overline{x}_{-i})},$$

Base case result

- Peyton Young: *under random scheduling the "all A" state the uniquely stochastically stable state.*
- Model of emergence of standards: gold vs. silver, driving on the left vs. right.
- Unrealistic feature: random norm adoption. No account of norm diffusion.

Properties of random schedulers

A random scheduler is:

(i) uniform: probability of getting scheduled is same.

(ii) non-adaptive: who gets scheduled does not depend on the past.

(a) who gets scheduled does not depend on *who got scheduled in the past.*

(b) who gets scheduled does not depend on the *past outcome of game-playing.*

(iii) fair: in $\theta(n \log n)$ steps all nodes get scheduled with probability 1.

Adversarial analysis

- *allow nonuniformity (drop (i)):* similar result to the one for baseline case.
- *allow adaptiveness (drop (ii a+b)):* can be just as fair as random scheduler and prevent stabilization.
- *only drop ii (b):* assume "social network of influences" (not necessarily the same as the game playing one). Scheduler: *random walk on this network*. Result again similar to the one for random scheduler.
- by now what was easy to show for random scheduling is quite nontrivial mathematically.

Making the result more robust

- *Time* until convention emerges: important !
- *Peyton-Young (based on Morris)* . Provably *small-world like structure implies $\theta(n)$ convergence time* for random scheduling.
- *Not true* for model with contagion.
- Instead of $\theta(n)$: new graph parameter related to *hitting time*.
- The time component of Peyton-Young's result: now true for new parameter.

Example II: PD with Pavlov dynamics

- n agents, situated at the nodes of a graph G.
- Each agent has a label from the set $\{0, 1\}$.
- At time zero the labels are chosen either uniformly at random, or according a fixed (but otherwise arbitrary) global configuration.
- At each step two of the players, i, j, that are connected by an edge update their labels from $X(i)$, $X(j)$ to $X(i) + X(j)$ (mod 2).

Base case result

- $(0,0) \to (0,0)$
- $(1,1) \to (0,0)$
- $(0,1) \to (1,1)$
- random scheduling: "all zero" unique fixed point, reached with probability one *red for all graphs with no isolated vertices*
- Convergence time (G. et al. 2002): exponential on complete graph, star graphs, $O(n \log n)$ on a cycle.
- Nonreversible Markov chain. Correlation: network structure $\to$ convergence time really nontrivial. More results (Mossel and Roch, arxiv.org/math.PR October 2005)

Genealogy of the model

- Shoham and Tennenholtz (AIJ 1994) "Colearning", distributed coordination model.
- Kittock (SFI proceedings 1994) experiments, this dynamics.
- Axelrod: Pavlov dynamics for IPD.
- Sidowsky "minimal social situation", Thibaut and Kelley, "mutual fate control" (1959).
- Coleman n player MSS (2005).

Types of scheduler and issues

- an edge-daemon is able to choose *both* players of the interacting pair.
- node-daemons choose only one of the players. The other player: random among the neighbors.
- *fairness*
- *adaptiveness*

Adversarial scheduling: results

- *Edge daemons are too strong*. One can preclude stabilization on "almost all" graphs, even for a non-adaptive daemon.
- *Nonadaptive node daemons:* similar to random schedulers.
- *Adaptive node daemons:* similar to random schedulers on almost all graph topologies (in random graph sense).

Convergence time

No mathematical results for convergence times adversarial scheduling. Convergence time seems consistent with the $O(n \log n)$ convergence time for random schedulers.

$\pi \vert n$	4	8	16	32	64	128
id	2.486	4.225	6.401	8.33	10.498	13.135
p3	2.469	4.039	5.807	7.662	9.639	11.718
$\pi \vert n$	256	512	1024			
id	16.091	17.954	20.331			
p3	14.323	16.054	19.826			

What about simulations ?

- other mathematical model (omitted) Schelling's segregation model (Peyton Young).
- model checking: technique used for hardware verification. Search for "bad events".
- scheduler: automaton. "Bad event": formula in temporal logic. Techniques from automata theory (Vardi and Wolper).
- More robust: model checking for interactive Markov chains (Herrmans).
- **LONG TERM:** adapting model checking MC to agent systems.

Conclusions

- Adversarial analysis is surprisingly feasible ...
- ... leads to robust results ...
- ... and could be used for agent-based simulations as well.

Theoretical results:

- with M.V. Marathe (VBI), S.S. Ravi (SUNY Albany CS).
- submissions to Games and Economic Behavior, Theoretical Computer Science. Available on request.

UNDERSTANDING INSURGENCY BY USING AGENT-BASED COMPUTATIONAL EXPERIMENTATION: CASE STUDY OF INDONESIA

A.R. CHATURVEDI,* Purdue University, West Lafayette, IN
D. DOLK, Naval Post Graduate School, Monterey, CA
R. CHATURVEDI, M. MULPURI, D. LENGACHER, S. MELLEMA,
P. PODDAR, C. FOONG, and B. ARMSTRONG, Simulex, Inc., West Lafayette, IN

ABSTRACT

Intra-state conflict is becoming an endemic feature of the post-Cold-War era, increasingly challenging international stability and security. Specifically, protracted violent conflict in the form of insurgency is being predicted as the most likely form of future warfare. This highlights the necessity of understanding the conditions under which tensions emerge within a state and converge toward violent conflict. In this paper, we use agent-based modeling as an integrative tool to understand the conditions that favor the emergence, duration, and intensity of insurgency. We present a Virtual International System developed in the Synthetic Environment for Analysis and Simulation (SEAS-VIS) to analyze insurgency in a strife-torn region of the world. SEAS-VIS provides an environment in which to conduct computational experimentation as a way to begin to understand the largely qualitative aspects of insurgency. The theoretical models used in building SEAS-VIS agents are calibrated from open-source data and validated against published real-world incidents. We then use the validated SEAS-VIS to analyze dynamic interrelationships among grievances, level of resources, and organizational capacity to mobilize members toward social action.

Keywords: Insurgency, agent-based computational experimentation, SEAS-VIS

INTRODUCTION

Low-intensity, protracted civil conflict in the guise of ethnic, religious, regional, or linguistic differences is increasingly becoming an endemic feature of the post-Cold-War era, threatening the territorial integrity and stability of various countries in the present international system. New and fast-evolving trends have contributed extensively to this growing global security threat. Globalization, especially with regard to travel and the speed of information interchange, is facilitating cooperative aggression by like-minded but far-flung individuals and groups. Messages posted on the Internet sites by radical groups to spread their ideology, mobilize for specific causes, generate funds, claim responsibilities for recent attacks, and divulge the technical know-how of weapons construction are all becoming a common phenomenon. Similarly, privatization of weapons is not only facilitating the ease of weapons acquisition but also putting the potential of macro-terrorist acts into the hands of small groups or even individuals (Victoroff 2005). According to Fearon and Laitin, the number of total dead from civil conflict (16.2 million) between 1945 and 1999 far outnumbers those from inter-state conflict

* *Corresponding author address*: Professor Alok R. Chaturvedi, Krannert School of Management, Purdue University, West Lafayette, IN 47907; e-mail: alok@purdue.edu.

(3.33 million). Compounding the problem is human suffering, as more and more people get displaced as a result of endemic violence and economic devastation. This necessitates the need for research communities as well as policy makers to understand the conditions under which tensions emerge within a state and converge toward violent conflict.

Insurgency is a dynamic, adaptive, and nonlinear form of low-intensity warfare. It is defined as "a technology of military conflict characterized by small, lightly armed bands practicing guerrilla warfare from rural base areas…that can be harnessed to diverse political agendas, motivations, and grievances" (Fearon and Laitin 2003). We complement this definition by adding the urban base as a chosen tactical area of operation for present-day insurgencies, since the urban terrain lends itself to anonymity, camouflage, public attention, recruiting and logistical support, and extensive media coverage. One of the key strategies of insurgency is to prolong the fight against the dominant power through asymmetric means (mobile conventional war) in order to discredit and delegitimize the government. Interrelated strategy is to maintain the precarious balance between creating dissension (through terror tactics to decrease support for the dominant power) and increasing sympathy amongst the masses for the rebel cause and/or for possible recruitment. For its production and maintenance, insurgency focuses on coordination at several levels: low-key political organization focused on recruitment and infrastructure; continuous procurement of resources for maintenance functions, such as recruiting and training; and information dissemination to generate a level of popular support. Propaganda, bombings, kidnappings, assassination, and assaults on key infrastructures are some of the known tactics that insurgents employ to create havoc and insecurity.

While conspicuous attempts are being made to synthesize explanations of insurgency, few methodological tools are available that fully integrate the theories and strategies at various levels of a socio-political system: individual, group, national, and international. This paper uses Synthetic Environment for Analysis and Simulation-Virtual International System (SEAS-VIS) (Chaturvedi et al. 2004), an agent-based system, to study insurgency in Indonesia. In recent years significant research has emerged using agent-based modeling as a technique to elucidate the causes of protracted civil conflicts. Notable studies are on ethnic mobilization (Bhavani and Backer 2000; Srbljinovic et al. 2003; Cederman 2005), emergence of ethnocentrism (Axelrod and Hammond 2003), and emergence of secessionism (Lustick 2004), to name a few.

One of the advantages of using agent-based modeling is that it overcomes some of the difficulties associated with the real world. One of the difficulties pertains to the scarcity of comparable and generalizable cases of insurgency, which, in turn, are context- and time-dependent. Second, the enormity of variables and interaction effects and the immense difficulty in gathering relevant data pose daunting challenges to scholars and policy makers and can, at best, lead to only a partial understanding of insurgency and of its mitigation. Finally, real-life cases are serious risks to the local implementers, who often lack the necessary information or the optimal solutions for conflict resolution (Lustick et al. 2004).

Computational experimentation methodology presents an innovative way of analyzing protracted conflicts. In this approach, one re-creates the environment on the basis of theoretical models of behaviors and calibrates them to fit the situation at hand. If the theoretical models are robust enough, then the situations when re-created can give revealing insights into the situations under investigation. Obviously, there will always be a lack of data and deep understanding of the flow of information, the interaction between the key actors, and the cascading effects of events leading up to the conflict. An agent-based synthetic environment allows us to fill in the gaps through experimentation with the solitary and collective behaviors of individuals, groups, organizations, and institutions. Specifically, agent-based modeling can have immense usage in the social sciences that are still concerned with how macro-level phenomena emerge from micro-level actions. According to Sawyer, the "emergence of macro from micro is perhaps the most interesting feature of artificial societies. In the artificial societies…macro-structural phenomena emerge, attain equilibrium, and remain stable over time. Thus, artificial societies provide sociologists with a tool to explore the micro-to-macro transition" (Sawyer 2003, page 333). As Cederman (2004, page 6) aptly portrays:

> "agent-based models constitute artificial and indeed simplified worlds in which the plausibility and consistency of well-specified causal mechanism can be evaluated in a context that is more complex than that of standard, rationalistic modeling tools, but still much simpler than the real world. Serving as a stepping stone between micro and macro analysis, such models can help untangle interacting mechanisms that together generate the phenomenon to be explained. This perspective defines a generative research strategy that starts from such patterns and moves backward in the search for candidate mechanisms that could generate observed outcomes."

