ANL/DIS/TM-60

Proceedings of the Workshop on

Simulation of Social Agents: Architectures and Institutions

October 6-7, 2000
The University of Chicago

Edited by David Sallach and Thomas Wolsko

Co-sponsored by

Social Science Research Computation
The University of Chicago

and

Decision and Information Sciences Division
Argonne National Laboratory

June 2001

ISBN 0-9679168-2-8

Available by request from

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/

CONTENTS

Cultural and Normative Dynamics

SATURDAY, OCTOBER 7, 2000

Politics

Environmental Processes

Energy and Infrastructures

Closing Panel

Foreword

Welcome to the *Proceedings* of the second annual workshop on the simulation of social agents, co-sponsored by Argonne National Laboratory (Decision and Information Sciences Division) and The University of Chicago (Social Science Research Computing).

The articles in this collection provide a view of the state-of-the-craft of social agent simulation in the early twenty-first century. They reveal that agent simulation is being applied to an impressive range of research topics — financial markets, environmental impact, housing segregation, international conflict, the causes of energy shortages, the dynamics of status symbols. Clearly, agent simulation provides a methodology through which a vast array of issues can be framed and addressed.

Nevertheless, simulation models do not and cannot emulate whole societies in all of their multifaceted dynamics. The processes represented within current simulations are selected to represent a particular research focus. In the years to come, the representation of such social entities and processes will be at the heart of an active dialogue between the substantive social sciences and the emerging epistemic community of agent simulation modelers. This is an important dialogue, and one to which we hope this workshop series will continue to contribute.

David Sallach, Director
Social Science Research Computing
The University of Chicago

Thomas Wolsko, Director
Decision and Information Sciences Division
Argonne National Laboratory

Acknowledgments

ORGANIZING COMMITTEE

Victor Lofgreen coordinated the logistics and managed the audiotaping of the workshop. Jewel Carter also helped with logistics. Kathy Ruffatto handled registration and administration, and Bob Baker managed the workshop web site. Jane Andrew assembled and copyedited the proceedings, including presentations, abstracts, papers, and discussions, with help from Curt A. Strating and others from Alanwood Enterprises, Inc. (for transcription of the discussions) and Linda Graf and others in Argonne's Document Processing Center (for word processing).

ORGANIZING COMMITTEE

David Sallach, University of Chicago (co-chair)
Thomas Wolsko, Argonne National Laboratory (co-chair)
Victor Lofgreen, University of Chicago
Charles Macal, Argonne National Laboratory
Michael North, Argonne National Laboratory
John Padgett, University of Chicago
Randall Picker, University of Chicago

Opening Presentation

COMPUTATIONAL SOCIAL SCIENCE: AGENTS, INTERACTION, AND DYNAMICS

K.M. CARLEY, Carnegie Mellon University*

ABSTRACT

Computational techniques are an important tool in the social theorists' and methodologists' toolkit. In particular, multi-agent computational models are particularly valuable for the development of social theory. They enable the researcher to examine the relations among groups and individual agents and to witness the way in which these relations enable, constrain, and affect agent and group behavior, as well as emergent phenomena. During the past 40 years, computational analysis more generally has played a role in revolutionizing social theory. Important advances in the area of cognition, interaction, chance, and adaptation have been made, leading to new paradigms such as information processing theory. Fundamental results related to bounded rationality, satisficing, competency traps, emergent order, and learning clashes have emerged. Today, multi-agent models are enabling social scientists to ask fundamental questions about the nature of coordination, mechanisms for facilitating or inhibiting change, and the effect of scale and technology on social behavior. Results and findings are illustrated using a variety of classical and current models.

[edited transcript of presentation follows]

INTRODUCTION

A couple of years ago I had the opportunity to go to a naval war game. Now, that may not sound like a really cool thing to most of you, but it was actually quite interesting. But the first thing that happened was that my portable computer got knocked off a desk onto the concrete and exploded all over the place. The military guy who was in charge turned to me and said, "Thank God. One less agent I have to get a security clearance for."

You may think that's a dumb remark, but it was true, in a sense. We were looking at new technologies in the military, and we had a lot of agent-based models. And a lot of them, of course, were stored on my machine. So it really did have to go through not a security clearance, exactly, but a virus check. One might begin to wonder about the extent to which agent models truly are social agents. That's the underlying theme of what I'm going to be talking about.

It has become increasingly clear that computational techniques in general are an extremely useful tool for doing social theorizing and for doing methodological development

* *Corresponding author address*: Kathleen M. Carley, Carnegie Mellon University, Porter Hall 219a, Pittsburgh, PA 15213; e-mail: kathleen.carley@cmu.edu.

[*slide 2*].[1] In particular, the new multi-agent techniques — and don't ask me what an agent is — are particularly valuable. They enable the researcher to do a lot of things. If you read any paper in this area, you'll see people talking about how agents allow you to examine relations among groups and individuals, let you witness the way relations enable or constrain individual behavior, and, of course, let you do emergent phenomena most readily.

Some of the key areas where they've been particularly useful in the social sciences are studies of the fundamental nature of coordination and communication, facilitating or inhibiting change among agents, the effect of scale (like moving from worlds of five to worlds of twenty to worlds of thousands), and the effect, of course, of information technology or telecommunications technology.

Now, over the past 40 years of computer models in this area, we've seen important advances in areas such as the cognitive sciences, interaction, chance, chaos, and adaptation [*slide 3*]. Fundamental results out there spring from ideas as simple as the old notion of bounded rationality, as well as notions of satisficing, competency traps, emergent order, learning clashes, and so on. The computational techniques have been responsible in large measure for the evolution of new paradigms, including the information processing paradigm and now, of course, the neo-information processing paradigm. The reason for this influence, many people will argue, is that the techniques allow you to ask new questions and to explore fundamental assumptions and to break the assumptions and try new ones. Next I'm going to go through some of the basic findings.

NATURE OF MODELS

The first thing I thought it might be instructive to look at is a set of models that are out there in the workplace [*slide 4*]. Clearly, this isn't all the models that exist in the field. It's not even close. But these models have appeared in multiple papers, and each of them has some claim to fame. Most are from the organizations area. What I've done here is take each of the models and show the extent to which they model various characteristics, like learning, multi-agency, features of tasks, features of organizational structure, resources, and so on.

Two things should be readily apparent in this table. First, a lot of attention has been given to multi-agent-type models. Many of these models have multiple agents in some form. The second thing, which is even more readily apparent, is that in addition to paying attention to individual agents, they've paid attention to resources, that is, the kinds of things the agents can manipulate, play with, change, adapt, alter, etc. They've paid a fair amount of attention to knowledge, but not as much, as you can see by the fact that not many deal with learning.

These models, plus the countless models I haven't put up there, have been used in a variety of ways. The list I've got here is not exhaustive [*slide 5*]. Nor are these strict alternatives in a logical sense. But these are some of the modes in which people are using computational models and have used them in the past 40 years.

[1] Slide numbers refer to the presentation slides that are reproduced beginning on page 15. Slides 10, 16, 20, and 23-28 are not referenced, as they have been omitted from this publication.

Models as theory. The first mode is probably the most profound and the most irritating to noncomputational people, and that's the notion of the model as a theory. The argument dates back to early computer science. It holds that the model does the task it seeks to explain; therefore, it is a theory. As such, one needs to test it in new ways. An example coming out of that approach is the social Turing test that Newell and I worked on, which extends the Turing tests to a set of agents and says, "Well, if you're using the simulation model in an experiment and it generates socially realistic behavior, then it meets the social Turing test." There's a little more to it, but that's about it.

Models as agents. In the second mode, models are basically used as agents. Ever since the very beginning of computational modeling techniques in the social sciences, we've had agent-like models. We've had models running around doing things, like playing in tournaments. We've had them taking part in experiments, substituting for humans in the lab, and so on. Not always were these terribly sophisticated, but models have played the role of individual agents in individual ways.

Models as virtual worlds. Many models create virtual worlds. The most familiar example to a lot of people is the A-Life work, particularly the work on Sugarscape at Brookings. The idea here is that you're doing social theorizing from the ground up: you're building worlds out of tons of these little agents interacting, talking to each other, communicating, and building emerging social phenomena out of their interactions.

Models as empirically grounded theories. Another mode follows a tradition that's extremely strong within the social sciences compared to other areas where computational techniques are used. In this mode, one builds empirical data into the simulation model, not only as parameters, but also as input data and so on, right as an integral part of the model. An example of that is the work by David Heise and his group on affect control theory and their model of social exchange as a result of affective changes in personality. Built into the model are the massive data sets collected from multiple cultures from multiple countries on the way individuals respond in various emotional settings, the emotional rating of various words, roles, positions, and situations, etc. They use that data to predict how these things will behave over time.

The thing I love about this approach is that a lot of the data embedded in simulation models today is ethnographic in nature. That is, if I were going to build model of a particular process at GM, the first thing I'd want is a very detailed ethnographic study of that group, because that would give me the details I need to build a simulation model.

Models as hypothesis generators. The models that you have nowadays are often so large that their responsiveness — that is, their performance outcomes relative to all the different kinds of input data — is just really large, and you can't always analyze the whole thing. So you run virtual experiments — like a human laboratory experiment, but with a computer — and statistically analyze the results. The end result is a set of hypotheses that can then be tested in other settings or compared against other models.

One of the values of that strategy is that it allows you to generate hypotheses from your basic theory in a way that is more systematic than verbal theorizing. In fact, when some researchers have translated existing verbal theories into simulations, they've found that the predictions and the hypotheses coming out of the verbal theories were actually incompatible with the fundamental assumption. It could not be regenerated. So it's a useful technique.

COMPUTATIONAL ANALYSIS

Multi-agent models. The computational approach a lot of people are using is to build multi-agent models. In these models, the agents can be models of humans or of artificial agents, like web-bots or robots or avatars or databases, and these agents tend to be diverse [*slide 6*]. That is, you may have multiple types of agents in a model, or they may be diverse simply by virtue of having access to different knowledge or, in the case of Sugarscape, sugar, or whatever. So you've got agent heterogeneity, which is really important, because a lot of the outcomes that are derived are really the result of having a set of heterogeneous rather than homogeneous agents.

Socio-information processing models. Social action is often derived through interaction among the agents, not by individual outcomes. And a lot of the models — in fact, I would say the vast bulk of them, particularly in organizations — are socio-information processing models. That is, the agents are information processors with the ability to collect, analyze, distribute, communicate, or do whatever they want with this information, but they're also "socio" in the sense that they usually have a little bit of a model of the other [agents].

Because they're socio-information processing agents, they have a set of constraints that really dictate behavior. So you're really talking about constrained behavior. The oldest constraint, of course, is cognitive constraints on the way you think, analyze, process data, etc., but the other ones are, of course, social constraints (who you talk to, who's in your network) and technological constraints (what technology you have access to for communication, for travel, for moving).

Emergent behavior. Out of these models you get emergent behavior. You can get it at the individual, group, organization, or population level or at multiple levels at once. You can also talk about the co-evolution of those systems. Emergent behavior is often studied by detecting patterns, and one of the difficulties in this area is doing pattern detection. That is, the ideas that people have and the models they're generating at this point are outstripping our ability methodologically to extract patterns, especially patterns that are dynamic over time, particularly for large-scale groups.

Equilibrium. Another feature of these models is that you rarely run them at equilibrium. You're usually not interested in questions of equilibrium; you're concerned with questions of dynamics and of change. In fact, a lot of the modelers right now will argue there's no such thing as an equilibrium; that's just a false assumption. These models also often tend to be empirically grounded at some level. At least they're based on empirical findings. The parameters may be input from other data sets, and so on. One of the implications is that you need multi-level validation and multiple types of validation.

CRITICAL FINDINGS IN FORTY YEARS OF RESEARCH

With that as background on the field in general, what I want to go through now is a series of critical findings from the past 40 years in this area. I think that over this time period a lot of really nice and really important results have come out, but they've come out of this department over here and that discipline over there. But if we put them all together, they're really suggesting a whole new way of thinking about human action and social behavior [*slide 7*]. So computational analysis has generated this entirely new way of thinking about it. The work that's been done in this area is strongly interdisciplinary. Many of the critical findings are the result of teams of people working together, anywhere from four to 30 people in a team, and, as I said, from many disciplines.

Cognitive Constraints

One of the oldest findings is that, in terms of cognition, there are constraints, and they really do matter [*slide 8*]. Now, in the old paradigm — because we're talking about the 1950s, basically — the view of the world is that you have rational actors. But look at one of the first models, the "garbage can" model of Cohen, March, and Olson. They said, "Look, these actors aren't rational." Human beings are not that way. They're boundedly rational at best. They have limits on their ability to process information, and moreover, what they process is in part affected by what they're interested in. A similar argument is made in the behavioral theory of the firm, which is a much more detailed, elaborate model by Cyert and March. And in those two models alone, they were able to demonstrate quite conclusively that if you use boundedly rational actors, the fundamental results that you might get from assuming rational actors just go by the wayside, and behavior looks totally different.

So, for example, the results might show that in organizations most decisions, rather than being made by finding the right decision, are actually made by oversight, by accident, because you happen to be in the right place at the right time — and that data actually reflects the world. In other words, they showed that cognitive and social limitations really affect the social outcomes.

Architecture

The second finding, which is a further refinement of the first, was the idea that in terms of cognition, the actual architecture matters [*slide 9*]. The term "architecture" from cognitive science refers to the actual hard wiring in your brain for how you process and handle information. If you're talking about a computer program, it's the actual code for how it handles information.

In terms of architecture, it's not that there was one paradigm; there were a whole bunch of views out there. There's the strong structural position, coming out of sociology: you have networks; they really matter; if you replace the human beings with rocks or rats, it doesn't matter — you'll get the same effect. Then there's the rational actor view that says agent differences just don't matter; they're not relevant. And there's the decision-making view that says outcomes are based on decision processes and the information gathering has no impact. But in fact the simulation work in this area said, "Well, that's not true."

What our models are suggesting is that agent variation exists, and the variation comes at two levels. First, it comes in terms of information processing capabilities. That is, if I build a model with certain capabilities, with a certain way of handling information, and then I build another one, they may generate fundamentally different results. Experiments with humans and rats show that the same situation will generate fundamentally different results because the cognitive architecture matters. The second reason why it matters is that agents are heterogeneous not just in terms of their cognitive architecture but also in terms of the knowledge that they have, and differences in knowledge also matter.

Technology

The findings about cognitive architecture opened a door to another finding, which is that we can treat any kind of telecommunication technology. Any information processing technology itself adds an agent with information processing capabilities that are distinct from those of humans. When you put humans and these artificial agents together, you will get distinctive differences in the outcomes because the exact way they are processing the information is different. Here again, it's not just that constraints matter, as we just saw, but now it's that the *exact* constraints totally affect the outcomes.

Another way of looking at this is to say that agents don't have to be people. Here's an example. We're looking at information diffusion, namely, the time for information to diffuse relative to how professional a group is. The group is more professional if it has more information than the average group or average society *and* if that information is select and special to that group. So there is a set of knowledge that they know and that is more or less exclusive to their group.

A good example of a professional society would be physics. [At first] you have a world with just human beings who can communicate information. [Then you look at] what happens when you add additional agents, which are their web pages, and how that affects the diffusion of information. [We see] that it takes the longest for information to diffuse (the highest point) when the new idea originates in a very professional group and it's going to a very unprofessional group. Think of this as physics going to the general public; it takes a long time. But when you introduce web pages, the whole surface changes. And now where did information move the fastest? Physics to the general public. Some people say this is a good explanation for what happened with the cold fusion story.

Interactions

So we've talked about cognition. Another huge area where there has been a lot of research is interaction [*slide 11*]. The earliest work in this area was associated with games. In particular, it was associated with the Prisoner's Dilemma game. The basic issue was the type of social intelligence needed to get groups to cooperate. You all know of Axelrod's tournaments, where people came in and played the Prisoner's Dilemma, and they sent in a computer program. Much to his and everyone else's surprise, a simple tit-for-tat strategy won in the first tournament.

That was very important because it showed the value of the reciprocity norm. That result has since inspired a lot of research, including work showing that interesting social orders, like cooperation, can emerge even with zero-intelligence actors; that the order in which the individuals interact matters to the outcome; that you can generate different behavior if you allow people to alter the choices.

This one particular finding, that interaction matters, led to a huge range of other findings in this area of interaction. But the fundamental idea is that through simulation, through these agent models, we were able to see, in fact, that interaction becomes a critical determinant of social outcome. So it is not just cognition; it is interaction, too.

Networks

More recently, people have been working on networks. The argument is that networks matter [*slide 12*]. I'm sure you have all heard about the "small world" phenomenon. The basic idea there is that we have a set of people who can interact. Think of the old rumor game. They are communicating information, and let's say the red guy there thinks up the piece of information, starts the rumor. It goes around the world and comes back; in this case, it takes about six links, because it has to go from person to person to person. One of the ideas that came out of Watts and Strogatz and others is simply that judiciously placed links between people, random links even, that cross these boundaries will speed up the flow of information. Lots of people have showed results like this.

The more fundamental thing that came out of this whole line of work is that the exact pattern of the network matters: that by changing what the fundamental underlying network looks like, you can dramatically change the rate of information flow, who gets what information, and the kind of social outcomes you get. And this is just an example of a change speeding things up.

So first we said cognition matters. Then we found out that it's the exact nature of cognition that matters. We found from simulation that interaction matters; then we found that the exact pattern of interaction matters.

Chaos

Chaos has become a buzzword in organization theory [*slide 13*]. In the old paradigm, chaos is the natural state of the world, and the order is imposed. To get order, you have to have power, and you have to exercise that power. For example, the managers of an organization might dictate its culture. The new view of the world coming out of the simulation work is that order emerges naturally, and it emerges under a couple of very simple assumptions. One is that the agents have some cognizance of each other, that they have very simple models of each other. A lot of people — Kephart, Padgett, and many others — have work that speaks to this basic finding. The idea again is that knowledge of others enables realistic outcomes, and in particular it enables order. By "order" we're talking about everything from distribution of size of firm effects, to distribution of organization effects, to distribution of number of people in your networks, to overall patterns of coordination, and so on. To get these kind of ordering effects, your knowledge has to include a little bit of knowledge of others.

Chance and Paths

Another result coming out of this work is that chance, or the path you are taking, matters [*slide 14*]. The old paradigm included a variety of views. People would argue that your starting conditions should be irrelevant; that chance is irrelevant; that there is a single end state, and so on. Some of the work in simulation, however, suggested that this is just not the case. It suggests, in contrast, that most social systems are complex, that the systems are nonlinear, and hat therefore they exhibit path-dependency effects — so the order in which things happens, the starting conditions, and minor variations in starting conditions can have wildly different outcomes.

I like to think of this as a lost continent of research, because this is an old finding. It really dates back to some of the early systems work. Cyberneticists were saying this a long time ago. It was forgotten for a while, but now it's the big news again. If we put this result in the social context we are talking about, what it means is that it matters *who* you know *when*, it matters *what* you know *when*, it matters in what *order* you learn it. And those continue to matter over time. So it's not just having knowledge of others; the order in which you get that knowledge also matters.

Adaptation and Learning

The next finding is in the area of adaptation [*slide 15*]. Here a finding coming out of simulation is that if you're going to get adaptation, which we define as the ability to maintain or improve performance, you really need to avoid learning clashes. So far, we have been talking about agents that know stuff and are interacting, but we haven't really talked much about their ability to learn. There are many old paradigms that talk about organizational learning. For example, people talked about such things as, "Well, learning is the average of these experiences of the individuals in the company." That's experiential learning: how much you do something is important. In this case, turnover would be important, because since it's the average, it matters who's there when you calculate it. Or there is the argument coming out of population ecology that "organizations don't learn, they evolve." Or there is the argument coming out of formalization work, that learning is really imbedded in the storage procedures, databases, etc., within the company.

The newer view says, yes, all of those are right. It is not either/or; they are all right. In fact, we now have this ecology of learning mechanisms going on in the social world. There are many kinds of learning: experiential, structural, expectation-based, and so on. The psychologists say, "We've been telling you that for the past three decades." What that means from a social organizational process is that because there are many kinds of learning and because the learning is occurring not just at the individual level but at the synthetic agent or group level, those different kinds of learning can clash with each other. If they do clash, you won't get adaptation. So adaptations in part result from the absence of clashes.

You also need meta learning to respond to change processes, such as innovation, new technologies, and so on, and in many cases that means you need to have meta learning to learn how to trade off between exploration (that is, group or structural-level learning) and exploitation (that is, individual learning). A variety of models exist in this area, like the work by March and Leventhal on the exploration/exploitation model, our work on ORGAHEAD, and so on. But the basic finding from all of these is that learning clashes can impede adaptation.

Here is an example that shows how this effect came out in some simulations. We have a response surface showing a hypothetical organizational performance as a function of the organizational design. For simplicity, the only elements of the design I'm looking at here are size and density. The organization learns strategically and figures out what it wants to do and then decides what to do with it. This strategic learning moves the organization through this response surface by causing it to change its size and density by hiring and firing people, teaching them to talk to each other, putting them in change management seminars, whatever. The only thing is, the individuals themselves are all the while going through experiential learning, which has an S form for most human beings for most problems.

Now, let's imagine a particular organization that wants to improve its performance. How does it do that? Well, it decreases its size and tries to get people to talk together more, so it increases density, maybe by doing more team projects. Their actual performance, however, may not be this high. It may be much lower, because this is only the maximum possible performance. But why? Well, by making those changes — downsizing, getting rid of people, getting people to interact more, making them talk to each other — what are they doing? They are getting rid of the lessons of experience. They are obviating the value of old learning, and so their whole performance as a group is going to degrade.

We also have the performance over time of set of simulated organizations. Two things you will note. First, most of them are improving, on average. Second, the occupancy of first place changes sporadically. Now and then one of them accidentally really plummets. So if you explore why they plummeted, you will find in every one of those cases that it was because of a learning clash between one or more types of learning.

Transactive Memory

The next finding is in the area of transactive memory [*slide 17*]. This actually came out of computer science directly. Vegner and others said, "Look, in computer systems, a lot of the intelligence relies on the referential information and the ability of the computers to communicate with each other. It lies in the connections." He built a computer model of this and said, "Hey, maybe this is also true of human beings." Since then, there have been a number of experimental studies and field studies in transactive memory that have shown that in fact transactive memory is important, and it does improve performance in small groups, teams, organizations, and so on.

Transactive memory is your knowledge about who knows *who* and who knows *what*. The idea is that if you have models in which each of your agents not only has knowledge of others, but also has knowledge of who others know and what others know, they actually do better and the performance of the entire group improves. This is actually a very valuable thing for the social scientists to have found out, because it is a very important way of tying a historically important idea called social capital into our understanding of organizational performance.

Systems

The next finding is that that there are strong interactions between cognition and between interaction [*slide 18*]. In the old paradigm, you could talk about strict structuralism — that rocks and rats don't make a difference in how a network behaves, or there was psychological individualism — that is, individuals operate psychologically and don't care about what anyone else is doing. The new arguments are different. They say, "Look, not only is there an interaction between structure and people and cognition, there's actually more: what's really happening is that we have kind of soft boundaries between agency and structure."

To explain this idea, you have to understand the notion of the synthetic agent, which is an agent that's created out of other agents. So, for example, a group is a synthetic agent because it includes the individuals within it, as well as their knowledge of each other. Groups have the ability to learn, and they inherit all the properties of individuals in that sense. When you have synthetic agents along with human agents in a group, the boundary between what is an agent and what is a structure changes somewhat. As another example, consider that you as an agent have a certain set of knowledge at your disposal and a certain way in which you forget it and a certain way in which you use it. Let's imagine that you as an agent are busy working for Ford and you are painting those little racing stripes on the side of the cars. So the tools at your disposal include paint and the paintbrushes, and the car is completely distinct from you, the paintbrush is completely distinct from you, and there is a clear, physical boundary between you and everything else that's going on.

Now, however, Ford says, “Robots are going to be much better at painting these lines straight than human beings.” So they replace the agent, the human being, with a new agent, which is the robot. But the brush and the paint can are built into the robot. So part of the old way — picking up the brush, dipping it in the paint — is gone. The boundary has been muted between agents. That’s what we mean by mutable boundaries. The whole agent world is filled with discussions of mutable boundaries and how they can affect the outcome of group performance.

Some of the models in this area, like the e-commerce models, IBIZA, multi-agent models, or Soar docking with ELM, have shown that when you look at the interaction between a structure and agents, you see a huge interaction that totally influences the outcome. So you have to take both of them into account simultaneously.

One example comes out of some earlier work where we used Soar models and ELM (experiential learning model) agents and humans. We let all of them do a radar task where they had to determine for a bunch of planes going through space whether the planes were hostile or not. It turns out that if the agents are on a team or if they are in a hierarchy, the type of agent that performs best — the human, the Soar agent, the ELM agent — totally depends on both the task and the cognition. It is a complex nonlinear interaction between them.

SYNTHESIS AND NEW DIRECTIONS

So let’s put all these things together. The findings from the past 40 years include that cognition really matters and that these limitations, the constraints on cognition and social action, affect the outcomes [*slide 19*]. Moreover, we found that the exact pattern of the information processing capabilities matters. It matters if your agents have interrupt schemes or not, for example. Interaction is critical to determining social outcomes, but it is not just interaction in general; it is the exact pattern in which those interactions occur. Knowledge about others enables more realistic outcomes and allows social order to emerge, and the order in which agents get that knowledge is also critical. The knowledge of others also has to include your knowledge of whom others interact with and what those others know in order to get effective performance or to get realistically human behavior out of these systems. Finally, when you move up to the group level, you get learning clashes that can impede adaptation and you get an interaction between cognition and structure. In a sense, that is a lot for 40 years, but what it really suggests is that we are currently seeing a new paradigm evolving in this area.

Now, let me illustrate what is happening because of this new paradigm. One thing is that people are now building new agent-based models that not only have increasingly smart and increasingly humanlike and increasingly cool agents, but these agents are put into networks. In fact, one of the big issues right now in robotics is how to build and think about and manage the networks, the communications patterns, between our robotic agents. If you go to any of the multi-agent conferences, they will talk about this.

Nonhuman Networks

The key thing here is that [social] networks are not exclusively among people [*slide 21*]. The network paradigm has permeated the area in terms of thinking about networks connecting agents to agents, but also networks connecting people to knowledge or agents to knowledge (the knowledge network), agents to tasks (the assignment network), agents to organizations (the work network), and so on. We have networks at every level connecting agents, knowledge, tasks, organizations, and so on. And that network-based approach is allowing modelers in this area to begin to say, “Oh, here’s how my modeling hooks up to yours, because we are both representing our data in the same way.” So there is an increasing amount of data sharing and model sharing.

Dynamic Networks

One difficulty, and the thing I haven’t talked too much about so far, is that these little networks that we are talking about are dynamic [*slide 22*]. They start off in one way and then over time they keep changing, and eventually what you end up with is a totally new network. What you’ve got is a set of agents — little smiley-faced blue guys. They’re in two different groups, with some interaction between them. Each of them has some knowledge. (I didn’t show tasks.) And they are interactive. When they interact, they learn from each other, that is they communicate and send new information, which leads to the growth of new information at those purple explosions, and new people come into the company. And of course the groups reformulate and change. So, for example, this guy who was connected over here learns new things that makes him more like this group, so eventually he might move over to this side, taking his knowledge with him.

One of the big issues in this area is group emergence; that is, how do I know when a group emerged, how can I track it, how can I understand it, how can I see it? That’s an unsolved problem. Simon has been telling me for years that we need to solve this. Yes, it’s true. It’s not solved.

Technology and Networks

Another problem in this area is dealing with technology. I told you we can treat technology as an agent. That’s fine. However, the big issue is not just treating technology as an agent; it’s the fact that when we talk about technology, we are talking about new information technology giving new access to more people, more information, more of the time. So the issue is how to represent this situation in such a way that we can look at matters of scale. This question is particularly important if you’re trying to model, say, the whole Internet, for example. So people are starting to use other techniques for looking at really large-scale agent networks.

Here’s another example of a result from this area. This is a transactive memory model that looks at organizational performance over time when you’ve got worlds that have just people or people plus databases or people plus avatars in the company. One of the things you see happening is that when you include databases in these systems, they slowly decrease the amount of know-who, whereas if you have avatars, people know more about who’s out there, and avatars also decrease what people know. More importantly, the other thing that technology does is change the fundamental way in which information is shared, the extent of that sharing, and the extent to which people interact. So, for example, if you put a Lotus Notes database system in your company, one of the big impacts will be that it will actually reduce the overall level of shared knowledge in the company, delay people interacting with each other, and lead to less interaction overall. These changes may not be that great from the perspective of living in the company, but they may have no impact on corporate performance.

Applications for Enhanced Networks

I want to close by extending that point in the organization setting and telling you about a real-world application. For this application, we went out and collected network data from a real company — it happened to be the Navy. We plugged all this into the system and used it to predict things about organizational adaptation.

One of the first predictions was that to be a really successful team, what they should do is spend more time retuning or retasking people and engaging in redesign tasks — that is, changing who is reporting to whom and who is doing what. If they spend time hiring and firing their personnel, they're probably going to be low performers. That was a generic finding coming out of this work. We then went and collected data on an actual platoon who were playing this one particular war game, plugged it into the model, and predicted who would be their emergent leaders. We got that exactly right, because we predicted that their emergent leaders would be those who had a high workload and who preferred to shed tasks rather than accept new personnel, which, of course, they did.

But we noticed something very interesting in that experience. In their after-action report the participants would talk to each other about why they did what they did. We just asked them, "Well, do you know what each other is doing?" Let's say I was the guy who was trying to shed a task. I'm Bravo. I've got a high workload, I'm trying to give you my task. You're not necessarily accepting it. So you're trying to do it, but you're not accepting it. So afterwards, we go to the other guys, Alpha and Charlie, and say, "Well, what do you think Bravo was trying to do?" They say, "I don't know. He was just sloughing off. I have no idea why he was asking me to go and do this task." No idea. They have no common view of the world.

Then we went to the admiral and said, "We understand that you're running the simulation for us, but could you have these guys tell each other what they're doing?" So the next time, they began by telling each other what they were doing. The next time the guy who was the emergent leader was trying to shed tasks, other people accepted them, because they knew what he was trying to do. So we thought, "This is really interesting." It fit our model, but we also saw something that we didn't know would happen, namely, that there seemed to be a trade-off between having a common view of the world — a common operational picture, a common understanding of what others were doing — and the performance of the group. We designed these groups to have no communication and to have optimal performance, and they worked. They were high-performance teams, they had no communication, they were great. They also had no view of what the others were doing.

So then we went back to the simulation and said, "Will this observation that we just had come out of the model?" The results of simulations that were done after the fact showed that the factors that led to high performance were dramatically different than those that led to having a common operational picture or having high adaptivity. My point about this is that the new work in this area is integrating data much more heavily. There's a real synergy in the new models because of the network perspective in going back and forth between real data and simulations.

This new approach is being used in a whole bunch of areas; these are just some of them [*slide 29*]. Each of you has many, many others. But the point I want to make here is that in this area there are applications both at the research level and at the totally applied level. I mean, how much more applied can you get than knowledge management, human resources policy evaluation, and looking at bioterrorism — bioterrorism isn't up there; that's a new project. And of course there are many universities, as you know, where this is being done. This is just some of them with big programs in the area, and for the graduate students here, I'll be glad to share this list with you.

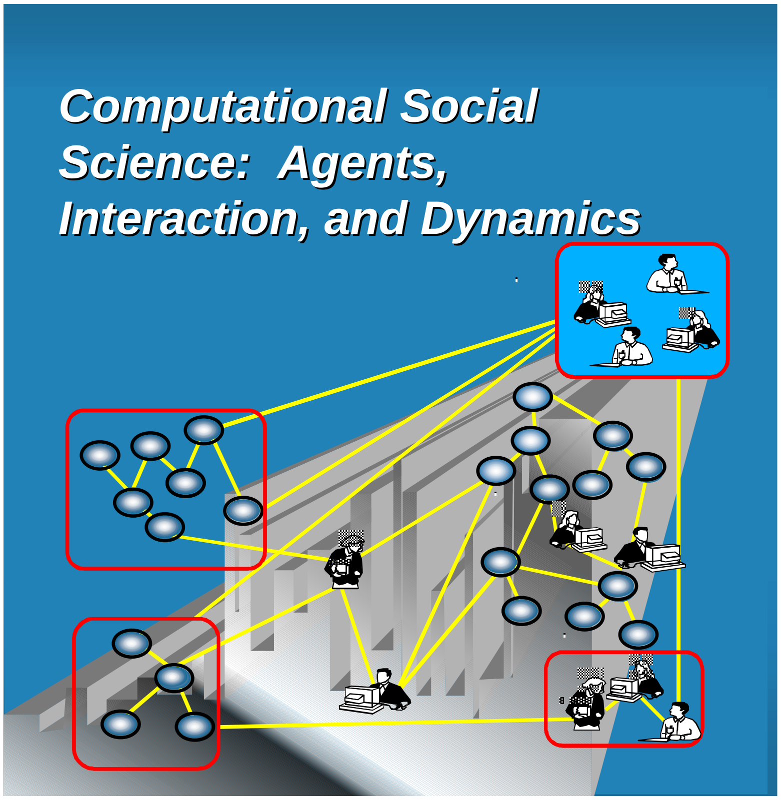

Carnegie Mellon

CASOS

Center for Computational Analysis of Social and Organizational Systems

Kathleen M. Carley

Director & Professor
Carnegie Mellon University
Porter Hall 219a
Pittsburgh, PA 15213
Tel: 1-412-268-3225
Fax: 1-412-268-6938
Email:
kathleen.carley@cmu.edu

2000

Computational Techniques

- Important tool for social theorists and methodologists.
- Multi-agent models particularly valuable.
- Enable the researcher to:
  - Examine the relations among groups and individual agents.
  - Witness the way in which these relations enable, constrain, and affect agent and group behavior.
  - Examine emergent phenomena.
- Key areas:
  - Nature of coordination.
  - Facilitating or inhibiting change.
  - Effect of scale.
  - Effect of technology.

2

The following slides have been omitted: slides 10, 16, 20, and 23-28.

40 Years of Computational Theory

- Computational analysis aided in revolutionizing social theory.
- Important advances in the areas of:
  - Cognition.
  - Interaction.
  - Chance.
  - Adaptation.
- Fundamental results:
  - Bounded rationality.
  - Satisficing.
  - Competency traps.
  - Emergent order.
  - Learning clashes.
- New paradigms such as information processing theory.

Question existing paradigms
Explore non-standard assumptions

Illustrative Models: Black-High, Star-Medium, and White-Low

	Agent		Task			Structure		Resource	
	Learning	*Multi-Agent*	*Types*	*Reassignment*	*Rework*	*Types*	*Dynamic*	*Kind*	*Amount*
VDT	❍	◗	❍	◗	●	❍	❍	◗	●
TAEMS	❍	●	●	❍	❍	◗	❍	●	●
STEAM	◗	●	◗	❍	❍	❍	◗	◗	●
COMIT	●	◗	◗	◗	❍	❍	❍	●	●
TacAir-Soar	◗	❍	❍	❍	❍	❍	❍	❍	❍
Orga-head	●	◗	❍	●	❍	●	●	◗	●
Garbage Can	❍	◗	❍	❍	❍	◗	❍	❍	●
NK	❍	◗	❍	❍	❍	❍	❍	❍	◗
Construct	◗	◗	❍	❍	❍	◗	●	◗	●
ELM	◗	◗	❍	●	❍	●	●	◗	●

The Many Faces of Models

- Models as theory
  - The model does the task it seeks to explain
  - Social Turing test
- Models as agents
  - Tournament participants
  - Substituting for humans in the lab
- Models as virtual world
  - Social theory from the ground up
  - A-life
- Models as empirically grounded theory
  - Empirical data as integral part of model
  - May be ethnographic
- Models as hypothesis generators
  - Response surface analysis
  - Systematic generation of hypotheses

5

Computational Analysis Approach

- Multi-agent models
  - Agents are humans or artificial
  - Collections of diverse agents
  - Heterogeneity of agents
  - Social action derived through interaction
- Socio-information processing approach
  - Social linkages, constraints
  - Cognitive capabilities, constraints
  - Technological changes, constraints
- Emergent behavior
  - Individual, group, organization, population level
  - Patterns
  - May not be equilibrium
- Empirically grounded
  - Models based on empirical findings
  - Parameters, mechanisms, input may be from other data sets
  - Multi-level validation

6

The Nature of the Social Agent

- Fundamental question
- Computational analysis generated many of the puzzle pieces
- Interdisciplinary
- Teams of scientists

7

Cognition: Constraints Matter

- One of earliest social science computational models
- Old paradigm - *rational actors*
- New paradigm - ***Boundedly rational actors***
- Models:
  - A behavioral theory of the firm - Cyert and March
  - Garbage can model - Cohen, March and Olsen

Cognitive and social limitations affect outcomes

8

Cognition: Architecture Matters

- Alternative paradigms - agents are agents
  - Strong structural position
  - Rational actor - agent differences not relevant
  - Decision making - outcomes based on decision process, not information process
- New paradigm - agent variation
  - Agents are heterogeneous
  - Technology as agent
  - Information processing capabilities matter

Exact information processing capabilities affect outcomes

Reciprocity - Interaction Matters

- Game - Prisoner's Dilemma
- Issue - what type of social intelligence is needed to get groups to cooperate?
- Tournament - Axelrod
- Outcome
  - Tit-for-tat
  - Reciprocity norm

0 intelligence
Order effects
Alternate norms
Alternate choices

Interaction is a critical determinant of social outcomes

Networks Matter

Watts and Strogatz
Young
Markosky
Many others

Small world phenomena

The exact pattern of the network matters

Chaos: Knowledge of Others Matters

- View - Leviathan - chaos as the natural state
- Old paradigm - order is imposed
  - Requires power
  - Requires the exercise of power
  - E.g., organizations - managers dictate culture
- New paradigm - **order emerges**
  - Agents need to be cognizant of others
  - Agents need to have models of others

Kephart
Padgett
Many others

Knowledge of others enables realistic outcomes

Chance: Path Matters

- Old paradigm
  - Starting conditions are irrelevant
  - Chance is irrelevant
  - Single end state
- New paradigm
  - Complex systems
  - Social systems as non-linear
  - Path dependence
  - Minor variations in starting conditions can have wildly different outcomes

The Lost Continent
Forrester
Sterman
Many many many others

Who you know, what you know, what you learn when is critical

14

Adaptation: Maintain or Improve Performance

- Old paradigms - organization learning?
  - Average of agents: turnover, experiential
  - Organizations don't learn - they evolve
  - Formalization - procedures, databases, etc.
- New paradigm - learning ecologies
  - Many kinds of learning
  - Change processes - innovation, technology, legislation
  - Adaptation
    - Absence of clashes
    - Meta learning
    - Trading off exploitation and exploration
- Models
  - EE - March, Levinthal
  - ORGAHEAD - Carley

Learning clashes can impede adaptation

15

Transactive Memory

- Idea:
  - In computer system, intelligence lies in part in connections and referential data
  - Maybe this is true of human cognition
- Transactive memory
  - Your knowledge of
    - Who knows who
    - Who knows what
- Models
  - TM model - Wegner
  - Construct-o - Carley

Who knows who knows what and who knows who affects outcome, TM improves performance

17

Interactions: Systems Matter

- Old paradigm
  - Strict structuralism
  - Psychological individualism
- New paradigm
  - Soft boundaries between agency and structure
  - Synthetic agents
- Models
  - E-commerce, IBIZA - Krishnan
  - Multi-agent soar docking with ELM – Carley, Lin, Prietula

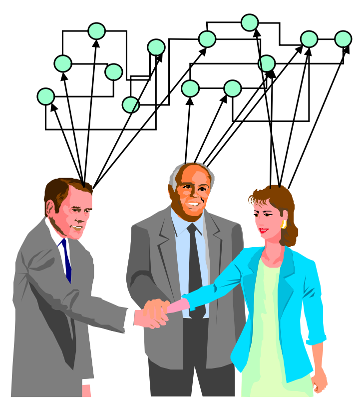

Interaction between cognition and structure influences social outcomes

18

Putting the Pieces Together

- Cognitive and social limitations affect outcomes
- Exact information processing capabilities affect outcomes
- Interaction is a critical determinant of social outcomes
- The exact pattern of the network matters
- Knowledge of others enables realistic outcomes
- Who you know, what you know, what you learn when is critical
- Who knows who knows what and who knows who affects outcome, TM improves performance
- Learning clashes can impede adaptation
- Interaction between cognition and structure influences social outcomes

New Paradigm

19

A Meta-network Representation of Organizations

	Agents	Knowledge	Tasks	Organizations
Agents	**Interaction Network** *Who knows who*	**Knowledge Network** *Who knows what*	**Assignment Network** *Who does which tasks*	**Work Network** *Who works where*
Knowledge		**Information Network** *What informs what*	**Needs Network** *What is needed to do which task*	**Competency Network** *What knowledge is where*
Tasks			**Precedence Network** *Which task precedes which*	**Market Network** *What tasks are done where*
Organizations				**Inter-Organizational Network** *Which organizations interact*

21

Dynamic Networks

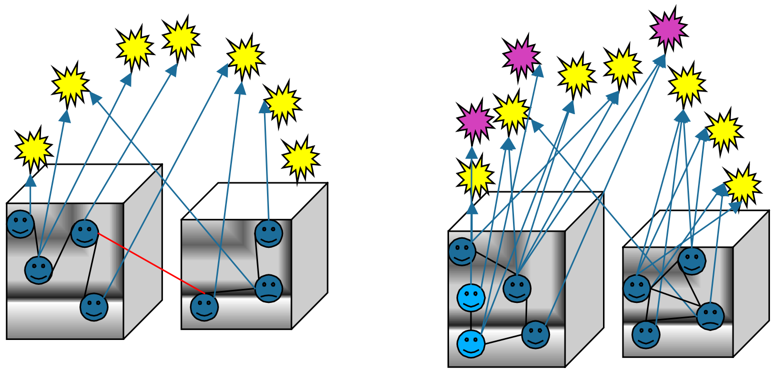

Application Areas

- Team design
- Organizational architecture evaluation
- Team flexibility and robustness analysis
- Tech transfer schemes
- HR policy evaluation
- Knowledge management
- Interaction of training and design
- Vulnerability analysis and information security
- Impact of IT
- Social network evolution and illicit drug/information transmission
- Risk management vs. information diffusion vs. early disease/virus detection

Academic Interest

- Universities and educational/research units forming programs
  - Stanford
  - Michigan
  - MIT
  - University of Texas - Austin
  - University of Texas - Dallas
  - John Hopkins
  - Cornell
  - Duke
  - University of Rochester
  - University of Arizona
  - UCLA
  - Penn State
  - University of Chicago
  - University of Illinois Urbana Champaign
  - Santa Fe Institute
  - England
    - University of Manchester
    - University of Surrey
    - London School of Economics
  - The Netherlands
    - Groningen
    - University of Amsterdam
  - Italy
    - University Trento
    - University Degli
    - University of Bologna
  - France
    - Insead
  - Japan
    - Kyoto University
    - University of Tsukuba

DISCUSSION: OPENING PRESENTATION*

T. WOLSKO, Argonne National Laboratory, Moderator

Michael North: I'm Michael North from Argonne National Laboratory. One of the things we've been talking about here that I think is very interesting is the idea of increasing specificity, in the sense that people are going from a general notion of finding an optimal solution to realizing that the optimal, if it exists at all, depends very much on the exact situation. I think the Lotus Notes example is particularly relevant to what we're talking about. Clearly, in a lot of cases, as you're seeing, it may make things worse. But if you have certain types of organizations — ones with high turnover, for instance — Lotus Notes may actually help. Could you elaborate on some of the things that you've seen change in people's thinking about moving from a general optimum to very contact-specific optimum?

Kathleen Carley: One of the main things I've seen happening is that a lot of the models today are much larger, much more detailed and specific than earlier models. That makes it difficult to publish these models in the journals; there are just too many pages of codes. So we have some issues in the field about how to archive these models — how can we share them? And that's led to a whole new research area on what is the right infrastructure for model sharing.

Because these details really matter and because there are so many findings out now showing that as you change, say, the cognitive architecture or the interaction, you change the results, you see people paying more attention to trying to dock or compare their models one against another. We don't have something like meta-analysis in psychology, yet there is more of this kind of discussion. In some sense, it's like this area is out of its infancy but is still in its adolescence.

North: Yes, that's one of the things I was going to note, too, in terms of repeatability. That's a hallmark of the scientific method. But how do we try to set up systems that are repeatable, given that not only are the problems so specific — you can slightly vary things and get a new problem space — but also that all the tools we're using are so chaotic and so complex in how they depend on those initial conditions.

Carley: Well, one of the answers that a lot of universities are looking at is open-source code. And there's another approach, which says that it doesn't really matter if your models are identical. We do a type of meta-analysis, and we get the same basic finding from a set of models, and that finding is very robust.

So that would be another possibility. The third thing, though, is some new work going on by logicians who are actually trying, on an experimental level, to show that, given an experiment that's completed — a virtual experiment on the computer — can you then back-prove that it occurred logically from its underlying assumptions? That's very state-of-the-art.

* [Editor's note: The discussion sessions were recorded with the speakers' knowledge and transcribed. The transcripts were edited for continuity and ease of reading; every effort was made to identify speakers and interpret comments accurately.]

North: Then this experimentation just seeds ideas to a certain extent.

Claudio Cioffi-Revilla: Claudio Cioffi from the University of Colorado. Could you elaborate on the notion you mentioned that "how groups emerge is an unsolved problem." Surely you meant there may be new *solution*s to how groups emerge? There's a game theory of how groups emerge; there's a stochastic theory. What did you mean?

Carley: I'm glad you asked that. What I'm talking about is not the process by which groups emerge. We have lots of theories, lots of work on that topic. The problem is, in an agent-based world, I want to be able to see a set of agents and say, "Aha, that's a conglomerate; there is a group." Some work has been done on recognizing when a set of agents is a group using clustering techniques, divisionary techniques, and so on. But if a mechanism is not predefined, what there isn't right now are good automated techniques for saying, "That agent belongs to this group and not that group."

It's not the *process* of groups evolving, it's recognizing that what I have is a group. A lot of the best work right now still involves drawing up a diagram and saying, "Yes, I can circle it; that's a group." Or you recognize a group, then backtrack, analyze it statistically, and say, "Okay, it kind of fits with my cluster analysis, it kind of feels right, so it's a group." So it's coming from that level, not from theory.

Jonathan Bendor: Jon Bendor of Stanford University. Going back to your remarks about the importance of starting positions and the anything-can-happen point of view. I'd like to generate a discussion about that, because I come from a somewhat different research tradition. In the research tradition that I belong to, assuming anything can happen is generally a weakness of a research program, because it reduces the empirical content of theories. If anything can happen, you look at a particular phenomenon and ask, "Why did this happen?" and you answer, "Well, give me the right starting position and I'll give you that outcome." Then the theory might well be vacuous — it doesn't exclude anything from happening. And a theory that doesn't exclude anything from happening is unfalsifiable.

It seems to me that in quite a few of these theories, however, if you work at them a little, you can show that there will generally be a unique limiting distribution to them. That doesn't mean that you reach a static kind of equilibrium — the individual agents or the individual populations of agents don't eventually stop changing. What happens, though, is that you reach a probabilistic equilibrium, in the sense that a Markov chain, for example, generally reaches a unique limiting distribution — and that's a lot of the empirical content of the theory. So you're saying, "Probabilistically, we know that it's going to go to a steady state, and we're going to try to characterize some of the properties of that steady state." Along the way, of course, the path matters. It's going to look different if you start here versus somewhere else, but in the long run you are going to reach a probabilistic steady state.

Carley: You're overinterpreting what I said. The issue is not that given your different starting positions you can end up anywhere, it's that given different starting positions you may end up in different ending conditions. Most of these models don't have the feature that you can end up anywhere. They do have sets of outcomes at some point. But you will get two very different ending conditions — not all possible endings, but different ones, depending on where you start.

Bendor: Well, what is meant by an ending condition?

Carley: Take one of our organizational adaptation models. If you run this model for several time periods, you see a behavior pattern where you get a probabilistic average of always hitting the same areas — just like what you were talking about. But there will be two different probabilities. If you look at the data, they will bifurcate. One set of organizations will end up with a high average performance level, another set will end up with a low average performance level — and where they started will affect where they end up.

North: In the work we're doing, end states and these long-run issues do matter to some extent, but we're actually much more concerned with the *process*, the path. The fact that these distributions may exist doesn't make a whole lot of difference to us. We care about the path of the market through the space. We also feel that in many cases the long-run equilibrium, or statistical equilibrium, will never really occur, because by the time you get that far out the whole environment has changed. So for us, the existence of these distributions is nice, but they may not actually matter.

Bendor: Two comments on that. First, there are results on the rate of convergence, and often it is surprisingly fast. And back to your first point, if you know what the long-run steady-state distribution of the system is, that gives you some information about the path.

North: The problem is that just knowing the starting and ending points is not enough. The path could do anything in between. A great example is the military.

Bendor: Anything with high probability?

North: We're trying to say that low-probability events can make a big difference here. A good example is a large military force fighting a small military force — there have been studies that indicate the large military force can lose. It's low probability, but it matters to the loser.

John Bower: John Bower from London Business School. I wonder if you'd comment a bit more on the techniques that people are starting to use to analyze the output from agent models. I get the strong feeling, in fact a slightly queasy feeling, that we often try to convince people about the output and the quality of the models by saying, "It quacks like a duck, so it's definitely a duck."

Carley: If you're in artificial intelligence, that's exactly the test that's used, because you only need to show proof of concept. But let me back up a second. A whole variety of new techniques are being developed for analyzing output from agent models. Some of these are standard statistical techniques, but there are also new techniques for network-based models. A new set of techniques derived from time-series analysis lets you look at complex data over time. There's also work coming out of engineering; many people use Fourier transforms, for example. But it is still the case for much of the work out there — and this is something we have a lot of trouble with in the journals — that people will run a model once and present their results as definitive, when in fact the model has stochastic elements that have to be demonstrated over hundreds of runs.

The other factor is that we have very large models, with huge numbers of agents in them. They actually break most standard statistical packages. If you're looking at networks of agents, the models will break all the network statistical packages. And if you want to visualize the results, they'll break all the visualization packages. One of the big research areas being funded by NSF right now is in designing analytical software that can handle these large models.

Robert Reynolds: Bob Reynolds, Wayne State University. I noticed that in your summary of results you didn't specifically address agent organization and problem-solving, though you talked about the infrastructure in which agents interact. For a particular problem, or an environment that poses a sequence of problems, how do the agents adapt?

Carley: There's some work in this area, but not a lot that has led to coherent findings. Some findings have shown that agent structures can adapt to a task environment. There's work that shows that when agents are structured in a way that's perfect for the task, they will perform better. We know how to design teams for tasks now; that's a pretty resolved area at this point. And we know about methods of adaptation. But in terms of what types of adaptations are needed for what types of tasks — other than some findings about teams, hierarchies, and volatility of tasks — we're still learning.

Artificial Markets

FINANCIAL MARKET EFFICIENCY IN A COEVOLUTIONARY ENVIRONMENT

B. LeBARON, Brandeis University*

ABSTRACT

What does the evolutionary interaction between different types of investors look like? This paper discusses research trying to understand the evolutionary dynamics between agents using differing lengths of past data to make decisions on portfolio choice. Computer simulations of a simple agent-based model show that agents taking a long perspective on past data have a difficult time dominating shorter perspective agents. The resulting dynamics replicate many features of actual markets. Furthermore, strategies become more homogeneous near sharp price declines, suggesting a liquidity-based explanation for market crashes and excess volatility.

INTRODUCTION

Traditional equilibrium models for financial markets often rely on the process of evolution, either explicitly or implicitly. These models generally assume that asset prices are the result of market participants holding common rational beliefs about the behavior of economic variables in the future and acting in a well-prescribed fashion on these beliefs. These economic worlds are simple, tractable, and generally at odds with the empirical evidence. Observed financial markets often appear too volatile and too predictable to be explained as the outcome of well-learned rational beliefs and strategies. Furthermore, the existence of large trading volume in financial markets adds important questions concerning heterogeneity across participants, since if all people agreed on asset valuations, trade would be unnecessary in most situations. This paper uses an agent-based model of a simple financial market to explore the evolutionary aspects of market dynamics, with the goal of understanding the barriers to market efficiency that cannot be eliminated through evolution and learning alone.

Most modern analysis of financial markets includes two crucial assumptions: markets are populated with rational agents, and they are in some kind of stationary equilibrium. Together these two assumptions yield tractable, testable restrictions for well-crafted theories. The first assumption can be relaxed somewhat in that irrational players may appear at times, but their suboptimal strategies will be driven out of the market.[1] Weakening the second assumption causes

* *Corresponding author address:* Blake LeBaron, Graduate School of International Economics and Finance, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453-2728; e-mail: blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research and an external faculty member at the Santa Fe Institute.

[1] Friedman (1953) provides the most often cited arguments for an evolutionary foundation for assuming rational behavior. In his case, for flexible exchange rates, he specifically comments on how less-than-rational speculators will be driven out of the market. Recent research on noise trading (DeLong, Shleifer, Summers & Waldmann 1991) and evolution (Blume & Easley 1990) has begun to suggest some flaws in the evolutionary argument for rationality in a financial setting.

more difficulty, since it is closely linked to the rationality assumption. Out of equilibrium, it becomes difficult to judge rational versus irrational strategies, as the economic landscape is in a continuing state of change. The biological term for this is coevolution. Strategies evolve against this current set of strategies in the population and not some well-defined fixed fitness norm.[2] In such a world, rationality can only be judged relative to the current population and not some well-defined fixed target that players should want to attain. It would be convenient to argue that these out-of-equilibrium dynamics can be ignored. However, this leaves open the critical question of how markets reach equilibrium in the first place.

In order to analyze out-of-equilibrium dynamics as well as convergence properties, markets will be populated with boundedly rational learning agents.[3] These are relatively simple agents, trading and learning about price dynamics as they go along. This facilitates the analysis of overall market dynamics and convergence properties when agents are faced with the same situations seen by ordinary people. It is important to realize that boundedly rational does not necessarily translate into stupid agents. They are often faced with situations where being completely rational may be computationally intractable, involving the beliefs of all the other market participants along with their dynamic decision-making processes. The only option available is to follow simple rules of thumb that are empirically tested and adjusted over time through learning.

This paper focuses specifically on one type of rationality, the appropriate use of past information. If financial time series were completely stationary, then more historical data would always yield better investment decisions. However, in practice, it appears that many market participants chose to ignore some past information to focus on the present. Recent arguments about a "new economy" are a good example of this. This behavior might indeed be rational if markets have changed, yielding the historical data irrelevant and giving those following it the survival chances of dinosaurs.

An evolutionary struggle between short- and long-horizon investors will be explored in an attempt to assess when and if the long-horizon types will evolutionarily dominate the market. The setting will be one with a completely stationary dividend process. In such a world, it might seem obvious that the long-horizon investors should dominate, but this is not necessarily the case. Prices move endogenously according to the traders' strategies and can even move in such a way as to enhance the strategies of short-term traders. It is also important to realize that in such a market, it is not clear what is a rational or irrational strategy without the guidance of a market safely in equilibrium. Given a turbulent market of short-run investors, it may be individually rational to become one of them, as opposed to taking the more difficult path of sticking to a long-run perspective.

---

[2] A good biological example is to think about evolution against predators versus evolution against climate. In the later case, one is probably safe in assuming a fixed fitness landscape, but in the former, this landscape is an ever-changing target.

[3] The concept of boundedly rational agents was introduced by Simon (1969). Recent applications in macroeconomics are summarized in Sargent (1993) and Sargent (1999). Often the argument for bounded rationality rests on actual bounds on computing power in the brain. However, bounded rationality might be argued for in terms of robustness. In a complex financial world, all strategies will be incomplete in some aspect, so it can be argued that simpler strategies may do better in terms of avoiding some really big mistakes. This is a little like arguing that you don't want to get "too smart for your own good."

MARKET DESCRIPTION

The market simulations used here are part of the class of economic models referred to as "agent based." Models of this type consist of large numbers of interacting agents, each acting independently of the others, often with active learning and adaptation.[4] Agent-based markets share many features: many interacting individuals, evolutionary dynamics, learning, and bounded rationality. However, the key distinguishing feature is that heterogeneity itself is endogenous. Markets can move through periods that support a diverse population of beliefs and other periods where these beliefs and strategies might collapse down to a very small set.

The market is a very simple one with a single equitylike security paying a random dividend each period and available in a fixed supply of one. This dividend follows a stochastic growth process that is calibrated to aggregate dividend series for the United States. There is a risk-free asset that is available in infinite supply paying a constant real interest rate of one percent per year. Portfolios are rebalanced and trades are made at a monthly frequency. Also, prices are determined and dividends are paid each month, which can be thought of as the basic unit of time in the market. Therefore, this is more of an experiment concerned with longer-term macroeconomic behavior as opposed to the minute by minute dynamics of day trading.

The basic actors are a set of 1,000 agents. These agents adjust their portfolios and trade independently. They have well-defined objectives in terms of optimal portfolio allocations.[5] However, they differ in one key respect: they have different views about how much past data are relevant in making their decisions. Some may take a long-horizon perspective using an equivalent of the past 20 years of data, while others view that only the past year or two of data are important. For them, previous data have become irrelevant in the investment decision-making process. As trades and time go by, the agents accumulate wealth and consume. Evolution takes place by eliminating agents with the lowest amount of wealth and replacing them with new ones. In this way, a "survival of the fittest" dynamic is imposed on the population of trading agents. As mentioned previously, it is important to understand how much pressure this puts on the market to move to a homogeneous rational outcome.

An important piece of the market is given by the trading strategies themselves. These can either be thought of as rules of thumb followed by the investors or as institutional managers that dynamically adjust their clients' portfolios. Since there are only two assets in the market, these strategies can only be market timing strategies if they deviate at all from simple buy and hold strategies. The strategies convert a subset of current market information into a recommended portfolio allocation. The allocation gives a fraction of savings to put in the risky asset.[6] The

---

[4] Examples of this include Cont & Bouchaud (2000), Epstein & Axtell (1997), Palmer, Arthur, Holland, LeBaron & Tayler (1994), Arthur, Holland, LeBaron, Palmer & Tayler (1997), Kim & Markowitz (1989), Levy, Levy & Solomon (1994), Lux (1997), Tay & Linn (2001). Also, the website maintained by Leigh Tesfatsion at www.econ.iastate.edu/tesfatsi/ace.htm is an important source for agent-based research in economics. Finally, the site at www.brandeis.edu/~blebaron/acf summarizes agent-based research in finance, and a survey of some of the early research can be found in LeBaron (2000). Some commentary on the construction of agent-based models is given in LeBaron (2001a).

[5] The agents have logarithmic preferences over expected future consumption. Their time rate of discount is set to 0.95 per year. This is a well-understood optimization problem in economics and finance (Merton 1969 and Samuelson 1969). It yields a consumption value that is a constant fraction of wealth and an investment strategy that should maximize expected log returns.

[6] Before making asset allocation decisions, agents consume a certain fraction of wealth, leaving the rest for investment. Also, they are not allowed to borrow or to sell shares short.

market information includes past returns, dividend yields, and two moving average technical indicators. The rules can be built off this information in any combination.[7] It is also important to realize that rules can ignore any piece of information too. In many ways, the ability to ignore superfluous information is part of what is being tested.

Two further issues related to trading strategies remain. First, agents must decide which rules to use. At any given time, there is a set of 250 active rules (investment advisors). Agents choose the rule that has performed the best in terms of their objective function. In doing so, they make their decision based on past data using only what they feel is relevant. In other words, if the agent believes that the only the past two years of data are important (a relatively short horizon type) this is the range over which available strategies will be evaluated.[8]

The second crucial issue is how rules learn and adapt over time. If an investment advisor has at least one agent signed up for its services, it will continue to exist with no change in its interpretation of market information into a dynamic strategy. If the advisor finds itself with no customers, it will be eliminated and replaced with a new advisor. The new advisor is created from the current population of active advisors using a genetic algorithm. This gives an interesting evolutionary dynamic to trading strategies. Those that are being used survive, and those that aren't are eliminated. The genetic algorithm tries to bring useful features of the current active strategies in to future ones. Success is determined purely on whether anyone is using a given strategy.[9]

Trading takes place each period. Agents all enter the market equipped with a chosen rule and their current portfolio positions. This gives a well-defined function for shares as a function of any given market price. Therefore, in principle, the market could be cleared by a Walrasian auctioneer operating each period. This is essentially what is done. A numerical procedure is used to find a price that sets the demand for shares of the risky asset equal to the fixed supply in the market of one share.[10]

COMPUTATIONAL EXPERIMENTS

Benchmark Runs

The following sections provide example runs of the market. In all cases, the dividend series follows a random walk that is roughly calibrated to postwar U.S. aggregate dividends

---

[7] The actual structure of the rules is given by an artificial neural network and is detailed in LeBaron (forthcoming 2001b). These provide a flexible functional form for turning the data inputs into actual portfolio weights.

[8] To introduce further heterogeneity in trader behavior, it is assumed that traders only evaluate a small number of rules each period. They do not do a complete search over all possible advisors. Specifically, their current advisor is compared to one chosen randomly from the upper half of the advisor distribution measured using the agent's own view of how much history to use.

[9] It would be difficult to use any other fitness measure, such as expected returns, since agents don't view these from a common perspective. They are estimating returns over differing time horizons.

[10] A more realistic market might consider the actual market microstructure. However, since this market will be viewed as a fairly long-range (monthly) pricing series, the temporary equilibrium assumption does not seem unreasonable.

with an annual growth rate of 2 percent and an annual standard deviation of 6 percent.[11] The risk-free rate of interest is fixed to a constant 1 percent per year. All these rates are adjusted to monthly frequency, which is the benchmark time horizon for trading and dividend/interest payments.

The key variable of interest in these experiments is the horizon length of the agents. This represents the distance they look into the past. Two different experiments will be considered. The first, referred to as *all horizon*, uses agents drawn randomly from 5 to 250 months in length. This allows for a diverse population with many different investment horizons competing against each other. The second experiment, referred to as *long horizon*, loads the market with a set of agents using a relatively long time horizon. In this case, agents are drawn from a distribution between 220 and 250 months. This loads the market with only long-horizon investors. The objective is to see if this group performs differently from the diverse investor horizon case.

The market is run for 10,000 periods, which, in calibrated time, is actually over 800 years. Figure 1 shows a run for the first 3,000 periods for a set of all horizon agents. The figure displays both price and volume. The figure shows three different phases for the market. In the first, the market is slowly adjusting, and the price is catching up to get on the correct exponential growth path. In the second phase, the market appears almost cyclical in its fluctuations about the constant growth trend. There appear to be smooth cycles that don't look very reminiscant of actual markets. In the final phase, after period 1500, the market begins to look more normal, with some run-ups in the price followed by some sudden crashes. Trading volume also increases during this later period too.

Figure 2 presents the same information for the long-horizon investors. Here we see a very different picture. After an initial adjustment phase, the price maintains a relatively steady growth rate, and trading volume drops to near zero.[12] This picture shows a glimpse of a market more closely resembling a traditional efficient market equilibrium. Further results will reinforce these early pictures.

In Figure 3, a detailed picture of the two price series is compared. This figure displays prices taken from the final 1,000 periods (roughly 80 years) of a 10,000-period run. This is done to capture behavior after all agents have settled down in their learning periods. The first two figures demonstrate that during the early periods, behavior might not reflect the final outcomes in the different market experiments. This picture very clearly repeats the message of the early figures. The top panel, which corresponds to the case where traders are using diverse horizon lengths, shows a market price that occasionally takes some large swings, often ending in dramatic crashes. In the lower panel, where traders are only long horizon, the picture shows a much smoother price dynamic with fewer large moves. The price still does indeed move, which reflects the fact that the fundamental dividend series is a somewhat volatile random walk. However, the amplification of volatility from the lower to the upper panel is clear.

[11] See Campbell (1969) for a summary of values for several different countries.

[12] Investors enter the market with no innate knowledge of how it works, or how prices move with the fundamental, so it is sensible that some learning must take place for a short time.

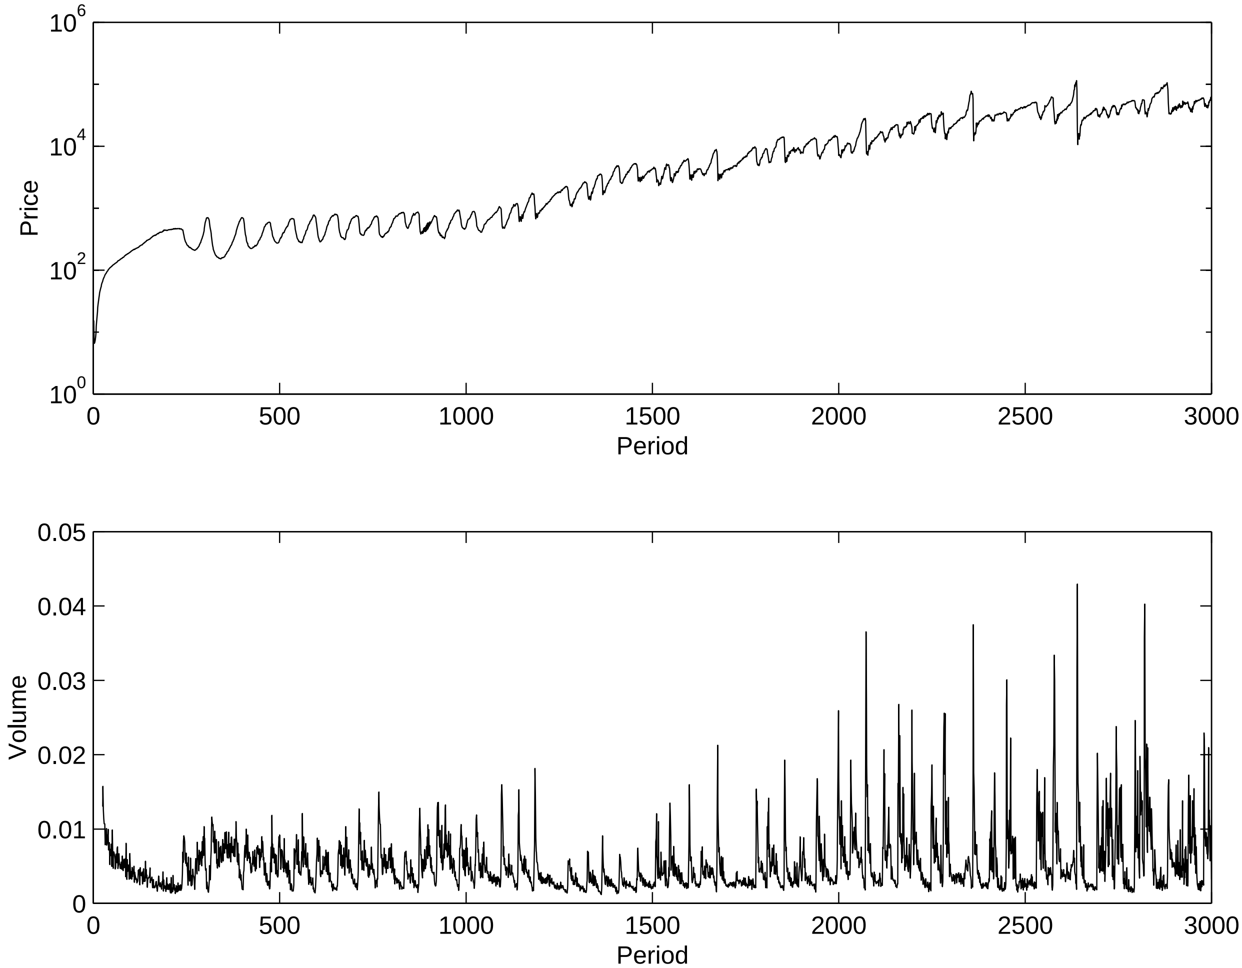

FIGURE 1 Price/Volume: All Horizons

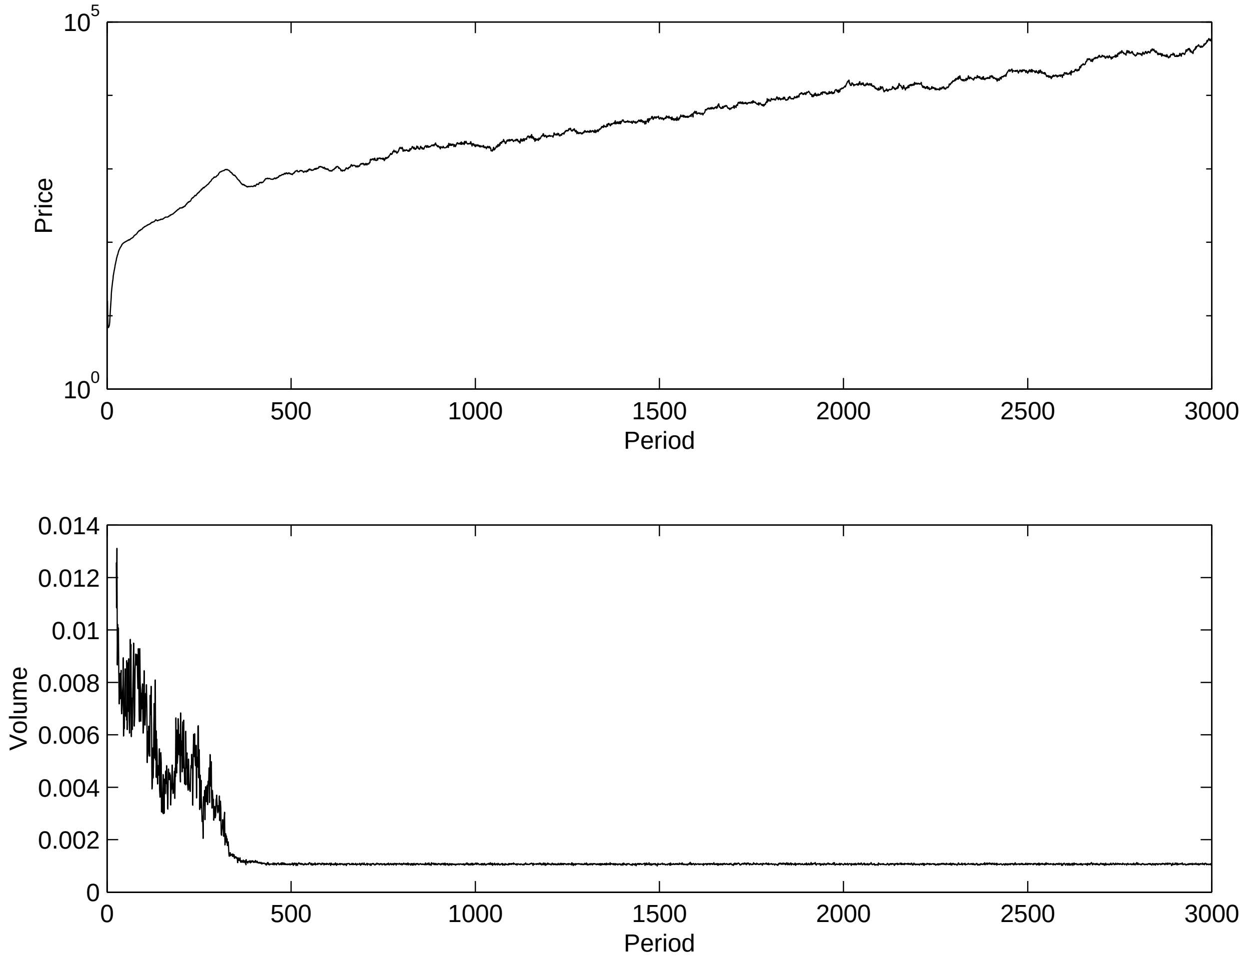

FIGURE 2 Price/Volume: Long Horizons

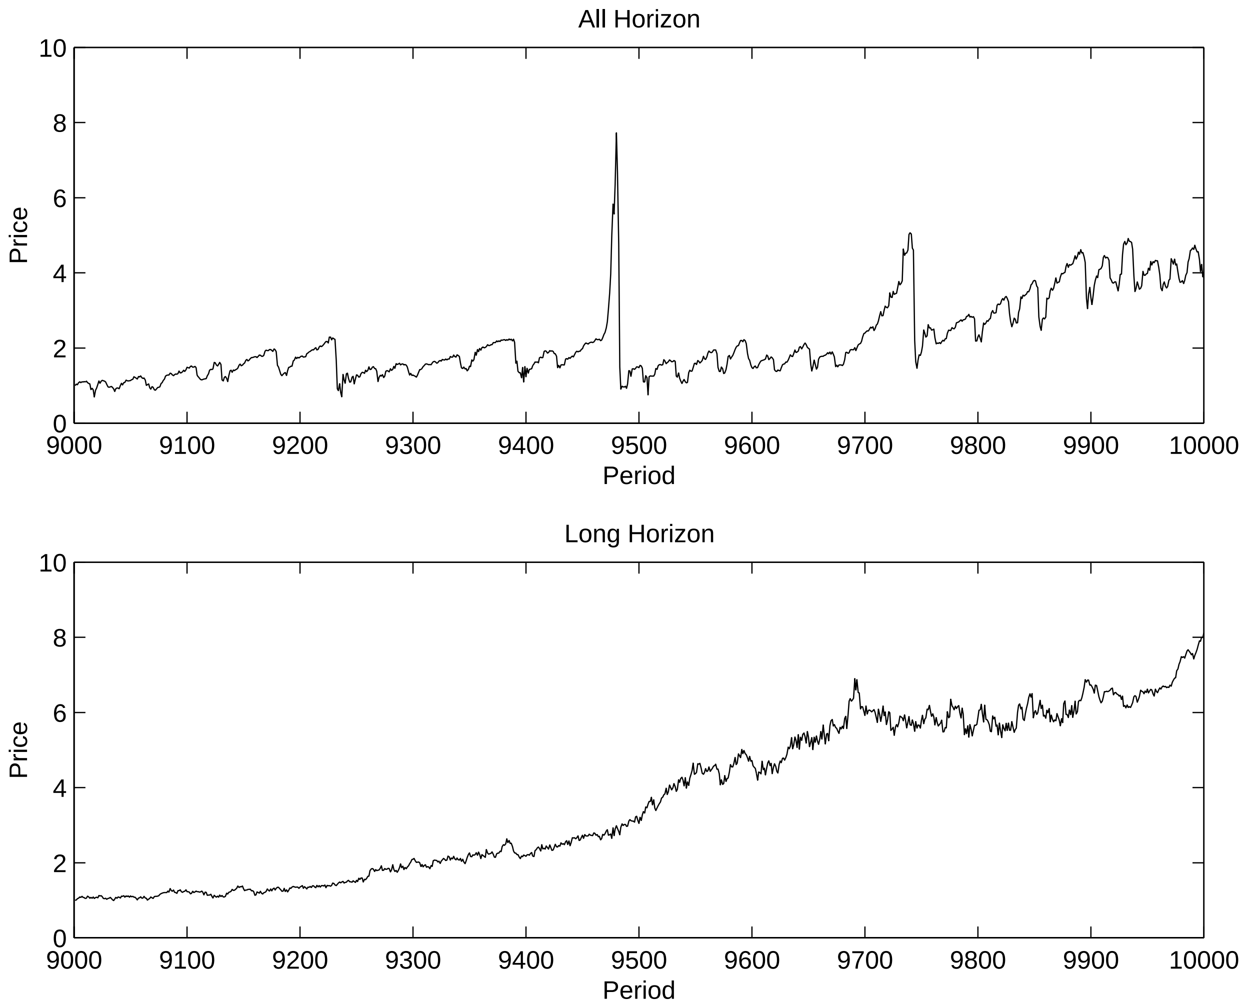

FIGURE 3 Price Comparison

RETURNS, VOLUME, AND VOLATILITY

Figure 4 compares returns across the two different computer experiments. The plot shows monthly returns inclusive of dividends. The return comparisons in Figure 4 confirm the earlier plots showing a much more variable return in the all-horizon case. Visually, returns also appear to exhibit some very large moves, with a few months yielding returns over 50 percent. Also, there appears to be some clumping in that periods of large movements are grouped together. None of these features are present in the lower panel, which corresponds to the long-horizon case. The computer-generated market yields a fairly homogenous set of returns with few, if any, large moves.

Figure 5 presents more evidence on the return distributions. Here, return histograms are compared with normal distributions. In the bottom panel, it is clear that the normal distribution provides a relatively good fit to the computer-generated returns. In the upper panel, the unusual aspects of the large moves in the all-horizon case become clear. The histogram is more peaked and yields several very unusual observations that are well outside the normal range. This replicates the sorts of distributions that are often observed in actual markets.