Further, this new tool allows us to integrate and evaluate various existing theories, paradigms, and courses of actions in a single holistic framework. This "third way of doing science," as eloquently stated by Axelrod (2003), is a "virtual" interactive system that creates artificial autonomous agents that mimic the behavior patterns of their counterpart in the real world. These autonomous agents "have control over their own behavior and can act without the intervention of humans or other systems" (Sawyer 2003). They can interact with other agents within the virtual environment and are able to communicate, negotiate, and cooperate with each other. Agent-based simulations allow the following (Buodriga and Obaidat 2004; Chaturvedi et al. 2005; Sawyer 2003):

- Virtual experimentation, in which consequences of decisions can be measured and analyzed;
- Integration of multiple theories from various specialized disciplines, for a comprehensive understanding of underlying phenomena;
- Creation of representation of agents with multiple decision strategies, both rational and nonrational;
- Modeling of heterogeneous actors who can modify their behavior during the course of the simulation; and
- Facilitation of a seamless and interchangeable integration of human and software agents.

In the following sections, we begin by briefly describing some of the key theoretical premises that elucidate conditions favoring insurgency — premises that we use for the model development and computational experimentation. This is followed by a description of SEAS-VIS. We then present the configuration of a small-scale artificial society within the VIS concept to study the impact of external shocks (tsunami) and the critical role of organizational mobilization on the level of insurgency.

DETERMINANTS OF REBELLION: MOTIVE, RESOURCES, AND ORGANIZATION

Research on conditions favoring rebellion is rich and varied. Theories explicating rebellion are a combination of economic, political, symbolic, and psychological factors essential to the development of conflict. We draw upon three well-established concepts — grievance, resources, and mobilization — as explanations of protracted rebellion.

Deprivation Theory

One of the dominant perspectives in the study of intra-state conflict is the deprivation model, which examines the range of discriminations and disparities experienced by minority groups as contributing factors of rebellion. Deprivation can be in the form of psychological (perceived inequality), economic (resource inequality), political (repression, lack of political rights or representation), or social inequity (inequality such as group domination and suppression) (Gurr 1970, 2000; Schmid 1983; Harmon 2000; Krueger and Maleckova 2002; Duckitt et al. 2002; Post et al. 2003; Besancon 2005). Two underlying assumptions characterize this concept. First, rebellion may be due to an increase in the gap between expectation and outcome (i.e., a gap between the valued things and opportunities that people think that they are entitled to and the things and the opportunities they actually get). Second, there is continuous and crucial comparison with other people or groups. It is the sense that one's group is not doing as well as other groups. Thus, deprivation is a "psychological process in which judgment is made relative to one's own expectations. The aggregation of these individual perceptions and frustrations leads to a social movement intent on violent political change" (Regan and Norton 2005).

Rebel-Resources Theory

A more recent body of research analyzes extensively the connection between natural resources and likelihood of conflict (Collier and Hoeffler 2004; Fearon 2004; Weinstein 2005; Regan and Norton 2005; Humphreys 2005; Lujala et al. 2005). In an influential paper, Collier and Hoeffler (2004) suggest that states that depended extensively on natural resources for capital generations were more prone to civil violence. Natural resources (especially lootable resources) can provide finances to rebel groups and increase the prospects of their success and decrease prospects for peace, since continued conflict may be more profitable for the rebels than an outbreak of peace (Addison et al. 2002; Ross 2004). Notable cases are in Sierra Leone, Congo, and Angola, where rebels used revenues from diamonds and other natural riches to finance their conflict against the government. When natural resources are concentrated in one area of the country, insurgent groups may be motivated by the assumption that seceding may be prosperous. Resources are also used as selective incentives to overcome the uphill battle of convincing and motivating individuals to rebel (Weinstein 2005).

One of the corollaries to the natural resource hypotheses is the existence of weak and natural-resource-dependent economies as being more violence-prone and a fertile and conducive environment for the development of insurgencies (Fearon and Laitin 2003; Collier and Hoeffler 2004). Especially in developing economies, the legitimacy of the government is much weaker as regimes are narrowly based, come to power by force, and remain in power by suppression. Maintaining power for these regimes also requires effective control over natural resources, since they remain vulnerable to many different groups who would like to gain control of the state through that very means. Regimes, as well as subnational groups, may also advance their interest by seeking outside support that is often granted, since abundant natural resources within a country may be an incentive for third parties, such as states or corporations, to engage in or foster civil conflict. A classic case in point is the competition between the United States and France over oil in Chad and their subsequent interference in Chadian politics that has "made and broken political leaders, has incited violence, and has shaped political agendas" (Humphreys 2005). There is extensive research on the international linkages that provide groups with popular encouragement through information warfare or resources, such as weapons, money, and training, as necessary to prolonging insurgency (Byman et al. 2001; Lobell and Mauceri 2004). Where the central authority has weakened or collapsed, predatory outside groups can take advantage of this situation in order to capture the spoils. Ethno-religious groups with affinities in neighboring states or otherwise can solicit support, resulting in diffusion of crisis.

Organizational Mobilization Theory

A contending perspective argues that organizations are the core protagonists of action and activism as a result of the mobilizational capacity of groups and organizations. Thus, deprivation is a necessary but not a sufficient explanation for rebellion (Tilly 1978; Tarrow 1994; Lichbach 1998). For example, advocates of resource mobilization theory focus on what compels aggrieved people to participate in social movements. They contend that organizations possess certain materials and resources that they use to generate actions that lone individuals are rarely capable of. These resources are generated by continuously participating in "supply chain" activities, such as resource procurement, accumulation, and recruitment of new members in order to sustain themselves. These resources, in turn, are directed toward activities that meet organizational goals. Organizations also provide their members with a sense of identity, existence, boundedness, coherence, agency, and mission that together may propel individuals toward violent behavior and provide justification for the same (McCarthy and Zald 1976; Tilly 1978; Jenkins 1983; Tarrow 1994; Klandermans 1984; Lichbach 1998; Brubaker 2004).

SEAS-VIS MODEL DESCRIPTION

We create a multi-level artificial society within the SEAS agent-based computational experimentation environment to test the three intra-state strife theories. We implement diverse social science theories in two distinct ways. (1) Certain fundamental or experimentally developed theories are explicitly encoded in the agents. Examples included well-being (Diener et al. 1993; Diener and Fujita 1995; Diener and Suh 1998; Diener and Lucas 1999; Peterson 1999; Diener et al. 1999; Kahneman et al. 1999), set point theories from psychology (Suh et al. 1996; Lucas et al. 2003; Lyubomirsky et al. 2005), and production and consumption theories from micro economics, etc. (2) Certain theories that represent emergent behaviors are observed and validated on the basis of the calibration of the primitives. Examples of such theories include sociological theories, such as social networks, and macro-economic theories, such as gross national product (GNP) and unemployment.

We configure Indonesia, a multi-ethnic virtual state within the SEAS-VIS, mirroring its counterpart in the real world. Virtual Indonesia is represented by four primitive constructs: individuals, organizations, institutions, and infrastructures (IOII). These four primitives are used to model higher-order constructs, such as geographical entities (nations, provinces, cities), political systems (type of government, political parties/factions), the military (soldiers, institutions), economic system (formal and informal structures), social systems (institutions, groups), information systems (print, broadcast, internet), and critical infrastructures (banking, oil and gas, electricity, telecommunications, transportation), as shown in Figure 1.

Political and social systems of a state are modeled as a multi-agent system representing the human elements. Individual citizen agents are constructed as a proportional representation of the societal makeup of a real nation. Each individual agent consists of a set of fundamental constructs: traits, well-being, sensors, goals, and actions. For example, a citizen agent is encoded with static traits (e.g., race, ethnicity, income, education, religion, gender, and nationalism) and dynamic traits (e.g., religious, political, societal, and violence orientations). We use Kahneman's (1999) concepts of subjective well-being, which refers to a person's assessment of his perceived state of happiness or well-being. The agent's well-being consists of eight needs: basic, political, financial, security, religious, educational, health, and freedom of movement. Traits and well-being together determine the set of basic goals for a class of agents. An agent uses its "sensors" to sense the environment and listen to messages from his/her leader(s), the media, and other members of the society. On the basis of the sensed information, each agent can autonomously choose from its repertoire of configurable action sets or adjust its goals. Traits, well-being, and goals determine the available actions each agent can take. For example, an agent can migrate to a different location (geography) to seek a better job to satisfy its financial well-being. Traits, well-being, sensors, and actions together determine the behavior of the agent.

We identify each agent's *desire* for each need. These desires are initially populated for each citizen on the basis of the socioeconomic class of the citizen. Further, we also identify

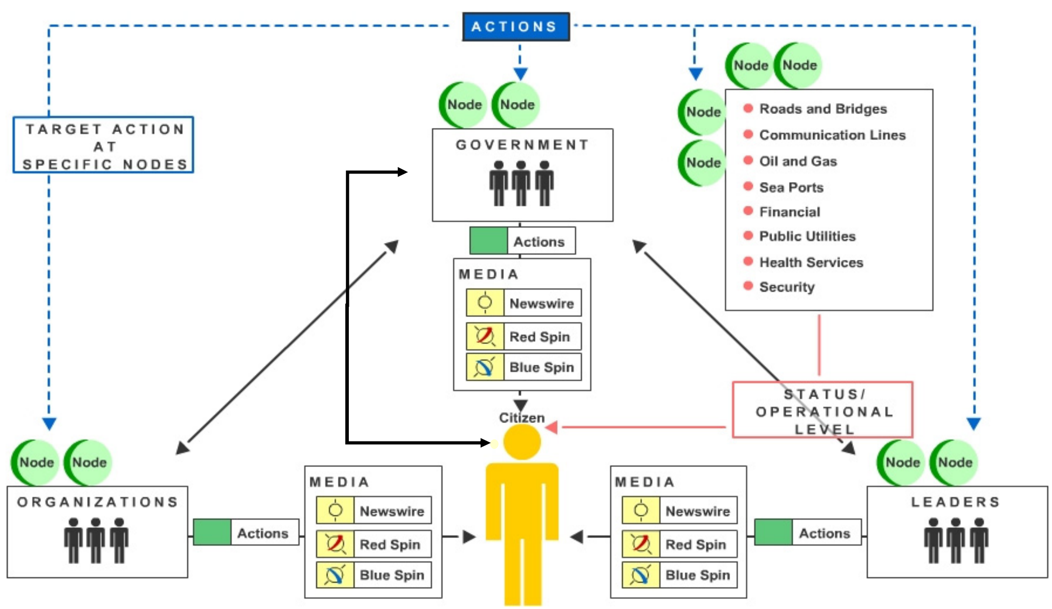

FIGURE 1 Schematic of SEAS-VIS

weights that identify the relative importance of the fulfillment of each need to the citizen. Each citizen forms a perception of the level of fulfillment of each need from several information sources, such as social groups, leaders, organizations, and the media. Each agent then identifies the deprivation of each need as the gap between the perception of a need and his/her desires for the need. By weighting the deprivation of each need, each citizen identifies the overall deprivation.