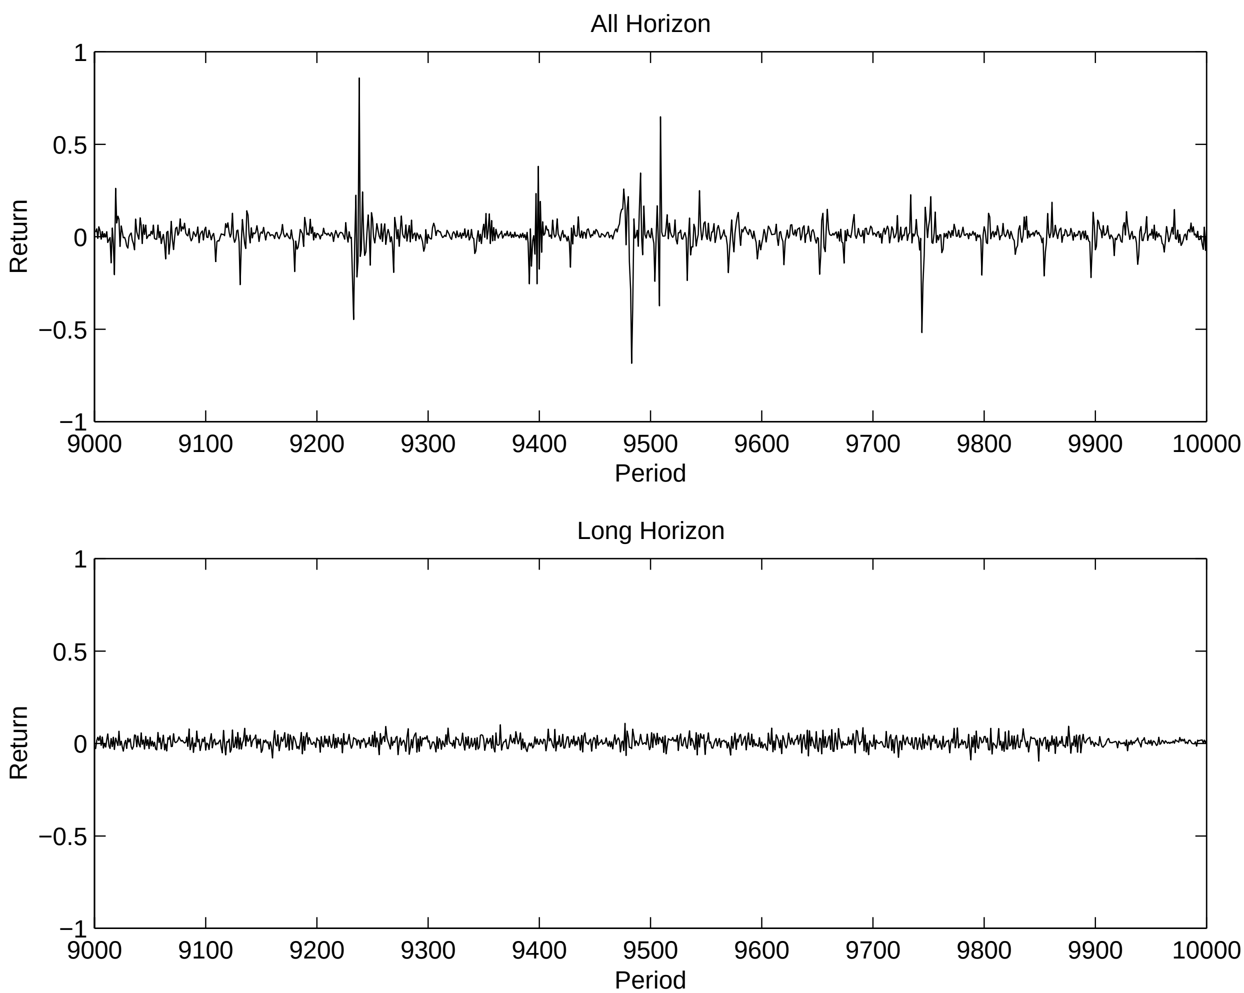

FIGURE 4 Return Comparison

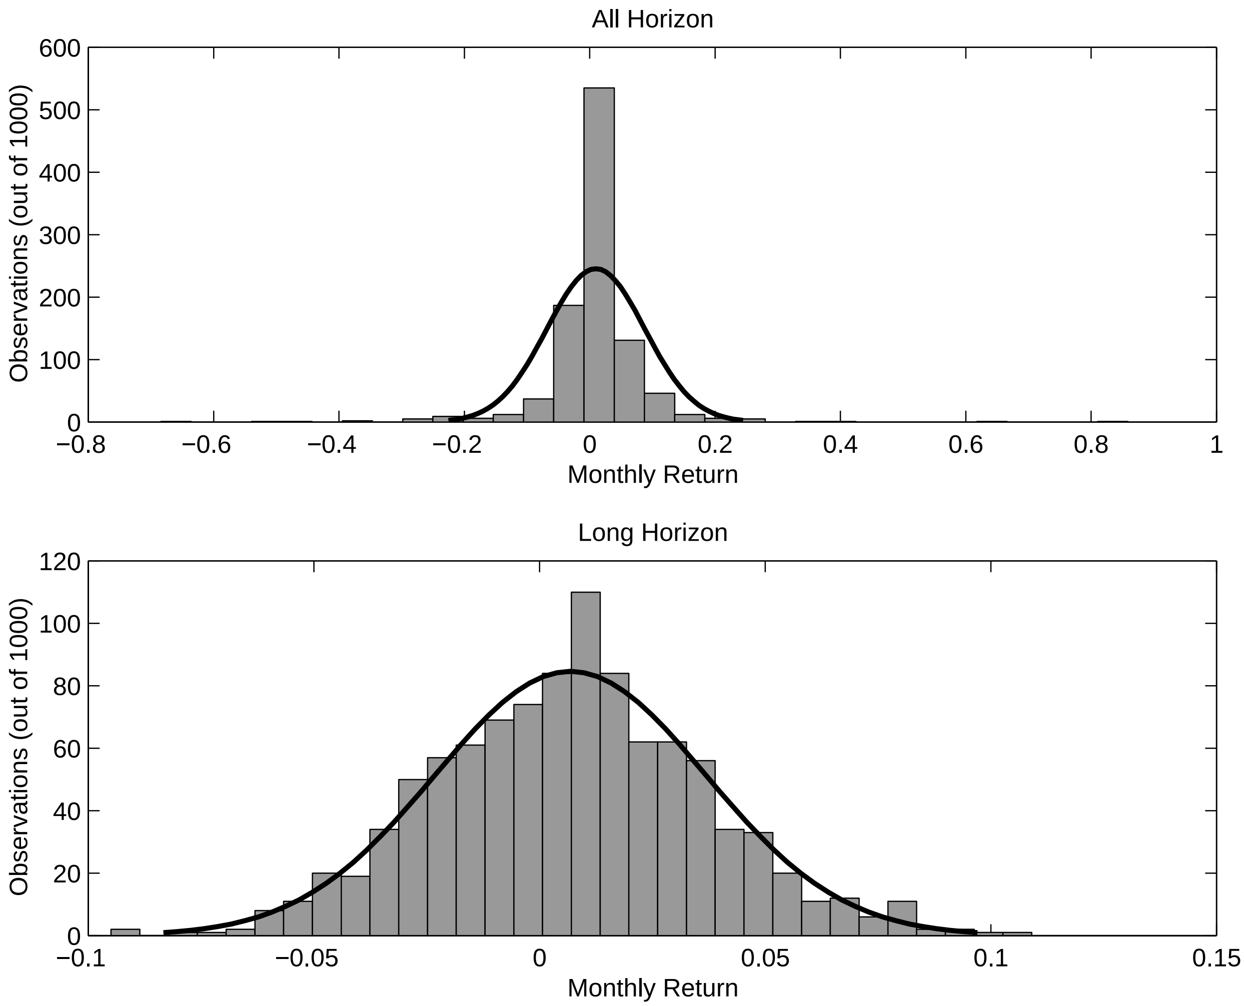

FIGURE 5 Return Distributions

Trading volume for the final 1,000 periods is shown in Figure 6. This plot shows a very strong distinction between the two cases, with very large and fluctuating volume in the all-horizon case and nearly zero volume in the long-horizon case. Obviously, trading volume is an important part of actual financial markets. The lower panel is merely reflecting the fact that the set of agents is in close agreement for asset valuations. In this case, they have no interest in trading with each other, while in the all-horizon case, differences of opinion and the desire to trade do not disappear.

Table 1 summarizes some of the results on the equity returns and trading volume. The mean return is larger in the all-horizon case than the long-horizon case. The standard deviation shows a large increase in moving from the long-horizon to the all-horizon case. Neither case demonstrates any significant skewness in returns. The most dramatic difference appears in the kurtosis estimates. These measure the thick-tailed aspects of the return distribution, which have been displayed in Figure 5. For a normally distributed random variable, kurtosis is 3. For the long-horizon case, it is very close to 3, with an estimate of 3.2. The all-horizon case shows clear evidence for leptokurtosis or “fat tails,” with an estimated kurtosis of 32. The volume numbers

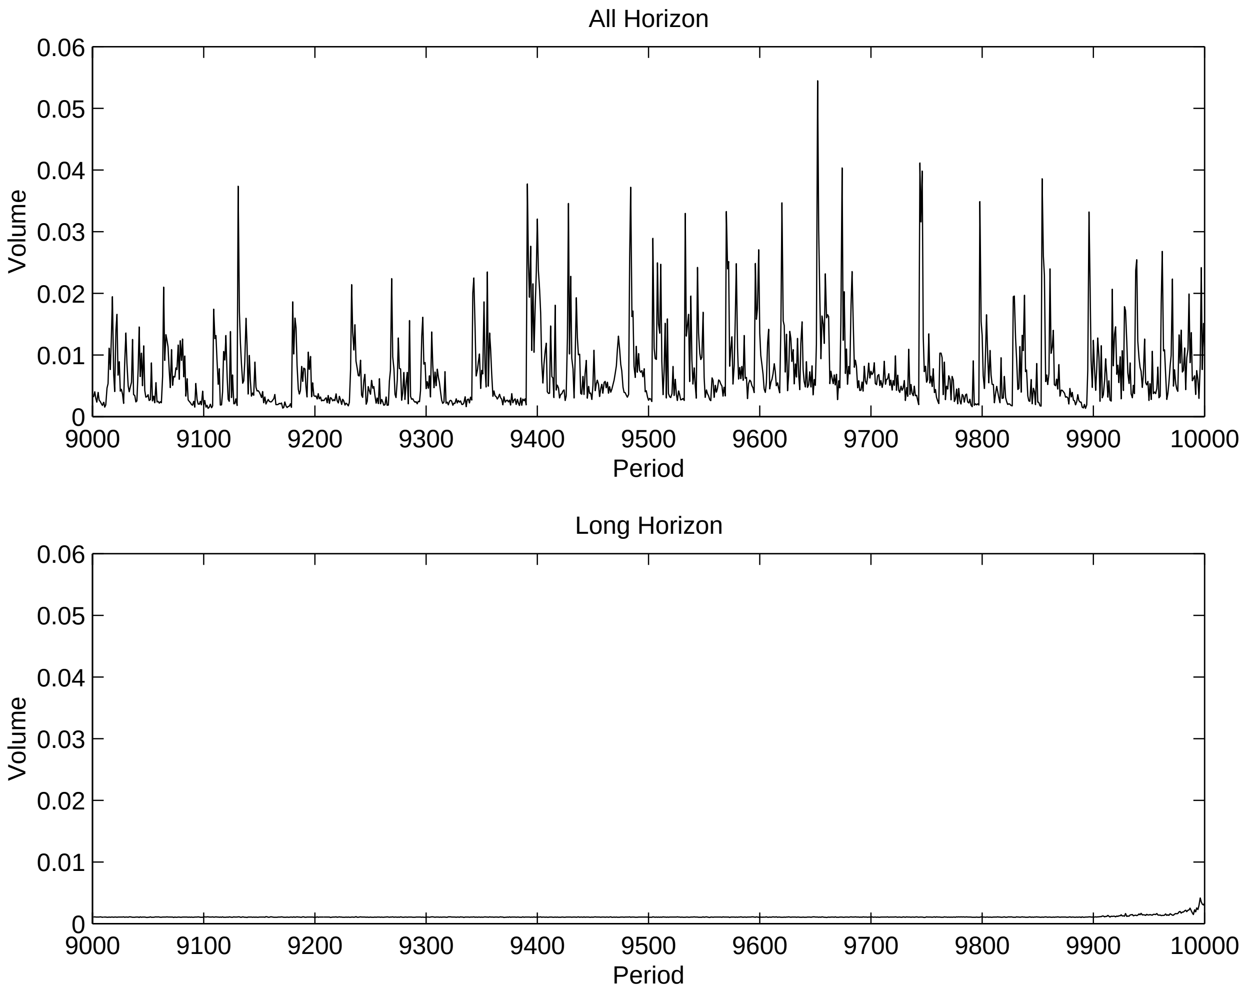

FIGURE 6 Trading Volume

TABLE 1 Excess Return Summary Statistics

	Mean	Std	Skewness	Kurtosis	Volume
All Horizon	10.6	27.1	0.3	32.6	8.7
Long Horizon	6.2	10.4	0.1	3.2	1.3
S&P	8.9	19.6	0.5	12.9	[15,78]

Summary statistics: Mean and Std are the annualized mean and standard deviation of the returns series, inclusive of dividends. Skewness and kurtosis are estimated at the monthly horizon. Values for the S&P are the total return less the 30-day T-bill rate monthly from January 1926 through June 1998. Trading volume reflects the percentage turnover at an annual rate. The value corresponding to the line S&P is the range of NYSE reported values from 1958 through 1999. It is taken from *NYSE Fact Book 1999* (2000).

report turnover at an annual rate. The all-horizon case displays trading volume that is nearly 7 times the value for the long-horizon case. Comparison numbers for the S&P are also presented. On means and standard deviations, the S&P values are between the all-horizon and long-horizon cases. The all-horizon market is actually generating more volatility and more large moves than in actual monthly data. In terms of trading volume, the NYSE clearly generates more turnover per year than the market simulations.[13]

Returns and trading volume in actual markets display several interesting dynamic features that have been hinted at in some of the earlier figures. First, returns are close to uncorrelated. Second, volatility, or the absolute value of returns, is positively correlated. In other words, large moves in either direction tend to follow large moves. Finally, trading volume is also positively correlated. Figure 7 displays the autocorrelations for returns, absolute returns, and trading volume for the long-horizon case. As in real financial data, returns show little or no correlations,

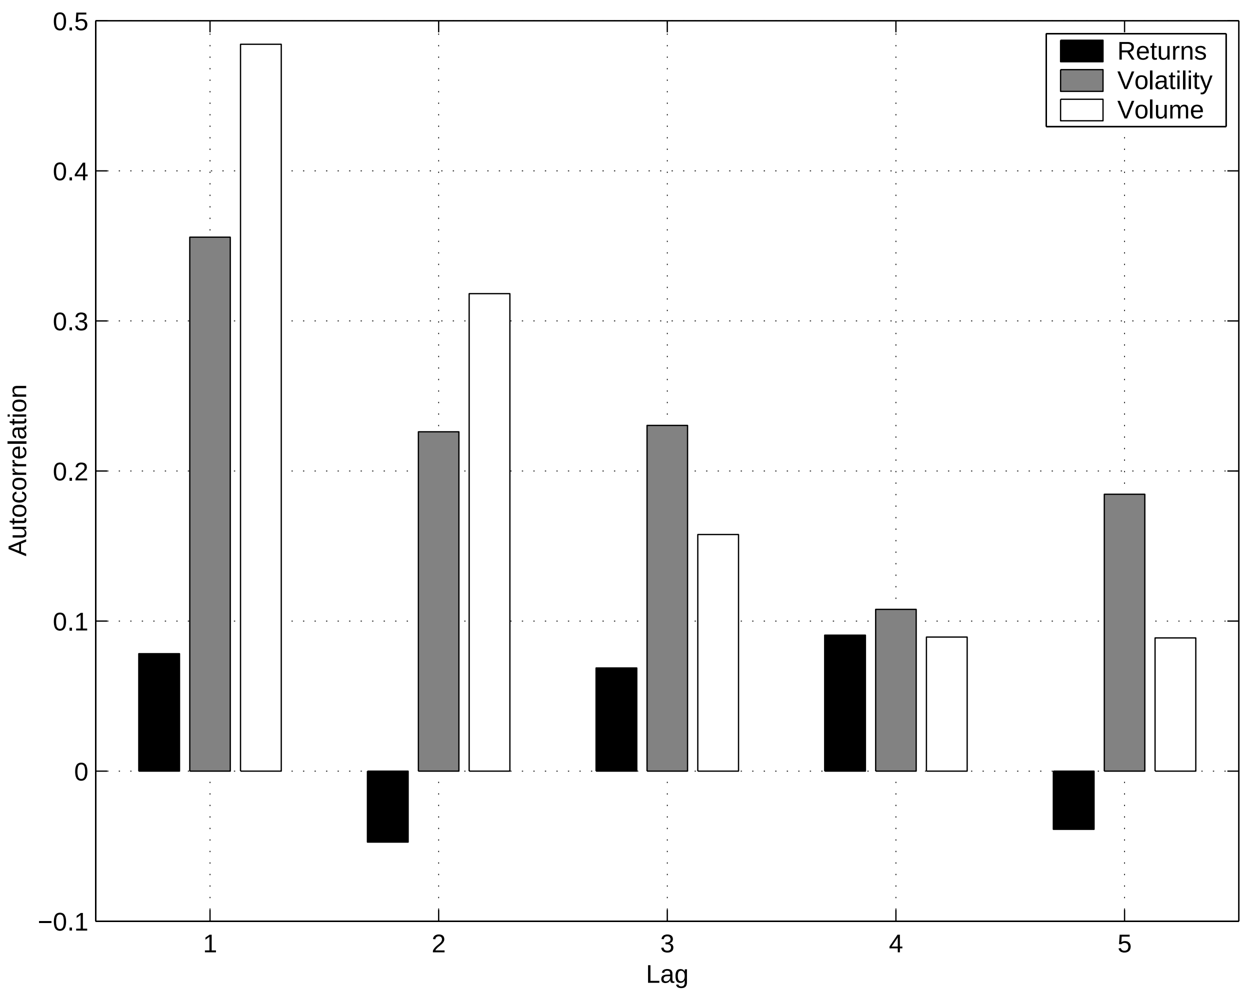

FIGURE 7 All Horizons: Return, Volatility, and Volume Autocorrelations

[13] This number is presented more for information purposes. This two asset market is difficult to compare with the NYSE in terms of trading volume, since the latter obviously has many more opportunities for trade.

positive or negative. However, both volume and volatility display strong positive correlations, going out several periods, as would be the case with actual financial time series. Figure 8 displays the autocorrelations for the market populated with long-horizon investors only. This displays a picture quite different from actual markets. Returns show a small amount of negative correlation at one lag, and then zero after that. The volatility and volume series show only negligible correlation compared with those from the all-horizon experiment. These figures clearly show a very different dynamic in the two different cases, with the all-horizon case showing a picture that more accurately reflects real markets.

Dividend Yields

Another feature of real markets is that they appear to deviate from accepted fundamental models of valuation, yielding some predictability from classical ratios such as price/earnings and dividend yields.[14] In these simple simulated markets, this a very easy and direct experiment as to

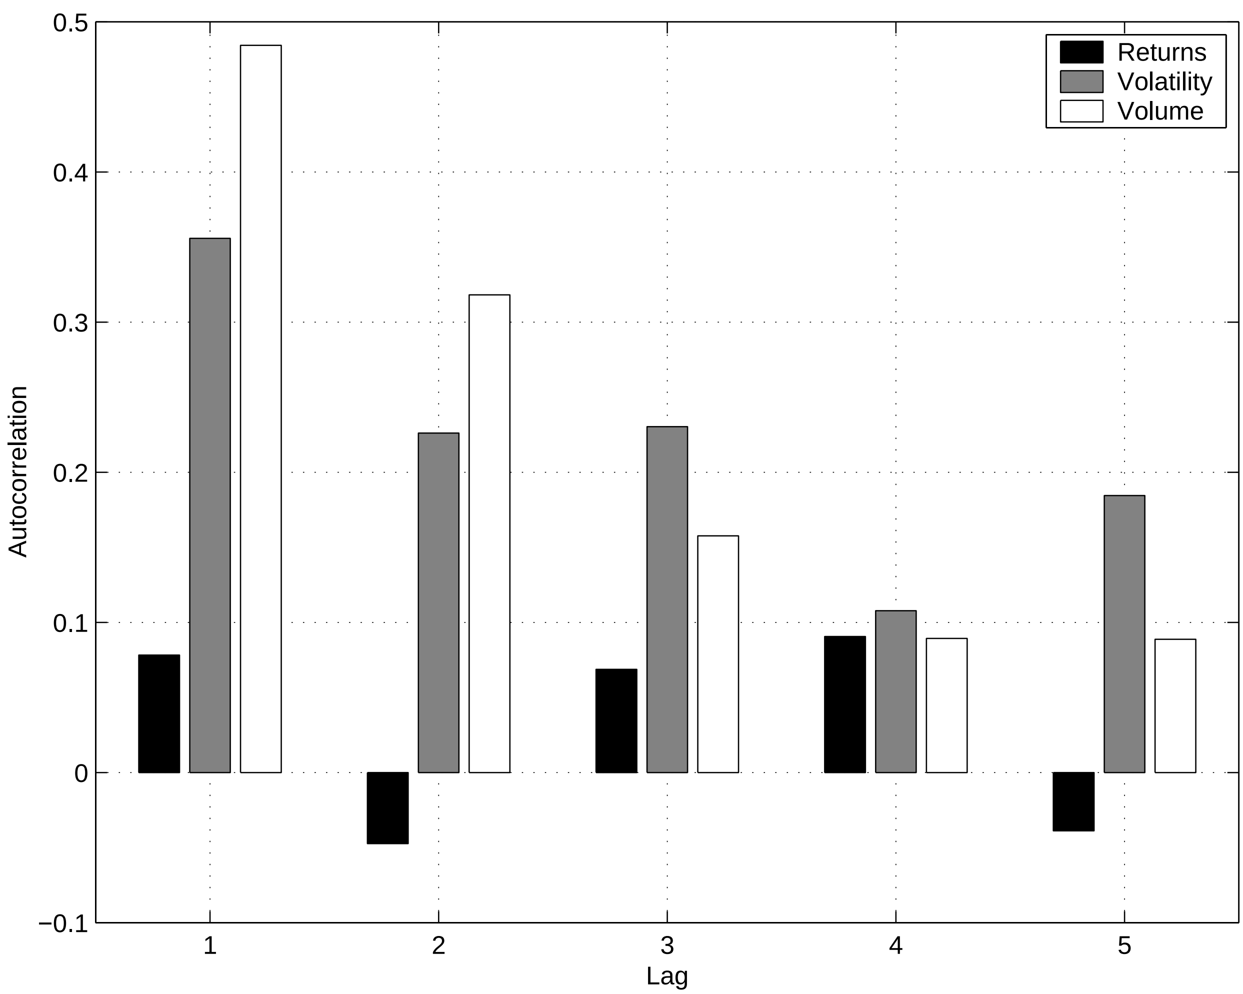

FIGURE 8 Long Horizon: Return, Volatility, and Volume Autocorrelations

[14] See, for example, Campbell & Shiller (1988), and for a recent commentary on valuation ratios and today's markets, see Campbell & Shiller (2001).

how well these markets are doing in terms of reflecting fundamental valuation. Since the equity asset only reflects a stochastic dividend stream, the dividend yield is the ratio of choice for valuation. In a constant growth situation, both the price and dividend will be growing at the same rate, and the ratio should be constant. Figure 9 compares the dividend yields in the two cases. When investors are long horizon in nature, the dividend yield is nearly flat, with only small variation around a value near 5 percent.[15] The situation is very different from the all-horizon case. It is clear that the dividend yield takes some very wide swings, going as high as 12 and as low as 2 percent. It is far from a stable series. This compares much more favorably to dividend yields and price/earnings ratios in actual financial series.

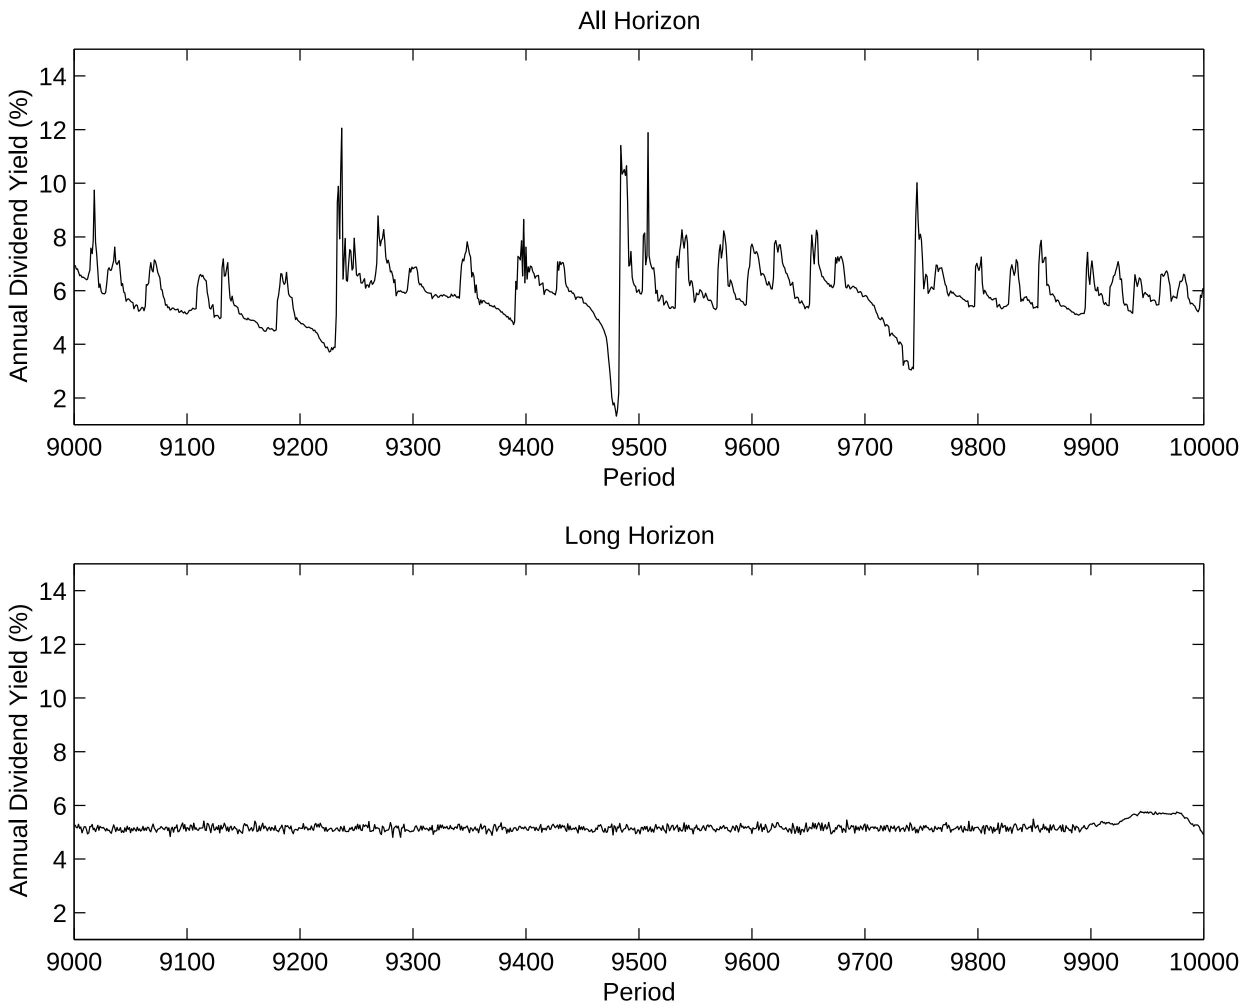

FIGURE 9 Dividend Yields

[15] The monthly dividend yield is annualized by multiplying by 12. The historical average for the S&P is close to 5 percent.

Crash Dynamics

The last two figures present an initial picture of what may be behind the large price changes in this model.[16] They take a snapshot of a short time series of prices, including some large crash periods, and compare these with two other related series.

Two measures will be used that try to assess some aspect of trader heterogeneity. As mentioned earlier, a key aspect of agent-based models is that the actual level of heterogeneity in the market is endogenous. It is possible this may be a precursor to market instability. First, the most obvious magnitude to check around crashes is trading volume. In Figure 10, several crashes are plotted along with the trading volume series. There is a weak indication that trading volume increases greatly after crashes, but it doesn't appear to show a very strong pattern before any of the large price drops in the figure, so it would be difficult to blame crashes on trading volume.

In Figure 11, the same price series is plotted against another measure of agent activity, rule dispersion. This variable measures the fraction of the rules that are currently being used by

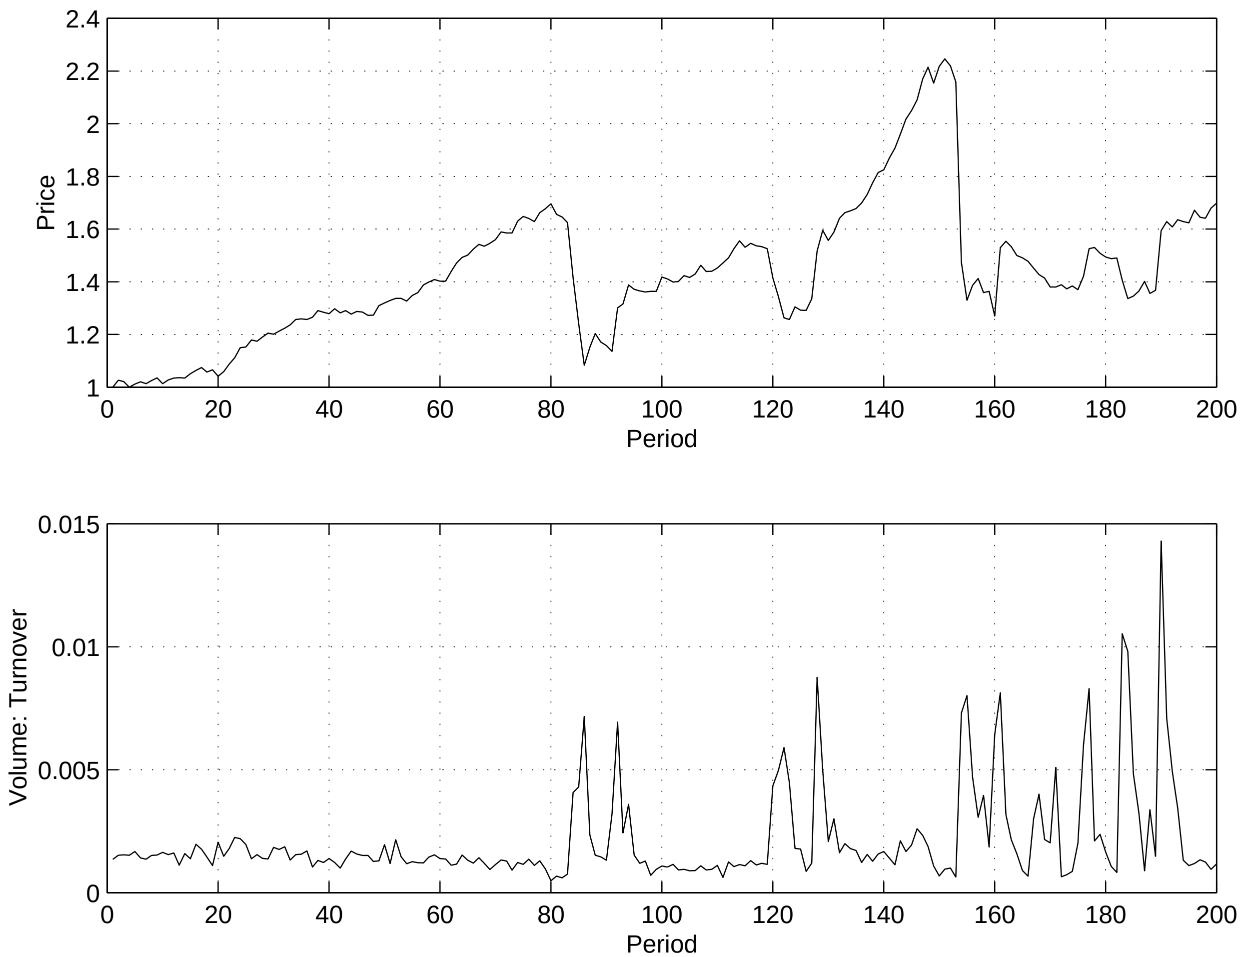

FIGURE 10 Crashes and Volume

[16] For a much more detailed statistical analysis, see LeBaron (2001c).

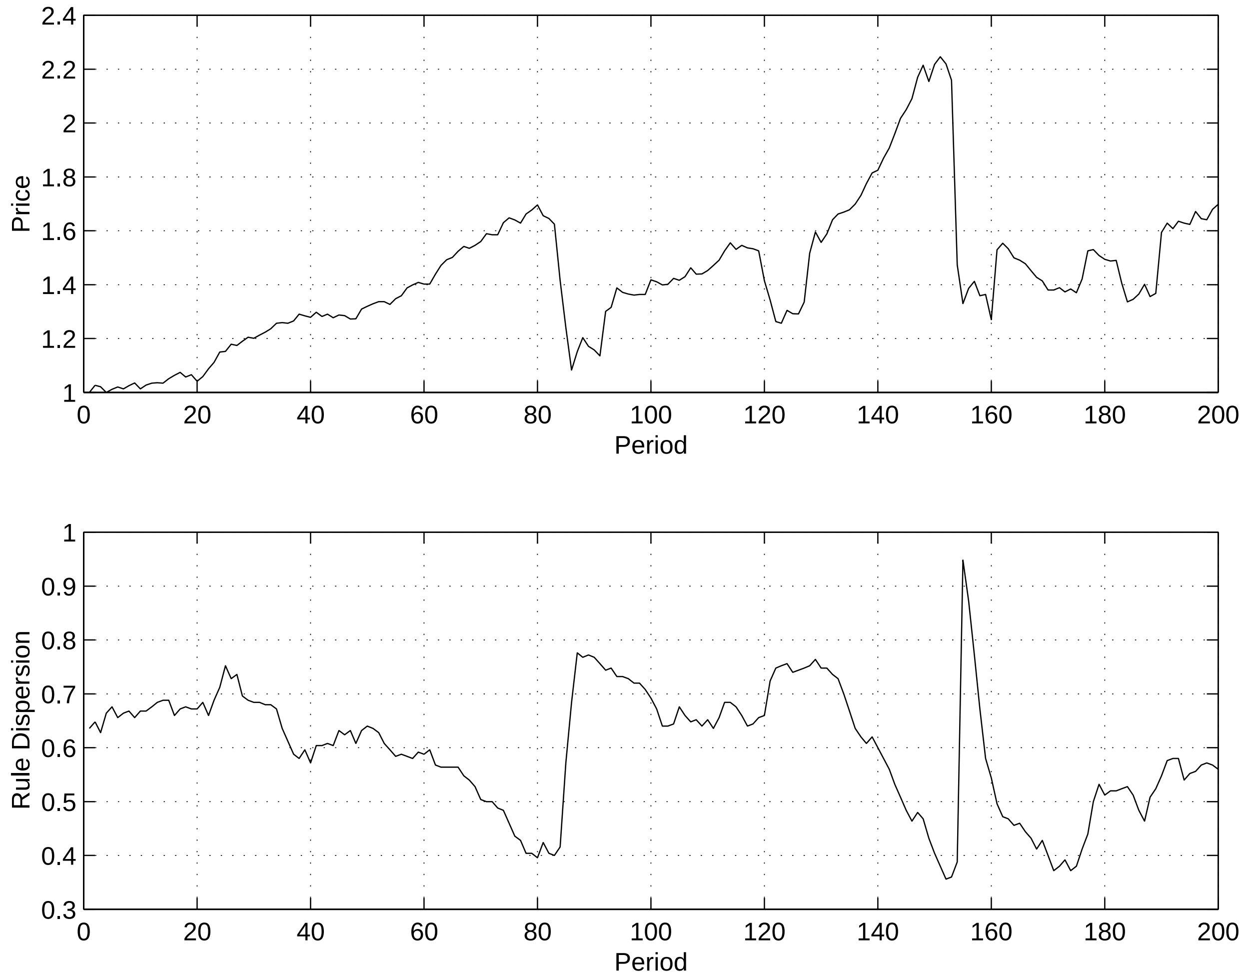

FIGURE 11 Crashes and Dispersion

agents in their active trading strategies. This comparison shows a very dramatic pattern occurring before both large crashes. Rule dispersion begins to fall long before the crash and reaches historically low levels at or near the crash date.

Armed with only this small picture of the overall dynamics, it is difficult to confirm an exact cause for market crashes. However, a simple story is starting to emerge. During the run-up to a crash, population diversity falls. Agents begin to use very similar trading strategies as their common good performance begins to self-reinforce. This makes the population very brittle, in that a small reduction in the demand for shares could have a strong destabilizing impact on the market. The economic mechanism here is clear. Traders have a hard time finding anyone to sell to in a falling market since everyone else is following very similar strategies. In the Walrasian setup used here, this forces the price to drop by a large magnitude to clear the market. The population homogeneity translates into a reduction in market liquidity.[17]

[17] This overall dynamic has some interesting parallels to the problems encountered by Long Term Capital Management. This hedge found it difficult to reduce its positions since many other traders had similar trades in place. See Lowenstein (2000) for a detailed description.

SUMMARY AND CONCLUSIONS

The results in this paper can be summarized along two dimensions. Its most important result is to provide a counter example to the argument that evolutionarily less rational strategies should be driven out of the market. Its second result is in generating time series that appear reasonable for fitting certain difficult-to-replicate features from actual markets.

As a counter example to evolutionary arguments about market efficiency, this model calls into question the basic structure of this argument. Who exactly is “less rational” in a world of heterogeneous agent investors? Will it be clearly obvious to investors to take a long-run perspective in a market dominated by short-run investors? In many ways, these computer experiments may simply be demonstrating that it is very difficult to go against the flow of the current market, even if you feel you must be right. Eventually, your performance relative to others will induce you to take a different (shorter) run perspective, and possibly further add to market instability and deviations from fundamental values.

It is important to realize that the evolutionary experiments implied in Friedman (1953) are quite a bit simpler than the actual market population dynamics. These arguments suggest a well-defined population of super-rational traders that is getting invaded by less rational types. In such a situation, the rational types may have already defined the environment, since their numbers allow them to dominate the pricing of traded assets. In other words, they have defined the rational world. In such a situation, the evolutionary argument is correct, and it should be difficult for the short-run types to stage a successful invasion of the market. The problem is that the more rational types often do not get to start with the luxury of having dominated the market. They must take over in a sea of noise from a heterogeneous population of less-than-rational types. To further complicate matters, this population may be changing its character over time, so the question of optimality gets quite murky.

Several time series features of financial data have proved to be difficult to replicate in standard models of asset pricing. This agent-based simulation easily handles many of these empirical regularities. These include “fat tailed” return distributions, increased volatility and trading volume, volatility and volume persistence, and highly persistent and variable dividend yields. Although the model may not match all of these entirely, it is still impressive that it can attack such a wide range of financial puzzles.

The model also takes a stand on the causes for large moves or market crashes. In the agent-based laboratory, the evidence suggests that crashes are caused by a move toward generally homogeneous markets. As agent strategies become aligned, market liquidity falls, and the market dynamics become brittle as traders cannot find counter parties for their trades. Hopefully, this will lead to further testable hypotheses concerning large price moves and to more general models of market liquidity.

The technology of agent-based financial markets is still in its infancy, but it appears to be a promising route for increasing our understanding of the dynamics of financial markets. It provides new insights for trying to understand financial theories in a world that may be far from the equilibrium that the theories were designed to describe. Features of financial series, such as large moves and excess volatility, may become easier to understand when viewed from a multiagent evolutionary perspective. There is a long way to go in terms of model fitting and assessment, but, for the moment, agent-based theories should take their place as a viable alternative to more traditional financial theories.

REFERENCES

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf and D. Lane, eds, "The Economy as an Evolving Complex System II," Addison-Wesley, Reading, MA, pp. 15-44.

Blume, L. and Easley, D. (1990), "Evolution and market behavior," *Journal of Economic Theory* **58**, 9-40.

Campbell, J. Y. (1999), Asset prices, consumption, and the business cycle, in J. B. Taylor and M. Woodford, eds, "Handbook of Macroeconomics," North Holland, pp. 1231-1303.

Campbell, J. Y. and Shiller, R. (1988), "The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies* **1**, 195-227.

Campbell, J. Y. and Shiller, R. J. (2001), Valuation ratios and the long-run stock market outlook: An update, Technical Report 8221, National Bureau of Economic Research.

Cont, R. and Bouchaud, J. P. (2000), "Herd behavior and aggregate fluctuations in financial markets," *Macroeconomic Dynamics* **4**, 170-196.

DeLong, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. (1991), "The survival of noise traders in financial markets," *Journal of Business* **64**, 1-19.

Epstein, J. M. and Axtell, R. (1997), *Growing Artificial Societies*, MIT Press, Cambridge, MA.

Friedman, M. (1953), The case for flexible exchange rates, *in* "Essays in positive economics," University of Chicago Press, Chicago, IL.

Kim, G. and Markowitz, H. (1989), "Investment rules, margin, and market volatility," *Journal of Portfolio Management* **16**(1), 45-52.

LeBaron, B. (2000), "Agent based computational finance: Suggested readings and early research," *Journal of Economic Dynamics and Control* **24**(5-7), 679-702.

LeBaron, B. (2001a), "A builder's guide to agent-based financial markets," *Quantitative Finance* **1**(2), 254261.

LeBaron, B. (forthcoming 2001b), "Empirical regularities from interacting long and short memory investors in an agent based stock market," *IEEE Transactions on Evolutionary Computation*.

LeBaron, B. (2001c), Volatility magnification and persistence in an agent based financial market, Technical report, Graduate School of International Economics and Finance, Brandeis University, Waltham, MA.

Levy, M., Levy, H. and Solomon, S. (1994), "A microscopic model of the stock market: cycles, booms, and crashes," *Economics Letters* **45**, 103-111.

Lowenstein, R. (2000), *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House, New York, NY.

Lux, T. (1997), "Time variation of second moments from a noise trader/infection model," *Journal of Economic Dynamics and Control* **22**, 1-38.

Merton, R. (1969), "Lifetime portfolio selection under uncertainty: The continuous time case," *Review of Economics and Statistics* **51**, 247-257.

NYSE Fact Book 1999 (2000), New York Stock Exchange, New York, NY.

Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. and Tayler, P. (1994), "Artificial economic life: A simple model of a stock market," *Physica D* **75**, 264-274.

Samuelson, P. (1969), "Lifetime portfolio selection by dynamic stochastic programming," *Review of Economics and Statistics* **51**, 239-246.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.

Simon, H. (1969), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

Tay, N. S. P. and Linn, S. C. (2001), "Fuzzy inductive reasoning, expectation formation and the behavior of security prices," *Journal of Economic Dynamics and Control* **25**, 321-362.

PRICE EFFICIENCY AND RISK SHARING IN TWO INTER-DEALER MARKETS: AN AGENT-BASED APPROACH

J. YANG, Bank of Canada*

ABSTRACT**

This paper investigates the effect of institutional arrangements on market performance using agent-based market simulation. The direct inter-dealer market is modeled as a decentralized dealership market and the inter-dealer broker market, is modeled as an order book market with double auction algorithm. On the simulated markets, six dealers trade with customers first and then with each other on a direct inter-dealer market and an electronic broker market, respectively, in two experiments. Market quality is investigated based on the results of the simulations. We find that given the same level of post-trade transparency, price discovery is relatively faster in an electronic broker market than in a direct dealer market, while the direct dealer market provides for greater opportunity for risk sharing. Furthermore, post-trade transparency increases price efficiency in both inter-dealer markets.

INTRODUCTION

This paper addresses the issue of institutional design in an inter-dealer market. We focus on inter-dealer trading and compare the effects of two distinct inter-dealer markets on market performance. These two markets are the direct inter-dealer market and the brokered inter-dealer market.

We directly compare the effects of two market structures on price efficiency and inventory control. The two markets differ only in the way trading is organized. In a direct inter-dealer market, trades are bilateral, simultaneous, and decentralized. Trades can occur at different prices at the same time. In an electronic broker market, trades are continuous and centralized and can only occur at one price at the same point in time, which is either at the best available bid or ask.

The experimental approach to financial markets is a growing field. A few experimental studies explore market performance in the context of the dealership market. Flood et al. (1999) examine the effects of pretrade transparency on price discovery. Their key comparison is between fully public price queuing (a pretrade transparent market) and bilateral quoting (a pretrade opaque market). They find that market liquidity is lower in the opaque market, but price discovery is faster in the pretrade opaque market. Thus, a trade-off exists between price

---

* *Corresponding author address*: Jing Yang, Bank of Canada, Financial Markets Department, 234 Wellington St., Ottawa, K1A 0G9; e-mail: jyang@bank-banque-canada.ca.

** This is a shortened version of Yang (2001). The author is grateful to all the participants of the Agent 2000 workshop at The University of Chicago and CEF2000 in Barcelona for their comments. Any errors that remain are the author's responsibility.

efficiency and liquidity. Bloomfield and O'Hara (1999) investigate the effects of both pretrade transparency and post-trade transparency on market performance. They find that post-trade transparency increases price efficiency but also increases the bid/ask spread. They find that pretrade transparency has no discernible effects on market performance. Sharing similar interests with these two studies, we attempt to use an agent-based approach in this paper, rather than a human-subject-based experiment, to study questions related to market institutional design.

THE MODEL AND THE EXPERIMENTS

Consider an economy with two assets, one riskless and one with stochastic liquidation value ("true value"), representing foreign exchange (FX). Dealers trade with customers first, and then they go to one of the inter-dealer markets, either a decentralized direct inter-dealer market or a centralized inter-dealer broker market to unwind their undesired inventory. The risky asset trades in periods $t = 1, 2, \ldots,$ T. Each trading period contains a series of trading rounds, *round = 1, 2,...,* R. The full information price of FX at time t is denoted by F_t, which is composed of a series of increments (e.g., interest differentials) so that

$$F = \sum_{i=0}^{T} r_i \ .$$

The increments r_i are i.i.d. with mean zero. Each increment r_i is realized immediately after each trading in period t. Realizations of the increments represent the flow of public information over time. The value of FX at t is defined as

$$F_t = \sum_{i=0}^{t} r_i \ .$$

Three signals define each period's information environment prior to dealer i's quote. The first two signals are received simultaneously. Before the third signal arrives, the first two signals are:

$$S_t = F_t + \varepsilon_t \text{, and}$$

$$C_{i,t} = F_t + \omega_{i,t} \ . \tag{1}$$

The noise terms ε_t and ω_t are normally distributed about mean zero, are independent of one another and across periods, and have variance σ_ε^2 and σ_ω^2, respectively. At the outset of each period t, all dealers receive a public signal S_t about the full-information value F, and also a private signal $C_{i,t}$. One potential source of private signals at the FX dealer level is order flow from non-dealer customers. Each dealer has sole knowledge about his own customers' order flow. To the extent that this flow conveys useful information, it potentially can be exploited in inter-dealer trading.

At the end of each trading round, dealers observe market information that represents a signal of market-wide order flow. This is given by:

$$V_t = \sum_{i=1} T_{i,t} + \zeta_t \ . \tag{2}$$

This net order flow measures the difference in buy and sell orders, and $T_{i,t}$ is the order placed by dealer i. The noise term ζ_t is normally distributed with mean zero and variance σ_ζ^2, which represents the precision of the signal. The empirical analogue of V_t is the signed order-flow information communicated by broker intercoms. The signal V_t plays a very important role in this paper. We interpret the precision of this order flow, σ_ζ^2, as the degree of transparency in the market. In later experiments, we vary the value of the precision in this signal to vary the level of transparency on FX markets.

In this paper, we model two inter-dealer markets. The direct inter-dealer market is a decentralized, quote-driven market. The electronic broker market is a centralized, order-driven, and continuous double auction market. Dealers trade with customers, and then they go to the inter-dealer market and trade among themselves. The choice of inter-dealer market is exogenous.[1]

For the direct inter-dealer market, dealers have to call one another bilaterally to ask for a quote. Then dealer *j* lays off his undesired inventory at dealer *i*'s price. In this market, quoting is simultaneous and independent, which implies that dealer *i*'s quote is not conditioned on the other dealers' quotes. Dealer *i* can only observe dealer *j*'s quote, not any other's quote.[2] This means that quote search is costly. Quotes are good for any size order. Dealers cannot refuse to quote, which implies that each dealer has an obligation to make a market. No trade information is revealed during the trading round. Only at the end of each trading round can each dealer observe a noisy signal about market-wide net order flow.

For the inter-dealer broker market, we choose a double auction setting that is closely related to Yang (2000).[3] In our simplified double auction market, dealers can either submit a bid/ask, or accept an existing bid/ask. A transaction occurs when an existing bid/ask is accepted. At the beginning of each trading round, we suppose a random permutation of the dealers, which determines the subsequent order of dealers. Initially, the dealers come to the market with their own price expectation, and they attempt to post or accept a bid (ask) order by comparing their price with the existing best ask (bid). They can only observe the best bid and the best ask price.

Let a be the best ask, b the best bid, $\hat{u}_{i,t}$ the expectation of the true price in this period, and γ the bid-ask spread. Then dealers trade with each other according to the following scenarios.

---

[1] An interesting future research topic is to make the choice between two inter-dealer markets endogenous to dealers. In reality, a common feature of dealership markets is that dealers have a choice between direct dealer trading or going to a broker market. Dealers are observed to switch between the markets, but the factors that determine such switching would be very interesting to investigate.

[2] To keep relatively similar information sets on both inter-dealer markets for comparability, we restrict the visibility of all the quotes in the direct inter-dealer market. In future research, this rule can be relaxed and each dealer can observe a few other dealers' quotes and trade with the lowest one, or not trade at all but just extract information.

[3] A few extensions in auction setting from Yang (2000) are first, that the assumption of fixed order size is relaxed, so dealers can submit an order at any size, and second, that partial execution of orders becomes possible.

Scenario 1. *If a best bid,* b*, and a best ask,* a*, exist in the market*

- If $\hat{u}_{i,t} > a$, he will submit a market order, buy $\delta(\hat{u}_{i,t} - a)$ units at this ask price;[4]

- If $\hat{u}_{i,t} < b$, he will submit a market order, sell $\delta(b - \hat{u}_{i,t})$ units at this bid price;

- If $b < \hat{u}_{i,t} < a$ and $\hat{u}_{i,t} < (a + b)/2$, he will post a buy order at a price of $(\hat{u}_{i,t} + \gamma)$ and a size of $\delta\gamma$.

- If $b < \hat{u}_{i,t} < a$ and $\hat{u}_{i,t} > (a + b)/2$, he will post a buy order at a price of $(\hat{u}_{i,t} - \gamma)$ and a size of $\delta\gamma$.

Scenario 2. If only the best ask, a*, exists*

- If $\hat{u}_{i,t} > a$, he will submit a market order, buy $\delta(\hat{u}_{i,t} - a)$ at this ask price;

- If $\hat{u}_{i,t} < a$, he will post a buy order at a price of $(\hat{u}_{i,t} - \gamma)$ with a size of $\delta\gamma$.

Scenario 3. If only the best bid, b*, exists*

- If $\hat{u}_{i,t} < b$, he will submit a market order, sell $\delta(b - \hat{u}_{i,t})$ at this bid price;

- If $\hat{u}_{i,t} > a$, he will post a sell order at a price of $(\hat{u}_{i,t} + \gamma)$ and a size of $\delta\gamma$.

Scenario 4. If no bid and ask exist,

- He will have an equal chance to post a buy or a sell order at price of $(\hat{u}_{i,t} - \gamma)$ or $(\hat{u}_{i,t} + \gamma)$, respectively, and a size of $\delta\gamma$.

Dealers use a prototypical inventory model to decide a quote, where price is linearly related to the dealer's current inventory,

$$P_{i,t} = \hat{u}_{i,t} - \alpha(I_{i,t} - I^*_{i,t}) + \gamma D_t \tag{3}$$

where $\hat{u}_{i,t}$ is the expectation of F_t conditional on the information available to dealer i at t; $I_{i,t}$ is dealer i's current inventory position; and $I^*_{i,t}$ is dealer i's desired inventory position. The

[4] The order size is decided by equation (5), given the current market price.

inventory-control effect, α, is generally a function of the relative interest rate and the cost to carry a position, which remains constant in our experiments. The term $(I_{i,t} - I^*_{i,t})$ captures the inventory imbalance. The higher is a positive (negative) inventory imbalance, the lower (higher) is the ask (bid) price a dealer is willing to post. Dealers use quotes as an important tool to perform inventory control. D_t is a direction indicator variable that equals 1 when the transaction price $P_{i,t}$ is the ask and -1 when $P_{i,t}$ is a bid. For a given expectation $\hat{u}_{i,t}$, γD_t picks up the bid-ask half-spread.

Dealer i's quote schedule is a function of the expectation of F_t at the time of quoting, which is denoted $u_{i,t}$. In turn, this expectation is conditioned on the signals S_t and $C_{i,t}$, which dealers observe at the beginning of the period, and which are updated subsequently with newly revealed market information V_t. Each dealer uses a *neural network* to find a good mapping from the three signals. Dealers update their expectations during each trading round. These expectations play a central role in determining dealer i's quote, which we describe next.

We simulate three different transparency settings for each of the two inter-dealer markets by varying the precision of post-trade order flow information. For the semi-opaque setting, the semi-transparent setting, and the perfectly transparent setting, the variance in order flow signal equals 4, 1, and 0, respectively.

RESULTS

In this section, we test three hypotheses concerning the effect of different market structures on price discovery and inventory control. The test of hypothesis 1 is that market structure has significant effect on the price discovery process. Following Bloomfield and O'Hara (1999), we use the change in the price error to measure price efficiency, which is calculated as the absolute value of the difference between the market prices and the theoretical prices. Table 1 gives a descriptive summary for the movement of price errors in the two inter-dealer markets. The results show price errors decline more rapidly in the electronic broker market. In Table 2, the results of a t-test on two sample mean of the convergence rate are reported. The first two P-values suggest to reject the null at the 5% and 10% significance level, respectively, which implies that prices on a dealership market and inter-dealer broker market converge at a different rate.

This suggests that electronic broker markets reveal information more rapidly and completely than direct dealer markets. Since the two inter-dealer markets only differ in the trading mechanism, and electronic broker market trades always happen at the best bid or ask price, we infer that the difference in price efficiency may benefit from the centralized price information, which reduces the searching cost.

The test of hypothesis 2 is that post-trade transparency increases price efficiency. To test this hypothesis, we vary the transparency level on both markets. Since the market-wide net order flow is the only public signal revealed after each trading round, we vary the noise term in this order flow signal to accommodate different post-trade transparency levels. For semi-opaque, semi-transparent and transparent markets, the variance of the noise in the order flow signals equals 4, 1, and 0 respectively. We present the results for the different transparency levels in Table 2. The values of decline from 0.64 under the semi-opaque setting to 0.61 under the

TABLE 1 Change in the Price Errors on Two Inter-Dealer Markets

Price Errors	Average of All Trading Rounds	Average of First 40 Trading Rounds	Average of Last 40 Trading Rounds
E-broker market	0.3838 (0.1170)	0.7385 (0.2774)	0.2588 (0.0290)
Direct inter-dealer market	0.6856 (0.1321)	1.0662 (0.1666)	0.568 (0.0151)

Notes: Every experiment is run for 10 periods and 20 trading rounds in each period. Overall, the same experiment is run for 25 times. The statistics shown in this table are the average over 25 runs. Numbers in parentheses are standard errors estimated using the 25 runs.

TABLE 2 Impact of Market Structure on Price Efficiency at Different Transparency Levels

Price Efficiency Measure	β_1	β_2	**H0:** $\beta_1 = \beta_2$
Semi-opaque setting	0.6412	0.7810	-2.5792
	(0.182)	(0.201)	[**0.006**]
Semi-transparent setting	0.6120	0.6531	-1.4837
	(0.081)	(0.012)	[**0.072**]
Perfect post-trade transparency	0.4206	0.4418	-0.4028
	(0.160)	(0.1812)	[0.334]

Notes: The results presented in this table are for the regressions from equation $Y_{i,t} = \mu_i + \beta_i Y_{i,t-1} + \nu_{i,t}$. The dependent variables are the price errors of the direct inter-dealer market or broker market. Numbers in parentheses are the standard errors of the estimates using the 25 runs. The last column is the statistic for a t-test, with the P-values reported in the brackets. Every experiment is run for 10 periods and 20 trading rounds in each period. The same experiment is repeated 25 times. The statistics shown in this table are averaged over the 25 runs.

semi-transparent setting, and then to 0.42 under the perfect transparency setting on the electronic broker market. A similar pattern appears in the direct inter-dealer market. This suggests post-trade transparency increases the price efficiency and accelerates the decline of price errors in both markets. This is evidence in support of our hypothesis 2. This result is consistent with the experimental findings of Bloomfield and O'Hara (1999), where they find trade disclosure increases the price discovery process.

The test of hypothesis 3 is that the direct dealer market provides a better inventory control opportunity. To investigate which inter-dealer market offers a better inventory control opportunity, we first examine whether the inventories exhibit mean reversion and then measure the adjustment speed of inventory imbalance corrections. A higher adjustment speed implies better risk sharing and inventory control opportunities for an inter-dealer market.

In inventory models of dealership markets,[5] the degree of competitiveness of a dealer's quotes depends on his relative inventory position as reflected in equation (4). The higher the positive (negative) inventory imbalance, the lower (higher) is the ask price that the dealer is willing to post. When a dealer has an extreme inventory position, he is able to post competitive quotes on one side of the market, and he stands a better chance of executing the public order flow in the desired direction. This results in a relatively quick reduction of his inventory imbalance. On the other hand, when a dealer's inventory is closer to the desired level, he is not able to post competitive prices and therefore stands a poor chance of executing the public order flow. As a result, his inventory takes a longer time to revert to the desired level. This implies that in competitive dealership markets, the relative inventories of the dealers should be mean reverting.

Let $IM_{i,t} = I_{i,t} - I^*_{i,t}$ denote the inventory position of market maker i relative to the desired inventory level at time t. We run the following regression,

$$\Delta IM_{i,t} = a + \theta IM_{i,t-1} + \tau_{i,t}\,, \tag{5}$$

where i denotes the dealer, and $IM_{i,t}$ represents the inventory imbalance of dealer i at the end of time t. The coefficient θ captures the intensity of mean reversion.

The results, which are reported in the second and third columns of Table 3, show that dealer's inventories in the direct dealer market have greater force of mean reversion (the absolute value of the coefficient). The average mean reversion coefficient for dealers' inventories in the electronic broker market is –0.064, and in the direct dealer market, –0.28.

The speed of adjustment of inventories is directly related to the mean reversion coefficient θ, which represents the fraction of the deviation between actual and desired inventories that is eliminated each day. A useful measure of adjustment speed is the inventory half-life, denoted by h. It is defined as the expected number of days required to reduce a desired inventory by 50%, where

$$h = \frac{\ln 2}{\ln(1-\theta).} \tag{6}$$

---

[5] See Ho and Stoll (1983), Madhavan and Smidt (1993), and Hansch, Naik, and Viswanathan (1998) for a theoretical inventory model and empirical tests.

The last two columns in Table 3 provide the estimates of the inventory half-life for six dealers in the two inter-dealer markets. The results show that it takes, on average, 12.94 trading rounds for an inventory imbalance to be reduced by 50% in the direct dealer market. In contrast, the inventory half-life, on average, is 22.09 trading rounds in the electronic broker market. This implies, on average, that dealers in the direct inter-dealer market are able to adjust their inventory imbalance faster, which, in turn, implies that dealers carry lower inventory-risk in this market.

CONCLUSION AND FUTURE RESEARCH

In this paper, we directly compare two different types of inter-dealer markets. Given a fixed level of market transparency, we find that there is a trade-off between price efficiency and inventory control. The centralized electronic broker market has greater price efficiency than a decentralized direct dealer market, but the direct dealer market provides greater market depth and a shorter inventory half-life. Post-trade transparency increases the price efficiency of both markets. We close with a discussion on further research. Our experiments investigate the effects of market structure on market performance. Having established substantial differences in market behavior, the natural question follows whether one type of inter-dealer market is likely to dominate another. When the choice of trading on certain inter-dealer markets become endogenous, which inter-dealer market prevails? Such a topic is well suited to exploration in an experimental market setting. We believe further research on this topic will provide more insights into inter-dealer markets.

TABLE 3 Inventory Mean Reversion Coefficients and Implied Half-Life

Dealers	θ on e-Broker Market	θ on Direct Dealer Market	Half-life on e-Broker Market	Half-life on Direct Dealer Market
Dealer 1	-0.0184	-0.0118	38.076	59.091
Dealer 2	-0.0134	-0.1085	52.105	6.7281
Dealer 3	-0.0595	-0.2304	11.989	3.3429
Dealer 4	-0.0460	-0.3458	15.410	2.3333
Dealer 5	-0.0646	-0.8960	11.070	0.0903
Dealer 6	-0.1851	-0.1216	4.0812	6.0388

Note: The results presented in this table are for the regressions based on equation (5) and from the calculation using equation (6). In both markets, the transparency parameter takes the value of 1, which implies that the results shown in this table are for semi-transparent markets. Every experiment is run for 10 periods and 20 trading rounds in each period. Overall, the same experiment is run 25 times. The statistics shown in this table are the average over 25 runs.

REFERENCES

Bloomfield, R., and O'Hara, M. (1999) Market transparency: Who wins and who loses? *The Review of Financial Studie*s, Spring 1999 Vol. 12 pp. 5-35.

Flood, M., Huisman, R., and Mahieu, R. (1999) Quote disclosure and price discovery in multiple-dealer financial markets, *The Review of Financial Studie*s, Spring 1999 Vol. 12 pp. 37-59.

Flood, M., (1994) Market structure and inefficiency in the foreign exchange market, *Journal of International Money and Finance* (1994), 13(2), pp. 131-158.

Garcia, R., and Gencay, R. (1999) Pricing and hedging derivative securities with neural networks and homogeneity hint, *Journal of Econometric*s, forthcoming.

Glosten, L. (1994) Is the electronic open limit order book inevitable? *The Journal of Financ*e, Vol. XLIX, No. 4, September 1994.

Hansch, O., Naik, N., and Viswanathan, S. (1998) Do inventories matter in dealership markets? Evidence from the London Stock Exchange, *The Journal of Financ*e, Vol. LIII, No. 5, 1998.

Hertz, J., Krogh, A., and Palmer, R. (1991) *An Introduction to the Theory of Neural Computation*, Addison-Wesley Publishing Company, Reading, Mass.

Ho, T., and Stoll, H. (1983) The dynamics of dealer markets under competition, *Journal of Finance* 38 (1983) pp. 1053-1074.

Hutchinson, J., Lo, A., and Poggio, T. (1994) A nonparametric approach to pricing and hedging derivative securities via learning networks, *Journal of Finance* 49(3), pp. 851-889.

LeBaron, B., Arthur & Palmer (1999) Time series properties of an artificial stock market, *Journal of Economic Dynamics and Control,* 23 (1999) pp. 1487-1516.

LeBaron, B. (forthcoming 1999) Agent based computational finance: Suggested reading and early research, *Journal of Economic Dynamics and Control.*

Lyons, R. (1997) A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* Vol. 42 (1997) pp. 275-298.

Lyons, R. (1995) Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economic*s, Vol. 39 (1995) pp. 321-351.

Lamoureux, C., and Schnitzlein, C. (1997) When it's not the only game in town: The effect of bilateral search on the quality of a dealer market, *Journal of Financ*e, Vol. LII, No. 2, June 1997.

Madhavan, A. (1992) Trading mechanisms in securities markets, *Journal of Financ*e, Vol. XLVII, No. 2, June 1992.

Madhavan, A., and Smidt, S. (1993) An analysis of changes in specialist inventories and quotations, *Journal of Financ*e, Vol. XLVIII, No. 5, December 1993.

O'Hara (1995) *Market Microstructure Theor*y, Blackwell Publishers, Cambridge, USA.

Yang, J. (2000) The efficiency of an artificial double auction stock market with neural learning agents, in Chen, S.H. (ed.), *Evolutionary Computation in Economics and Financ*e, Springer-Verlag (forthcoming).

Yang, J. (2001) Price Efficiency and Inventory Control in Inter-dealer Trading, Bank of Canada, internal report.

THE COMPARISON BETWEEN SOCIAL LEARNING AND INDIVIDUAL LEARNING: AN APPROACH BASED ON GENETIC PROGRAMMING

C.-H. YEH, I-Shou University, Taiwan
S.-H. CHEN, National Chengchi University, Taiwan*

EXTENDED ABSTRACT

Economic simulations have been widely employed in the study of economics. The learning behavior of economic agents is also modeled by the computational framework. Different learning algorithms are used to *simulate* different style of learning behavior. In the literature, they can be classified into social learning and individual learning. In social learning, traders learn from other traders' experience, while in individual learning, they learn from their own experience. The implications between social learning and individual learning have been stressed. However, what's more important is that the simulation result is not only influenced by *how we learn*, but also *what we can learn*. In other words, both of the learning styles and potential knowledge space contribute to the outcome. In terms of the techniques of evolutionary computation, the potential knowledge space is related to the representation. Therefore, how we represent the knowledge is one of the most important steps in the economic simulations. According to Lucas (1986),

> In general terms, we view or model an individual as a collection of decision rules (rules that dictate the action to be taken in given situations) and a set of preferences used to evaluate the outcomes arising from particular situation-action combinations. These decision rules are continuously under review and revision; new decision rules are tried and tested against experience, and rules that produce desirable outcomes supplant those that do not. (p. 217)

From the viewpoint of representation, if a decision rule can *hopefully* be written and implemented as a computer program, and since every program in terms of its input-output structure can be understood as a function, then, based on the LISP programming language, every function can be represented as a LISP S-Expression, and hence a parse tree. This representation of decision rules is exactly what genetic programming does.

In the past few years, genetic algorithms (GAs) and genetic programming (GP) have frequently been used to model economic agents. There are two ways to implement GAs and GP, that is, single-population GAs (GP) and multi-population GAs (GP), which are distinguished from social learning and individual learning. Vriend (2000) pointed out their difference and consequences based on the framework of genetic algorithms. However, according to the traditional implementation of GAs, pre-specified domain knowledge is required, which makes the result more parochial. In this paper, multi-population GP is employed. The topic used to investigate the influence of social learning and individual learning is the *artificial stock market*.

---

* *Corresponding author address*: Chia-Hsuan Yeh, Department of Information Management, I-Shou University, Kaohsiung County, Taiwan 84008; e-mail: spockyeh@ms38.bitnet.net. Research support from NSC grants No. 89-2415-H-214-002 is gratefully acknowledged.

The basic framework of the artificial stock market considered in this paper is the standard asset pricing model employed in Grossman and Stiglitz (1980). The dynamics of the market is determined by an interaction of many heterogeneous agents. In this market, there are two assets available for traders to invest. One is the riskless interest-bearing asset called *money*, and the other is the risky asset known as the *stock*. At each period, each of them has to make a decision about how many shares of stock he should hold based on his forecast about the future (the sum of price and dividends in the next period) in order to maximize the one-period expected utility. Their forecasts are formed by genetic programming. The control parameters of this genetic programming are shown in Table 1.

At the evaluation date *t*, each trader has to make a decision. Should he change his mind (the strategy used in the previous period)? This psychological activity is modeled by two probabilities which describe the intensity of peer pressure and self-realization respectively. If the trader decides to change to a new idea, then he will go to the business school (in social learning) or think about it (in individual learning) to get useful strategies. Once he gets a new idea, he will compare the new idea with his old one used in the previous period. If the new idea outperforms his old idea, he will adopt the new one. Otherwise, he will go to the business school or think about it once again until either he succeeds or he fails for a pre-specified time.

Understanding the difference between social learning and individual learning more precisely based on the representation of GP, we consider three different scenarios, Markets A, B, and C. The markets B and C are distinguished by the different number of ideas in each trader's mind. Market A is the case of social learning, which is the same one used in Yeh and Chen (2000). In this paper, the influence of the number of ideas for each trader is discussed. In principle, the traders are more adaptive when they have more ideas in mind. Therefore, they have more chances to discover the patterns of price dynamics, so their survivability is also higher. However, they may cause the market to become more complicated beyond control. Their survivability is then reduced. Which of these is the most possible outcome is a question that needs to be studied.

Based on the experimental design described above, the difference between social learning and individual learning is then investigated. In this paper, we focus on

- their chance in discovering the fundamental price,
- the market efficiency,
- their chance in generating exotic behavior.

Keywords: Social Learning, Individual Learning, Genetic Programming, Artificial Stock Market

REFERENCES

Chen, S.-H., and C.-H. Yeh, 2000, "Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market," *Journal of Economic Dynamics and Control*, forthcoming.

Grossman, S.J., and J., Stiglitz, 1980, "On the Impossibility of Informationally Efficiency Markets," *American Economic Review*, 70, pp. 393-408.

Lucas, R.E., 1986, "Adaptive Behavior and Economic Theory," in Hogarth, R.M., and M.W. Reder (eds.), *Rational Choice: The Contrast between Economics and Psychology*, University of Chicago Press, pp. 217-242.

Vriend, N., 2000, "An Illustration of the Essential Difference between Individual and Social Learning, and Its Consequence for Computational Analysis," *Journal of Economic Dynamics and Control*, Vol. 24, No. 1, pp. 1-19.

Yeh, C.-H., and S.-H. Chen, 2000, "Toward an Integration of Social Learning and Individual Learning in Agent-Based Computational Stock Markets: The Approach Based on Population Genetic Programming," presented at the Sixth International Conference on Computing in Economics and Finance (CEF '2000). Barcelona, Catalonia, Spain. July 6-8, 2000.

TABLE 1 Parameters of the Stock Market

The Stock Market	
Shares of stock (H)	100
Initial money supply (M_1)	100
Interest rate (r)	0.1
Stochastic process (D_t)	Uniform distribution, U(5.01,14.99)
Parameters of Genetic Programming	
Function set	{+,-, ×,%,Sin,Cos,Exp,Rlog,Abs,√}
Terminal set	$\{P_t,P_{t-1},\ldots,P_{t-10},P_t+D_t,\ldots,P_{t-10}+D_{t-10}\}$
Selection scheme	Tournament selection
Probability of creating a tree by reproduction	0.10
Probability of creating a tree by immigration	0.20
Probability of creating a tree by crossover	0.35
Probability of creating a tree by mutation	0.35
Probability of mutation	0.3
Probability of leaf selection under crossover	0.5
Mutation scheme	Tree mutation
Replacement scheme	(1+1) Strategy
Maximum depth of tree	17
Maximum number in the domain of Exp	1700
Number of generations	4000
Traders	
Number of traders (N)	100
Number of ideas for each trader	1 (A, Social Learning), 10 (B), 25 (C)
Degree of RRA (λ)	0.5
Criterion of fitness (Traders)	Increments in wealth (Income)

DISCUSSION: ARTIFICIAL MARKETS

L.P. HANSEN, University of Chicago, Moderator

Michael North [to Blake LeBaron]: One thing you mentioned is that homogeneity tends to lead to crashes. I'm wondering if that's related to the amplification effect, where a gain in the system goes way up when you've got homogeneity because everyone is essentially making the same decision.

Blake LeBaron: There's definitely a dynamics of the crash, where you have homogeneous behavior, but what's almost more interesting is the buildup to the crash, which is where the agents are locking in on the same thing. In effect, as those horizons click by, agents are finding it harder and harder to swing the other way.

North: That's what I was referring to in terms of amplification.

LeBaron: It poses an interesting question: is that a good model for what happens in a crash? You carefully monitor various things during the crash. Unfortunately, one of the things you can't measure in the real world is heterogeneity. I'm thinking a lot about proxies for that, but I haven't settled on anything yet.

Peyton Young: I don't understand the distribution of memory lines in your model. In your slides, it didn't appear that there were any substantial differences in the returns for the short, medium, and long memories at various times. They fluctuated fairly wildly. But were the ones with poor performance being deselected? Did long-memory agents doing poorly at a certain interval get deselected through the genetic algorithm? I didn't get a sense that was true.

LeBaron: There is a uniform distribution across memory links going in. In many runs I've actually locked that population of memory link agents. What happens is that wealth shifts; it sloshes around between agents. This change takes the place of evolution over the agents — when agents aren't doing well, you can say either that they're going away or that their wealth is effectively going to zero. Because of the constant relative risk-aversion preferences, they each have the same impact on the price, and so you think of it as evolution across those agents. I have run it both with and without evolving the agents and it hasn't made any difference. The main points are the constant relative risk-aversion preferences and the fact that it's merely wealth that's shifting.

Eventually, when I consider issues such as bankruptcy, borrowing, and lending — lending is not in the model now — I will have to come to grips with how the model should handle agents who are going bankrupt. It will mean taking them out in some clean way — cleaning up bankruptcies in these models is messy, as in the real world.

Claudio Cioffi-Revilla: [What about] the distribution of downturns?

LeBaron: What we know definitively is that stock returns reported at relatively high frequencies — daily, weekly and monthly — are not Gaussian. There are many large moves. What happens beyond that is the subject of controversy. Some researchers give what they believe is evidence of power laws. I’m personally not yet convinced they exist. The problem is, of course, there is just so little data in the tail. So we don’t know much more than that large moves occur a lot more frequently than they would in a Gaussian distribution. Beyond that, there is some controversy about asymmetry. Are there more large drops than gains? There seems to be some evidence of that. But in all of these cases I think the empirical literature is still undecided.

Researchers pretty well agree on two other facts, though: first, that there is convergence to Gaussianeity at the very long horizon, and second, that second moments exist in the tails, which is related to power-law scaling in the tails — which is an important issue. Whether there are moments beyond second moments is unclear.

There’s a set of controversial issues about power laws in financial data. Researchers also argue over power laws and the persistence of volatility. It appears to be a very long-range persistence that is connected to power laws. I also think that issue is somewhat controversial, and it somewhat overlaps with possibly very low frequency regime changes — things like the Great Depression. Those things stand out, but also impair your power law identification. As an empirical issue in agent models, this is important: are we seeing power laws? We need to perform well-defined hypothesis testing on this issue.

The Evolution of Norms

CONVENTIONAL CONTRACTS

P. YOUNG, The Brookings Institution*

[Reprinted with permission from *The Review of Economic Studies*.]

* *Corresponding author address*: Peyton Young, Economic Studies, The Brookings Institution, 1775 Massachusetts Ave. NW, Washington, DC 20036-2188; e-mail: pyoung@brook.edu.

DISCUSSION: THE EVOLUTION OF NORMS

Daniel Diermeier [to Peyton Young]: I'm Daniel Diermeier from Kellogg. Two questions. First, if you don't have an error, if it's a pure best response, then every norm is an attracting state, a stable state?

Peyton Young: Right, including inefficient norms.

Diermeier: And second, with respect to the empirical motivation: since you get different distributions across different populations —for example, in agricultural contracts — would you say that the set of feasible contracts is different in these populations? And is there empirical support for that? Say two-thirds/one-third or half/half: would that be driven by the feasible set in the Kolayez-Modinsky formulation?

Young: Yes. When we take this to data, the model has to be modified slightly to incorporate *a priori* focal effects. Half/half has a natural *a priori* focal property to it, so there may be some psychological bias toward certain — fractions, let's say. There's a way to rework the model to allow these built-in psychological biases toward particular contracts. In other words, they're not going to be generated endogenously. You build in the biases, and you ask, "What's the solution in an evolutionary sense?" and you can get answers.

Jonathan Bendor: Jon Bendor, Stanford. It's a very intriguing result, such a sharp result, with relatively spare assumptions. Can you give us some intuition for why, say, Proposition 1 holds?

Young: Yes. There's a natural intuitive argument, though probabilists in the audience will immediately see that it's not good enough. The idea is that diversity is being generated by idiosyncratic shocks. So if you're in a norm where the unattached state is not that unattractive compared to the current norm, it takes less idiosyncratic shock to move such a person into trying an experiment, into demanding something better.

It's actually very natural. People who are not doing very well vis-à-vis the unattached state don't have that much to lose. Even though rationally they shouldn't be demanding something better, we build in irrationality by calling it a random variable. I didn't assume a normal distribution for this variable, but you could. And so its variance would push a person into trying something else with higher probability than if, say, the norm were in the equal-division state, because then the random shock has to be much, much larger for it to be reasonable to try something else. That's effectively what's driving this result.

Michael North: Are you saying that people who have relatively little to lose are willing to take bigger risks?

Young: That's one possible interpretation. Now, it might not be a bigger risk per se; it might simply be that, given there's a certain amount of idiosyncratic variation in each stratum of the population, there will be a larger group that is predisposed to deviate from the norm. So it isn't taking a risk, necessarily; it's simply that this heterogeneity exists. We can interpret this epsilon as a crude measure of some heterogeneity that persists.

John Bower: I just want to throw an empirical problem at you. My father happens to be a farmer, and he happens to rent his land out to a factory that makes potato chips. Now, every farmer that rents his land out to this factory uses the same contract. It has nothing to do with norms — the factory's the only game in town; it sets the contract. It's also the same in the North Sea with the oil market. There are four big players. They set the contracts; they're the people with the dominant bargaining power. Norms don't apply. Isn't that what's going on here?

Young: No, but I'm delighted that you highlighted this point. In Illinois agriculture, for example, there are no equivalents of your potato chip factory, that is, not for the whole state. It's a highly decentralized market. There are hundreds of thousands of farmers with acreage between one and four or five hundred acres.

I choose these examples because you would expect heterogeneity in them. In your example you would not. Our theorem applies when it is a truly decentralized bargaining situation. So I don't think that you've given a counterexample; you've merely said there are contracting situations where one dominant party is contracting against a whole lot of other small parties. This model is not meant to address that situation.

Bower: Yes. Just one last comment. The point I was making was that when we see this, we shouldn't only ascribe it to the setting of norms. There are many cases where norm-setting is not past the gate.

Ian Lustick: Ian Lustick, University of Pennsylvania. I have a question about the concept of norm that you're using. You say it is essentially absolute uniformity. [Young: Yes.] But it's also the state that you are explaining or that you are predicting in your model. [Young: Yes.] That's an odd meaning of the word "norm" in political science or in most other areas, where we want norm to be doing something as an independent variable. It's not clear to me whether you insist on absolute uniformity. Something can be a norm even if there's a cloud of idiosyncrasy around it. So is there any effect that the uniformity or quasi-uniformity has on behavior, or is it a pure resolution of behavior, a reflection of behavior?

Young: Well, this is a deep question. What are the motivations of individuals for entering into particular kinds of contracts? I think this is a very interesting area. I think it's also an area that I'm not equipped to study. It would take someone like a rural sociologist or an ethnographer.

In the model, we're asserting that expectations at the individual level arise by observing the demands made by others in the population. As we are using the term, a "norm" is a kind of purified concept, something where all diversity has been driven out; yet, in running the model we'll never see that. We discuss the norm because it's that state around which probability is concentrated. Norms and near-norms are what survive when you run the dynamic process. But to define "near-norm" gets to be awkward, so I didn't do it.

Maurits van der Veen: Maurits van der Veen from the University of Pennsylvania. How reliant is your model on having exactly two types of actors? What if there were, say, three or four types, so you wouldn't know exactly which type you would see in the next round?

Young: If the bargaining is bilateral but there are more than two types, I'm confident that you could run these kinds of results with no difficulty. A trilateral bargaining situation would have some complications; there are some technical problems. The result may not survive without some additional assumptions.

Cultural and Normative Dynamics

SIMULATION OF THE LEARNING OF NORMS

H. VERHAGEN, The DECIDE Research Group,
Stockholm University, and the Royal Institute of Technology [*]

ABSTRACT

Multiagent system research tries to obtain predictability of social systems whilst preserving autonomy at the level of the individual agents. In this paper, social theory is called upon to propose a solution to this version of the micro-macro problem. The use of norms and learning of norms is one possible solution. An implementation of norms and normative learning is proposed and evaluated using simulation studies and measures for the internalizing and spreading of norms. The results of the simulations provide food for thought and further research. Modeling the norm set of the group in accordance with the individual mindset in absence of other information (i.e., at the beginning of the group-forming process) proves to be a more fruitful starting point than a random set.

INTRODUCTION

Multiagent system designers face the problem of how to ensure efficiency at the level of the multiagent system whilst respecting individual autonomy. This is in a sense the inverse of the problem of sociology, which tries to explain how social cohesion is possible in a world where individuals become more autonomous. Multiagent system development is like social engineering, with a focus on the reduction of behavior variance whilst maintaining the agent's autonomy. In this sense, multiagent systems research is focused on solutions to the micro-macro problem. I will discuss some theories from the social sciences that may help multiagent systems research in balancing individual autonomy and system efficiency. Issues of autonomy will not be discussed; for these, the reader is referred to previous writings, e.g., Verhagen (2000).

One possible solution to the problem of combining social level efficiency with autonomous agents is the use of central control (thus limiting the individual's autonomy severely). In human social systems such as organizations this is realized via bureaucracy. In multiagent systems it is the central coordinator that plays this role. This solution works only when the social system's environment has a low rate of change (including changes in the set of individuals included in the social system) since central control has as one of its main characteristics a low rate of adaptability (see Carley and Gasser [2000] for a discussion on the impossibility of an optimal organizational structure). When flexibility is of the essence, other solutions are called for. An intermediate solution is internalized control, e.g., the use of social laws (Shoham and Tennenholtz 1992). Structural coordination as proposed in Ossowski (1999) is another example of an intermediate solution that is only suitable for closed systems (or at least systems in which the behavior of new members has to conform to preconceived rules). Open

[*] C*orresponding author address:* Harko Verhagen, Electrum 230, SE-16440 Kista, Sweden; e-mail: verhagen@dsv.su.se; http://www.dsv.su.se/~verhagen/welcome.html.

systems (Gasser 1991) (with respect to the composition of the social system and the environment in which it is to function) require the most flexible solution. This involves a set of norms and learning at all levels, including the level of norms, based on reflecting upon the results of actions. During the lifetime of the system, norms evolve (or possible even emerge) to adapt to changes in the circumstances in the physical and social world.

NORMS IN SOCIAL THEORY

Habermas (1984) tries to synthesize almost all schools of social theory of the twentieth century into one framework. In the process, Habermas distinguishes four action models. Each action model makes presumptions about the kind of world the agents live in, which has consequences for the possibilities modes of rational action in that model. The model of interest for my present purposes is the normative action model. Habermas identifies the use of norms in human action patterns as normatively regulated action. “The central concept of complying with a norm means fulfilling a generalized expectation of behavior. The latter does not have the cognitive sense of expecting a predicted event, but the normative sense that members are entitled to expect a certain behavior. This normative model of action lies behind the role theory that is widespread in sociology” (Habermas 1984, p. 85).