Over long periods of time, citizens adjust their desires for each need relative to their perception of that need. A citizen could be influenced to adjust desires by organizations and leaders through coercion or persuasion. Each citizen's desires are also influenced by the desires of other citizens in his social network. The increase in perception of needs of a few citizens could lead to higher desires not only in those citizens but also other individuals in their social groups. Such an adjustment of desires across social groups whose perceptions do not change could lead to a higher sense of deprivation in the citizens of the social groups.

Citizens adjust their *weights* as certain needs become more significant because of conditions in the environment. Citizens focus on needs that they are most deprived of and attach less significance to those needs that are fulfilled. Organizations, leaders, and the media could influence a citizen into adjusting weights by attaching significance to certain issues. Citizens are also influenced or coerced by their social groups in the needs to which they attach the most significance.

The leader agent is encoded with influence levels that reflect his/her power within the group, organization, or institution. A leader agent is categorized as social, religious, and/or political and has a repertoire that is larger than that of citizen agents and includes additional traits, such as power base, ideology, and his/her stance on economic, political, and social policies. Leader agents are able to affect the political and social climate of the synthetic environment and impose their stances upon citizens and organizations to promote their respective goals. The goal of leader agents is to set the agenda of the organization or institution in which they reside and persuade the citizen/member agents to make decisions that favor those positions.

Clusters of agents form groups, organizations, or institutions. They differ from individuals with regard to the rules that govern their behavior and intent. Groups are either informal or formal. Formal groups' rules of engagement are published and are relatively static, while those of informal groups are only known to their members and continuously evolve on the basis of interactions among the environment, leader, and members.

An organization is composed of a structured group of artificial human citizen and leader agents. Citizens that subscribe to an organization make up the member population, and the combined behaviors and interactions of members and leaders results in the behavior for the organization. Organizational leadership constantly seeks maintenance and growth of the organizational membership by providing tangible and intangible benefits, and citizens subscribe on the basis of a perceived level of benefit that is received from the organization. Leaders attempt to influence the organization to align with their ideologies by framing issues and attitude sharing. Members also influence each other's attitudes through the formation of intra-group social networks that emerge from levels of affinity between members. In addition, through inter-organization networks, attitudes and resources may be shared between organizations. Through these internal and external interactions, organizations cause significant changes in perception and attitude change and become core protagonists of activism in the model. The interaction between organizations and other entities is shown in Figure 2.

Organizational deprivation is modeled in terms of the well-being and attitudes of leaders and members. As well-being decreases and attitudes become more hostile, an organization may choose to mobilize to action in the form of a demonstration, a riot, or an attack. The course of action depends upon the ideology of the organization and the extremity of the unrest. An organization that is more willing to use violence to achieve its means may be more prone toward rioting or suicide-bombing attacks, whereas an organization that subscribes to nonviolent means may choose to arrange a demonstration.

An organization exercises its power through the control over its resources and its ability to procure and maintain its resource base. Organizational networks, member recruitment, and member maintenance are primary sources of resource procurement and maintenance. A higher level of control over these resources contributes to a higher level of effectiveness when organizations are mobilized to action.

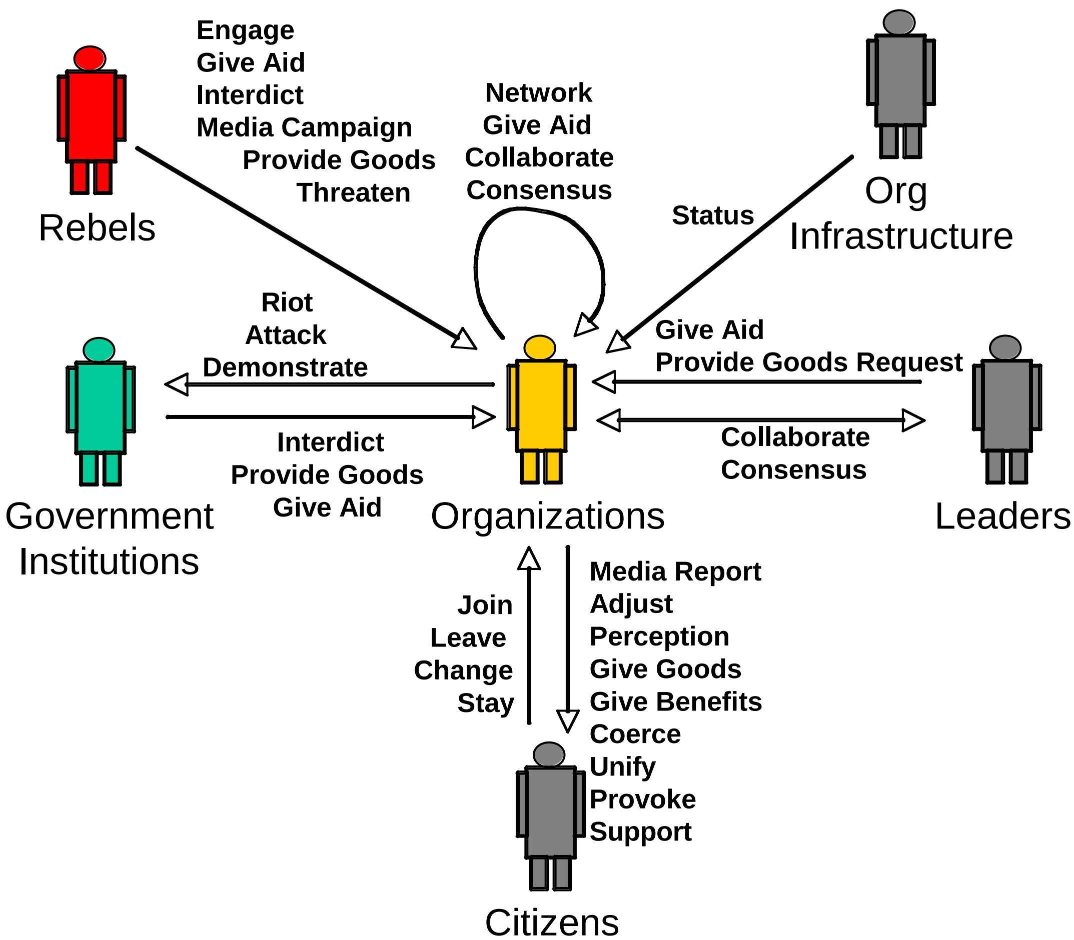

FIGURE 2 Interactions between organizations and other entities

Agenda setting is a significant contributor to organizational activism in the model. Every organization sets an agenda based on its ideology and goals that directs mobilization. This agenda is adaptive and dynamic as a result of intra-organizational and inter-organizational information sharing. Internally, leaders influence the attitudes and perceptions of well-being of their members and other leaders, and members share like information within their own social networks. Externally, organizational networks impact perceptions and attitudes through the interactions among leaders as well as member social networks across organizations.

Within our model, the media also play a significant role in providing information to members in the form of reports on well-being and attitudes. Media organizations consist of television, radio, newspapers, and magazines. They make choices about what information to cover, which people to cover, what statements to report, what story elements to emphasize, and how to report the information. The media is able to set the agenda for domestic policies as well as foreign policy issues. Incidents are framed on well-being components and formalized in a media report. For example, if the media's agenda is to arouse public against the government and if basic needs are below a certain threshold level, then the media frames it as government being responsible for the dire conditions of the people. Citizens subscribe to a media organization on the basis of their ideological bent. Subscription to a particular media is dependent upon the congruence of the ideology of the media with the ideology of the citizens subscribing to it. Media organizations are primarily focused toward framing the issues for their audiences in such a way that they increase their viewership as well as their influence. When the media infrastructure agents are reduced in their capacity to report, then the media conglomerates are also decreased in their ability to spin reports.

We model institutions as "governmental entities," such as the army, police, legislature, courts, executive, bureaucracy, and political parties — entities that are able to formulate policies that are legally binding and that have more discretionary resources. We also consider institutions as structures that are products of individual choices or preferences, the later, in turn, being constrained by the institutional structures (i.e., an interactive process). The government institution agents represent the leadership and various branches of the government. Institutions are like formal organizations with an additional power to influence the behaviors of members and nonmembers.

Examples of traits, well-being, sensors, and actions of different classes of agents are given in Table 1.

INSURGENCY INDICATOR

Epstein (2002) and Cederman (2004) have modeled civil violence wherein a central authority seeks to suppress unorganized rebellion. By building upon these models, and using the three well-established concepts of grievances, control over resource, and mobilization, we define a metric called *Insurgency Indicator* to indicate the overall level of insurgency against the government in a region. At the aggregate level, *Insurgency Indicator* is observed as the ratio of the number of mobilized citizens to the total population. It is given as:

Insurgency Indicator, S = total number of mobilized citizens/total population.

TABLE 1 Traits, goals, sensors, and actions for each entity type

IOI Categorization	Entity Type	Traits	Goals	Sensors	Actions
Individual	Citizens	Age Income Education Ethnicity Religion Ideology	Maintain and enhance personal well-being	Leaders Organizations/media Institutions	Demonstrate Riot Join organizations Leave organizations
	Leaders	Type Power oriented vs. affiliation oriented Responsive vs. ideologue Ethnicity Race Income Education Attitude towards group, state	Maintain and enhance personal influence Maintain and enhance the influence of their organization Maintain and enhance well-being of their members	Followers' well-being Organizational power base Control over resources	External Consensus Collaborate Internal Set agenda Unify Coerce
Organization	Informal groups Formal organizations Networks	Type Political, religious, social, economic, media Size Control over resources Ideology Ethnicity Nationalism Religion	Survive Maintain Increase membership Seek influence	Member well-being Other organizations	Demonstrations Riots Attacks Set agenda Collaborate Unify Seek consensus Coerce
Institution	Government	Type Political Military Economic Spatial Central Provincial Local Power Resource Competence Nationalism	Policy implementation Policy adjudication Policy enforcement Policy formulation Influence policies	Population's well-being component Public's confidence/legitimacy Public's trust Resource availability Other institution's actions Incoming actions DIME actions	Collaborate Unify Coerce Enforce Respond Prepare Recover Reconstruct Attack Ally Defend Aid Coerce Trade

Each individual agent evaluates its position at multiple levels in order to determine its intention to mobilize and join the rebellion. This intention of a citizen agent to rebel depends upon its grievance and its perceived net risk in acting against the grievance.