This view is also in agreement with Tuomela (1995). Tuomela distinguishes two kinds of social norms (meaning community norms), viz. rules (r-norms) and proper social norms (s-norms). Rules are norms created by an authority structure and are always based on agreement-making. Proper social norms are based on mutual belief. Rules can be formal, in which case they are connected to formal sanctions, or informal, in which case the sanctions are also informal. Proper social norms consist of conventions, which apply to a large group such as a whole society or socioeconomic class, and group-specific norms. The sanctions connected to both types of proper social norms are social sanctions and may include punishment by others and expulsion from the group. Aside from these norms, Tuomela also describes personal norms and also potential social norms (these are norms that are normally widely obeyed but are not in their essence based on “social responsiveness” and that in principle could be personal only). These potential social norms contain among others moral and prudential norms (m-norms and p-norms, respectively). The reasons for accepting norms differ as to the kind of norms:

- Rules are obeyed since they are agreed upon.
- Proper social norms are obeyed since others expect one to obey.
- Moral norms are obeyed because of one’s conscience.
- Prudential norms are obeyed because it is the rational thing to do.

The motivational power of all types of norms depends on the norm being a subject’s reason for action. In other words, norms need to be “internalized” and “accepted.”

NORMS IN ARTIFICIAL AGENT SOCIETIES

The use of norms in artificial agents is a fairly recent development in multiagent systems research (c.f., e.g., Shoham and Tennenholtz 1992, Verhagen and Smit 1997, Boman 1999). Multiagent systems research uses different definitions of norms. In Conte and Castelfranchi (1995), the following views on norms in multiagent system research are described:

- Norms as constraints on behavior.
- Norms as ends (or goals).
- Norms as obligations.

Most research on norms in multiagent systems focuses on norms as constraints on behavior via social laws (c.f., e.g., Briggs and Cook 1995, Mali 1996, Shoham and Tennenholtz 1992). These social laws are designed off-line[1] and agents are not allowed to deviate from the social laws (except in the work by Briggs; see below). In this sense the social laws are even more strict than the r-norms Tuomela (1995) describes, which come closest to these social laws. The social laws are designed to avoid problems caused by interacting autonomous selfish agents, thus improving cooperation and coordination by constraining the agents' action choices. This view on norms is based on the view of norms as developed within game theoretical research such as Ullman-Margalit (1977). In Briggs and Cook (1995), agents may choose less restrictive sets of social laws if they cannot find a solution under a set of social laws, thus introducing a possibility for deviation. This approach is close to the approach in Boman (1999), where sets of norms are used by an artificial agent decision support system (pronouncer) to reorder decision trees, with the agent having the possibility to refrain from using the reordered decision tree. The reasons behind this are not further developed in Boman (1999), in contrast to Briggs and Cook (1995). However, the title of Briggs and Cook's article ("Flexible Social Laws") is deceptive: it is not the laws that are flexible but the way they are applied. The laws do not change, it is the agent who decides to apply them or not. The agent is only allowed to deviate from a social law if it cannot act. Thus, the authors deny that not acting can be a choice and disconnect the choice of applying a social law from more realistic reasons other than the possibility to act.

Work on cognitive grounded norms is conducted in the group around Castelfranchi and Conte (c.f., e.g., Conte and Castelfranchi 1995, Conte et al. 1999a, Conte et al. 1999b) or in research inspired by their work (c.f., e.g., Saam and Harrer 1999). In Conte et al. (1999a), norms are seen as indispensable for fully autonomous agents. The capacity for norm acceptance is taken to depend upon the ability to recognize norms, normative authorities, and on solving conflicts among norms. Since normative authorities are only of importance in the case of r-norms, the agents should also be able to recognize group members to be able to deal with s-norms Tuomela (1995). In Tuomela (1995), a theory solving conflicts among norms of different categories is developed that can complement the research described in Conte et al. (1999a). The origins of norms is not clarified in Conte et al. (1999a). However, the possibility of norm deviation is an important addition to multiagent systems research on norms.

[1] In Shoham and Tennenholtz (1997), social laws and conventions are not designed off-line but emerge at runtime. Social conventions limit the agent's set of choices to exactly one. The agents are not allowed to deviate from the social laws or conventions. Furthermore, a central authority forces agents to comply.

Norms are viewed in this article as internalized generalized expectations of behavior. The agent expects itself to follow the norms (in this sense, norms steer the agent's behavior). Other members of the group to which the norms apply are also expected to behave according to the norms. Norms can thus be used to predict the behavior of fellow group members. Deviation from norms is possible but is followed by rebuke from the other group members (see, e.g., Gilbert 1989, 1996).

The modeling of norms and the learning of norms is partially based on Boman (1999), which presented a general model for artificial decision making constrained by norms. In this model, agents adhere to norms via local adaptation of behavior or via groups exercising their right to disqualify action options. The adaptation of behavior consists of an internalization of group norms, or more precisely a synchronization of the individual behavior dispositions to those of the group. The learning of norms constitutes to an individual behavior pattern endorsed by the group and is thus the basis for socially intelligent behavior. The assessments in the information frames gradually evolve, in order for the agent to act in accordance with the norms of its group. The group norms, or social constraints, are not merely the union of the local information frames of its members, but rather develop interactively, as do the local information frames.

The use of norms as a mechanism for behavior prediction and control assumes at least two components: a theory of acceptance of norms (which is the focus of Conte et al. [1999a]) and a mechanism for the spreading and internalizing of norms. I will focus on the second component and in particular test the possible use of communication of normative advice as a way of spreading norms. As for the acceptance of norms, to reduce the complexity of the research topic, in this article I presume that if an agent is part of a group, it blindly accepts the norms of that group. The degree to which the group norms are applied in the agent's decision making is dependent upon the degree of autonomy the agent has with respect to the group, with autonomy meaning the freedom of choosing to not comply with the norm. In general the group does not condone this behavior, and sanctions may be applied. However, I will not discuss these mechanisms in the current article. I will study the influence of normative comments on previous choices. This contrasts with the norm-spreading mechanism in Boman (1999), where normative advice is sought before an agent makes its choice.

I will introduce the simulation model developed to test the usability of these concepts in multiagent systems and discuss the results obtained so far. After this I will indicate possible topics for further research.

SIMULATION OF THE SPREADING AND INTERNALIZING OF NORMS

The simulation model consists of several agents roaming a two-dimensional space. The agents form a group, with one of the agents acting as the leader. Every spot in the two-dimensional space may contain nothing, one piece of resource A, one piece of resource B, or one piece of both resources. The agent has a choice to do nothing, move to another spot, or take resource one or resource two (if available). Combining the number of content alternatives with the choice alternatives and outcome alternatives (whether the chosen alternative is realized or not) gives 20 combination alternatives in total. These alternatives are summed up in a so-called decision tree.

A decision tree is a general structure for summing up alternatives, probabilities of outcomes, and utility values of those outcomes. For example, if an agent finds itself in a spot with one item of resource A and one item for resource B, the agent has the following choices:

- Consume resource A.
- Consume resource B.
- Move to another spot.
- Do nothing.

Suppose the agent has almost full capabilities in consuming the resources (e.g., probability is 0.9), the probability of moving to another spot being successfully executed is less high (e.g., 0.8), and the probability for doing nothing is equal to one, and for all probabilities goes that they add up to one for all alternatives. Suppose also that consuming resource A has a utility value of 0.8 (and 0.3 for its counterpart), while consuming resource B has a utility of 0.4 (thus expressing that the agent prefers resource A over B), moving (M) has a utility of 0.6 (thus being preferred over resource B), and doing nothing (N) has a utility of 0.1. Given these conditions, the numerical representation of the decision tree for this situation is:

$$T_{AB} = (A, 0.9, 0.8, 0.1, 0.3,\ \ B, 0.9, 0.4, 0.1, 0.3,\ \ M, 0.8, 0.6, 0.2, 0.4,\ \ N, 1, 0.1, 0, 0)$$

Choosing one of the four alternative actions is solved via calculating the estimated outcome of each of the four alternatives and picking the one with the highest outcome. Estimated outcome is the product of the probability of the alternative working out times its utility minus the product of the probability of its counterpart times its utility.

Thus in this situation, the following four estimated outcomes can be calculated:

- Consume resource A has an estimated outcome of $(0.9 \times 0.8) - (0.1 \times 0.3) = 0.69$.
- Consume resource B has an estimated outcome of $(0.9 \times 0.4) - (0.1 \times 0.3) = 0.33$.
- Moving to another spot has an estimated outcome of $(0.8 \times 0.6) - (0.2 \times 0.4) = 0.4$.
- Doing nothing has an estimated outcome of $(1 \times 0.1) - (0 \times 0) = 0.1$.

Consequently, the agent chooses to try to consume resource A.

Description of Decision Making Model

Every agent has a private decision tree containing its personal evaluations (self model) and a group decision tree containing the evaluations the agent presumes for the group (group model). The group model expresses the agent’s interpretation of the norms the group holds. When faced with a decision situation, the agent combines both trees to one single tree that can be evaluated in the manner as described above. In this combination of decision trees, the degree of autonomy of an agent relative to the group determines what weight the group model and self model are given respectively. The resulting decision tree is thus calculated as follows:

for each consequence c in $T_{s,g}$

$$\frac{c_n = c_s \times a + c_g (1-a)}{2}$$

FIGURE 1 Function for creating the mixed decision tree: T = decision tree, s = self model, g = group model, n = new decision tree, c = consequence, a = autonomy value.

After making its choice, the agent announces its choice situation and choice to all other group members and executes its choice. The other group members will send feedback to the agent, consisting of their self model for that situation. The group model is updated every time n messages are communicated by the other agents. This memory size can also be varied. The feedback of the leader of the group weighs heavier than the feedback of the other agents. This is set via the leadership value, with a leadership value of ten expressing that the leader's feedback is interpreted ten times. The group model is updated via:

$$g(s,a,c) = \frac{\sum_{n=1}^{n} f(s,a,c)}{n}$$

FIGURE 2 Function for changing group model: g = group model, s = situation, a = alternative, c = consequence, n = amount of feedback messages considered, f = feedback chunk.

The agent's self model is updated based on the outcome of the choices (i.e., feedback from the environment).

Implementation of the Model

The simulation model is implemented in Java, with each agent being a separate thread. The agents communicate with the environment and each other through a router programmed using JATLite. Varying the settings for the agents requires editing of some data files and Java code files and (re)compiling these. All simulation runs were run for 100 minutes (during test runs this proved to be adequate for the system to reach equilibrium) and repeated 6 times for each setting. During the simulation run, a log file is kept for each agent and is updated every minute with the agent's self model and coalition model at that point in time. These are read into Excel and analyzed offline.

Simulation Setups

The following simulations were run: the leadership factor varied between 1, 2, 5, and 10; autonomy (on a scale from 0 to 1) had a value of 0.0, 0.4, or 0.8; and the initial group model was either set to a default model (equal for all agents) or equal to their self model (different for all agents). This gives six simulation setups in total. The following three hypothesis were formulated:

1. The higher the degree of autonomy, the lower the predictability of behavior will be.

2. The higher the leadership value, the higher the predictability of behavior will be.

3. If the personal decision tree equals the initial group decision tree, the predictability of behavior will be higher compared to an initial random group decision tree.

The variance of behavior can be measured in several ways. One way is by determining the difference between an agent's self model and its group model, expressing the internalizing of norms. This is calculated as:

$$\frac{\sum_{j=1}^{m}\left(\frac{\sum_{i=1}^{n}\left|s_i - g_i\right|}{n}\right)_{\mathbf{j}}}{m}$$

FIGURE 3 Norm internalization measure:
n = amount of alternatives, *m* = amount of agents,
s = self model value for alternative *i*, *g* = group model
for alternative *i*.

Another measure is the differences in the group models over the agents, expressing the spreading of norms. The spreading of norms is calculated via:

$$\frac{\sum_{j=1}^{m}\sqrt{\frac{\sum_{i=1}^{n}(g_i - \overline{g})}{n}}}{m}$$

FIGURE 4 Norm spreading measure: *n* =
amount of alternatives, *m* = amount of
agents, *g* = group model for alternative *i*,
$\overline{g}$ is the mean group model over the
agents.

A higher variance of behavior is thus a higher norm-spreading factor (the agents are closer to each other with respect to their vision on the norms of the group) and a lower norm-internalizing factor (the agents have a higher difference between their self model and their group model).

In figures 5 and 6, the first hypothesis is tested. The hypothesis that a higher autonomy value results in lower behavior predictability holds for the spreading of norms (Figure 5), but not for the internalizing of norms (Figure 6). The graphs testing hypothesis 2 (a higher leadership value leads to a higher predictability of behavior) show the same result (Figures 7 and 8). Varying the strategy of choosing an initial group model, as formulated in hypothesis 3 and tested in Figures 9 and 10, shows a different result. Here, the hypothesis holds for nearly all simulations, the only exception being the norm spreading in the case in which the autonomy factor is 0.8 and the leadership factor is 4.

DISCUSSION AND FUTURE RESEARCH

The formulated hypotheses did not hold for all measures and simulations run. Several factors may play a role here. One possible cause is that not all situations occur during the simulation. Since the normbases are only updated for the situations that occur, some utilities do not change during the entire simulation. A second explanation may be that the variance between different runs with the same setting could be greater than the difference between runs with different settings. A third possible explanation is that a norm spreading and norm internalizing are measures of learning on different levels. Further simulation experiments will be conducted to draw conclusions about why the norm spreading factor does not follow with the first two hypotheses. Another topic for future research will be the formation of groups.

ACKNOWLEDGMENTS

The author wishes to thank Magnus Boman and Johan Kummeneje for the inspiring discussion and work on joint articles, without which this article would not have existed. Special thanks to Stefan Fägersten for developing the simulation model as part of his master’s thesis work.

This work was sponsored in part by NUTEK through the project Agent-Based Systems: Methods and Algorithms, part of the PROMODIS program.

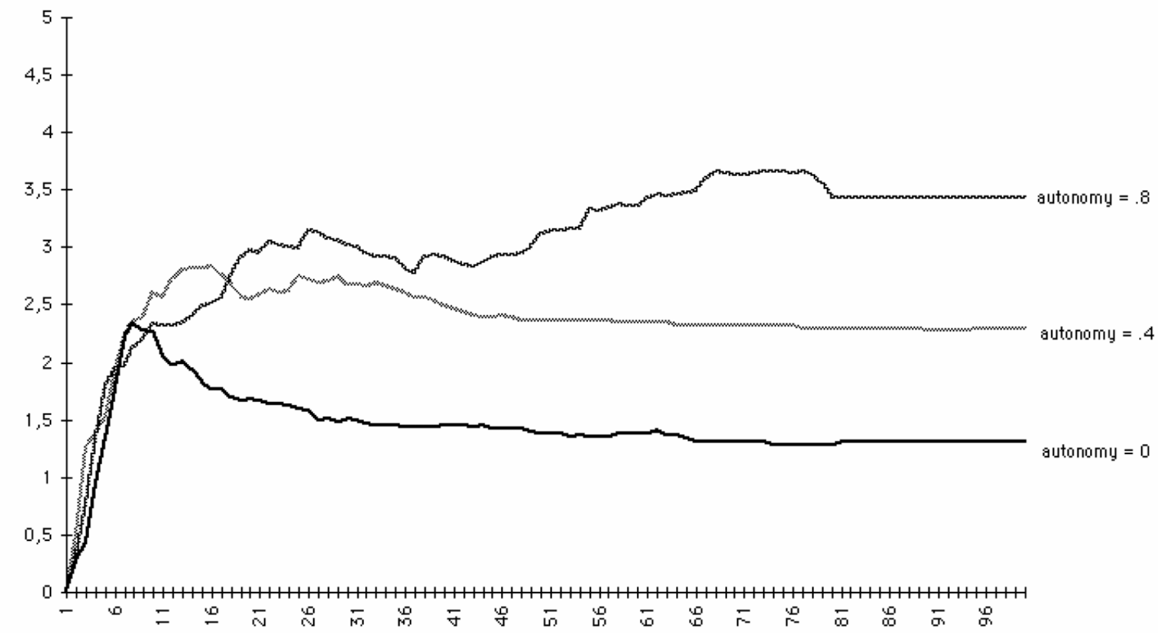

FIGURE 5 Norm-spreading factor varying autonomy with leadership factor = 4 and initial group model = default model

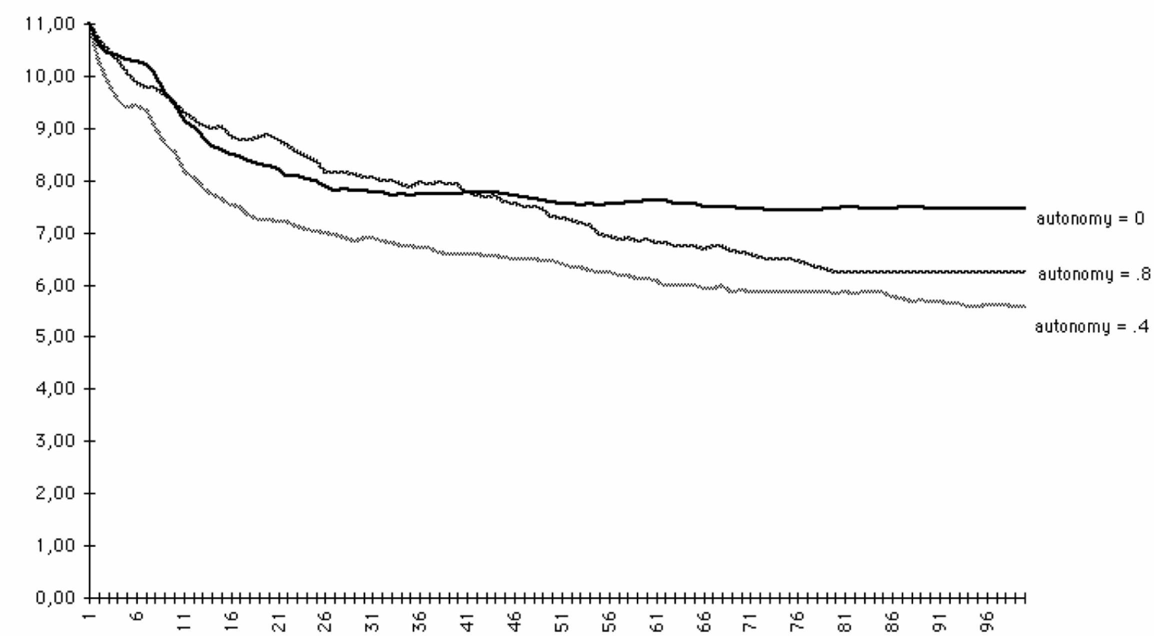

FIGURE 6 Norm-internalizing factor varying autonomy with leadership factor = 4 and initial group model = default model

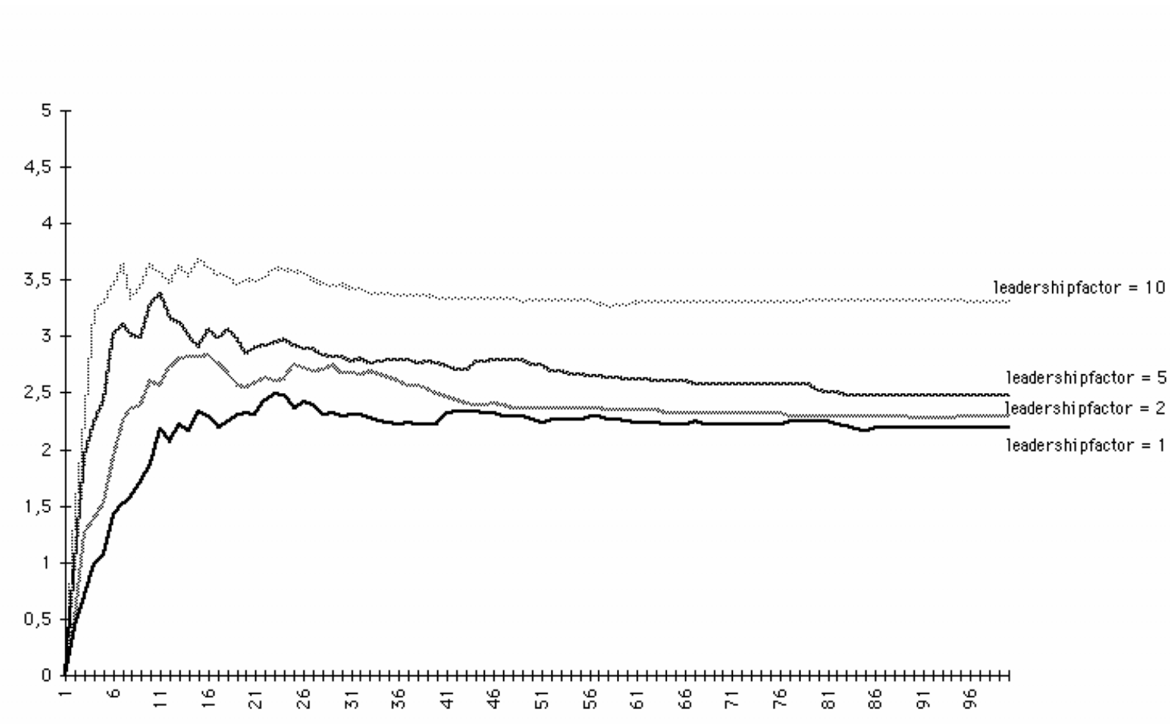

FIGURE 7 Norm-spreading factor varying leadership factor with autonomy = 0.4 and initial group model = default model

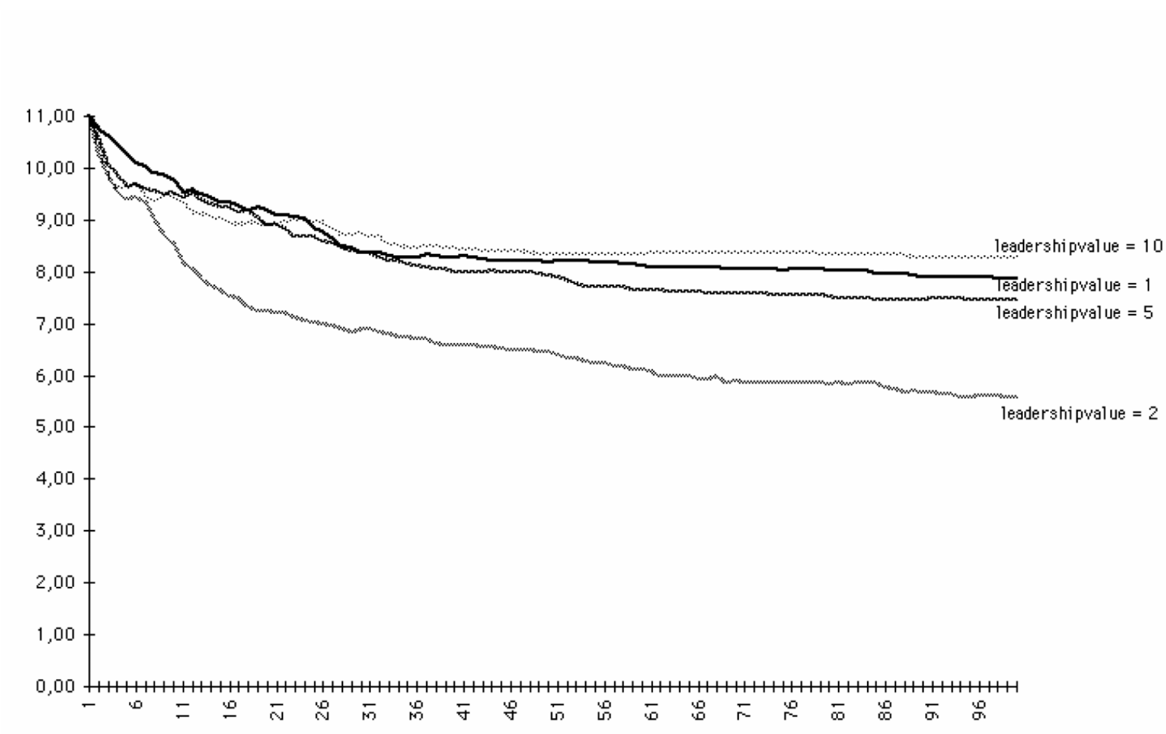

FIGURE 8 Norm-internalizing factor varying leadership factor with autonomy = 0.4 and initial group model = default model

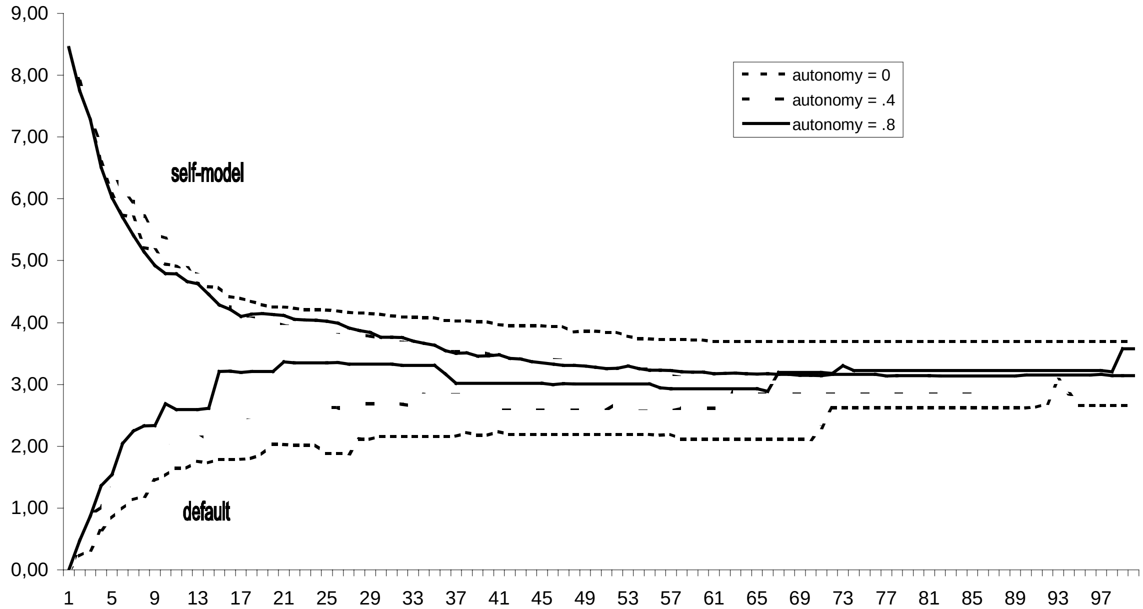

FIGURE 9 Norm-spreading factor varying autonomy and group model with leadership factor = 4

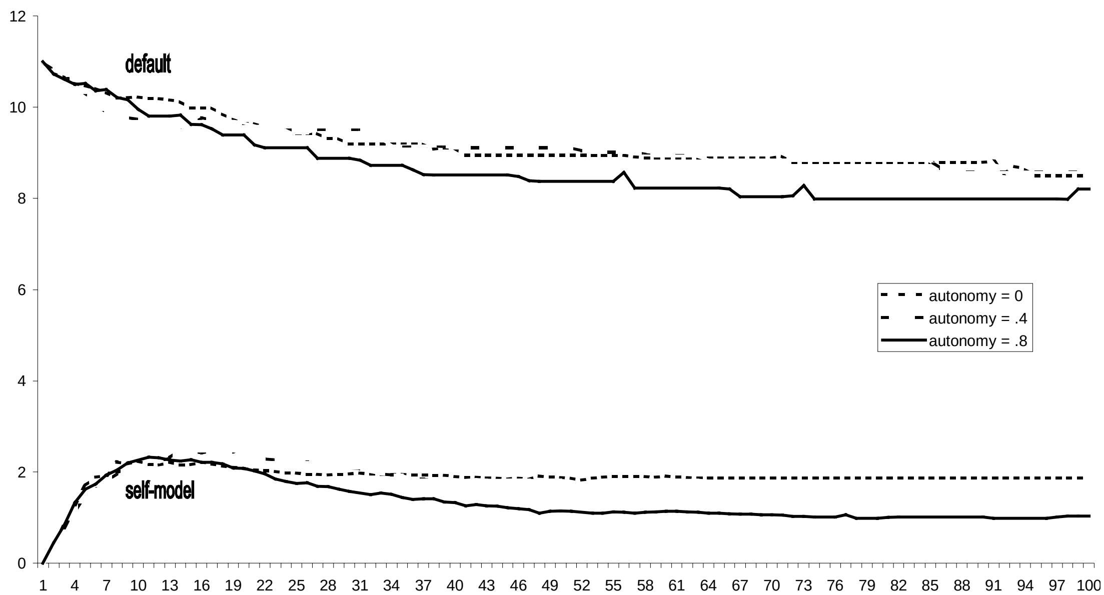

FIGURE 10 Norm-internalizing factor varying autonomy and group model with leadership factor = 4

REFERENCES

Boman, M., 1999, "Norms in Artificial Decision Making," *Artificial Intelligence and Law*, 7 (1):17-35.

Briggs, W., and Cook, D., 1995, "Flexible Social Laws," pp. 688-693 in *Proceedings of the 1995 International Joint Conferences on Artificial Intelligence*, Morgan Kaufmann.

Carley, K.M., and Gasser, L., 1999, "Computational Organization Theory," pp. 299 - 330 in G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press.

Conte, R., and Castelfranchi, C., 1995, *Cognitive and Social Action*, UCL Press, London.

Conte, R., Castelfranchi, C., and Dignum, F., 1999a, "Autonomous Norm-acceptance," pp. 99-112 in *Intelligent Agents V: Agent Theories, Architectures, and Languages*, Springer-Verlag, Berlin.

Conte, R., Falcone, R., and Sartor, G., 1999b, "Introduction: Agents and Norms: How to Fill the Gap?" *Artificial Intelligence and Law*, 7 (1):1-15.

Gasser, L., 1991, "Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems Semantics," *Artificial Intelligence*, 47 (1-3):107-138.

Gilbert, M., 1989, *On Social Facts*, Routledge.

Gilbert, M., 1996, *Living Together: Rationality, Sociality and Obligation*, Rowman and Littlefield.

Habermas, J., 1984, *The Theory of Communicative Action, Volume One, Reason and the Rationalization of Society*, Beacon Press, Boston. Translation McCarthy, originally published as *Theorie des Kommunikativen Handelns, Band 1 : Handlungsrationalität und gesellschaftliche Rationaliserung*, 1981, Suhrkampf.

Mali, A.D., 1996, "Social Laws for Agent Modeling," pp. 53-60 in M. Tambe and P. Gmytrasiewicz, editors, *Agent Modeling Papers from the AAAI Workshop*, AAAI Press.

Ossowski, S., 1999, *Co-ordination in Artificial Agent Societies*, Springer-Verlag, Berlin.

Saam, N.J., and Harrer, A., 1999, "Simulating Norms, Social Inequality, and Functional Change in Artificial Societies," *Journal of Artificial Societies and Social Simulation*, 2(1), electronic journal. On-line at http://www.soc.surrey.ac.uk/JASSS/2/1/2.html.

Shoham,Y., Tennenholtz, M., 1992, "On the Synthesis of Useful Social Laws for Artificial Agent Societies (Preliminary Report)," pp. 276-281 in *Proceedings of the National Conference on Artificial Intelligence*, San Jose, CA.

Shoham, Y., and Tennenholtz, M., 1997, "On the Emergence of Social Conventions: Modeling, Analysis, and Simulations," *Artificial Intelligence*, 94 (1-2):139-166.

Tuomela, R., 1995, *The Importance of Us: A Philosophical Study of Basic Social Norms*, Stanford University Press.

Ullman-Margalit, E., 1977, *The Emergence of Norms*, Clarendon Press.

Verhagen, H.J.E., and Smit, R.A., 1997, "Multiagent Systems as Simulation Tools for Social Theory Testing," presented at the International Conference on Computer Simulation in the Social Sciences, Sienna, Italy; available in Verhagen, H., 1998, *Agents and Sociality*, Ph.D. thesis, Stockholm University.

Verhagen, H.J.E., 2000, *Norm Autonomous Agents*, Ph.D. thesis, Department of System and Computer Sciences, The Royal Institute of Technology and Stockholm University, Sweden.

EMPIRICALLY TESTING A COMPUTATIONAL MODEL: THE EXAMPLE OF HOUSING SEGREGATION

R.H. SANDER, School of Law, UCLA*
D. SCHREIBER, Department of Political Science, UCLA
J. DOHERTY, Empirical Research Group, UCLA

ABSTRACT

Thomas Schelling's famous model of housing segregation started with a few coins on an eight-by-eight grid and some very simple assumptions about individual preferences. Using the SWARM programming environment, we have extended Schelling's concept to examine the contemporary debate about the nature and causes of housing segregation. We begin with basic preferences functions derived from empirical data on neighborhood racial composition and add a variety of putative factors in housing decisions. The result is a computational model of racial housing segregation that provides insight into empirical patterns of segregation and desegregation in the late twentieth century.

INTRODUCTION

Thirty years ago, the economist Thomas Schelling advanced a theory to explain the persistence of racial segregation in an environment of growing racial tolerance (Schelling, 1971, 1978). Schelling posited a simple model that made a straightforward point: if the racial makeup of one's neighbors is a decisive factor in choosing housing, then the collective interaction of individual preferences will tend to produce segregation, even if many individuals tolerate or even prefer integration.

Schelling's use of "micromotives" to explain "macro" phenomena has become a familiar concept, but for many years it did not advance very far as a practical tool for studying segregation. Mathematically, it is much easier to analyze the aggregate behavior of individuals in market models, in which everyone is engaged in the same transaction, than in Schelling's "tipping" model, in which individuals react to their local environment rather than an aggregated market. But the spread and increased accessibility of computational modeling makes it possible to experiment with much more complex models, incorporating both neighborhood-based racial preferences and market-based housing choices into a more complete representation of urban forces.

In this paper, we describe our development of a computational model of housing segregation over the past year. We particularly focus here on a problem that has often bedeviled model-builders: how to "test" models with empirical data. We believe our experiments show some progress towards the goal of rigorously testing alternate theories of segregation.

---

* *Corresponding author address:* R.H. Sander, UCLA School of Law, Box 951476, Los Angeles, CA 90095-1476; e-mail: Sander@law.ucla.edu.

The Paradox of Modern Housing Segregation

Housing segregation — particularly of African Americans — continues to be a dominant feature of most American cities, and it has been linked to a wide range of urban ills (Massey and Denton, 1994). Scholars in the field generally agree that income differences between blacks and whites explain only a small fraction of current segregation levels (Muth, 1986; Clark, 1986); but beyond this, there is no consensus and little in the way of convincing evidence to demonstrate why housing segregation has remained so high.

As an historical matter, it is obvious and widely agreed that black housing segregation came about through organized, mostly private efforts to ghettoize blacks in the early twentieth century — particularly the years between the world wars (Sander, 1988). One can chart the rise of segregation through the index of dissimilarity, which measures the proportion of one group (e.g., blacks) who would have to move to different neighborhoods to achieve the same metropolitan distribution as a second group (e.g., whites). Where an index of 1.00 equals complete apartheid, black/white segregation rose in American cities from an average of 0.6 in 1910 to 0.90 in 1940.

Intense levels of housing discrimination persisted throughout metropolitan America during the 1940s and 1950s, and it is thus unsurprising that segregation remained extremely high into the 1960s. But in the late 1960s, things changed rather dramatically. The civil rights movement brought a striking change in the answers whites gave to survey questions about integration; many more whites expressed marked levels of tolerance. The same movement brought the Civil Rights Act of 1968, which outlawed a wide range of discriminatory conduct in housing markets. Moreover, some real changes in American cities followed in the 1970s. Measured levels of housing discrimination fell sharply by 1977 (Wienk et al; 1979). Black migration to previously white housing accelerated (Sander, 1998). Although inner-city poverty remained acute, the black middle class grew and in many ways flourished.

Given these changes, many observers predicted rapid declines in housing segregation. Generally, however, the declines were almost imperceptible. The average black/white index of dissimilarity across a range of cities fell from .88 to .81. Most of the decline was concentrated in a fraction of cities; in many of the nation's largest cities, the black/white index had fallen only a couple of points by 1990 (Sander, 1998).

This, then, is the central paradox of segregation: why have levels of black/white separation changed so little when so many related factors have changed a lot? Related to the paradox are a number of smaller but important puzzles: why has housing in predominantly black areas "flipped," from being more expensive than white housing prior to 1970 to being less expensive afterwards? Why are Latinos, who experience levels of discrimination similar to those experienced by blacks, and who have about the same incomes as blacks, much less segregated from whites? And why have a few areas (e.g., Santa Clara County in California or Seattle, Washington) experienced sharp declines in segregation?

Scholars of segregation have not been able to provide robust theories that account for these related puzzles. While an extended discussion of the literature is beyond the scope of this paper, we believe that there are two central problems: a paucity of attempts to rigorously specify testable hypotheses derived from alternative theories of segregation and the difficulties noted earlier in operationalizing Schelling's insights. We turned to computational modeling as a way of overcoming both problems.

Building a Computational Model of Segregation

Although Schelling described his segregation model as a "thought experiment" and focused on simple numerical examples, it lends itself so readily to computational modeling that its has been a standard computational demonstration for many years. In these models, agents are arrayed on an open grid, with each interior agent surrounded by eight squares. Each agent is either black or white, and they are either "happy" or "unhappy," depending on whether no more than some critical threshold of their neighbors are persons of the other race. Some random squares are vacant, and in a series of rounds, each agent who is unhappy moves to a vacant square. The surprising report, as we noted earlier, is that for a wide range of preference levels, integrated neighborhoods will "tip" towards one group or another, leading the outnumbered group to flee and thus producing segregation.

We wanted to bring this simple model closer to the real world in two ways. First, rather than have agents be either "happy" or "unhappy," we wanted them to evaluate and compare a wide range of possible states. Second, we wanted racial preferences to be only one of many factors agents used to compare neighborhoods. We achieved both goals by creating a multivariate utility model through which agents are periodically asked to compare a range of neighborhoods, evaluate the overall utility they would achieve at each location (considering several variables), and decide if they would like to move.

Since we sought to explore the evolution of segregation after the civil rights revolution, we modeled our schematic city on a prototypical large American metropolis in 1970.[1] The city is biracial (with "reds" and "blues" standing in for "whites" and "blacks"), and the minority group is systematically clustered near the center of the city (that is, the initial index of dissimilarity is 1.00). In the models described in this paper, the population of the city is 2,375 (a number that seems to us, so far, large enough to capture the key dynamics of the simulation). In each period, a randomly selected tenth of the population is given the option of moving to one of five alternative locations (which are randomly selected from the entire city) or remaining at the agent's current location. The agent compares the sites using five different criteria: housing cost, distance of the new site from the present location, local discrimination, the racial makeup of the immediate neighborhood (the eight adjacent agents using Moore's definition), and the racial makeup of the surrounding community. Each factor is closely related to a competing theory of housing segregation. We operationalized each variable as follows.

Housing cost: Five percent of the city's cells are vacant at any time, and the 2,500 cells in the city at large are divided into 25 communities (demographically analogous to census tracts) each containing 100 cells. Popular communities have fewer vacant cells than unpopular ones. The housing price then becomes a simple inverse function of the vacancy rate. In actual runs of the model, neighborhood vacancies range from zero to over 20%.

Distance from present home: We calculate the distance of each of the five sites to which an agent might move from the agent's present location, using the Pythagorean theorem. We assume that, other things being equal, utility monotonically declines with the distance an agent

[1] The federal Fair Housing Act went fully into effect on January 1, 1970, and the decennial census conducted in April 1970 thus captures nicely the state of segregation as society embarked on its experiment in housing desegregation. Please note that though we use the term "city" to describe our model, we envision it as a representation of a metropolitan area, in which the movement choices of agents are limited to the region of the model.

moves, because we also (simplistically) assume that an agent's relatives, job, friends, and church are close to its current location.

Discrimination: In our model, some of the majority-group (red) agents discriminate against minority-group (blue) agents. We capture this by allowing a specified proportion of red agents to impose a utility cost on adjacent squares if they are occupied by blues.

Racial composition of neighbors: Survey data show substantial variation in what blacks say is their "ideal" neighborhood racial mix; the same is true of whites. We therefore created six different utility functions (three for each race) in which each possible neighborhood racial makeup is characterized by a unique utility for the agents. In evaluating moves, the agent calculates the racial makeup of its Moore neighborhood and considers the associated utility.

Racial composition of community: In much the same fashion, agents evaluate the overall racial makeup of their current and alternative communities. Again, we define the community in tracts of 100 agents. Agents use the same utility functions they apply to their Moore neighbors, but to larger groups of neighboring agents.

With these combined factors, we generate an overall utility function like this:

$$Utility = w_1 * neigh\_pref + w_2 * tract\_pref + w_3 * \left(\frac{occupancy_i}{totalhousing_i}\right) + w_4 * distance + w_5 * discrimination$$

Note that each element of the equation has an associated weight. We can thus investigate a wide variety of model specifications by changing only a few values. For instance, we can set w_1 to 1.0 and all the others to 0.0 and run a model akin to Schelling's original (only the race of immediate neighbors matters). The model's flexibility turns out to be crucial in generating empirically testable predictions.

THE BASIC MODEL

We sought, as noted, to create a simulation patterned after a prototypical city in 1970. In our basic model, the minority-group blues, comprising 20% of the overall population, are clustered in an inner-city ghetto. As in most 1970 cities, the black district is more crowded and has higher housing prices (controlling for quality) than similar white neighborhoods (thus, the model starts with no vacancies in the ghetto). We used real-world survey data from the 1960s and 1970s to model the racial preferences of reds and blues (Pettigrew, 1973; Farley et al., 1979) bearing in mind that survey data do not necessarily measure true preferences. We tried to capture the diversity of real-world preferences by assigning one-third of the agents of each race a set of preferences that roughly described the most tolerant third, least tolerant third, or middle third of real-world preferences within each race. Finally, we experimented with varying levels of discrimination to simulate that measured in the first large-scale studies of discrimination after passage of the Fair Housing Act (Wienk et al., 1979).

From this starting point, we run the model through a sequence of time periods. In each period, one-tenth of the agents compare their current location with four randomly selected alternatives; agents move to whatever location from the available options maximizes their utility. A variety of programmed metrics monitor, over time, various indicia of the city's condition: the level of red/blue segregation, the proportion of agents choosing to move; the utility of the average red or blue agent, and so on.

As a general matter, this "basic model" behaves in a way that broadly mirrors the typical evolution of large American cities after 1970. A significant number of the most tolerant blacks take advantage of lower discrimination levels to move into white neighborhoods; a smaller number of the most tolerant whites migrate towards the inner city. As the black population expands into white areas (mostly areas adjacent to the existing ghetto), the price of housing in mostly black areas falls relative to housing costs in mostly white areas. Segregation falls modestly, but significantly.

PUSHING THE MODEL: PRIMITIVE EMPIRICAL TESTS

We were delighted to create a complex computational model that could mimic the real world in some simple ways. But this achievement, by itself, did little to solve the puzzles of housing segregation with which we began. How could we reliably determine which elements of our model drove particular results? How could we use the model to compare theories?[2]

A related problem is the choice of parameter weights. Although we could model individual variables (e.g., racial preferences and housing prices) with close attention to real-world data, the *relative* weights given to the variables in our model were intrinsically arbitrary. One might, through the empirical housing literature, get some idea of how distance interacts with price, but how much relative weight might migrants give to price, racial makeup, and discrimination?

It is of course possible to get an intuitive feeling about model dynamics by changing parameters: if one increases the frequency of discrimination, for example, the runs of the model tend to produce less desegregation. The challenge is to make this experimentation systematic enough to produce clear and measurable results.

As of this writing, we have made modest but interesting progress on this front. Essentially, we have adopted the method of systematically varying parameters, measuring specific output variables, comparing input variation with output variation, and then comparing model results with empirical, demographic data. Here we give, very briefly, two examples of this technique.

Income Inequality and Segregation

The most common lay explanation of housing segregation blames income inequality — the inability of blacks to afford housing in white communities. Most segregation scholars discount this explanation, pointing to the substantial overlap in black and white income distributions and a range of data showing that different economic clusters of blacks experience

[2] This is a common curse of computational modeling: because the models are designed to illustrate the complexity of particular kinds of interactive phenomena, it is difficult to rigorously derive and "prove" particular results from the model. In general, if proofs could be derived from the model specifications, the computational model itself would be superfluous.

nearly identical levels of segregation (Muth, 1986; Farley, 1986). Nonetheless, there has been surprisingly little theoretical or empirical work explaining exactly why this should be so.

Within our model, it is possible to vary "income" across agents by simply varying the weight different agents give to housing cost. Our assumption is that affluent agents can make housing choices with less regard to housing cost than poor agents. Using this idea, we created three different simulation environments: (1) a world where all agents have the same income; (2) a world where agent incomes vary along a uniform distribution, with the maximum income three times the minimum income; and (3) a world like the second one, except that minority-group (blue) incomes vary across a narrower range and are, on average, about 20% lower than red incomes. We ran simulations of each world 100 times and measured the index of dissimilarity after each run. The results are shown in Figure 1.

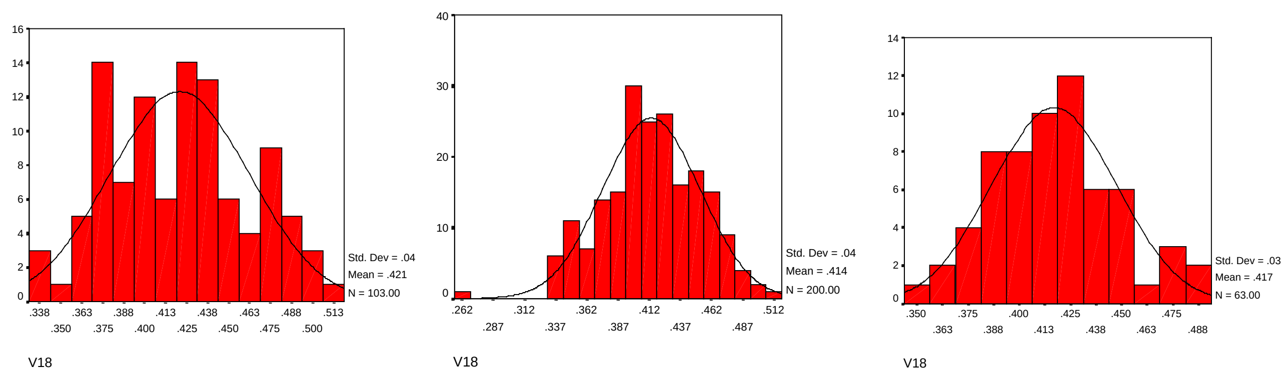

Housing cost = 50
mean dissimilarity = .42

25 < Housing Cost < 75
mean dissimilarity = .41

Blue/Red economic disparity
mean dissimilarity = .42

FIGURE 1 Economic differences do not drive segregation in the model

The basic result is clear. In all three models, the average dissimilarity level (over 100 model runs) that exists after 100 time periods is nearly identical (0.41 or 0.42). The income configuration does not seem to matter. This is so even when the housing price (and implicitly, income) variable is itself playing a critical role in the model.

Neighborhood Identity and Segregation

Another theoretical view of segregation emphasizes the role of neighborhood boundaries (Sander, 1998). It suggests that urban residents are sensitive not only to the racial makeup of their immediate neighborhood, but also of broader boundaries — of ethnic enclaves, school districts, and suburbs — that have crucial significance for some aspect of their social or political life. These boundaries are more likely than other locations to be "defended" by some majority-group members through discrimination and to be avoided by minority-group members.

In our model, agents consider the racial composition of their immediate Moore neighborhood and their broader (tract) community. We explored the role of neighborhood boundaries by systematically varying the relative weight that agents gave these two definitions of community. As Figure 2 illustrates, we found that the potential desegregation in the model city declines dramatically when the tract weight passes some threshold. At that critical point, tracts as a whole tend to become segregated and remain so. As in many real-world cities, an ordinary street can become a long-term placeholder between white and black communities if the boundary denotes some well-established distinction by which one of the separated communities defines itself.

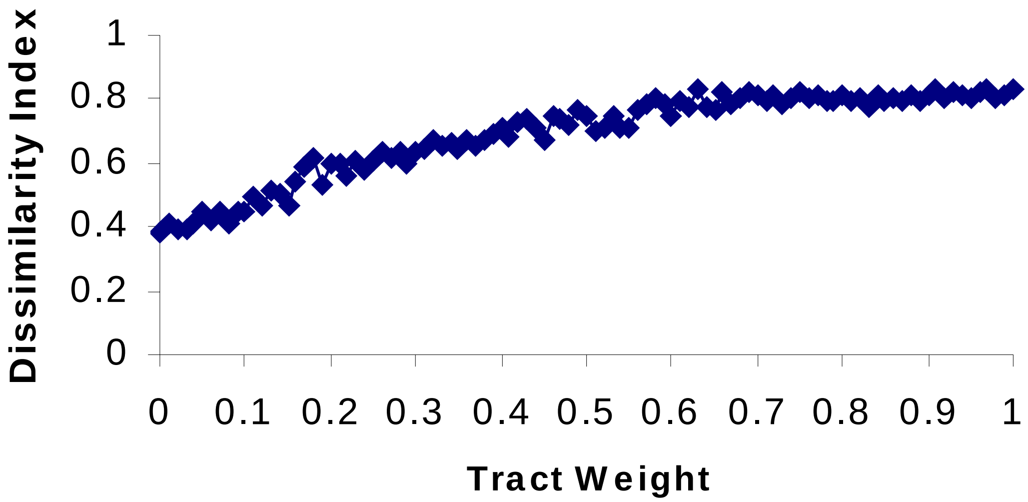

FIGURE 2 Stronger geographic borders increase segregation in our model

CONCLUSION

The computational modeling of housing segregation is still in its infancy. However, we believe that our first steps have been fruitful. A utility-based model seems to work effectively to combine many real-world factors shaping housing and neighborhood choice. A manageable agent space and iterated time sequences produce computational runs that strongly resemble actual urban dynamics. And systematic variation and outcome measurement allow us to develop data that can be directly analyzed and compared with empirical measurements of the real world. We are hopeful that, in time, these methods will make it possible to test robustly and modify competing theories of housing segregation.

REFERENCES

Clark, W.A.V., 1991, “Residential Preferences and Neighborhood Racial Segregation: A Test of the Schelling Segregation Model,” *Demography* 28:1.

Farley, Reynolds, et al., 1979, “Barriers to the Racial Integration of Neighborhoods: The Detroit Case,” *Annals of the American Academy* 441:97.

Farley, Reynolds, 1986, “The Residential Segregation of Blacks from Whites: Trends, Causes, and Consequences,” *Issues in Housing Discrimination*, U.S. Commission on Civil Rights, pp. 14-28.

Massey, Douglas, and Denton, Nancy, 1994, *American Apartheid*.

Muth, Richard F., 1986, "The Causes of Housing Segregation," *Issues in Housing Discrimination*, U.S. Commission on Civil Rights, pp. 3-13.

Pettigrew, Thomas, 1973, "Attitudes on Race and Housing: A Social-Psychological View," *Segregation in Residential Areas*, National Academy of Sciences, p. 25.

Sander, Richard H., 1988, "Individual Rights and Demographic Realities: The Problem of Fair Housing," *Northwestern University Law Review* 82:874-928.

Sander, Richard H., 1998, "Housing Segregation and Housing Integration: The Diverging Paths of Urban America," *University of Miami Law Review* 52(4):977-1010.

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

Schelling, Thomas C., 1978, *Micromotives and Macrobehavior*. New York: W.W. Norton.

Wienk, Reid, Simonson and Eggers, 1979, *Measuring Discrimination in American Housing Markets: The Housing Market Practices Survey*, U.S. Department of Housing and Urban Development.

DYNAMICS OF STATUS SYMBOLS IN HIERARCHICAL SOCIETIES

R. PEDONE and R. CONTE, National Research Council, Rome*

INTRODUCTION

Among both laymen and (a good share of) social scientists, it is commonly believed that social agents are strongly characterized by a built-in inclination to imitate or follow given others. The attractor-follower is considered one of the most important and frequent patterns of social life and is usually attributed a strong influence on the emergence of social institutions.

An alternative hypothesis about the emergence and maintenance of social hierarchies was suggested by a famous sociologist, Georg Simmel (*The Fashion*, 1957). In his view, agents are not only inclined to follow but also to distinguish themselves from others. He described "fashion" as a combined effect of both attitudes. What are the effects of the interplay between these two distinct inborn social attitudes? When social groups are ordered by rank, agents imitate symbols designating the higher hierarchical levels and abandon those designating the lower-level ones. As a consequence of this dynamics, the status symbols will spread rapidly through the population, proceeding downward from the higher to the lower levels of the social hierarchy. However, as soon as they spread, these symbols will be abandoned and replaced with new ones. Thereby, the "Simmel effect" allows for the social hierarchy to persist under the variability and instability of its status symbols.

A previous study (Pedone and Conte, 2000) aimed to simulate a spatial version of the Simmel effect. Two social attitudes were implemented — i.e., to minimize physical distance from higher-level agents and to maximize the distance from lower-level ones — in a population of autonomous agents scattered on a two-dimensional torus (i.e., a grid with virtually united edges). The underlying intuition was that to imitate social attractors gives only a partial account of the requirements for social differentiation and that agents are not only, nor mainly, characterized by a tendency to imitate, but also by an inclination to distinguish themselves.

PREVIOUS STUDY

The main objective of the study was to observe the emergence of hierarchically homogeneous clusters as an effect of the dynamics of the features designating the agents' status (status symbols). In particular, two rules were implemented: imitation of higher-level agents and avoidance of lower-level agents. While imitation alone was expected to increase convergence of the agents on the same symbols, social avoidance alone was expected to reduce such convergence. Only their combination was expected to produce a segregating effect, with agents belonging to the same hierarchical level sharing the same symbols of status.

The simulation model was based upon a spatial metaphor of the Simmel effect, in which the status symbol was identified with the agents' location on a two-dimensional grid representing

* *Corresponding author address*: Roberto Pedone, Institute of Psychology, National Research Council, Viale Marx 15, 00137 Rome, Italy; e-mail: conte@ip.rm.cnr.it.

the environment of the simulation. In the simulations, therefore, imitation was implemented as approaching higher-level agents, while social avoidance was implemented as moving away from lower-level agents. The effects of these rules have been observed in several experimental conditions, with or without intersection between followers and fugitive, with or without social barriers in the hierarchy that limit the agents' social perceptions, and whether imitation dominates avoidance or the other way around. Generally speaking, results confirmed our expectations. The Simmel effect takes place when both rules coexist. Furthermore, the effect is reinforced by but not conditioned to social barriers. In the condition of intersection, clusters form from all social levels rather than from the highest level only. Finally, the clustering is more visible when imitation dominates avoidance than with the opposite pattern, and this is consistent with the intuition that avoidance has a desegregating effect.

Results and Discussion

These findings seem to suggest that a simulation study of the Simmel effect is not only feasible but also informative. First, such a study could help to highlight minimal conditions for social segregation. Second, it questions the oversimplified view, held by social as well as multi-agent scientists, that a social agent is essentially ruled by imitation; instead, it calls for a more complex view of social agents, where imitation is combined with social avoidance. Third, it shows that a simulation study of the interplay between social and cultural processes is feasible and useful. Cultural and social processes interact without necessarily presenting the same qualitative effects. Indeed, cultural dynamics may coexist and even favor the maintenance of social status quo. Fourth, a simulation study of the Simmel effect has a memetic (Dawkins, 1976; Blackmore, 1999) side effect: it may help clarify some issues concerning the transmissibility of memes (i.e., units of culturally transmissible information). In particular, in the phenomenon under study, the fecundity of memes seems to ultimately lead to their extinction, a finding which is in apparent contrast with the memeticists' thesis that memes are self-replicating units.

PRESENT STUDY

A more realistic metaphor of the Simmel effect is under implementation. In particular, a clearly symbolic representation of status features is preferred to a spatial representation. Such a representation is more congruous with a study of status symbol dynamics. Two specific processes are under investigation:

- Social dynamics of status symbols, that is, how they move along the social hierarchy. Here, no substantial difference is expected from the previous study.

- Cultural dynamics: how symbols change (get corrupted) as long as they move on the social hierarchy, as a function of agents' individual and social characteristics.

Hypothesis

The cultural dimension of status symbol dynamics is certainly a worthwhile object of investigation in its own right. However, it is expected to have significant and complex interrelationships with the social dimension. In particular, our hypothesis is that corruption of status symbols is a direct function of social influence. The wider the influence of given symbols, the more they get corrupted and the sooner they will fade away. The more limited the range of influence of given symbols, the lesser their degradation and decay. In turn, social influence is a direct function of social status: the higher the status, the wider and stronger its influence (measured in terms of how many agents are likely to be influenced). Therefore, the higher the status, the more its symbols will be corrupted and decay. In principle, the more complex the social hierarchy, and the stronger and faster the social and cultural dynamics of status symbols (i.e., the more easily corrupted they will get, the sooner they will decay, and the earlier they will be replaced). In this sense, the degree of corruption of a given status symbol is expected to be an indicator of the social distance it went though.

Model Discussion

The final paper will present a set of experimental simulations in which agents distributed on a social hierarchy are characterized by complex symbols designating their status. Agents' perception of others' status symbols varies as a function of both their personal and social endowments, with lower-status agents suffering from more severe limitations in perception. Agents' capacity to reproduce others' status symbols, instead, varies only with their status, with lower-status agents suffering from more severe limitations in reproductive capacity. Simulations will be run to observe social and cultural dynamics of status symbols in more or less "complex societies" (i.e., number and distributions of social hierarchy levels).

REFERENCES

Blackmore, S. (1999) *The Meme Machine*, Oxford University Press.

Dawkins, R. (1976) *The Selfish Gene*, Oxford University Press.

Pedone, R., and Conte, R. (2000) "The Simmel Effect: Imitation and Avoidance in Social Hierarchies." In S. Moss and P. Davidsson (Eds.), *Multi-agent Based Simulation*, Springer.

Simmel, G. (1957) "Fashion," *American Journal of Sociology* 62 (May):541-558.

THE SOCIO-GENETIC SOLUTION: A NEW LOOK AT LANGUAGE GENESIS

T. SATTERFIELD, University of Michigan*

ABSTRACT

Creole languages have long been a point of contention in linguistic circles. Broadly defined, creoles are grammars resulting from an amalgamation of two or more languages, such as when speakers of differing mother tongues find a need for rudimentary communication during economic or social transactions. Creolization occurs if the "invented" system becomes the native language of the speech community. There are several hypotheses for how linguistic properties and social contact each came to bear on the formation of creole languages in past centuries; however, until recently no reliable method for testing these complex interactions existed. By building agent-based "societies," it becomes possible to examine premises of linguistic theory and to reconstruct historical contexts, with an eye to isolating patterns and factors that are most relevant to the acquisition and transmission of languages. Implemented in Swarm 2.1.1, the current model consists of a multiagent population drawn from historical records of Surinamese sugar cane plantations (Arends 1995). Each agent is endowed with a demographic profile and linguistic parameters. Linguistic features include a set of genetic constraints stipulating the environmental conditions required for successful analysis and acquisition of any language. Three experiments using the Swarm model are described. The results provide viable motivation for advancing a "socio-genetic" solution for the emergence of prototypical creole languages.

INTRODUCTION

This paper describes a model that simulates a language formation process known as creolization, using a computer program that tracks "speaker-agents" who enter and emerge from the learning environment. Throughout history, creole languages have arisen only in restricted social contexts, during trade transactions, or under circumstances of upheaval such as slavery. In such cases, speakers of diverse languages must bridge a communication gap when there is little opportunity or will to learn others' language(s). The situation results in the use of a simplified communication system known as a *pidgin*.[1] In time, this newly formed code may nativize (become the native tongue in a speech community), at which point we say that the pidgin has *creolized.* The resultant language is called a *creole.*

* *Corresponding author address*: Teresa Satterfield, Department of Linguistics, University of Michigan, 1076 Frieze Bldg., Ann Arbor, MI 48109-1255; e-mail: tsatter@umich.edu.

1 I use the term "*pidgin*" not in the derogatory sense attributed to it in many standard dictionaries, but in the technical linguistic sense, as a type of "auxiliary" language or strategy that is called upon by speakers in addition to their own native languages to enable communication to take place.

Where the controversy lies, and where motivation for agent-based modeling begins, is that, unlike pidgins, which confine themselves to a few simple structures, the creole language regularizes and expands into complex grammatical forms. These syntactic innovations are akin to "normally" transmitted native languages, yet unlike the ambient mother tongues or pidgin elements that function as the creole's primary linguistic input. Moreover, some argue that creoles worldwide pattern uniformly, such that among themselves they share far-reaching structural similarities, which make them distinguishable from "conventional" languages.

Since pidgins never constitute a mother tongue for any speaker, how is a full and nativized creole language acquired given "difficult" environments? Most researchers agree that both "nature" (the biological capacity for humans to acquire language) and "nurture" (external features shaping human development) play a part; however, this consensus is not integrated into research programs. The majority of creolization investigations make strong arguments for either the genetic component of language acquisition or the social aspect. Below, I sketch the dichotomy that holds between these frameworks.

The *Language Bioprogram Hypothesis (LBH)* (Bickerton 1984, 1992*)*, at its most extreme, depicts creolization as a "bet-hedging" reflex related solely to biological (first) language acquisition processes. Given a "chaotic" linguistic environment, innate "blueprints" yield the most generic structures possible for natural language. To the extent that children generate creoles as mother tongues only in language contact situations, there is a basis for the related claim that creole languages emerge abruptly within one generation of speakers.

A contrasting viewpoint can be labeled the *Social Context Hypothesis (SCH)* (Thomason and Kaufman 1988). Framed in the premise that creoles are indicators of human interaction in unique and varying communicative contexts, they argue against a bioprogram, stating that children possess only enough *a priori* linguistic knowledge for environmental and social conditions to impact acquisition in important and interesting ways. The SCH questions "universal" creole similarities, using comparative historical data as evidence that creoles contain broad structural differences. Finally, the SCH points to nativization as a gradual process requiring generations of participation from the community's adult speakers.

Researchers tacitly assume that mechanisms formulated in the LBH and the SCH cannot function bilaterally. The proposed model aims for a fuller understanding of the creolization process. Object-oriented programming provides an optimal system for simulating complex and dynamical aspects, such as those in language formation. Following Epstein and Axtell (1996), the proposed setting for language genesis functions as an *artificial society*, "cultivating" creolized syntactic structures *in silico*. The primary goal is to discover the role of certain innate and socially based mechanisms in generating prototypical creole grammars.

Overview

This model simulates real-world multilingual plantations that supplied optimal conditions for growth of pidgins and continued development of existing contact languages. While plantation-based slave labor tended to lead to creole genesis, it is not known if these languages necessarily originated from pidgin stages. Nor is creolized grammar an inevitable outcome of the plantation system. We will focus on one known context which spawned and perpetuated creole language.

The ethnolinguistic and demographic parameters used in the model derive from statistics on Sranan Tongo (Surinam Tongue) (Arends 1995, Migge 2000). Sranan is the native language of 200,000 inhabitants of modern Surinam and functions as *lingua franca* for the majority population. Sranan's earliest histories document mid-17th century English settlers who set up small farms. A Dutch invasion drove out the English by 1680. Until 1690, the colony experienced accelerated growth as Dutch planters expanded farms into large-scale sugar plantations and imported increasing numbers of African slaves. By 1750, every planter-master owned 45 to 60 slaves. Plantation numbers were routinely decimated from slaves' lowered life expectancy, low birth rates, and escape. Records speculate that male slaves outnumbered female counterparts 2 to 1. Children were 15% of the total slave cohort; approximately 40% of these infants died before age five. In short, the population was sustained through the constant influx of new slave labor rather than from natural growth.

Plantations functioned as strict hierarchical organizations. In Surinam, a large social distance existed between different groups on the plantation, causing considerable reduction in social networks. Social boundaries included European-African; adult-child; and elite slaves, including overseers and house slaves, versus field hands, and to a lesser extent, skilled African laborers. Ethnolinguistically, imported populations contributing to the formation of Sranan link to language families on the western African coast. Within 70 years, the population was linguistically homogeneous, with three individual languages as most prominent (Arends 1995).

Implementation of Demographic and Linguistic Data

Conforming closely to the previous statistics, each inhabitant of the artificial society is assigned a demographic profile designating age, race, health, death, cultural identity, and social status. Planters have the highest social standing. Slaves receive a status index based on their function in the plantation. Overseers have high status among slaves, while field hands and infirm slaves receive the lowest index. Upon reaching age 12, the child's status is assigned through random "inheritance" of one of his parents' indices.[2] The status index is essential, as it drives movement and subsequent language transmission, following the SCH.[3] Slaves' language variables are randomly assigned tags constituting different A(African)-*language* families. The planter language variable is one E(European-based)-*languag*e.

Per LBH specifications, individuals have innate linguistic capacities, characterized by a lexicon, represented as a list of values corresponding to each native and non-native word stored by the agent; a storage unit for accumulating new vocabulary items; lists of derivational morphemes (building blocks of word formation) and grammatical morphemes (tense and plurality affixes on lexical items, etc.). A morphology storage unit accumulates morphemes as a native language. Sentence structures, or syntactic constraints, are represented as links mapping words and morphology in various ratios depending on language specified. A-languages are mapped three morphemes to one word, E-languages have one morpheme per word.

---

[2] Offspring can only receive planter status if both parents are planters; if one parent is a planter, the child receives slave index.

[3] Given sociolinguistic studies (Gumperz 1971), socio-economic status appears to play a role in language transmission.

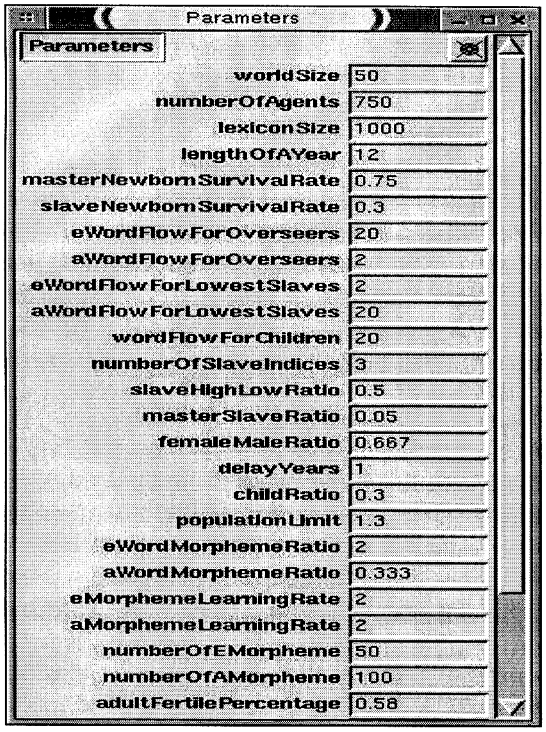

FIGURE 1 Model parameters

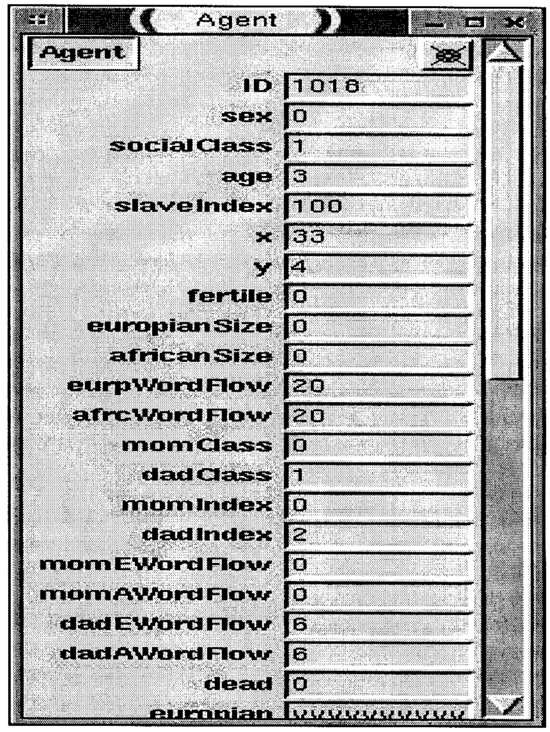

FIGURE 2 Agent profile

Agents process linguistic input via bounded cognitive resources *(e.g.*, memory and time). Following the LBH, children and adults possess different computational limitations for acquisition and storage of lexical and grammatical information.

DESCRIPTION OF THE SWARM MODEL

Swarm experiments generally involve three components: agents, environment space, and rules. Basic creolized systems are hypothesized to emerge as by-products of rule-governed interaction between agents operating in the established social environment. Distributions of innate linguistic knowledge, lexicons, and demographics are entered to initialize the population of agents. The Temporary Memory Buffer prohibits adult learners from acquiring an unchecked amount of vocabulary, on analogy with the restricted computational resources of adult learners. Child-agents enter the environment with no linguistic affiliation but possess inherent capacities to store and generate any language. Children acquire both words and grammatical morphemes at higher rates than adults, following the LBH. Steps performed by the model are listed in Table 1.

TABLE 1 Language-Learning Rules for Agents

1. Initialize **adult** (>15.0 years) Agent's distribution of language features and demographic features.
2. **Language Contact.** Initialize Random Walk *(1* chronological year = *12* cycles). Select unoccupied space within *X* Von Neumann neighbor-agents.
3. **Language Learning.** Talk to *X* neighbors. If input language = [+*African*], form probability of encountering lexical and morphological information based on population ratio of African languages: Language *A*, Language *B*, Language *C*, Other African Languages.

   A. Adult:

   If neighbor has [> status index], look at *X* words in his lexicon. If you do not find those words in your lexicon, select those *X* items and copy into your Temporary Memory Buffer. Talk to *X* neighbors during this step. Consult Temporary Memory Buffer. Add new vocabulary item(s) to your lexicon only if you encounter *>2.0* instances of those new words.

   B. Child:

   If neighbor has [> status index], or if he is a child *>5.0* years, look at *X* words in his lexicon and *Y* morphemes in the morphology listing. If you do not find those words and morphemes in your lexicon, select items and copy them into your respective lexical and morphological storage units. Talk to *X* neighbors during this step. Add new vocabulary item(s) and morphemes to storage areas immediately when you encounter *1.0* instance of new word *W* and new morpheme M.
4. Language learning for current cycle completed.

SAMPLE RUNS AND RESULTS

Three experiments examine the roles of "nature" and "nurture" in creole genesis. For each test, the number of "years" of contact is given, as well as the linguistic profile of the community and an average of individual language patterns. Population and age distributions were monitored. Primary experimental variables were child language acquisition mechanisms (LBH) and rates of social contact (SCH).

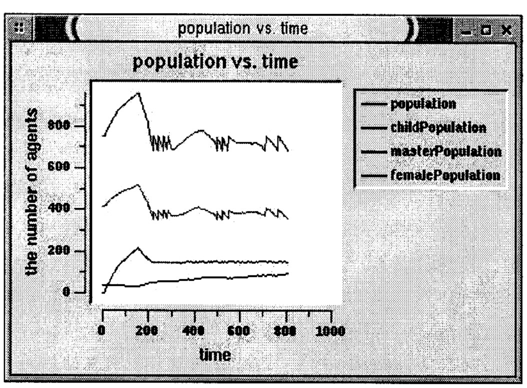

FIGURE 3 Population distribution

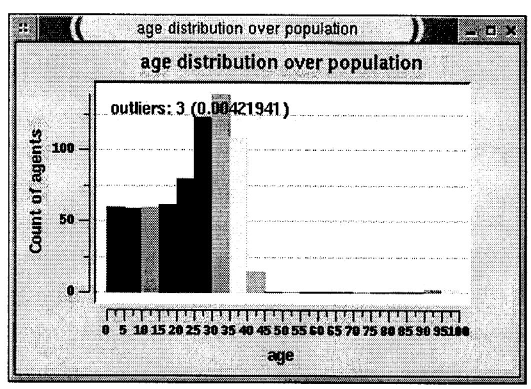

FIGURE 4 Age distribution

Experiment 1

All variables were kept constant in Experiment 1; subsequently, elements of the SCH (translated as high-frequency and high-quality social contact) and the LBH (translated as biologically endowed linguistic mechanisms) are assumed maximally operative. The end result indicates a setting composed of a 15% minimum of children and unlimited contact (per appropriate social hierarchies) between neighbor-agents. Figures 5-7 display results (12 cycles on time axis = 1 year).

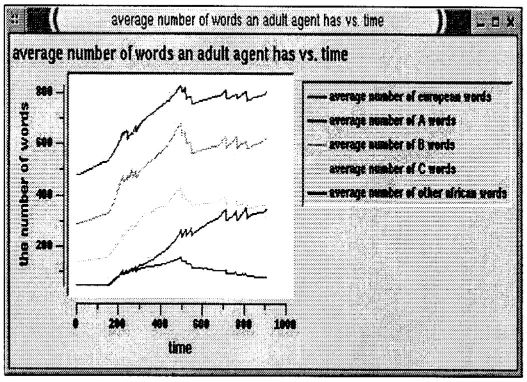

FIGURE 5 Adult lexical storage

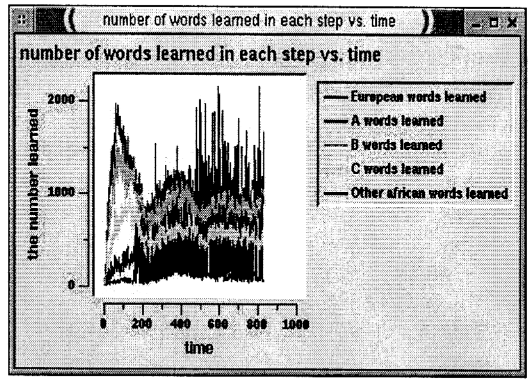

FIGURE 6 Lexical items acquired

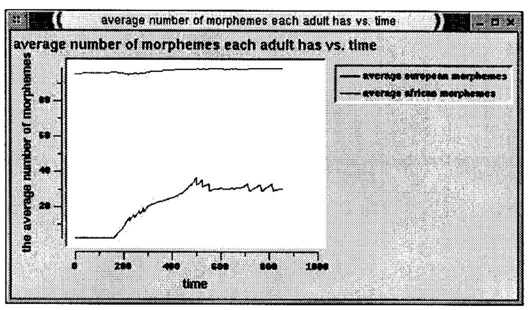

FIGURE 7 Adult morphemes

Experiment 2

Experiment 2 manipulates the social contact variable to examine the impact of the SCHs. Instead of unlimited access to linguistic information from all surrounding neighbors of appropriate status, agents interact with one neighbor per cycle. Results may predict whether creole genesis is possible with minimal agent interaction. Results are shown in Figures 8-9.

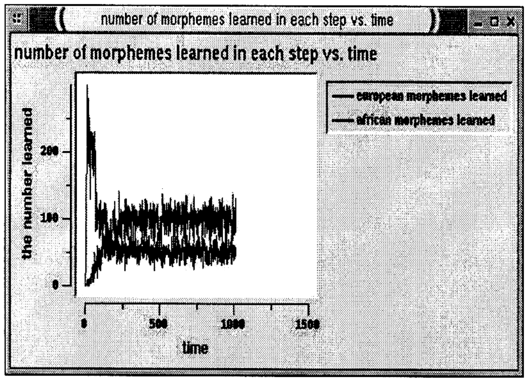

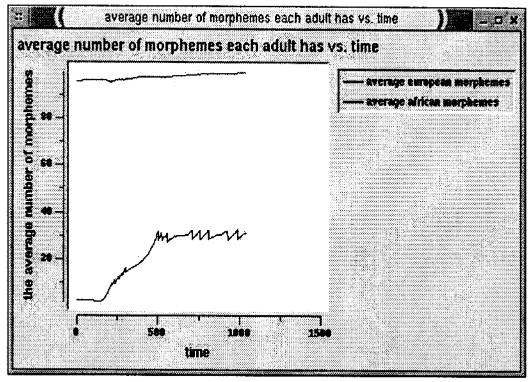

FIGURE 8 Morphemes acquired (no SCH) **FIGURE 9** Adult morphemes (no SCH)

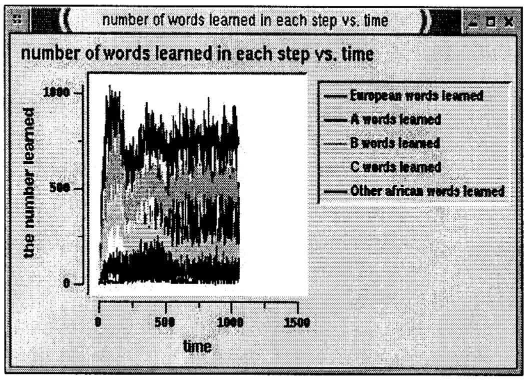

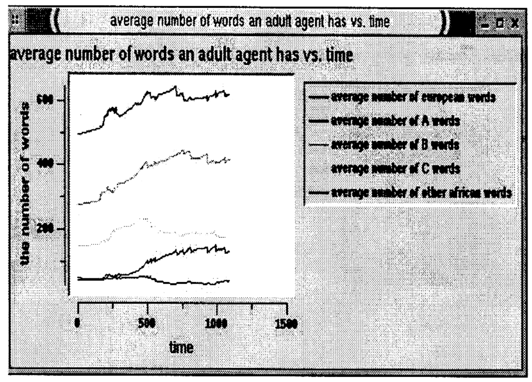

FIGURE 10 Lexical items acquired (no SCH) **FIGURE 11** Adult lexical storage (no SCH)

Experiment 3

Experiment 3 manipulates the child-driven language variable to investigate the role of the LBH. In this scenario, no children are born into the population. Adults interact, with numbers maintained through periodical importation of slaves into the contact setting. Findings may predict whether creole genesis is possible in the absence of child learners. Results are presented in Figures 12 and 13.

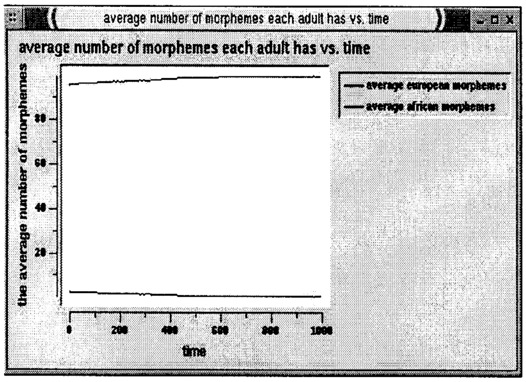

FIGURE 12 Adult morphemes (no LBH)

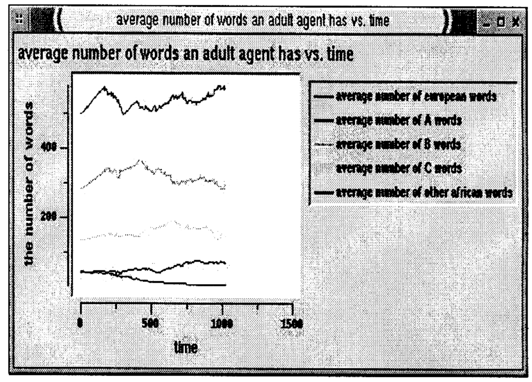

FIGURE 13 Adult lexical storage (no LBH)

DISCUSSION

Experiment 1 produces dynamic characteristics expected with language change. After three generations (approximately 70 years), Figure 5 shows how E-language words come to replace certain A-language varieties, as documented in real-world plantation history. In line with historical slave language pidgins, Figure 6 exhibits initial trends of an abrupt overlapping of A-language acquired lexical items. The presence of children shortly after cycle 100 brings an "averaging" effect for the acquisition rate. An explosion of E-language vocabulary emerges in part from the fact that children require less frequency and memory to retain new vocabulary items but, apparently, also from the influx of slaves from Africa, seen shortly after cycle 200 and again at cycle 450. Stored morphemes of the respective languages are distributed, as shown in Figure 7. Since morphology is acquired entirely as a function of first language, gradual storage of increased non-African morphology might arise if European planter children, after age eight, have more contact with European adults than do other groups. E-language morphemes acquired in late childhood could still be easily transmitted to younger children. From the data, speakers in this scenario possess a diminishing A-language vocabulary, with substantive, but separate, knowledge of both A- and E-language grammatical properties derived from childhood interactions.