Each agent evaluates its personal grievance against its government. This is measured as a function of the agent's subjective well-being and its perception of its government's legitimacy. Therefore, high deprivation may be either counterbalanced by a high legitimacy or bolstered by a low one in producing a grievance against the government. It follows from the previous description of deprivation and organization models that the grievance of a citizen therefore depends on his or her base desires, perceived reality through media reports, and the actions/attitudes of organizations, leaders, and the government.

An agent's net risk in addressing its political grievance is the product of its level of risk aversion and perceived incarceration or punishment. The risk propensity of an agent reflects diminishing returns to increasing gains and losses as held by prospect theory, so that agents decreasingly become risk acceptant when faced with increasing erosion in well-being. The perceived probability of incarceration increases with repression and enforcement, while it decreases with the number of citizens already mobilized against the government.

Therefore, a citizen's intention to join the insurgency is determined as follows:

Intention to Rebel, I = f {grievance, risk propensity},

Grievance, G = f {subjective well-being, legitimacy},

Subjective Well-being, W = f {basic needs, political needs, financial needs, security needs, religious needs, educational needs, health needs, and freedom of movement needs},

Legitimacy, L = f {government actions; media, organization, and leader attitudes}, and

Risk Propensity, R = f {media, organization, and leader actions}.

ILLUSTRATIVE SCENARIO: INSURGENCY IN ACEH, INDONESIA

We create Virtual Indonesia (VI) within SEAS-VIS. VI consists of political, military, social, economic, information, and infrastructure entities or nodes. The behaviors of these nodes were mined from open source data (Polity IV, Indonesia Public Opinion Survey 2005, CIA Fact Book, Worldpress.org, *Europa Magazine*, etc.). We model behaviors of a total of 474,073 agents. Included in this count are 473,500 citizen agents, 9 named leaders, 9 named organizations, 9 media organizations, 14 sectors, and 406 critical infrastructure nodes. The interactions between these nodes are emergent. Individuals, organizations, and institutions modeled in VI are given in Table 2.

Our experimental setting consists of the six phases outlined below. In these phases, we observe how the insurgency indicator fluctuates over time during the period December 2004 and August 2005 and use that as a basis for prediction until February 2006. We explain the reasons behind these fluctuations based on citizen agents' well-being and the roles of the media and organizations in mobilizing them to rebel against the government.

A. Pre-tsunami: We calibrate our experimental scenario for Aceh, Indonesia, where there is a pre-existing active secessionist movement led by GAM and its leader Hasan Di Tiro. We insert Tsunami as an external shock to the system at the end of this phase.

TABLE 2 Named agents and agent classes

Citizen (473,500)	Named Leaders (9)	Named Organizations (9)	Media (9)	Sectors (14)
Javanese	Yusuf Kalla (VP)	Golkar	Jakarta Post	Oil
Acehnese	Megawati	PDI P	Indonesia Times	Gas
Sundanese	Sukarnoputri	PPP	Jakarta Times	Power
Batak	Hamzah Haz	GAM	Jaringan Islam	Education
Minangakabau	Husan di Tiro	NU	Liberal	Financial services
Banjarese	Hashim Mujadi	Muhamadiya	Sinar Harapan	Agriculture
Bantanese	Amien Rais	Jemmah Islamiyah	Voice of Islam	Water
Madurese	Abu Bakar Bashyir	MMI	Radio Republik	Manufacturing
Buginese	(Leader of MMI)	Democrat Party	Indonesia	Military industrial
Betawi	Susilo Yudhoyono (P)		Televisi Republik	Transportation
Chinese			Indonesia	Telecommunication
Malay				Government services
Other				Labor
				Capital

B. Immediate aftermath of tsunami: We insert our best approximation of response to the calamity by the local government and the international community.

C. Post-tsunami recovery.

D. Intermediate aftermath of tsunami.

E. Local government permits greater freedom to citizen and media while interdicting organizations.

F. Prediction of the outcome of government policies on insurgency indicator.

Figure 3 shows the insurgency indicator for Indonesia and three specific provinces, while Figures 4 and 5 illustrate the impacts of deprivation, changes in resource levels (flow of aid), the media, and organization mobilization on the indicator. Using these results, we discuss here the fluctuations of insurgency in Aceh.

Insurgency in Aceh rises sharply in phase B. This spike is explained by the acute deprivation following the Tsumani and increased political grievance resulting from the delay in relief from the government. As media attitude is mostly pro-rebel and anti-government, citizens perceive the government relief effort to be ineffective in meeting their needs and blame it for their hardship. Hence, citizens are mobilized by hostile organizations like GAM.

In phases C and D, flow of international aid leads to reduction in deprivation and a positive shift in media attitude (anti-government, centrist, and moderate right media conglomerates). Furthermore, opposition groups and leaders, such as PDI-P, GAM, Golkar,