Experiment 2 examines creolization under limited social interaction during 1,000 cycles. Figure 8 shows that the initial short-lived rise of A-language morpheme acquisition, gives way to E-language morphology acquired at slightly lower rates than A-language forms. The learning pattern is relatively robust despite limited contact; however, morphologies of each language do not overlap or combine as the learner acquires stable quantities of A- and E-language morphemes. Figure 9 displays results identical to Figure 7 for morphological storage in an adult speaker. Again, the agent has maximum knowledge of both A- and E-language grammatical properties. In the face of restricted interaction, Figure 10 shows that lexical acquisition is 50% less than rates obtained in Experiment 1. A-language (precisely, African language A) forms are acquired at a prominent rate throughout the test, overlapping with E-language activity in the final cycles. Figure 11 tracks adult lexical inventory. Similar to Experiment 1, E-language words gradually replace certain A-language varieties; however, A-languages comprise the majority of words stored. An adult's profile in this scenario corresponds to a small and receding A-language vocabulary (slightly more A-dominant than Experiment 1), with substantive, but separate, knowledge of both A- and E-language grammatical properties.

Experiment 3 reflects a contact scenario in the absence of child language. Figure 12 shows the adult speaker with a full storage of A-language morphemes; however, contrasting with Figures 7 and 9 involving children, Figure 12 has no E-language morphology present. Figure 13 signals that A-language lexicons completely dominate E-language vocabulary, the latter falling to zero within 1000 cycles. Empirically, A-language words are stored at the same quantities found for Experiment 2 (See Figure 11), with a slight exception for the highest A-language. Speakers in Experiment 3, as adults, are the least influenced by contact, possessing a small A-language-only vocabulary, with full knowledge of solely A-language grammatical properties.

Statistically speaking, Experiment 1 represents the most fertile conditions for hypothetical creole genesis, since active child populations and relatively unrestricted contact trigger the highest incidents of change. Patterns in Experiment 2 did not differ extensively from Experiment 1. Limited contact appears to affect most dramatically the number of lexical items acquired and stored, while having minimal influence on the acquisition of morphology. Experiment 3 deviates greatly from the projected creolization outcome, as the lack of children equated to zero feature transmission across language groups.

CONCLUSION

In attempting to build from formulations of the LBH and SCH, this model offers experiments which duplicate the complex nature of creolization, simultaneously following social and (biological) linguistic development in a plantation setting. Preliminary evidence from the simulations demonstrates that genetic and social mechanisms must function bilaterally if a prototypical creole grammar is to potentially emerge. Overall, our findings suggest that innately specified capacities play a crucial role in organizing disparate input and in generating richer structures in *child* language. However, without significant amounts of social interaction, agents cannot attain native-speaker quantities of lexical inventories, nor do extensive nativized mixtures occur. Notably, striking similarities in morphological acquisition patterns of Experiments 1 and 2 favor the LBH claim that extralinguistic factors and environment do not greatly influence grammatical acquisition. Insofar as this is the case, such results may contradict the conservative “bet-hedging” premise of the LBH, since speakers redundantly retained native levels of morphology in both language groups. Finally, no experiment validated the related LBH claim that creole languages develop within a single generation of speakers. With roughly four generations of agents (1,000 runs), few mixed language forms were materialized. Given these factors, we tentatively advance a socio-genetic solution for analyzing questions related to creole genesis.

Motivated by theoretical and practical concerns, modifications are presently being considered, with an eye to providing more explanatory models of language contact. For example, a syntactic formation involving fixed templates of *X* words to *Y* morpheme affixes is not suggested by current linguistic theory, although this strategy efficiently eliminates extensive searches for language family tags. Second, we posited rigid constraints on adult acquisition of non-native words and morphology, whereas the youngest children had few restrictions for storing any morphological or lexical item. It is clear that such stipulations, along with issues of linguistic input, must be coordinated with recent findings in first- and second-language acquisition theory. Lastly, demographic factors should be fine-tuned to better pinpoint those “critical mass” conditions necessary for triggering individual and population-wide creolization over time.

ACKNOWLEDGMENTS

I would like to thank Wei Chen and Rick Riolo at the Center for Study of Complex Systems (CSCS), University of Michigan, for valuable assistance on this project.

REFERENCES

Arends, J., ed., 1995, "Demographic factors in the formation of Sranan," in J. Arends, ed., *The Early Stages of Creolization,* John Benjamins, Amsterdam, pp. 233-285.

Arends, J., P. Muysken, and N. Smith, eds., 1994, *Pidgins and Creoles: An Introduction,* John Benjamins, Amsterdam.

Bickerton, D., 1999, "How to acquire language without positive evidence: What acquisitionists can learn from creoles," in M. DeGraff, ed., *Language Creation and Language Change,* MIT Press, Cambridge, MA, pp. 49-74.

Bickerton, D., 1992, "The sociohistorical matrix of creolization," *Journal of Pidgin and Creole Languages* 7:307-318.

Bickerton, D., 1988, "Creole languages and the bioprogram," in F. Newmeyer, ed., *Linguistics: The Cambridge Survey,* Cambridge University Press, Cambridge, pp. 267-284.

Bickerton, D., 1984, "The language bioprogram hypothesis," *Behavioral and Brain Sciences,* 7:212-218.

Bradshaw, J., ed., 1997, *Software Agent*s, MIT Press, Cambridge, MA.

Chomsky, N., 1995, *The Minimalist Program,* MIT Press, Cambridge, MA.

Chomsky, N., 1986, *Knowledge of Language: Its Nature, Origin, and Use,* Praeger, New York.

Epstein, J., and R. Axtell, 1996, *Growing Artificial Societies,* MIT Press/Brookings Institute, Cambridge, MA.

Gumperz, J., 1971, *Language in Social Groups,* Stanford University Press, Stanford, CA.

Holm, J., *Pidgins and Creoles, Volume 1: Theory and Structure,* Cambridge University Press, Cambridge.

Migge, B., 2000, "The origin of the syntax and semantics of property items in the Surinamese plantation creole," in J. McWhorter, ed., *Language Change and Language Contact in Pidgins and Creoles,* John Benjamins, Amsterdam, pp. 201-234.

Rickford, J., and J. McWhorter, 1997, "Language contact and language generation: Pidgins and creoles," in F. Coulmas, ed., *The Handbook of Sociolinguistics,* Blackwell, Oxford, pp. 238-256.

Satterfield, T., 1999, *Bilingual Selection of Syntactic Knowledge,* Kluwer, Dordrecht.

Sebba, M., 1998, *Contact Languages: Pidgins and Creoles,* St. Martin's Press, New York.

Siegel, J. , 1999, "Transfer constraints and substrate influence in Melanesian Pidgin," *Journal of Pidgin and Creole Languages* 14(1):1-44.

Seigel, J., 1987, *Language Contact in a Plantation Environment,* Cambridge University Press, Cambridge.

Thomason, S., 1997, "A typology of contact languages," in Spears, A., and D. Winford, eds., *The Structure and Status of Pidgins and Creoles,* John Benjamins, Amsterdam, pp. 71-88.

Thomason, S., and T. Kaufman, 1988, *Language Contact, Creolization, and Genetic Linguistics,* University of California Press, Berkeley, CA.

DISCUSSION: CULTURAL AND NORMATIVE DYNAMICS

R. PICKER, University of Chicago, Moderator

[Presentation by Verhagen]

Randal Picker: Have you played with group size at all?

Harko Verhagen: We didn't. I should explain that the student's main goal was to have a central measure called self-esteem. How did the agent proceed perceive itself to be? That would be influenced by the feedback from the group. If an agent received negative feedback, stating that he should have done something else, his self-esteem would drop. If his self-esteem became too low he would go to another group. But it did not really work that way, because due to the problems in the Jet Lite code, we couldn't switch on the fly from group to group. The members were different anyway, so they did not arrive at the same group decision tree in the end.

Michael North: You mentioned the decision trees — that members maintain their own decision trees representing their personal views, as well as a group decision tree representing what they think the group is thinking. Can you give further details on what that decision tree represents, versus the world they're working in?

Verhagen: The decision tree enumerates all 20 possible alternatives. For each alternative, it gives possible consequences, which are at most two: the probability per consequence and the utility per consequence.

North: When people are trying to discover what they believe the group norm is in each step, how is that reflected in their decision making? Do you see the perceived norm as essentially an attempt to compress models of many people into a single, simple model?

Verhagen: Yes, in a sense it's just getting the average of the group, so to speak, weighted by the leadership values.

North: There are other ways to think about this process, too, in the sense that a norm is collective wisdom, versus just a model of other people's thinking. Do you see your model touching on that at all, or do you see it as just dealing with one branch?

Verhagen: I do think we get into the aspect that you mention.

[Presentation by Sander, Schreiber, and Doherty]

North: My understanding of Schelling's original work includes the idea that you don't need to have a huge difference in preference or strength of preference to have a segregated outcome. You could want just one neighbor to be the same as you and end up with extreme segregation. Have you looked at preference thresholds — how low they can go and still result in segregation?

Richard Sander: No, and that's an important thing to do. The general thrust of our work is that the Schelling mechanism is clearly important, but it's not all-important, so we need a model that involves the other factors. But, no, we haven't tested yet to determine when pushing up the preferences has an overall weight and utility function. It starts producing heavier segregation then.

Darren Schreiber: I've explored that a bit, but not systematically. One of the interesting aspects of our model is that with one switch you can turn it right back into Schelling, and use it to monitor all of the dependent variables that we are looking at. So it really is a good test bench for this kind of question.

North: You mentioned that you are attempting to model income issues by weighing the housing costs function. Could that result in bracket creep? Because if you look at a change in housing, you could move to a slightly more expensive one, then slightly more expensive again.

Sander: There are two ways that we're fiddling with the housing function. One is by changing the importance that housing cost has in the overall model, which is one way to test the idea of whether the "black premium" drove integration in the '70s. If housing costs have too little weight, so that people aren't driven enough by that factor, then you lose the rapid mobility phenomenon that has been clearly observed. The other way is by varying the individual weights people assign to housing. In our model, within those ranges there is a continuous incrementation of weighting. I should also mention that we operationalize our model by giving a fraction of the agents five alternate houses to consider in each period. If one of them improves an agent's utility, he moves there.

Benjamin Schoepfle: I'm Benjamin Schoepfle from Argonne National Laboratory. I'm wondering how you define tracts. Are they defined relative to the agents or are they fixed *a priori*?

Sander: They are fixed *a priori*.

Schoepfle: Did you think about other important tract boundaries, such as highways and railroads, and school districts? I did a dissertation on school redistricting from an operations research perspective. It seems that these boundaries, especially school districts, are very powerful engines for driving segregation — and desegregation, but mostly segregation.

Sander: The tract boundaries are fixed *a priori*, but they are intended to symbolize those types of borders. The tract boundaries in the model are just arbitrary boxes, but their counterparts in the real world are whatever boundaries people think are important for defining what they want the racial makeup of their neighborhood to be.

Ian Lustick: Ian Lustick from the University of Pennsylvania. You mentioned that agents are given five housing options in your model. As I was listening to your presentation I was thinking of Doug Massey's book, *American Apartheid*, and that type of work, which portrays the issue of segregation as overdetermined by 10 or 12 powerful factors, including the failure to enforce the laws. Don't you create a problem for yourself in terms of verisimilitude by simply stipulating that everybody gets five choices? You didn't mention any of the literature; perhaps you could refer to that.

My other comment: You say you haven't fiddled with potential thresholds in the characteristics you give agents, but it would be a good question to look at. What would be the effect of including a small number of people with very strict preferences, as opposed to changing everybody's settings marginally?

Sander: On your first question: we give people five choices, but we can introduce into the model the level of discrimination they are going to encounter with each choice. So if you build in a high level of discrimination by reds against blues, then all the choices that blues have outside the blue area may be terribly unacceptable.

Generally, what we are finding in preliminary work is that the discrimination cost, or the racism cost, is not as significant as the preference functions in driving these neighborhood outcomes. In other words, even when you introduce a substantial level of discrimination, you find a significant number of moves, because you have some initial moves that come from people who have a strong desire for integration and are trying to get lower prices by moving out of the predominantly black area. Then their moves enable other people to move into these little sub-groupings.

On the second issue, almost all whites in the survey data tend to say that their highest level of utility is an all-white neighborhood, but different whites vary in terms of how diverse a neighborhood can be before they say it's not ideal. In the model, for the least tolerant reds, their utility goes down monotonically as the proportion of minorities increases. And that is fixed; the agents don't change as they change neighborhoods. It would be nice, although perhaps overly optimistic, to have the preferences themselves evolve as people experience different types of neighborhoods.

[Presentation by Satterfield]

Claudio Cioffi-Revilla: Claudio Cioffi from Colorado. How do we know that the calibration of the time scale that you have in this simulation in fact corresponds to 100 years, as you said? How do we know it's not 100 months or 100,000 years? Well, 100,000 years is probably extreme; but how can you demonstrate that these are in fact years?

Teresa Satterfield: That's a good point. You're right, I can't at this point say I know for a fact that my 12 runs are equivalent to a chronological year. But we tried to look at empirical information and see what the trends were over time, and coordinate those historical developments with our results. Hopefully, we'll refine that and get closer. I don't think you can ever mimic real-world behaviors 100 percent using these kinds of simulations, but I think they're good tools. For example, in theoretical linguistics, they allow us to see some of the gaps that no one has been able to even conjecture about in terms of these historical developments and language acquisition. Records just weren't that well preserved, so we're not sure how exactly languages were being used in the different populations. So even if they only give us an idea to start from, I think it's better than nothing.

Cioffi-Revilla: I think this would be very helpful, both for the sake of your own model, to demonstrate it, but also for those of us who work in different domains — because each domain seems to have its own tricks for solving the time-scale calibration problem. A comparison of these strategies, of these solutions, would be a very helpful methodological gain for the field. In political science, for example, we have a feel for the rate at which nations form and dissolve, we have a feel for the rate at which riots break out, we have a feel for how long wars last, and so forth — a lot of this is just plain empirical background.

Christopher Langton: Chris Langton from Santa Fe. Could you talk about what the theory says about children who grow up in a bilingual family? How do they *not* create an intermediate language?

Satterfield: It seems to be based on the input received. Social-contact languages, often called “creole” or “pidgin,” seem to result when children are bombarded with multilingual degenerate input, a simplified, conglomerate version of communication that the adults in their lives put together. For kids to become balanced bilinguals, where they master two or more languages, they appear to need balanced input coming in, and it has to be of a very substantial nature. It can’t be degenerate or simplified; it can’t be pidgin. As long as you get constant, consistent input from two different languages, you can distribute the linguistic knowledge into two grammars and can become truly bilingual.

Langton: So the quality of the language is very high and therefore they don’t need to construct their own.

Satterfield: Well, I’m saying that if you want to get two separate systems that don’t converge at some point, then you need to have very strict constraints for both languages, to the extent that the child will be able to maintain them as separate languages. What happens with the pidgins is that the children and the adults pick certain characteristics of each language and then simplify that amalgam to create this new form of language. We know that bilingual kids are able to differentiate and maintain separate languages because they are getting enough of the two languages as input.

John Padgett: John Padgett, University of Chicago. If I have understood your talk correctly, you have been focusing on the creolization process mostly from the angle of how individuals construct their own lists and recombine — speech production, you might say. The other side of it, though, is speech comprehension, how you compute what other people say. I would think that in these pidgin or creole situations, comprehension could be deeply problematic because other people are uttering these highly hybridized words that are not yet in your list. I’m wondering how you would operationalize this more collective convergence process. Aside from individuals generating these hybrids, how does the whole thing converge down into some common language?

Satterfield: This is something that’s still in progress. But one thing we know about these pidgin and hybrid languages is that they have a very simplified structure, maybe just subjects and verbs. So it seems to be something that’s canonical within the human realm. The lexicon for any pidgin is a maximum of 100 words, so we’re not talking about manipulating many words within the mental dictionary.

Politics

STUDYING PERFORMANCE AND LEARNING WITH ABIR: A RESEARCH NOTE

I.S. LUSTICK, D. MIODOWNIK, and
M. van der VEEN,
University of Pennsylvania*

The Agent-Based Identity Repertoire (ABIR) model is a two-dimensional array of square cells. Agents neither move nor die. In each time step, agents compare an "activated" attribute to the comparable attribute of their neighbors and note any general information available about changes in the attractiveness of the different ways in which this attribute can be activated. Changes in activation at the individual agent level yield patterns of change at the macro-level — patterns that then constrain the behavior of individuals. The model was originally designed, and is still being used, to refine, explore, and test propositions associated with "constructivist" theories of political identity.[1]

The model has also been used to study deliberative democracy.[2] With problems of democratic stability, opinion diversity, and flexibility in mind, we transformed the model into the "Agent-Based Argument Repertoire" (ABAR) model, using it to investigate the behavior of polities modeled as two-dimensional landscapes inhabited by citizens. At any particular time, each citizen articulates one and only one argument. All arguments understood by any agent are contained in the agent's repertoire of arguments. Each agent monitors arguments articulated in its eight-member Moore neighborhood and also registers signals available from the mass media to all citizens regarding current evidence as to the validity or attractiveness of the argument. A citizen determines whether to continue articulating the argument it was making based on a comparison of the "argument weights" present in its neighborhood for different arguments, taking into consideration the current signals regarding these arguments being sent by the mass media, and whether or not the argument is readily understood, i.e.,, present in the citizen's repertoire.

Among our findings was a striking pattern linked to the presence of relatively small percentages of "opinion leaders." We modeled opinion leaders as citizens

- Whose arguments were offered with twice as much persuasive power as those of ordinary citizens;
- Whose argument repertoires were larger;

---

* *Corresponding author address:* Ian S. Lustick, Dan Miodownik, and Maurtis van der Veen, University of Pennsylvania; e-mail: ilustick@sas.upenn.edu, donm@sas.upenn.edu, and maurtis@sas.upenn.edu.

1 "Agent-Based Modeling of Collective Identity: Testing Constructivist Theory," *Journal of Artificial Societies and Social Simulations*, Vol.3, no. 1 (January 2000). http://jasss.soc.surrey.ac.uk/3/1/1.html; Ian S. Lustick, Dan Miodownik, and Stacey Philbrick, "The Institutionalization of Identity: Micro Adaptation, Macro Effects, and Collective Consequences," Paper presented at the annual meeting of the American Political Science Association, August 2000, Washington, D.C. ABIR was developed as a result of intensive collaboration between Ian S. Lustick and Vladimir Dergachev. An executable copy of the program, along with a manual for its use, is downloadable at http://www.polisci.upenn.edu/profileil.html.

2 "Deliberative Democracy and Public Discourse: The Agent Based Argument Repertoire Model," *Complexity*, Vol. 5, no. 4 (March-April 2000) pp. 13-30 (with Dan Miodownik).

- Who were more sensitive to marginal changes in the argument weights of arguments they were not currently articulating;

- Who were more able to replace familiar arguments in their repertoire with unfamiliar arguments being made by those around them and/or favored by the media;

- Who updated their activated argument earlier.

We found that in ratios of just 5% or 10% of the collectivity, the presence of opinion leaders

- Promoted greater local agreement among citizens when lack of sophistication made agreement difficult;

- Reduced general levels of disagreement when lack of sophistication among citizens made such levels very high;

- Promoted discussion and the possibility of change when early clusters of agreement among sophisticated citizens tended to lock in adherence to arguments no longer valid;

- Protected diversity when risky and volatile environments represented by the mass media signals tended toward the elimination of minority viewpoints.

In our current work we have asked a different kind of question. We are interested in the performance of the landscape as a whole, over time — in its ability to respond as an aggregate to changing signals of what is required for success. To score well, these landscapes must overcome the challenge of becoming overly committed to early signals that are likely to prove false over time. Also, in contrast to our work on deliberative democracy, our attention here has shifted away from measurements of the landscape at a selected point in time (such as $t = 500$ or $t = 1000$). Here we are interested in the efficacy, or performance, of a landscape throughout its entire history, not just in its profile as a snapshot in a particular, final, or intermediate time period.

We see this research as speaking directly to the question of whether and to what extent learning can take place as an emergent property of complex processes of local and simple algorithmic adaptation. Such an approach stands in contrast to other agent-based models that study learning in terms of increasingly intelligent agents, variation in the patterns of communication and decision among networks of agents, or selection mechanisms for the replacement and creation of agents (such as genetic algorithms).

Here we briefly report the results of several of our experiments, which were intended to probe the analytic space and virtual empirical domain, as well as test our operationalizations of key variables and conditions.

In Experiment 1, we used ABAR to create polities with 2,400 agents across a 48×50 cylinder shape. These polities were populated by different mixes of ordinary citizens, opinion leaders, and innovators. Ordinary citizens had three fewer arguments in their repertoires than did opinion leaders or innovators. Innovators had the same size repertoires as entrepreneurs and also had the same updating, but innovators articulated their arguments with the persuasive power of an ordinary citizen, i.e., with an "argument weight" of "1" rather than "2." Together we refer to opinion leaders and innovators as "mobilizers."

We produced histories, or runs, of these polities, each with activated arguments and repertoire arguments assigned randomly to all citizens, for eleven different repertoire size conditions beginning with 1,4 (ordinary citizen repertoire = 1, mobilizer repertoire = 4) and ending with 11,14. We ran this series of twenty histories per repertoire size:

1. With no "mobilizers" present (no innovators and no opinion leaders).

2. With 5% of the citizens transformed into opinion leaders.

3. With 5% of the citizens transformed into innovators.

Twenty histories were produced for each condition. A total of 660 histories — 220 histories each for the no mobilizer, 5% opinion leader, and 5% innovator conditions — were produced. Relatively stable environmental settings were used: volatility[3] of .005 and a range[4] of -2,+1.

For each group of twenty histories, average performance scores over 500 time steps, starting at t = 0, were calculated for each condition. These scores were produced by summing scores of the landscape for each time step. Scores for each time step were produced by multiplying the number of citizens articulating each argument by the bias assigned to that argument in that time step. Thus polities (landscapes) scored high when they contained citizens who adjusted more fluidly to changing indications about the validity of different arguments. Performance over history scores were produced by summing the scores for each of 500 time steps. We also corrected for variation in the average validity score available to the citizens within different histories.[5] The data from this experiment are displayed in Figure 1.

We can see that in the small size repertoire conditions (Repertoire [Rep] of ordinary citizen = 1, 2, or 3), performance scores rose rapidly and then more or less leveled off. Of interest is that the presence of 5% opinion leaders or 5% innovators did not make a dramatic difference in performance levels. The presence of mobilizers seems to have enhanced performance only in the low repertoire (relatively unsophisticated) conditions. We may also observe that innovators helped polities in one mid-size repertoire condition (citizen Rep = 6) achieve the highest score of the experiment, but that in the two largest repertoire conditions (most sophisticated), the "no mobilizer" condition outperformed the other two. Examination of these results in combination with visual observation of the dynamics of these histories suggests that an important reason for the relative success of polities with opinion leaders, at small repertoire sizes, is that opinion leaders alert their relatively unsophisticated neighbors to opportunities for adopting more satisfying, or valid, arguments and do so in part by using their greater persuasive power to "teach" arguments to unsophisticated neighbors that those neighbors previously did not

3 The probability that the mass media signal about the validity of any particular argument would be eligible for change in any given time step.

4 The range within which validity signals could vary.

5 This was accomplished in each time step by subtracting from the raw performance score for the total number of agents in the landscape the average of all biases assigned to identities in that time step multiplied by the number of agents in the landscape.

understand (i.e., were not in their repertoire). In contrast, as the repertoire size of citizens becomes moderately sophisticated (rep conditions 4,7 to 7,10) innovators help their polities achieve higher levels of performance than do opinion leaders. It appears that this added neighborhood sophistication turns the greater persuasiveness of the opinion leaders into a liability since it tends to "anchor" their neighborhoods by producing belts of ordinary citizens whose mimicry of the opinion leaders ends up reducing their ability to exploit available arguments to adapt to changing conditions.

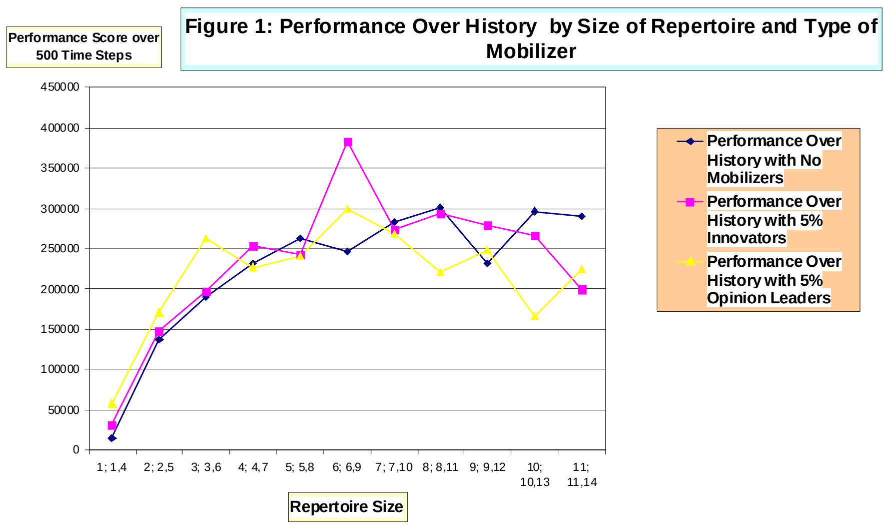

FIGURE 1 Performance over history by size of repertoire and type of mobilizer

In Experiment 2 we focused only on one repertoire condition: Rep = 8; 8,11[6] and changed the volatility and range of the environmental settings to moderately turbulent levels (.009 and –2,+2). Twenty histories were produced with no mobilizers, with 5% opinion leaders, and with 5% innovators. The results are displayed in Figure 2.

The average performance of these histories, under the three different mobilizer conditions, is reported in the three bars on the extreme left of the figure. Clearly, and somewhat surprisingly, the histories in which no mobilizers were present performed considerably better overall than histories in which either opinion leaders or mobilizers were present. These findings, combined with a visual examination of many of the histories, led us to note the prominence of cascades triggered within the early portions of many of the histories due to the juxtaposition of "+2" and "-2" validity signals for some competing arguments. Figure 3 displays a screen shot of ABAR showing the result of a typical cascade (in favor of the argument coded as light green) that occurred before t = 100. These cascades tended to leave one argument in such a dominant

[6] This notation should be read as follows: Rep(ertoire) = 8 (when only basic agents are present, their repertoire size is "8"); 8, 11 (when mobilizers are present, basic agents remain with repertoire size of "8" while mobilizers have a repertoire size of "11").

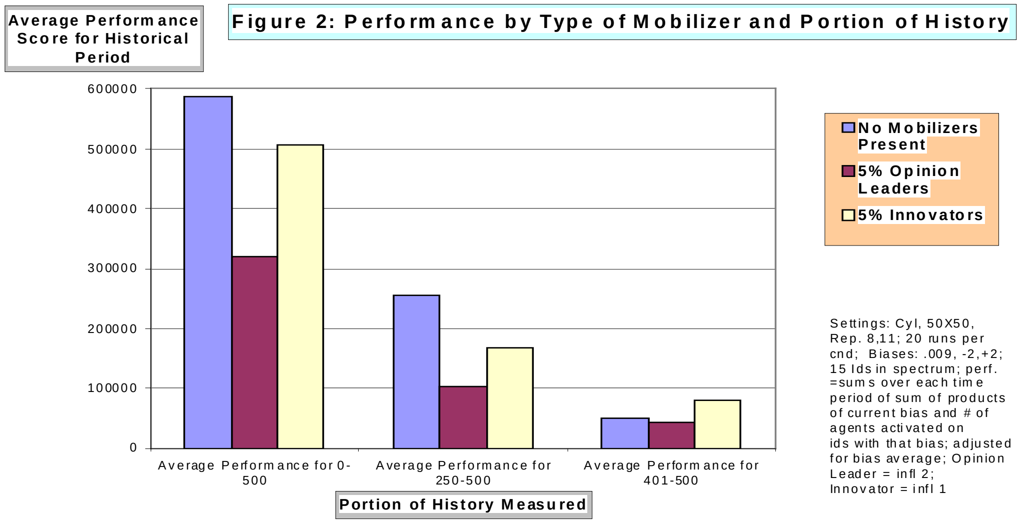

FIGURE 2 Performance by Type of Mobilizer and Portion of History

position within the polity that citizens articulating that argument could not readily adapt to other arguments even when the attractiveness of that argument decreased. In other words, for a large majority of citizens, neighborhood effects — pressures from all eight neighbors articulating the initially favored argument — outweighed media signals that the argument was no longer as valid. We hypothesized that only toward the end of these histories would the dynamism of the environment help other arguments break up the domination of the argument that benefited from the initial cascade, reasoning further that the contribution of the innovators and opinion leaders to the enhancement of performance would be manifest in the latter portions of these histories.

The data displayed in the columns on the right side of Figure 2 strongly support this analysis. While their absolute size is considerably smaller than the columns on the left side of the figure (because they represent not 100%, but only 20% of the total performance of the history), the performance score of the polities with no mobilizers is only barely greater than that of polities with 5% opinion leaders. Most striking, however, is that in the last 20% of these histories, the performance of histories with 5% innovators was considerably higher than performance scores in either of the other two conditions. The fact that polities with innovators out-performed polities with opinion leaders in each portion of the histories measured is added evidence that the “anchoring” impact of the opinion leaders’ persuasiveness interferes with the ability of the polity to “learn” about change in its environment.

Finally, we report the results of Experiment 3. Building on our analysis of the Experiment 2 data, we produced polities by running histories under relatively stable conditions (volatility = .005; range = -2,+1) to t = 500 without measuring their performance. Environmental conditions were then made more turbulent by increasing the volatility to .009 and the range to –2, +2 and histories from t = 500 to t = 1,000 were recorded. Our aim was to see if the performance of polities would be more sensitive to the presence of mobilizers under moderately turbulent conditions when multiple arguments, rather than one dominant argument, were well established in the polity.[7] The results are displayed in Figure 4.

---

[7] The results of additional runs, using 10% innovators, are not reported here because of their incommensurability.

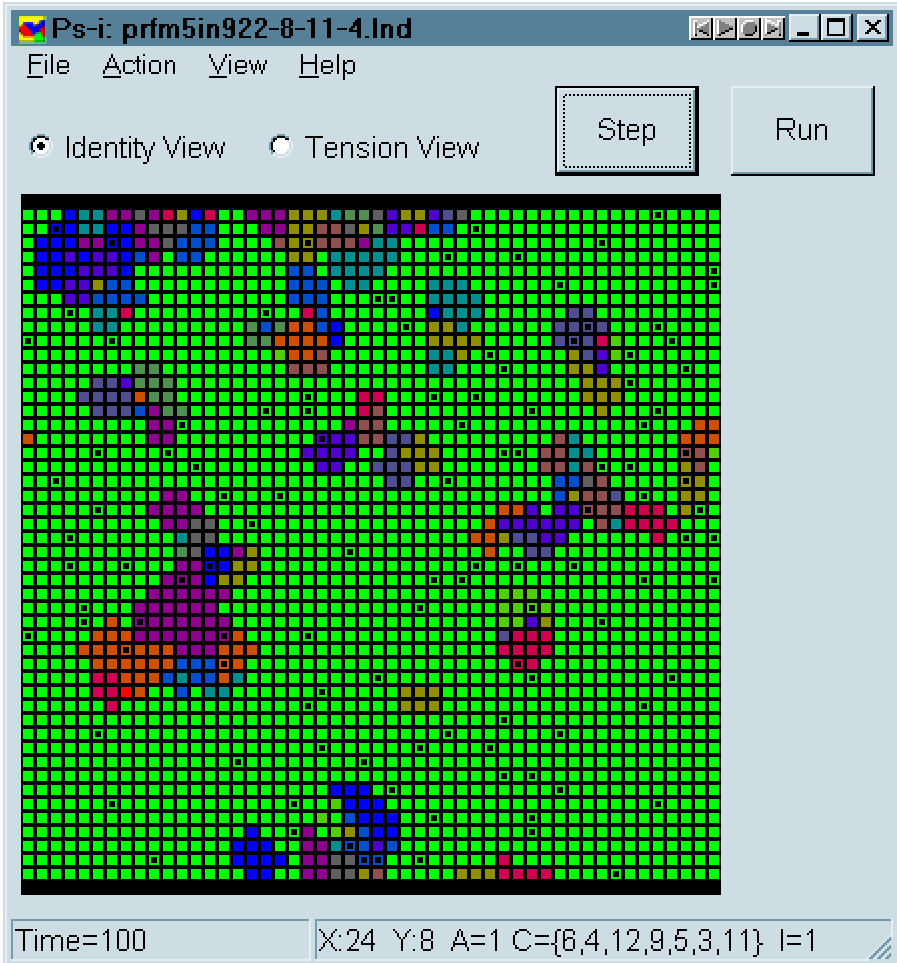

FIGURE 3 Cascade effect (favoring the argument coded as light green) that occurred early in the history

We can see, again, that the performance of polities increased dramatically as the sophistication of citizens increased from very low to moderate levels. And we can see that under these conditions the presence of 5% opinion leaders boosted the performance of the polities significantly. For reasons we do not yet understand, the performance patterns in both conditions converge and vary quite widely over conditions of greater sophistication. Clearly, this will require more carefully controlled work to explain, though we observe that when these two conditions (opinion leaders absent and present) are compared with one another across all repertoire sizes, the difference between them is significant in the predicted direction, i.e., greater performance with opinion leaders present, at the p = .001 level.

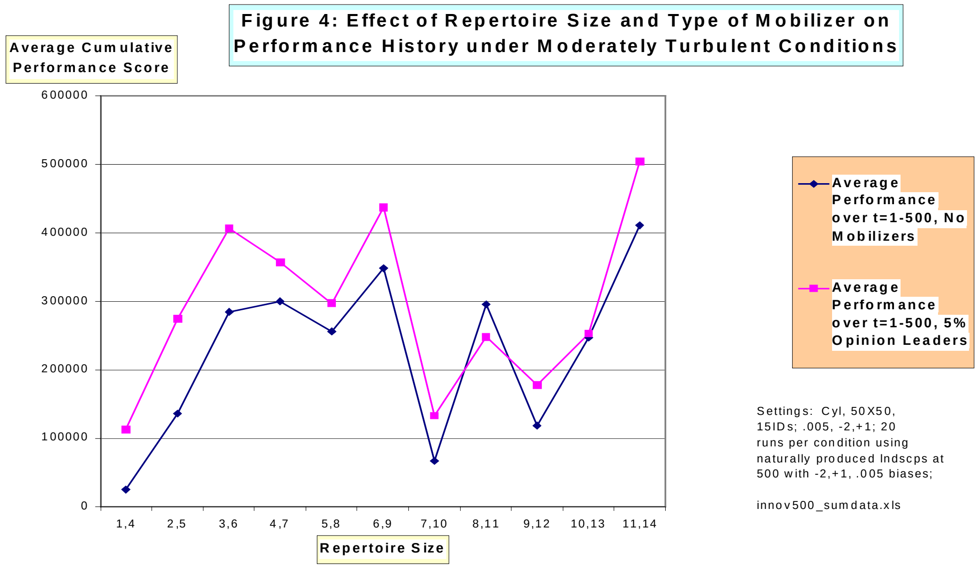

FIGURE 4 Effect of repertoire size and type of mobilizer on performance history under moderately turbulent conditions

ASPIRATION-BASED ADAPTATION IN GAMES

J. BENDOR, Stanford University*
D. DIERMEIER, Northwestern University
M. TING, University of North Carolina at Chapel Hill

I. INTRODUCTION

Behavioral game theory (Camerer 1990) attempts to synthesize classical game theory's focus on strategic interaction and behavioral decision theory's focus on the cognitive constraints of real decision makers. A key driver in the development of behavioral game theory has been the discovery of persistent anomalies in both field studies and laboratory experiments. Examples include cooperative behavior in the prisoner's dilemma (PD) and the rejection of strictly positive offers in Ultimatum games. Anomalies also confront theories of perfect rationality in allied fields. In political science, for example, classical game theoretic models predict that expected turnout in large electorates where citizens have strictly positive costs of voting will be very small. Generically, high expected turnout cannot be an equilibrium, for then the chance of being pivotal would be too small to warrant the effort of voting. Yet, millions of citizens vote in even the largest of elections.

In this paper we propose a behavioral game theoretic model of *aspiration-based adaptation* that can simultaneously address two such anomalies: cooperation in the finitely repeated PD and participation in large-scale elections such as the "paradox of voting" described above. Agents' behavior is based on two properties. First, they have aspiration levels which partition outcomes into two subsets: satisfactory and unsatisfactory (Simon 1955). Second, decision makers learn via trial-and-error, becoming more inclined to try actions that satisfice and less likely to try those that don't. Aspirations are dynamically endogenous: they adjust to an agent's experience (payoffs). We then turn this adaptive behavior loose in the context of two games: the prisoner's dilemma and an electoral participation (turnout) game. Regarding the latter, we show that agents who adapt in this manner turn out in substantial numbers even in large electorates and even if voting is strictly costly for *everyone*.

The paper is organized as follows. Section II presents the theoretical model, formalized as a finite-state Markov chain with stationary transition probabilities. It also provides an analytical result showing that this stochastic process is ergodic: it has a unique limiting distribution to which the system must converge, from any starting vector of initial values of the state variables.

* *Corresponding author address*: Jonathan Bendor, Graduate School of Business, 518 Memorial Way, Stanford University, Stanford, CA 94305-5015; e-mail: jonathan@gsb.stanford.edu.

The properties of this limiting distribution are then studied via a computational model. This model is described in section III; simulation results are reported in section IV. Conclusions are offered in section V.

II. THE MODEL

Most theoretical research on learning on learning in games (e.g., Fudenberg and Levine 1997) has focused on providing behavioral foundations for Nash equilibrium and other game-theoretic solution concepts. In contrast, our goal is to propose a behavioral model that can account for some of the observed behavioral anomalies implied by game-theoretic solution concepts. This motivates our focus on cooperation in the PD and electoral participation in the turnout problem. In each application we leave the game structure — and, hence, payoffs — alone. But instead of assuming fully rational actors with common knowledge about the game form, our players adapt their behavior over time as a consequence of past individual payoffs and aspirations.

Specifically, we assume that actors adapt according to a simple form of trial-and-error behavior which is consistent with basic axioms of reinforcement learning (Bush and Mosteller 1955): actions that are successful today are more likely to be used tomorrow; unsuccessful actions are less likely. This reinforcement learning is married to an aspiration level, a threshold that partitions all possible current payoffs into successful and unsuccessful ones, hence indicating which actions are coded as successes (and so worthy of reinforcement) and which as failures (and so inhibited). An actor's aspiration level itself adjusts to experience, reflecting prior payoffs.

We assume a finite set of actors N, where player i has a set Σ_i of finitely many actions. Let σ be an action profile with Σ denoting the set of pure action profiles. Actor i's action is denoted σ_i. Players interact at discrete time periods t according to the same (one-shot) game.

Payoffs. As in classical game-theory the payoff to each player depends on the action taken by each player — i.e, the action profile — but with a random component. Payoffs are thus modeled as $|\Sigma|$ non-degenerate (conditional) probability distributions with finite support. Note that this implies that for each σ at least two payoff values may occur with positive probability. To denote realized payoffs we use $\pi_{i,t}(\sigma)$ with corresponding (stationary) random variables $\Pi_{i,\sigma}$. We assume that the payoff realizations are mutually independent across agents and time.

Different games then correspond to different conditions on these random payoffs. For example, suppose that according to the classical game form some agent i prefers some σ to some other σ'. With stochastic payoffs, this assumption may be captured by a corresponding assumption on expectations or the supports of the random variables, such as $E(\Pi_{i,\sigma^*}) > E(\Pi_{i,\sigma^{*\prime}})$.

Propensities and Adjustments. The heart of the model is the learning behavior of each agent. As indicated earlier, the adaptive behavior combines reinforcement learning and endogenous aspirations. Thus in every period t, every actor i is endowed with a propensity (i.e., the probability) to realize each action $\sigma_i \in \Sigma_i$; call this $p_{i,t}(\sigma_i) \in [0,1]$ with $\Sigma_{\sigma_i \in \Sigma_i} p_{i,t}(\sigma_i) = 1$. Let $p_{i,t}$ denote actor i's vector of propensities. Each actor is also endowed with an aspiration level, call it $a_{i,t}$. Depending on $p_{i,t}$, an action σ_i is realized for each i. Given σ, the agent's payoff is realized as $\pi_{i,t}(\sigma)$. Realized payoffs are then compared to aspiration levels, which may lead to the adjustment of propensities or aspirations for the next period.