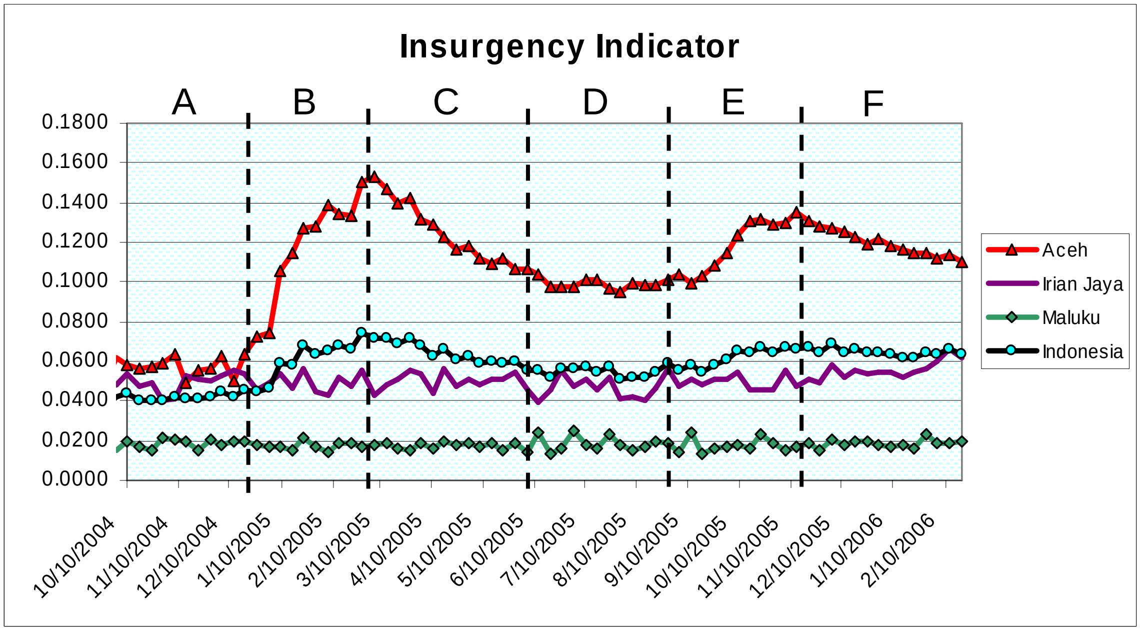

FIGURE 3 Insurgency indicator over time

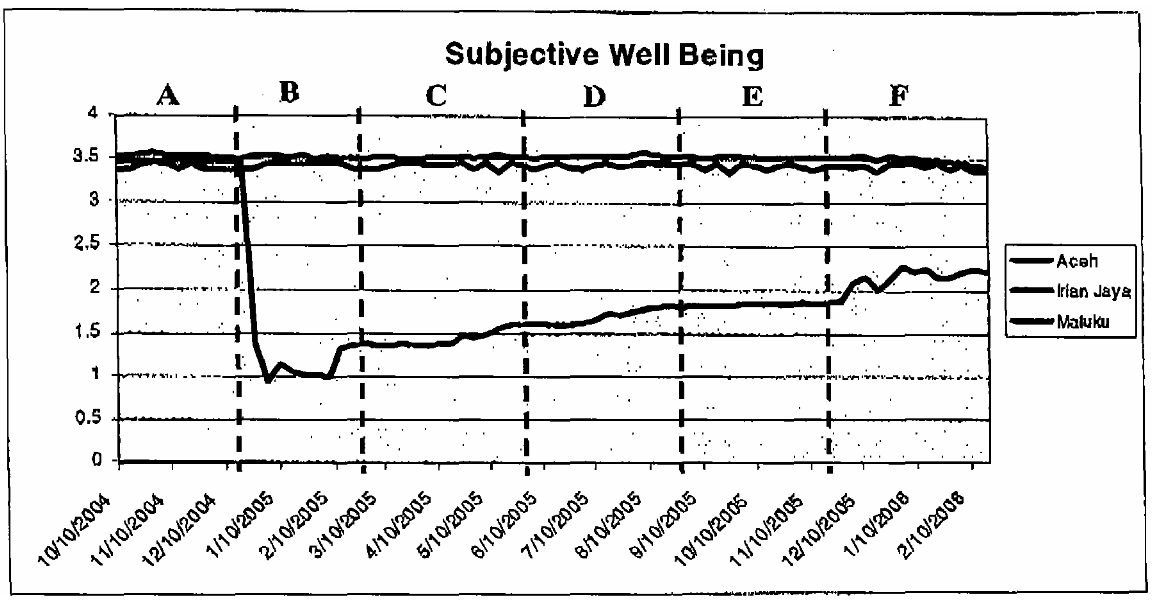

FIGURE 4 Impact of deprivation and grievance on well-being

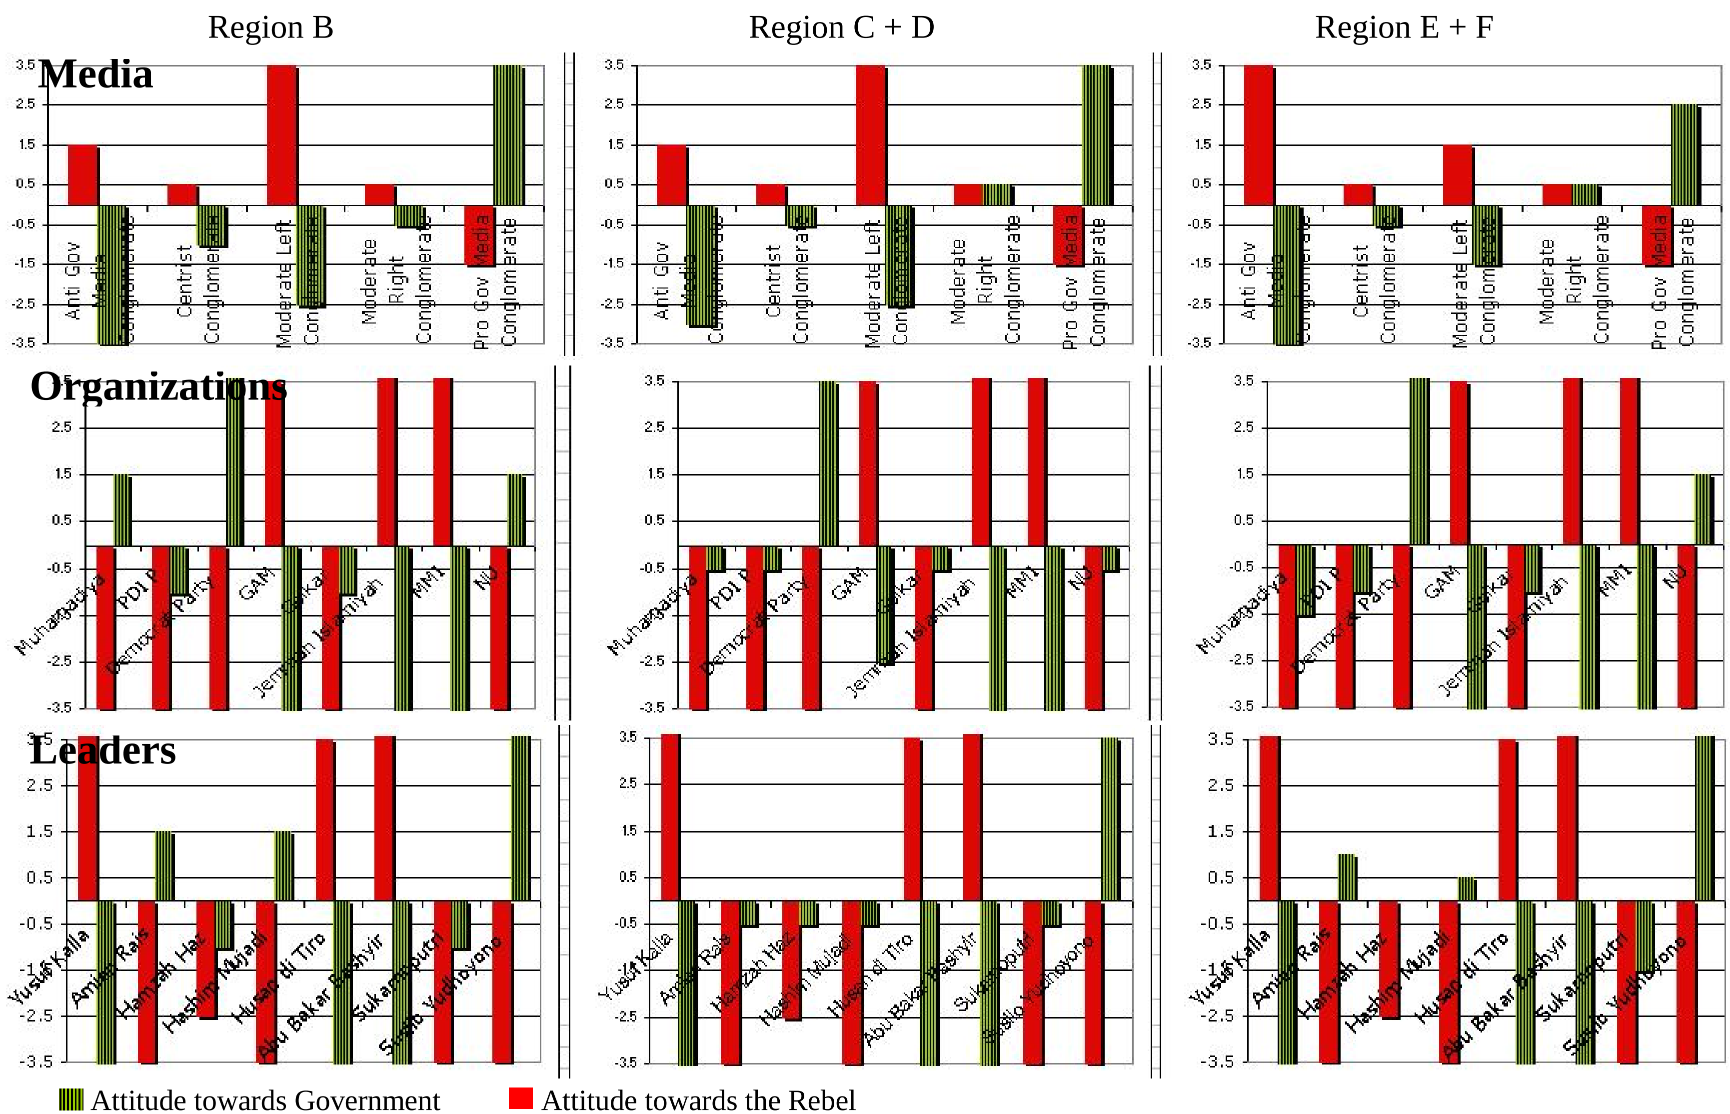

FIGURE 5 Role of media, organizations, and leaders in shaping citizen attitude and behavior, thereby determining the insurgency level

Hamzah Haz, and Sukarnoputri, soften their anti-government attitudes as they become more willing to cooperate. As a result, fewer citizens get mobilized, and insurgency is gradually mitigated.

In the predictive phase, with a government policy change, insurgency spikes initially (phase E) due to risk-acceptant citizens, media that is now free to favor the rebels (anti-government and pro-government media conglomerates), and organizations that respond negatively to being indicted (Muhammadiya, PDI-P, GAM, and Golkar). However, insurgency gets slightly alleviated in region F, primarily as a result of the continued aid and citizen freedom that increase aggregate well-being. In addition, the government has more support from leaders like Amien Rais, Hamzah Haz, and Hashim Mujadi, along with the organization NU.

CONCLUSION

In this paper, we present an agent-based simulation of intra-state conflict to understand the conditions that favor the emergence, duration, and intensity of insurgency. We present a Virtual International System developed in the Synthetic Environment for Analysis and Simulation (SEAS-VIS) to analyze insurgency in a strife-torn region of the world. SEAS-VIS provides an environment in which to conduct computational experimentation as a way to begin to understand the largely qualitative aspects of insurgency. We use theoretical models to configure SEAS-VIS agents, calibrate them from open-source data, and validate them against published real-world incidents. We model a total of 474,073 agents, with 473,500 citizen agents, 9 named leaders, 9 named organizations, 9 media organizations, 14 sectors, and 406 critical infrastructure nodes. We then use the validated SEAS-VIS computational experimentation environment to analyze dynamic interrelationships among grievances, level of resources, and organizational capacity to mobilize members toward social actions.

ACKNOWLEDGMENTS

This research is partially funded by grants from the National Science Foundation, Indiana State 21st Century Research and Technology Fund, and the U.S. Department of Defense.

REFERENCES

Abelman, N., and J. Lie, 1995, *Blue Dreams: Korean Americans and the Los Angeles Riots*, Cambridge, MA: Harvard University Press.

Addison, T., P. Le Billon, and M. Murshed, 2002, “Conflict in Africa: The Cost of Peaceful Behavior,” *Journal of African Economics* 11(3):365−386.

Axelrod, R., and R. Hammond, 2003, “The Evolution of Ethnocentric Behavior,” presented at the Midwest Political Science Convention, Chicago, IL.

Axelrod, R., 2003, “Advancing the Art of Simulation in the Social Sciences,” *Japanese Journal for Management Information System, Special Issue on Agent Based Modeling* 12:1−19.

Besancon, M., 2005, "Relative Resources: Inequality in Ethnic Wars, Revolutions, and Genocides," *Journal of Peace Research* 42(4):393−415.

Bhavani, R., and D. Backer, 2000, "Localized Ethnic Conflict and Genocide: Accounting for Differences in Rwanda and Burundi," *Journal of Conflict Resolution* 44:283−306, June.

Boudriga, N., and M. Obaidat, 2004, "Intelligent Agents on the Web: A Review," *IEEE*, 35−42.

Brass, P., 1997, *Theft of an Idol*: *Text and Context in the Representation of Collective Violence*, Princeton: Princeton University Press.

Buhaug, H., and S. Gates, 2002, "The Geography of Civil War," *Journal of Peace Research* 39(4):417−433.

Cederman, L.-E., 2004, "Articulating the Geo-cultural Logic of Nationalist Insurgency," presented at the Santa Fe Institute.

Cederman, L.-E., 2005, "Beyond Fractionalization: Mapping Ethnicity onto Nationalist Insurgencies," presented at Conference on Disaggregating the Study of Civil War and Transnational Violence, University of California Institute of Global Conflict and Cooperation, San Diego, CA.

Chaturvedi, A., D. Dolk, S. Mehta, and R. Ayer, 2005, "Artificial Labor Markets: An Experiment in Large Scale Agent Based Modeling," *European Journal of Operational Research.*

Collier, P., and A. Hoeffler, 2004, "Greed and Grievance in Civil War," *Oxford Economic Papers* 56(4):563−595.

Costa, L., and J. Matos, 2002, "Towards an Organizational Model of Attitude Change," *Computational and Mathematical Organization Theory* 8(4):315−335.

Diener, E., and F. Fujita, 1995, "Resources, Personal Strivings, and Subjective Well-being: A Nomothetic and Idiographic Approach," *Journal of Personality and Social Psychology* 68:926–935.

Diener, E., and R.E. Lucas, 1999, "Personality and Subjective Well-being," in D. Kahneman, E. Diener, and N. Schwartz (Eds.), *Well-being: The Foundations of Hedonic Psychology*, pp. 213–229, NY: Russell Sage Foundation.

Diener, E., and E.M. Suh, 1998, "Subjective Well-being and Age: An International Analysis," in K.W. Schaie and M.P. Lawton (Eds.), *Annual Review of Gerontology and Geriatrics: Focus on Emotion and Adult Development*, Vol. 17, pp. 304–324, NY: Springer.

Diener, E., E. Sandvik, L. Seidlitz, and M. Diener, 1993, "The Relationship between Income and Subjective Well-being: Relative or Absolute?" *Social Indicators Research* 28:195–223.

Diener, E., E.M. Suh, R.E. Lucas, and H.L. Smith, 1999, "Subjective Well-being: Three Decades of Progress," *Psychological Bulletin* 125:276–302.

Duckitt, J.C., I. Wagner, and I. Birum, 2002, "The Psychological Bases of Ideology and Prejudice: Testing a Dual Process Model," *Journal of Personality and Social Psychology* 83:75−93.

Epstein, J., 2002, "Modeling Civil Violence: An Agent-based Computational Approach," *Proceedings of the National Academy of Sciences of USA*, Vol. 99, No. 10, Suppl. 3.

Fearon, J., 2004, "Why Do Some Civil Wars Last So Much Longer Than Others?" *Journal of Peace Research* 41(3):275−301.

Fearon, J., and D. Laitin, 2003, "Ethnicity, Insurgency and Civil War," *The American Political Science Review* 97(1):75−90.

Gurr, T., 1970, *Why Men Rebel*, NJ: Princeton University Press.

Gurr, T., 2000, *People Versus States: Minorities at Risk in the New Century*, Washington, DC: U.S. Institute of Peace Press.

Humphreys, M., 2005, "Natural Resources, Conflict and Conflict: Uncovering the Mechanisms," *Journal of Conflict Resolution* 49(4):508−537.

Harmon, C.C., 2000, *Terrorism Today*, London: Frank Cass.

Jenkins, J.C., 1983, "Resource Mobilization Theory and the Study of Social Movements," *Annual Review of Sociology* 9:527−553.

Kahneman, D., E. Diener, and N. Schwarz, 1999, *Well-being: The Foundations of Hedonic Psychology*, NY: Russell Sage Foundation.

Klandermans, B., 1984, "Mobilization and Participation: Social Psychological Explanations of Resource Mobilization Theory," *American Sociological Review* 49:583−600.

Klandermans, B., and D. Oegema, 1987, "Potentials, Networks, Motivations, and Barriers: Steps toward Participation in Social Movements," *American Sociological Review* 5:519−531.

Krueger, A.B., and J. Maleckova, 2002, "Education, Poverty, Political Violence, and Terrorism: Is There a Connection?" Working Paper No. w9074, National Bureau of Economic Research; available at http://papers.nber.org/papers/w9074.

Le Billon, P., 2001, "Angola's Political Economy of War: Natural Resources and Armed Conflict," *Political Geography* 20(5):55−80.

Lichbach, M., 1998, *The Rebel's Dilemma*, Ann Arbor, MI: The University of Michigan Press.

Lucas, R.E., A.E. Clark, Y. Georgellis, and E. Diener, 2003, "Reexamining Adaptation and the Set Point Model of Happiness: Reactions to Changes in Marital Status," *Journal of Personality and Social Psychology* 84:527–539.

Lujala, P., N. Gleditsch, and E. Gilmore, 2005, "A Diamond Curse? Civil War and a Lootable Resource," *Journal of Conflict Resolution* 49(4):538−562.

Lustick, I., D. Miodownik, and R. Eidelson, 2004, "Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It?" *American Political Science Review* 98(2):209−229.

Lyubomirsky, S., K. Sheldon, and D. Schkade, 2005, "Pursuing Happiness: The Architecture of Sustainable Change," *Review of General Psychology* 9(2):111−131.

Masters, D., 2004, "Support and Nonsupport for Nationalist Rebellion: A Prospect Theory Approach," *Political Psychology* 25(5):703−726.

McCarthy, J.D., and M. Zald, 1976, "Resource Mobilization and Social Movement: A Partial Theory," *American Journal of Sociology* 82:1212−1241.

Post, J.M., E. Sprinzak, and L.M. Denny, 2003, "The Terrorists in Their Own Words: Interviews with Thirty-five Incarcerated Middle Eastern Terrorists," *Terrorism and Political Violence* 15:171−184

Regan, P., and D. Norton, 2005, "Greed, Grievance, and Mobilization in Civil Wars," *Journal of Conflict Resolution* 49(3):319−336.

Ross, M.L., 2004, "How Does Natural Resource Wealth Influence Civil War? Evidence from Thirteen Cases," *International Organization* 58(1):35−67.

Sawyer, K., 2003, Multiagent Systems and the Micro-Macro Link in Sociological Theory, *Sociological Methods & Research* 31:325−363.

Schmid, A., 1983, *Political Terrorism: A Research Guide to the Concepts, Theories, Databases, and Literature*, Amsterdam: North Holland.

Snow, D.A., and R. Benford, 1988, "Ideology, Frame Resonance, and Participant Mobilization," in B. Klandermans, H. Kreisi, and S. Tarrow (Eds.), *From Structure to Action: Comparing Movement Participation across Cultures*, *International Social Movement Research* 1:197−218, Greenwich, CT: JAI Press.

Snow, D., and R. Benford, 1992, "Master Frames and Cycles of Protest," in A.D. Morris and C. Mueller (Eds.), *Frontiers in Social Movement Theory*, pp. 133−155, New Haven, CT: Yale University Press.

Srbljinovic, A., D. Penzar, P. Rodik, and K. Kardov, 2003, "An Agent-based Model of Ethnic Mobilization," *Journal of Artificial Societies and Social Simulation* 6(1); available at http://jasss.soc.surrey.ac.uk/6/1/1.html.

Suh, E.M., E. Diener, and F. Fujita, 1996, "Events and Subjective Well-being: Only Recent Events Matter," *Journal of Personality and Social Psychology* 70:1091–1102.

Tarrow, S., 1994, *Power in Movement: Social Movements, Collective Action and Mass Politics in the Modern State*, Cambridge: Cambridge University Press.

Tilly, C., 1978, *From Mobilization to Revolution*, Reading, MA: Addison-Wesley.

Victoroff, J., 2005, "The Mind of the Terrorist: A Review and Critique of Psychological Approaches," *Journal of Conflict Resolution* 49(1):3−42.

Weinstein, J., 2005, "Resources and Information Problem in Rebel Recruitment," *Journal of Conflict Resolution* 49(4):598−624.

DATA SOURCES

48th UN Statistical Yearbook

2005 CIA World Factbook

World Trade Organization Statistical Database

BPS-Statistics Indonesia

www.worldpress.org

Human Development Report, 2004, United Nations Development Program (UNDP)

Annual Report on International Religious Freedom, 2004, U.S. Department of State

Public Opinion Survey: Indonesia 2005, International Foundation for Election Systems

ICG Asia Report No. 19, 2001, Communal Violence in Indonesia: Lessons from Kalimantan

The World Bank, 2004, Local Conflict in Indonesia: Incidence and Patterns, a World Bank publication of Social Development Notes in Conflict Prevention and Reconstruction

Asia Program Special Report No. 110, Piety and Pragmatism: Trends in Indonesian Islamic Politics

DISCUSSION

Computational Social Theory

(Evolutionary Models, Saturday, October 15, 2005, 3:15–4:45 p.m.)

Chair and Discussant: *Keven Ruby, The University of Chicago*

Keven Ruby: Hello, I'm Keven Ruby, and I will be trying to keep things on track. We'll begin this session with Saunders-Newton, Axtell, Roszman, and Frank talking about "Can Many 'Littles' Make a 'Much': One Approach for Transforming Underspecified Theory into Agency-oriented Rules and Behaviors."

Can Many "Littles" Make a "Much"? One Approach for Transforming Underspecified Theory into Agency-oriented Rules and Behaviors

Desmond Saunders-Newton: I'm with BAE Systems as well as University of Southern California. I'm joined by my colleague, Aaron Frank, who is also with BAE Systems. Our other collaborators in this particular effort are Rob Axtell, who's at Brookings, and Larry Roszman, who's also at BAE Systems.

Today, we want to focus less on some of the model results, and I'll tell you some reasons why. I want to think a lot about some of the interesting challenges we had with this particular project in terms of attempting to craft this very, very high level systems theory of cultural revolution into an agent-based system and the challenges that were associated with that. I'll tell you how we worked through this particular process and how we can actually think a little bit more effectively about finding a rigorous and more structured way to convert these oftentimes very elegant ways of thinking about the world into basically levels of agency.

[Presentation]

Saunders-Newton: This whole notion of formations, institution formations, and social coherence is interesting because you actually measure this and define the appropriate metrics. That was our goal for the six-month effort. As you can tell, six months is not a long time.

Unidentified Speaker: What do you mean when you say, "Game theory coming soon."

Saunders-Newton: Each of the first, second, and fourth building was actually instantiated in the six months. Game theory was coming soon. We're actually dealing with internal R&D now, but yes, that's correct.

[Presentation Continues]

Saunders-Newton: Thank you very much. Are there any questions?

Joanna Bryson: I have two short questions. One's very high level. You mentioned why a company like BAE might be interested in this area. Is that why BAE is doing this modeling, so they can sell planes only to stable countries? What's their interest?

Saunders-Newton: No, no, no. I'll give you a short answer and … decide … gives you more detail. For those who know, BAE Systems used to be British Aerospace, and once Maggie Thatcher released them, they started buying up crazy numbers of buildings. My group is actually called the Intelligence Innovation Division. We focus on the various intel communities, and we develop combinations of methodologies and advance computational tools to support their activities. I run the Social Computation and Complexity Directorate. So this is a business line for them, and we're working around that.

Bryson: Okay, great. Now I have a much more specific question. During your talk, you mentioned modeling just the elites. Of course, the problem then is having an emergent to lead and not wanting to get caught. At first, I was going to ask how you're going to have a mechanism for adding in new elites, but then I thought, well, maybe you should just have a pool of candidate elites. For example, the 50 governors, because who would have thought Bill Clinton would have come out of nowhere and become President of America? But it's not from nowhere because he was the governor of a state, right? But you don't normally model every state. And, again, Howard Dean from Vermont has an impact, but although many people have heard of him, that wouldn't necessarily be considered the top elites, right?

Saunders-Newton: Yes. Our initial efforts have been focused more on trying to capture the people we know now, and to the best that we know of, since we support a lot of the work of the intel community. They do track a lot of the emergent elites, so these people have some type of authority. We actually get to start earlier in the process than others may, but your point is well taken. If you were actually doing a totally open-source activity, it may be something of a challenge to capture all these things. You have to be open to it.

Ruby: Any other questions?

Charles Macal: Desmond, you suggested that your approach is based on a data-grounded approach relative to seemingly economic measure ….

Saunders-Newton: Data couple. I like data couple.

Macal: Okay. Well, there're data involved in it.

Saunders-Newton: Right.

Macal: I could see where economic, or some broader variables, would be available or could be measurable, but regarding the notion of the elites or the social aspects, what can you say about the data there, or how closely coupled is it potentially?

Saunders-Newton: In this particular community, ultimately you would never know. So the question is, how can you actually explore that space of possible data values? Can you ground the distributions so you can actually do something like a pliometric exploration of those values? Some of those you would just never know.

But there are also really cool social science methodologies, such as unintrusive observation, where people have used these kinds of proxy variations that capture things, much of which comes from the nationalistic inquiry community in terms of how they actually approach this work. That can be explored, and so that has not really enjoyed some of the discussions about using ethnographies as a driving, because some of those insights may be helpful to actually capture ways of distributing — finding what the right distributions are to fit with that.

Ultimately, though, your point's really good. What you want is something that's totally connected with the data, which came up often, occurred consistently with this client and past clients. I say, anyway, it's just black magic. This doesn't exist. But it doesn't mean that the model can't provide you with some insights on how to behave, and that's our general next push: what is the ultimate goal of the modeling process?

Macal: I have a comment. Regarding your last slide and your reference to J. Forrester, which you know I'm very sympathetic to. It was brought up earlier in Epstein's talk that some of the epidemiology models they're working with are at some level agent-based models, but at a larger scale they can be kicked up into differential equation models and make that transition. So these may be potentially complementary approaches.

Saunders-Newton: I think that's true. As a person who's actually stored … methods (most of my time was in methodologies), I actually think there are many roles: it's about finding the right use of the twos, that is, have twos where you can actually use two or three at the same time on the same problem, methodologically speaking.

Ruby: All right. Thank you very much.

Knowledge Swarms: Generating Emergent Social Structure in Dynamic Environments

Ruby: Next, we have Robert Reynolds with "Knowledge Swarms: Generating Emergent Social Structure in Dynamic Environments."

Reynolds: This work was done in conjunction with Bin Peng, who's just finished up her Ph.D. at Wayne State. I'm affiliated with the Computer Science Department. Xiangdong Che did some of the recent programming, and also the Museum of Anthropology, University of Michigan — that's where Kent Flannery is. I do a lot of work with Kent. So it's a great follow-up to this talk. In fact, if you have the CD-ROM of Agent 2003, I have an operationalization of his model at a higher level of agency; not at the detailed level, but basically at the level of a site, so it's a nice tie-in. Obviously, they read these papers. So my perspective is from the perspective of AI and knowledge systems, and I'm interested in the emergence of social intelligence and socially motivated problem solving.

[Presentation]

Ruby: Thank you very much. Do we have any questions?

Mengziao Zhu: Mengziao Zhu from University of Illinois, Urbana-Champaign. I really enjoyed your research and your talk, and I have read about your former research. From your presentation today, I have several questions for you.

First, do you use the homogenous agents in simulations? If you use the homogenous ones, what do you think is the mechanism for them to converge to the certain roles? Is it just randomly decided? And second, you mentioned knowledge models and knowledge resources in your presentation several times. Are they something similar to the strategies or the roles that are used in the system by the agents to behave or react to the environment and the evolving process of the knowledge sources? Are they something similar to the learning process in the AI area?

Reynolds: Yes, in fact, each of the knowledge sources is being adjusted in terms of learning. The interesting thing about homogenous agents — in fact, what we did is to effectively remove the agents' individual memories and put all of their experiences or a subset of their experiences into the belief space, and so they are completely stripped-down agents. So your question is great. Why would a role emerge in an agent that has no memory? Well, it's because the knowledge swarms have a memory, and that the size of, for example, the knowledge swarms on the roulette wheel is going to determine how frequently it's going to be used. For example, if you have one knowledge source that dominates everything else, then every time you spin the wheel for an agent, you'd get that knowledge source. So you'd get that one role coming back, even though the individual has no memory. It's an interesting way of looking at things, rather than starting at the level of the individual, starting at the level of their shared cognition and seeing how far you can go with that.

This takes us back to the notion of generative sufficiency. In fact, you can generate swarms at the individual level, even though individuals don't have any information about their past. It says something about the power of knowledge and the power of learning in directing the behavior of groups.

Ruby: I think we have time for one more question.

Ventkatesh Mysore: Ventkatesh Mysore from New York University. I have a question and a comment. Could you put your approach in the perspective of AI and mission learning techniques like descent and similar….

Reynolds: Well, the techniques that I use from traditional AI would be genetic algorithms — evolutionary computation. Certainly, we do inductive learning, and we have, but not in this example, used decision trees, and we use somatic networks. If it's available in AI, we can use it and represent it. Here, we go with some very simple inductive learning techniques and basically evolutionary computation.

Mysore: My comment is completely unrelated to the question. It has to do with an interesting example in biology that you might be able to exploit. It's called bacterial chemotaxis. If you have foot sores, the bacteria migrate toward the foot sores, so they have simple local computational routes that they use the chemical gradient to guide their motion toward the foot sores, and once in a while, they do a random tumbling to venture to a new direction. Maybe that is something you might be able to model.

Reynolds: Yes. Thanks. Sounds great, thank you.

Ruby: All right. Great.

Adversarial Analysis of Evolutionary Models and Multi-agent Systems (Toward Theoretical Foundations for Generative Social Sciences)

Ruby: Next, Gabriel Istrate will present "Robustness of Evolutionary Models and Multi-agent Simulations to Adversarial Scheduling."

Gabriel Istrate: Thank you. So, adversarial analysis of evolutionary models and multi-agent systems: I bet that at this hour this doesn't mean a lot to you. As we've seen, one of the main themes of the conference is the generative approach to social sciences. In the words of Josh, if you didn't grow it, you didn't explain it or it's more complicated.

[Presentation]

Ruby: Thank you very much. Are there any questions?

Bryson This isn't really a question. It's just a comment on your last slide. One of the things I really liked about Axtell's paper, "Why Agents?," is that he brought out the point that sometimes it's easier to work with agent-based modeling just because it's intuitive. And while it's probably worth getting some analysis at some point, I don't necessarily think it has to be the first approach. If you have multiple approaches, you should possibly go with the one you're most comfortable and competent at.

Istrate: I totally agree. As I said, I do that for a living.

Xinrong Lei: My name is Xinrong Lei. If there is a simulation, I don't know if the results are different because of the initial situation or because of the randomness of the schedule. In this case, according to your suggestion, should I first adjust the randomness of the schedule, or should I eliminate the influence of initial situation? Thanks.

Istrate: I think both are important. I don't know if there is any single answer to that. I wasn't trying to make the point that you should try to attempt what I did here. I just wanted to have a very modest point, which is, you wouldn't expect, given that you could talk about any kind of schedule, to have any sort of mathematical result, and to my surprise that was not the case.

Michael North: Following up on that point, I'm interested in what you did here and, particularly, how it might impact the agent modeling toolkits and things like that, because obviously scheduling is a big issue there. What do you think the prospects are for a more generalized form? The things you did are very interesting, but there's a broader question of developing or maintaining a toolkit. How would you adjust that so it could take these results into account?

Istrate: It was mentioned yesterday that there is an approach called model checking in the formal verification educational literature that basically specifies scenarios as temporologic formulas and schedulers as automata. One problem with that type of approach as applied to the agent literature is that it's almost like looking for contra-examples. Some recent literature has dealt with model checking for mark of chains and versions of it, and that seems more applicable to agent simulations. That's precisely the long-term goal that I was alluding to. Basically, that's kind of the root.

North: If there are other questions, I could follow up while we are waiting.

Macal: I've got questions.

North: Yes.

Macal: If I could just resummarize: the problem you're looking into is one in which the order of the choices of the agents is based on some strategic notion of interaction and a calculation on the part of the agents — "who goes next?" — based on some kind of Nash equilibrium or maximizing concept of utility, or at least an incremental improvement of utility for an agent, given its position. Some strategies could be adaptive in the sense that they're dependent on the state that the agents find themselves in. Others are not adaptive. They're independent of the state, maybe as a mark-off situation or whatever. So if that's a correct characterization, are you suggesting that agents can solve this complex selection problem, or in the instance of choosing which to move next, perhaps there's even a notion of bounded rationality that has to operate just to make the problem solvable given the agent perspective? Could you comment on that?

Istrate: Basically, I took apart one part of the dynamics and varied it. I tried to understand what is it in that part that makes the result hold for the random case and that could provide, for instance, if you want to talk about evolution, emergence of norms. Things like transmission via some sort of random walk mechanism are much more plausible than, "I'm just going to decide to update my strategy." So it was very modest in that sense, and I'm certainly not suggesting that agents could do that.

Macal: Okay.

North: Following up on the previous point, ideally, you'd want to find the one counterexample, the one bad event if you're presented with a scheduling problem like this. But even that could be, and I think is, extremely difficult given a complex problem. But is it possible, then, to consider general robustness, and so instead of one bad event, build this model and be faced with many, many bad events, or a large percentage of events being bad in some sense? Do you think that might weaken the problem and, therefore, make it more solid?

Istrate: I don't really have a good intuition about that, so I also started looking at what happens if I take one result and change the graph topology a bit. It's basically a beginning, an avenue, that could be investigated and not much more than that. I don't claim there is any sort of general guidance or ….

North: No, I think it's excellent work. I'm just interested in the implications

Understanding Insurgency by Using Agent-based Computational Experimentation: Case Study of Indonesia

Ruby: Okay. Now we have Alok Chaturvedi from Purdue University, who will talk about "Understanding Insurgency Using Agent-based Computational Experimentation."

Alok Chaturvedi: Thank you. I'll try to go very quickly. I'm the last speaker, so I won't hold you for long. I'm going to talk about a project that we have just started working on for the Joint Forces Command, and we will be using this for urban resolve, so we'll be supporting them in July.

[Presentation]

Unidentified Speaker: Are the data any good?

Chaturvedi: Once you have a common framework, you can do that. That's why I said, that continuous validation is the key. What we did was model 12 months before the tsunami, and then we continuously validated after the tsunami. The results are pretty amazing because if you — I'll come to that in a second.

[Presentation Continues]

Ruby: All right. Thank you very much. Are there questions?

Steven Wilcox: I noticed that you have hierarchal agents, and some of what you're talking about sounds like what is called generative practice theory by Gregg Courand and Michael Fehling. Do you know what that is, and do you see the resemblance?

Chaturvedi: We don't have a hierarchical agent *per se*. Okay? So many of those behaviors are emergent behaviors. It may look like we have a hierarchical agent, but we have population, and then it is all about who is communicating with whom. So what we have is a gigantic publish-and-subscribe architecture, if you may call it that, so that people are communicating. So we ask how does communication emerge as those things that are imposed by a certain structure in the society, so the organizations are formed, organizations are disbanded, institutions are formed, which are much, but, you know, like this is a fully integrated emergent type of society.

Ruby: The slides that you displayed, which I assume showed what the model does, were very complex. I'm wondering if you could, in more straightforward terms, talk briefly about the inputs that are creating or fueling the parameters of the model. What is the output, and what does this output mean? Is the output some sort of risk of insurgency or violence? And finally, what is it that your model does that one might not be able to do in following the media of, say, Indonesia in this given time period?

Chaturvedi: There are several things that we do, which models typically are going to do. Intervention is a big thing for us because we are working with the Department of Defense, and they want to take certain actions. It's not just Department of Defense, but it is interagency — coordination groups, as they call it — say the State Department or USAID. They're taking DIME actions, Diplomatic Information, Military and Economic. So those are the actions they're taking. The way we model is an end-sided game.

We've got what we call "glow," which is the coalition. We model red, which are the bad guys, so to speak. We've got green, which is the local government, and we have gray, which is our country axis — all the other countries in the region. All of them are interacting. We have a dynamic environment that is continuously running. They are all doing their own thing. The citizens are doing their own thing, and the organizations are doing their own thing.

So what do we observe? In this particular case, we're looking at propensity to do violence or to join the insurgency, or the intent to rebel, as we call it. But in different situations, we've got a full economic engine that is running, so it is generating all the economy, doing production, consumption and all those things — imports, exports — so you can look at that.

Now, if you're doing certain interventive strategies … let's say you have an economic sanction or you have a blockade and other things. It is going to reflect what is going on in the global economy. If there is an insurgency, one of the things that we are trying to model is what would happen if there is a major incident in Iraq. How is that communicated?

One of the critical things that we have is a story-telling model. We might have the same incidents observed by four different parties, and there are four different story lines, which we have. These story lines are going to their own networks, to their own subscribers, and then they are generating different types of behaviors and emotions. Again, that is going to feed back into the system. Now, you're going to observe what is going on in the political model, what is going on in the military model or economic or social and others so you can start observing. You can observe just about anything. You can look at what is going on in any infrastructure that was blown up, for example. What does it do? What happened yesterday, you know, the whole blackout? What impact did it have on society?

Ruby: Because you're pulling in data from the real world?

Chaturvedi: Yes.

Ruby: Or do you have an isolated universe that is somehow parallel to the ….

Chaturvedi: It is a parallel universe from the real world. We are building tools by which we are mining the real world, all the blogs and newspapers and others, and we put the data back into the system.

Saunders-Newton: I'll be so bold and presumptuous to expand upon Alok's response. Part of what drives this effort and a number of other efforts is this whole issue around effects-based operations. So the 7¢ story is basically about being consequent-aware. If you can actually think about the choice between using military force and using development dollars versus using diplomacy and using whatever of the national strengths that are available to you, can you model that to think about the possible consequences so you can think about second, third, or *n*'th order effects? That's the real issue about using simulations, so that's one of the outputs with a number of these efforts. Is it, and, again, Alok has this particular one, but there are a couple of others that are supporting some of these activities with this….

Reginald Tucker-Seeley: Reginald Tucker-Seeley from the Harvard School of Public Health, as well as the Dana-Farber Cancer Institute, so I've really enjoyed your cancer example. You're right; we're trying to go back and figure out what determinants yield the outcome that we're evaluating. It's a very difficult process.

You mentioned that the data sources were good in regard to validation. One thing that we note is that it's difficult to compare data from different sources across different scales, so could you possibly talk about how good the data are and the comparability of the data from several sources across different scales for your model?

Chaturvedi: One of the important things is how to have a consistent framework, so obviously we're looking at economics — political science, international relations. You have all these different areas, so people are describing things in very different ways. You have to come up with a common framework on how to model these things. That is one of the things we have done. And in going back to different skills, so you know, we have heard a lot of talk about ontology.

So let's say, there is a fire in this room. Over there is a door for all of us. Just consider that we are all agent simulations. For us, that door is for escaping, but the moment we open the door … say it is a totally different meaning for the fire model. The fire model has more oxygen, so the fire is getting more fuel to burn faster, right? It is all about how you translate the ontologies between different things, and especially, in this case, we may be running into the second time scale, whereas the fire model may be running into the millisecond or even the microsecond time scale.

So the thing is, we have another project, which we call simulation bridge, that is a shared reality engine. The way we operate over there is like here: we all are talking and listening, and we are sharing some reality with each other. Okay, so we're not really sharing whatever we have going on outside this room. That is how we model all these complex simulations so that we are only sharing the realities that we need to share and everything else is done offline. That is how we do the multi-scaling problems.

Larry Kuznar: Yes. It's an impressive compilation of many different kinds of information from many different parts of the world with different social systems, and I can appreciate that one of the problems you must have run into is having certain systems where it was hard, perhaps, for the base model to predict. Related to that is that a lot of calibration had to be involved in building the model and in adjusting it to these different databases. How long did that take? I mean, what was the scope of this project?

Chaturvedi: It's a huge problem. In fact, we have been working with the Joint Forces Command for over two years, almost three years now. We started very small — a very small region with very high level actions. Then we took one country or one city and expanded that and added more actions, more nodes. Right now, as you can see, we've got 12 million agents that we are running. I think, to the best of our knowledge, it is at least one or two orders of magnitude better than most of the things out there. The reason for that is that we approach this whole thing from the computer science perspective first. The first thing that we did was look at a good architecture to deal with this complex problem, where we can take multiple disciplines so that we'll be able to map things into our environment and solve some of the more critical scaling problem.

So once we did that, we started adding more and more models. We got a social science model and physics-based models. We are integrating with a lot of attrition models like JASAF and JWARS and others that are already attrition models out there. We are linking with two other

European models. One is Joanna and one is Eliance. One is a French system and one is the German system. But, again, I mean, you are absolutely right. For us, calibration is the key thing.

So we broke it down so that there are certain forward problems and there are certain inverse problems. Obviously, we said we are not going to try to solve those inverse problems, because, I mean, there would be a gazillion solutions to the same thing. Okay, so we said, “Let’s solve these forward problems and let’s calibrate at the inverse problems for our problem level.” So we modeled at the individual level, and we calibrated at the population level. So once you have bound those problems, it is a lot easier to, again, match the results from the real world. So that’s why continuous validation is our approach, and that is something that is working very well for us.

Ruby: All right. We actually have time for one more question.

Wilcox: What sort of mathematics or statistical methodology or economic methodology did you use to solve the inverse problems?

Chaturvedi: I am not saying we are solving the inverse problem. It is almost impossible to solve the inverse problem at this scale. Okay, it is absolutely impossible. It is intractable and there is no way trying — I mean, there is no reason trying to solve the inverse problem, okay?

So we are using statistics. We are using a whole bunch of different econometric matter to make sure that the synthetic data that we are getting and observe the data from the real world, and those are statistically — I mean, and solving all those T tests and other tests to show the results from there.

The other thing we are doing is that we are using bunches of different parametric and nonparametric methods to analyze the relationship going from actions to effects and from effects to actions because this is a pretty complex problem, and even making an attempt to solve the endless problem is futile. We can do that in smaller chunks, but not at the scale at which we are doing.

Ruby: All right. Thank you very much, and thank you to all the panelists this afternoon.

Macal: I’d like to thank Keven Ruby from The University of Chicago for chairing that session. Thank you, Keven.

As we close out this Agent 2005 session, I’d like to give a few thanks and acknowledgements. I’d like to thank the administrative people that are the AV recorders and the Program Committee of the Agent 2005, including David Sallach for computational social theory and Michael North for methods, toolkits, and more. I’d like to thank Michael North again for organizing and conducting the Repast training course earlier in the week.

Finally, I’d like to thank you all very much for your contributions, your active support, your attention, and your contributions to forwarding computational social science in the form of theory applications and toolkits and methods. We hope to see you next year.

LIST OF ATTENDEES

Alexandridis, Kostas
Purdue University

Alexeev, Alex
Indiana University

Anderson, Michael S.

Artel, Arsun
Illinois Institute of Technology

Bankes, Steven
Evolving Logic

Birkner, Irving
The University of Chicago

Boyd, Gale
Argonne National Laboratory

Bryson, Joanna
University of Bath

Buitrago, Corian

Bulleit, William
Michigan Tech

Calett, Dan

Carrie, Ana
Trinity College, University of Dublin

Cederman, Lars-Erik
Swiss Federal Institute of Technology at ETH

Chaturvedi, Alok
Purdue University

Chaturvedi, Rashmi
Simulex

Clemmons, Margaret
Argonne National Laboratory

Collier, Nick
Argonne National Laboratory and PantaRei Corp.

Davidson, Robbie
Argonne National Laboratory

Diermeier, Daniel
Northwestern University

Dunlop, Tina

Epstein, Joshua
The Brookings Institution

Fernandez, Luis
U.S. Environmental Protection Agency

Frank, Aaron
BAE Systems

Geller, Armando
Swiss Military Academy at ETH Zurich

Girardin, Luc
Swiss Federal Institute of Technology (ETH)

Griffin, William
Arizona State University

Gulyás, László
Loránd Eötvös University

Hoffer, Lee
Washington University School of Medicine

Howe, Tom
Argonne National Laboratory and The University of Chicago

Huang, He Qing

Istrate, Gabriel
Los Alamos National Laboratory

Jager, Wander
University of Groningen, the Netherlands

Johnston, Carl
George Mason University

Jung, Taehyun
Georgia Tech

Kennedy, Ryan
University of Notre Dame

Koehler, Matthew
The MITRE Corporation

Kuznar, Lawrence A.
Indiana Purdue University in Fort Wayne

Lauen, Doug
The University of Chicago

Lee, Doowan
The University of Chicago

Lei, Xinrong
University of Illinois at Chicago

Leveen, Sheen
Singapore Management University

Li, Zhian
Argonne National Laboratory

Liu, Chungshan

Lytinen, Steven
DePaul University

Macal, Charles
Argonne National Laboratory and
The University of Chicago

McBride, Mark
Miami University

Mellarkod, Veena S.
Texas Tech University and Argonne
National Laboratory

Mysore, Venkatesh
Courant Institute, New York University

North, Michael
Argonne National Laboratory and
The University of Chicago

Ozik, Jonathan
Argonne National Laboratory

Perk, Sinem
Illinois Institute of Technology

Perkins, Timothy

Rand, William
University of Michigan and Northwestern
University

Rewerts, Chris

Reynolds, Robert
Wayne State University

Ruby, Keven
The University of Chicago

Sallach, David
Argonne National Laboratory and
The University of Chicago

Saunders-Newton, Desmond
BAE Systems and the University of
Southern California

Schmidt, Shana
Arizona State University

Silbert, Neal

Snowdon, Anne
University of Windsor

Stephan, Craig
Ford Motor Company

Strout, Nathan

Suematsu, Yutaka Inti Leon
ATR International and Kyoto University

Sullivan, John
Ford Motor Company

Sydelko, Pam
Argonne National Laboratory

Tanzman, Ed
Argonne National Laboratory

Tatarynowicz, Adam

Tetiker, Mehmet Derya

Tucker-Seeley, Reginald
Harvard School of Public Health and Dana-Farber Cancer Institute

Upton, Steve
Referentia Systems, Inc.

Ussher, Leanne
Queens College, City University of New York

Vos, Jerry
Argonne National Laboratory and University of Illinois at Urbana-Champaign

Wander, Jager
University of Groningen

Wilcox, Steven
Northrop Grumman Corporation

Yin, Li
University of Buffalo

Younger, Stephen
Los Alamos National Laboratory and University of Hawaii

Zeidenberg, Matthew
University of Wisconsin

Zheng, Kangmin
Zhu, Mengxiao, University of Illinois at Champaign-Urbana