In our model we wish to capture agents that learn by trial-and-error; i.e., propensities and aspirations may adjust to payoff experience. However, since an actor's attention may be on other matters, these codings do not invariably lead to adjustments in propensities. Consistent with the spirit of bounded rationality, we allow for the possibility that humans are sometimes inertial: they do not invariably adapt or learn. Thus, with probability ε_p, an agent does not adjust her propensity in a current period. Similarly, with probability ε_a, an agent may not adjust her aspiration level. For simplicity we assume that ε_p and ε_a are mutually independent and i.i.d. across both agents and periods.

Propensities and aspirations may be adjusted randomly or deterministically. For each i, σ_i and t, we assume that there is a finite, time-invariant number of propensity levels (not necessarily the same for each individual i), denoted $p_i^1(\sigma_i)$, ..., $p_i^{l(i)}(\sigma_i)$ with $l(i) > 1$. So, $p_i^1(\sigma_i)$, = $p_i^{\min}(\sigma_i)$ and $p_i^{l(i)}(\sigma_i)$ = $p_i^{\max}(\sigma_i)$. To simplify the analysis we assume, for each i, that if some action receives i's maximal feasible propensity value in t, then all the other actions receive his minimal propensity value. (This is ensured, of course, if i has only two actions in the stage game, no matter what his maximal feasible propensity is. It also is ensured if i's maximal feasible propensity is 1.0, no matter how many actions there are.) To capture random propensity adjustment, we define for each i and σ_i a family of random variables $\{P_{i,t}(\sigma_i)\}_{t\in \mathrm{N}}$ with values drawn from $p_i^1(\sigma_i), ..., p_i^{l(i)}(\sigma_i)$. Propensity adjustment then corresponds to a (stochastic) dynamic process. By putting point mass on one of the possible realizations we can capture deterministic adjustments, e.g., the Bush-Mosteller rule.

As in the case of propensities, we assume that for each i and t, there is a finite, time-invariant number of aspiration levels (not necessarily the same for each individual i), denoted $a_i^1, \ldots, a_i^{m(i)}$ with $m(i) > 1$. Again, we allow for random aspiration adjustment (with point mass in the case of deterministic adjustments). So, formally for each i: $\{A_{i,t}\}_{t\in\mathrm{N}}$ is a family of (possibly degenerate) random variables with values drawn from $a_1, \ldots, a_{m(i)}$. We assume that $\{P_{i,t}(\sigma_i)\}_{t\in\mathrm{N}}$ and $\{A_{i,t}\}_{t\in\mathrm{N}}$ are mutually independent stationary processes.

When they adjust (i.e., when agents are not inertial), propensities as well as aspiration levels respond to experience, i.e., to the payoffs received in previous rounds. We consider a very general class of adjustment rules for both propensities and adjustments. These assumptions on propensity and aspiration adjustment formally define our concept of an Aspiration-Based Adjustment Rule (ABAR).

Propensities:

(P1) (positive feedback). For all i, t and action σ_i chosen by i in t:
- if $\pi_{i,t}(\sigma) \geq a_{i,t}$, then $\Pr(p_{i,t+1}(\sigma_i) \geq p_{i,t}(\sigma_i)) = 1$;
- if $\pi_{i,t}(\sigma) \geq a_{i,t}$ and $p_{i,t}(\sigma_i) < p_i^{l(i)}(\sigma_i)$, then $\Pr(p_{i,t+1}(\sigma_i) > p_{i,t}(\sigma_i)) = 1$.

(P2) (negative feedback-direct effect). For all i, t and action σ_i chosen by i in t:
- if $\pi_{i,t}(\sigma) < a_{i,t}$, then $\Pr(p_{i,t+1}(\sigma_i) \leq p_{i,t}(\sigma_i)) = 1$;
- if $\pi_{i,t}(\sigma) < a_{i,t}$ and $p_{i,t}(\sigma_i) > p_i^1(\sigma_i)$, then also $\Pr(p_{i,t+1}(\sigma_i) < p_{i,t}(\sigma_i)) = 1$.

(P3) (negative feedback-indirect effect). For all i, t and actions σ_i and $\sigma_i'(\neq \sigma_i)$: if i chose σ_i in t, then:
- if $\pi_{i,t}(\sigma) < a_{i,t}$, then $\Pr(p_{i,t+1}(\sigma_i') > 0) > 0$.

Aspirations:

(A1) For all i, t:
- if $\pi_{i,t}(\sigma) > a_{i,t}$, then $\Pr(\pi_{i,t}(\sigma) \geq a_{i,t+1} > a_{i,t}) = 1$.

(A2) For all i, t:
- if $\pi_{i,t}(\sigma) = a_{i,t}$, then $\Pr(a_{i,t+1} = a_{i,t}) = 1$.

(A3) For all i, t:
- if $\pi_{i,t}(\sigma) < a_{i,t}$, then $\Pr(\pi_{i,t}(\sigma) \leq a_{i,t+1} < a_{i,t}) = 1$.

Note that because aspirations can adjust to experience, eventually all agents' aspirations will be drawn toward the set of feasible payoffs. These assumptions formally capture the two central concepts of aspiration-based learning. The first feature is feedback: if the payoff associated with an action taken exceeds the aspiration level (i.e., if it is coded as a success), its propensity will increase (if the agent is not inertial and the current propensity isn't already maximal); if it is coded a failure, the agent's propensity to choose it in the future will decrease (if the agent is noninertial and the current propensity isn't already minimal). Note that in the case of games with only two actions, (P3) follows from (P2). The second feature is the assumption of endogenous aspirations: over time, aspirations adjust to payoffs received.

We can now describe a full cycle of learning. In each period t, an agent is endowed with a vector of propensity levels $p_{i,t}$ and an aspiration level $a_{i,t}$. Initially (i.e., for $t = 1$) these levels are assigned arbitrarily. Given the realized action of each agent, each agent receives a randomly drawn payoff that is conditional on the outcome of the election and her own action. This leads to a propensity adjustment with probability $1 - \varepsilon_p$, and to an adjustment of the agent's aspiration level with probability $1 - \varepsilon_a$. So, with probability $\varepsilon_a \varepsilon_p$ the agent is completely inertial.[1] Propensity adjustment occurs according to some adjustment process consistent with axioms (P1)–(P3). In the case of aspiration adjustment, axioms (Al)–(A3) must be satisfied. This cycle is depicted in Figure 1.

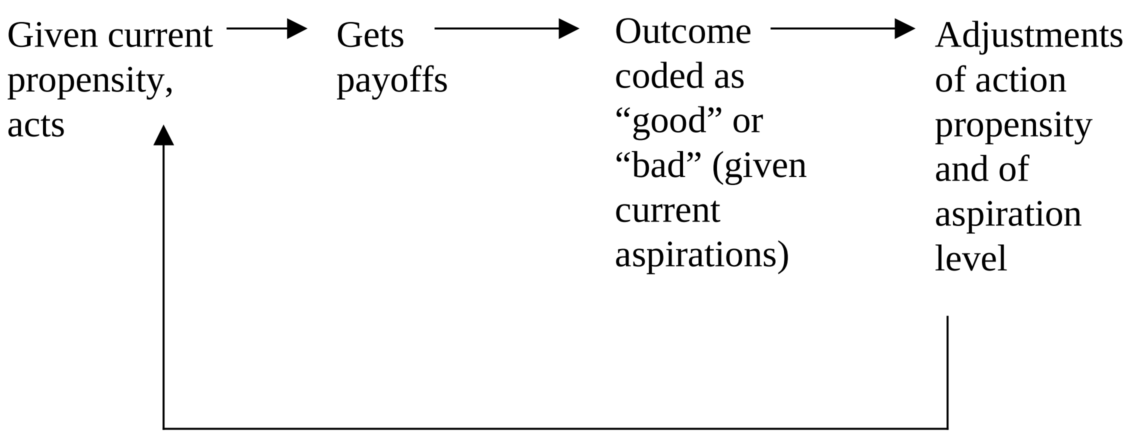

FIGURE 1 Full learning cycle

An example for an ABAR in a model with two actions is the well-known Bush-Mosteller model of adaptive learning. If an actor who takes action σ_i and happens to be noninertial in a given period codes the outcome as successful (i.e., if $\pi_{i,t}(\sigma) \geq a_{i,t}$), then

$$p_{i,t+1}(\sigma_i) = p_{i,t}(\sigma_i) + \alpha(1 - p_{i,t}(\sigma_i)),$$

where $\alpha \in (0,1]$ represents the speed of learning or adaptation, given a successful outcome. Similarly, if the outcome was coded as a failure, then

$$p_{i,t+1}(\sigma_i) = p_{i,t}(\sigma_i) - \beta p_{i,t}(\sigma_i),$$

[1] Note that an agent may be inertial with respect to only propensities or aspirations or both.

where $\beta \in (0,1)$.

Further, tomorrow's aspirations are a weighted average of today's aspiration level and today's payoff (Cyert and March 1963):

$$\alpha_{i,t+1} = \lambda a_{i,t} + (1 - \lambda)\pi_{i,t}(\sigma),$$

where $\lambda \in (0,1)$.

Because we assume a finite state space — and hence, finitely many propensity aspiration values — these transition rules are approximate; we assume that actual values of $p_{i,t}$ and $a_{i,t}$ are rounded to three digits for all i and t.[2] Thus, this combination of adjustment rules is a special case of an ABAR. For example, it specifies linear adjustment rules where ABARs in general are not restricted to a particular functional form. Moreover, Bush-Mosteller adjustments are deterministic, while the class of ABARs also includes probabilistic adjustment rules.

Our model defines a discrete-time, finite-state Markov process. That is, we have a family of random variables $\{X^t : t \in \mathbb{N}\}$ where X^t assumes values on the state space $S = \times_{i=1,\ldots,\mathrm{N}} S_i$ where S_i consists of elements of the form $(p_i, a_i) =: s_i$. Generic states are thus of the form $(p_i, a_i)_{i=1,\ldots,\mathrm{N}}$, denoted s. Note that given the independence assumptions on $\{P_{i,t}(\sigma_i)\}_{t\in\mathrm{N}}$ and $\{A_{i,t}\}_{t\in\mathrm{N}}$, $\Pr(X^t = s' \mid s) = \Pi_{i=1,\ldots,n}\Pr(X_i^t = s_i' \mid s)$ where $\{X_i^t\}$ is the (decomposed) family of random variables assuming values on S_i. Since transitions $\{P_{i,t}(\sigma_i)\}_{t\in\mathrm{N}}$ and $\{A_{i,t}\}_{t\in\mathrm{N}}$ are stationary, we have a stationary Markov process. Our goal is to study the long-run behavior of this process. Using our definition of ABARs, we are able to derive the following important result about this process. (See the appendix for the proof.)

Proposition 1: *For every game with finitely many strategies, the stochastic process converges to the unique limiting distribution independent of the initial state.*[3]

Proposition 1 provides theoretical foundations for our simulations. First, it implies that irrespective of the starting state the process will converge to the same limiting distribution. Thus,

[2] The rounding rule works as follows. To satisfy (P1)–(P3) and (Al)–(A3), we round up for all upwards adjustment, and round down for downwards adjustment (when feasible). For example, if shirking did not satisfice in period t and $p_{i,t}(S) - \beta p_{i,t}(S) = 0.306$, then $p_{i,t+1}(S) = 0.300$.

[3] Convergence to a unique limiting distribution is a robust property of the model that holds for many different specifications of randomness in the agents' behavior. It can be shown to hold if any of the following hold: (i) players tremble when choosing actions; (ii) propensities of zero and one are infeasible; (iii) payoffs are stochastic; (iv) players' states (propensities and aspirations) can tremble to "neighboring" states. (Bendor, Diermeier, and Ting 2000a).

our simulation results do not depend on the initial state.[4] Second, an alternative interpretation of the limiting distribution is that it also gives the long-run mean fraction of time that the process occupies a given state. Therefore, by considering a single run (for each parameter configuration), we can capture the limiting behavior of our process as if it were run for many different initial states. Note also that Proposition 1 does not require agents to use the same ABAR, only that they use some stationary ABAR.

In the remainder of the paper we apply this general model to the turnout game and the PD. In most games of interest (including these two games), it is difficult or impossible to derive the limiting distribution analytically. We therefore use simulation techniques.[5] Specifically, we use Bush-Mosteller adjustments with randomly perturbed payoffs. That is, a payoff $\pi_{i,t}$ is just the game's normal form payoff plus a random shock $\theta_{i,t}$, where $\theta_{i,t} \sim U[-\frac{\omega}{2}, \frac{\omega}{2}]$ and is i.i.d. across players and periods. The parameter ω therefore represents the size of the support of the shock. Thus, by specifying a game's normal form, our adaptive model is well-defined.

Let us first consider the turnout game. Here, each agent has two choices, to vote or stay home ("shirk"). We assume that the electorate is of finite size N and is divided into two blocs or factions, of n_D Democrats and n_R Republicans, with $n_D \geq n_R > 0$, and $n_D + n_R = N$. The faction that turns out more voters wins. Ties are broken by a (not necessarily fair) coin toss. Members of the winning faction earn a deterministic payoff of $b > 0$ for winning, whether or not they voted; losers get d. Voting is costly (whether a voter's faction won or lost). The private cost of voting is $c < b$. Payoffs are additive in the benefits and costs. Thus, winning voters get $b - c$; winning shirkers get b. Losing voters get $d - c$; losing shirkers get d. Unless stated otherwise, we normalize b to 1 and d to 0.

III. SIMULATION PROGRAM

Although it is simple, our theoretical model yields no tractable analytical solution. To establish some sense of the limiting distribution of agent states, we have therefore written a set of routines in ANSI C to simulate the adaptively rational behavior we have described. These routines are compatible with the GNU C compiler and most UNIX operating system configurations.

[4] This claim strictly speaking only holds "in the long run," i.e., if we let the simulation program run long enough.

[5] For partial analytical results in the turnout game see Bendor, Diermeier, and Ting (2000b).

As with our theoretical model, the computational model is readily adaptable to any normalform game. In all variants, it allows the user to specify the length and the adaptation parameters of each run. These include the number of simulations, the number of periods per simulation, the inertia probabilities ε_a and ε_p, the reinforcement and inhibition rates α and β, the aspiration updating rate λ, the size of the payoff support ω, a vector of initial propensities for each action for each type of player $p_{i,1}$, and initial aspirations $a_{i,1}$ for each player. The core routines are then mated with different "shells" to describe specific games. In a 2×2 game, for example, the shell additionally allows the user to specify the payoff matrix. In an election game, the user must specify each faction's size and payoff from winning or losing the election.[6]

Roughly speaking, the program works as follows. When a run begins, the program initializes a pseudorandom number generator with an integer representing the current time. This standard procedure effectively ensures that each run's random parameters are independent of those of other runs. It also allows runs to be replicated easily, since each stream of random numbers is described by a single integer, the random seed. The program then initializes a custom data structure that keeps track of state variables (i.e., propensities and aspirations, which may change over the course of a run) and statistics related to the history of play. These variables revert to their original values for each new simulation.

In each period, moves are realized in accordance with the underlying action propensities. After payoffs are revealed, the data structure is updated to reflect the (possibly) changed propensities and aspirations, as described above. All computers are finite-state machines; thus, any implementation of the updating mechanism would provide the discrete state space required by our model. However, a key feature of our routines is that they resolve the discrete state spaces internally. Thus, the state spaces are independent of the particular platform on which the program is run. To do this, we assume that the state space has a fixed number of decimal digits (three by default). Aspirations are then simply rounded to the nearest thousandth. Propensity updating is more complicated because of the constraint that $\Sigma_{s_i \in S_i} \; p_{i,t}(s_i) = 1$. When propensities are recalculated, all values are truncated (not rounded) to three digits. The excess thousandths are then redistributed to those action propensities which lost the most from the truncation. Thus, if a player has updated raw propensity values of 0.3338, 0.3337, and 0.3325, her propensities for the next period become 0.334, 0.334, and 0.332.

The data structure associated with each run can be used to recover statistics at different levels of the run. To reduce runtimes, only those statistics that are requested by the user are

[6] It is also possible to substitute in behavioral models other than the Bush-Mosteller; for example, by allowing a small sub-population of voters to vote if their likelihood of being pivotal is sufficiently high.

collected. The set of statistics that may be collected varies according to the game being investigated, but typical sets include the following: (i) the moves, payoffs, and adjusted propensities and aspirations after each period; (ii) average and cumulative propensities and aspirations for each simulation, with a histogram of propensities for each agent for each simulation; (iii) average propensities and aspirations across simulations for certain periods; and (iv) a histogram of final-period propensity and aspiration levels across simulations. For most runs, the final set of statistics is the most useful, as it provides the most direct sense of the limiting distribution of the process.

IV. APPLICATIONS

To illustrate our computational model at work, we present examples of two common games. The run reports associated with each game display their most important starting parameters, as well as histograms of final period propensities across simulations. While the program does not generate graphics on its own, the output files it creates are easily read into other programs for further processing.

The first example illustrates the well-known prisoner's dilemma (see Run 1 at end of paper).

The figure plots the number of players (out of 1,000) who choose "cooperate" at the end of each run as we vary the relative payoff from defecting. When c = 0.5, defection produces a high payoff, and 78% of players do so in period 1,000 of our simulations. When c = 1.5, defection becomes less appealing: only 35% of players do so. Agents are thus predominantly playing the *strictly dominated* (non-Nash) strategies of cooperating. Note that while classical game-theory predicts universal defection independent of the cost value (as long as the game is a PD) our model predicts widespread cooperation among agents. Moreover, the relative frequency of cooperation depends critically on the value of c.

The pattern may be explained by the role of endogenous aspirations. When c is low, the mutual defection payoff is much higher than 0, the payoff from unilateral cooperation. Further, the payoff from unilateral defection is much higher than the payoff from mutual cooperation. It is therefore more likely for players to experience mutual cooperation as unsatisfactory. Hence, they will begin *to mutually defect*, and continue doing so most of the time. When c is high, the payoff from mutual defection is not much bigger than 0, while the payoff from unilateral defection is only slightly better than $b/2$. Defection thus has only a small effect on aspirations, and cooperation will satisfice.

The second example (see Run 2) illustrates the turnout game, the most famous anomaly for the rational choice approach in the study of politics (e.g., Palfrey and Rosenthal 1985). In this game, two factions (D and R) elect an official. The factions sizes are n_D = 5,00l and n_R = 5,000. Each voter may either vote for her preferred candidate, or abstain. We assume that candidates and policy platforms are fixed; thus, members of the winning faction receive a payoff of 1, members of the losing receive 0, and all citizens who vote pay a cost of 0.25.

This result contrasts sharply with those of game-theoretic models of turnout (e.g., Palfrey and Rosenthal 1985) which predict very low turnout levels for all but the smallest electorates. As we show in Bendor, Diermeier, and Ting (2000b), turnout of about 50% may occur even in electorates of *a million* voters. Further, as in the PD, varying the payoffs leads to changes in the prediction. For example, higher costs decrease expected payoff (though not linearly), and turnout decreases as faction size become more unequal.[7]

To illustrate the independence from starting values in Proposition 1, consider Run 3. In all 1,000 simulations, vote propensities are in the neighborhood of 0.5 at the end of period 1,000. This occurs despite the initially low starting propensities of all voters.

V. CONCLUSION: MULTI-AGENT MODELLING

As we noted in the introduction, behavioral game theory tries to synthesize classical game theory and behavioral decision theory by uniting the strategic focus of the former with the cognitive focus of the latter. The present model is an example of how that marriage might be consummated. We represent the strategic relations among agents by different normal form games (here, by the PD or the turnout game). From the perspective of classical game theory, this is a completely standard representation. But our agents are far from game-theoretically rational. Instead, they adapt via simple trial-and-error, in a way that respects the constraints of bounded rationality. Thus, our model is built on two distinct pillars, one from each of the “parent” research programs. Though of very different intellectual lineages, these two pillars — the strategic environment of the normal form game and the agents’ adaptive processes — are consistent with each other. The marriage is possible.

It is also fruitful. For example, in the turnout game, the model reveals that citizens who adjust by aspiration-based adaptation vote in substantial even in large electorates. Thus, the union

[7] See Bendor, Diermeier, and Ting (2000b) for details.

of game theory and adaptive rationality proposes a solution to a long standing problem in political science, the so-called “paradox of voting.”

Finally, it worth emphasizing that our adaptive mechanism can be plugged into any normal form game whatsoever. These adjustment processes are completely separable from the kind of strategic environment that agents face. Thus, the domain of the marriage’s offspring is unlimited.

ACKNOWLEDGMENT

We would like to thank Sunil Kumar.

Run 1: Prisoner's Dilemma

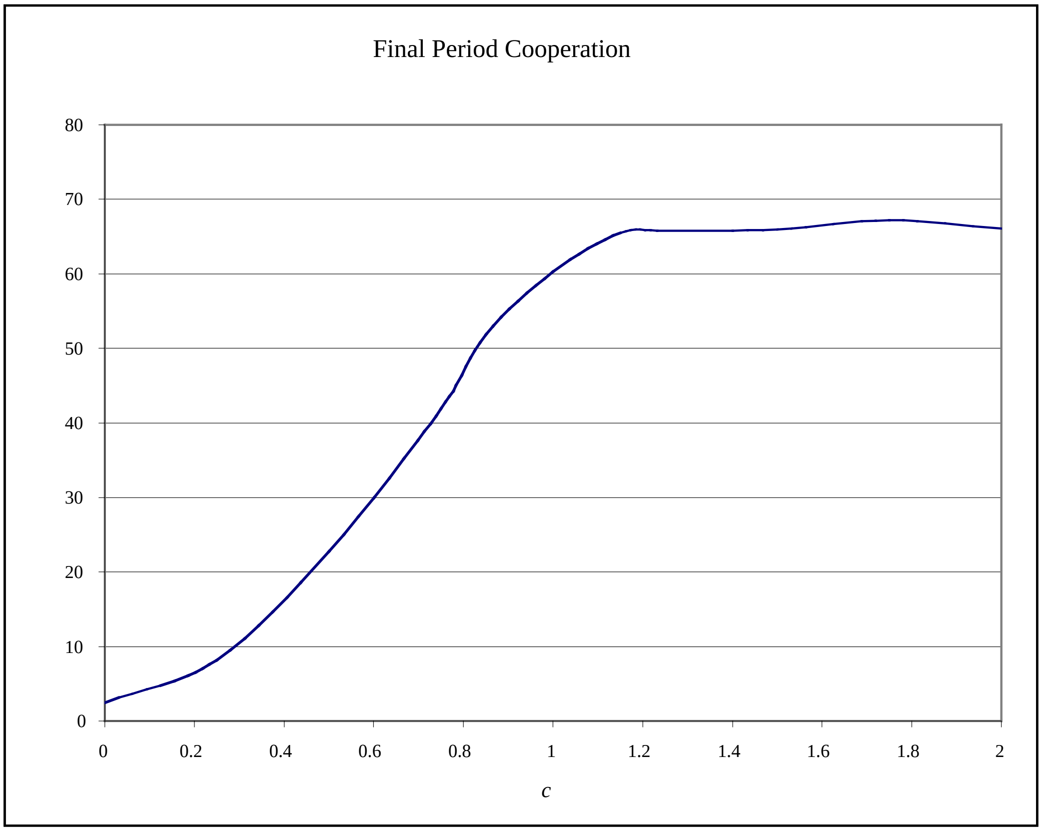

Starting Values: 1,000 Periods
5,000 Simulations

Payoff Matrix:

	C	**D**
C	2,2	0, 4–c
D	4–c, 0	2–c, 2–c

Run 2: The Turnout Game

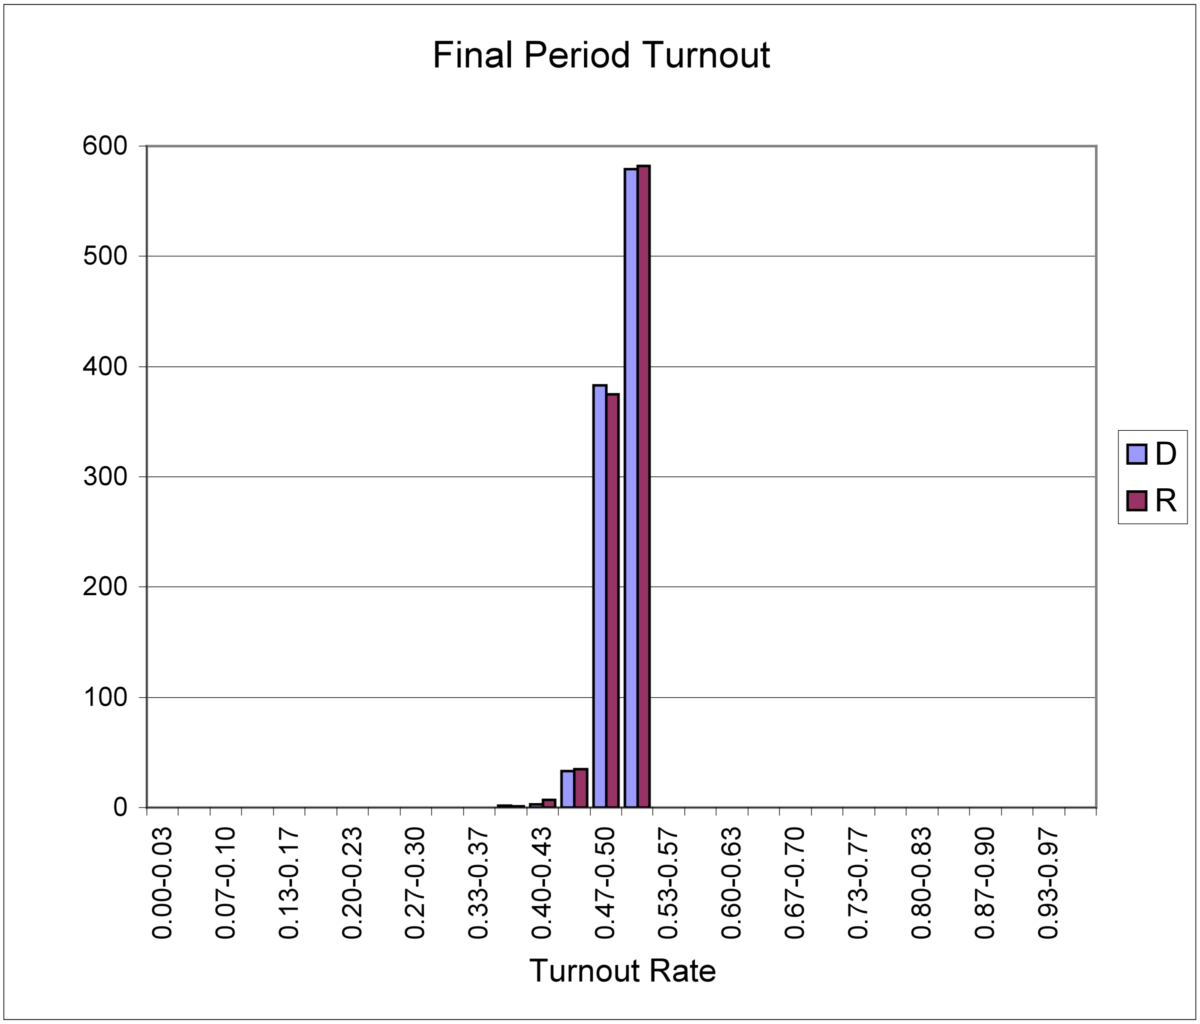

Starting Values: 1,000 Periods
1,000 Simulations

Faction	D	R
Population	5,001	5,000
b	1.0	1.0
c	0.25	0.25
Aspirations	0.5	0.5
Vote Propensities	0.5	0.5

Run 3: Turnout with Low Initial Propensities

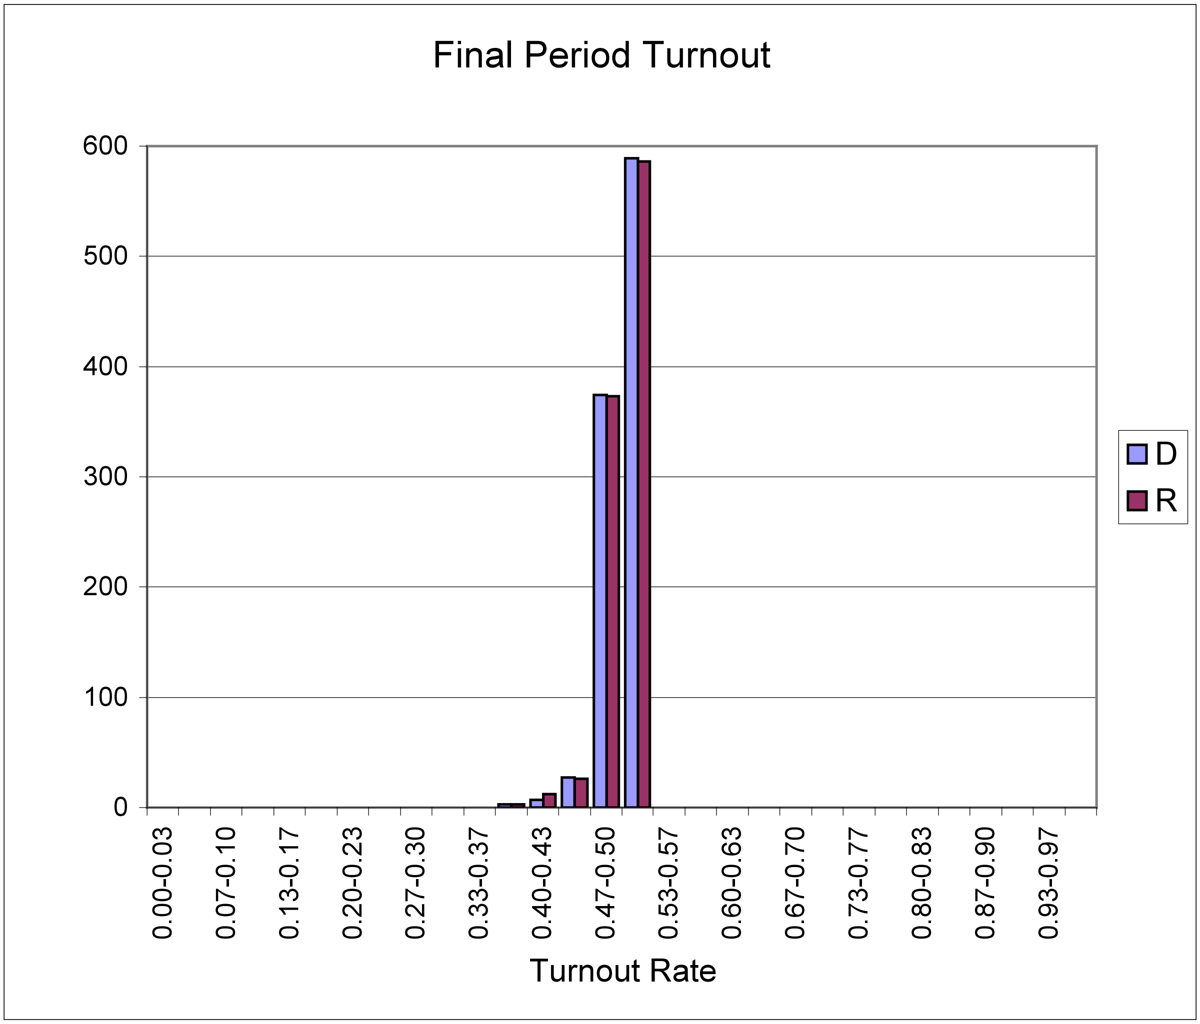

Starting Values: 1,000 Periods
1,000 Simulations

Faction	D	R
Population	5,001	5,000
b	1.0	1.0
c	0.25	0.25
Aspirations	0.5	0.5
Vote Propensities	0.01	0.01

APPENDIX

To show Proposition 1, we use standard results from the theory of Markov chains (e.g., Feller 1968). In particular, to demonstrate that our chain converges to a unique limiting distribution, we show that our chain must satisfy either of the following two conditions, each of which suffices for ergodicity: (1) our chain is aperiodic and irreducible, or (2) it can be partitioned into two nonempty sets — a closed, irreducible and aperiodic set of nontransient states and a set of transient states.

That condition (1) suffices for ergodicity is a standard result (ibid., p. 394). That condition (2) suffices follows from the conjunction of two well-known facts. First, for any chain that satisfies the partition of (2), with probability one it must eventually be absorbed into the set of nontransient states (Karlin and Taylor 1975, p. 90). Second, any closed set may be treated as a distinct new Markov chain, and analyzed independently (Feller 1968, p. 384). And because the closed set is aperiodic and irreducible, it is ergodic. Together, these two facts imply that if condition (2) holds, then with probability one our chain must be absorbed by a set of states that is an ergodic Markov chain; hence, the original (undecomposed) chain is itself ergodic.

Aperiodicity follows immediately from inertia, which holds for every state. So we need only show that (1′) our chain is irreducible, or (2′) it can be partitioned into two nonempty sets — a closed, irreducible set of nontransient states, and a set of transient ones. For this the following lemma is useful.

Lemma. *If a finite, stationary Markov has a state that is accessible from all states, then either it is irreducible or its set of states can be decomposed into two nonempty sets — a closed irreducible set of nontransient states, and a set of transient ones.*

Proof: Call the distinguished state s^*. Since s^* is accessible from all states, it must belong to every closed set of states in the Markov chain. By Theorem 3 of Feller (1968, p. 392), any Markov chain can be partitioned, in a unique way, into nonoverlapping sets T, C_1, C_2, ..., where T is composed of all transient states and each C is closed and irreducible. Since these closed sets are disjoint and s^* belongs to each of them, our chain must have only one closed set. Hence, either the set of transient states is empty, whence (1′) holds, or it is not, whence (2′) does.

The proof of the Proposition relies on the Lemma by identifying a state that is accessible from all other states. We use "wpp" for "with strictly positive probability." To denote an arbitrarily chosen state s at time t we write $s_i = (p_i(\sigma_i^1), ..., p_i(\sigma_i^j), ..., p_i(\sigma_i^{l(i)}); a_i) =:$

$(p_i^{-j}(\sigma_i^{-j}),\ p_i(\sigma_i^j);\ a_i)$. Our candidate for a state s^* that can be reached from every state is defined as follows: each actor i puts maximal propensity on action σ_i^1, and every i's aspiration level is the maximal possible payoff given $(\sigma_1^1, \ldots\ \sigma_i^1, \ldots, \sigma_N^1) =: \sigma^*$. So for each i we have $s_i^* = (p_i^{-1}(\sigma_i^{-1}), p_i^{\max}(\sigma_i^1); a_i = \pi_i^{\max}(\sigma^*))$.

The proof consists of two steps. In period t consider an arbitrary state s and agent i.

Step (1):

If $p_{i,t}(\sigma_i^1) > 0$, then, given inertia, for all $(p_{i,t})_{-i} =: p_{-i}$ some action profile σ of the form $\sigma = (\sigma_i^1, \sigma_{-i})$ occurs at $t+1$ wpp. Also (given non-degenerate random probabilities) wpp agent i receives payoff $\pi_i^{\min}(\sigma)$. If $a_{i,t} = \pi_i^{\min}(\sigma)$ then, given inertia, we are done. If $a_{i,t} > \pi_i^{\min}(\sigma)$ then by (A3) $a_{i,t+1} < a_{i,t}$ wpp. If $a_{i,t+1} = \pi_i^{\min}(\sigma)$, we are done. Otherwise, repeat this procedure while for all $t' > t$ keeping $p_{i,t'}$ "frozen" at $p_{i,t}(\sigma_i^1)$, and $(p_{i,t'})_{-i} = p_{-i}$. (This occurs wpp given inertia.) At some finite τ, $a_{i,t+\tau} = \pi_i^{\min}(\sigma)$, which yields wpp a state s' such that $s_i' = (p_i^{-1}(\sigma_i^{-1}), p_i(\sigma_i^1); a_i = \pi_i^{\min}(\sigma))$ with $s'_{k\neq i} = (p_{-i}, a_{-i,t+\tau})$ where $a_{-i,t+\tau}$ is some vector of feasible aspiration levels.

Suppose on the other hand that $p_i(\sigma_i^1) = 0$, then some profile $\sigma' = (\sigma_i^k, \sigma_{-i})$ must occur wpp (where $k \neq 1$). Using the same procedure as above, wpp agent i receives payoff $\pi_i^{\max}(\sigma')$ and for some finite τ, we have $a_{i,t+\tau} = \pi_i^{\max}(\sigma')$. Now again since at time $t+\tau$ we also have $p_i(\sigma_i^k) > 0$, $\sigma_i' = (\sigma_i^k, \sigma_{-i})$ occurs wpp such that i receives payoff $\pi_i^{\min}(\sigma')$. Since there are at least two possible payoff levels we have $\pi_i^{\min}(\sigma') < \pi_i^{\max}(\sigma') = a_{i,t+\tau}$. So, by (P3), wpp $p_{i,t+\tau+1}$, $(\sigma_i^1) > 0$ and we can use the same argument as above.

Step (2):

From step (1) we know that from any state s any state of the form $s_i' = (p_i^{-1}(\sigma_i^{-1}),\ p_i(\sigma_i^1);\ a_i = \pi_i^{\min}(\sigma))$ with $s'_{k\neq i} = (p_{-i}, a_{-i,t+\tau})$ and $p_i(\sigma_i^1) > 0$ can be reached wpp for each fixed i at some time t. Thus $\sigma_i = (\sigma_i^1, \sigma_{-i})$ occurs wpp such that i receives payoff $\pi_i^{\max}(\sigma)$. If $p_{i,t+1}(\sigma_i^1) = p_i^{\max}(\sigma_i^1)$, then we are done. If not, then by (P1) we have $p_{i,t+1}(\sigma_i^1) > p_{i,t}(\sigma_i^1)$. If $p_{i,t+1}(\sigma_i^1) = p_i^{\max}(\sigma_i^1)$ we are done. If not, repeat this procedure (while "freezing" the aspiration level at $\pi_i^{\min}(\sigma)$) until for some finite τ, $p_{i,t+\tau}(\sigma_i^1) = p_i^{\max}(\sigma_i^1)$, which yields a state s'' with $s_i'' = (p_i^{-1}(\sigma_i^{-1}),\ p_i^{\max}(\sigma_i^1);\ a_i = \pi_i^{\min}(\sigma))$. At this point freeze the agents' propensities, and repeat the action profile σ while letting i receive $\pi_i^{\max}(\sigma)$. By (A1), eventually i's aspiration must hit $\pi_i^{\max}(\sigma)$, yielding $s_i''' = (p_i^{-1}(\sigma_i^{-1}),\ p_i^{\max}(\sigma_i^{1});\ a_i = \pi_i^{\max}(\sigma))$.

Since i was chosen arbitrarily, steps (1) and (2) can now be used for some other $j \neq i$ with state s''' as the starting state. Repeating these steps for all actors yields s^* wpp.

REFERENCES

Bendor, Jonathan, Daniel Diermeier, and Michael Ting. 1999. "Recovering Behaviorism: Adaptively Rational Strategic Behavior with Endogenous Aspirations." Unpublished manuscript. GSB: Stanford University.

Bendor, Jonathan, Daniel Diermeier, and Michael Ting. 2000b. "Ergodicity in Models of Adaptive Behavior with Endogenous Aspirations — A technical note." Unpublished manuscript. GSB: Stanford University.

Bendor, Jonathan, Daniel Diermeier, and Michael Ting. 2000b. "A Behavioral Model of Turnout." Unpublished manuscript. GSB: Stanford University.

Bush, Robert, and Frederick Mosteller. 1955. *Stochastic Models of Learning*. New York: Wiley.

Camerer, Colin. 1990. "Behavioral Game Theory." In Robin Hogarth (ed.), *Insights in Decision Making*. Chicago: University of Chicago Press.

Cyert, Richard, and James March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.

Feller, William. 1968. *An Introduction to Probability Theory and its Applications* (third ed.). Vol. 1. New York: John Wiley.

Fudenberg, Drew, and David Levine. 1999. *The Theory of Learning in Games*. Cambridge, Mas.: MIT Press.

Karlin, Samuel, and Howard Taylor. 1975. *A First Course in Stochastic Processes* (second ed.). New York: Academic Press.

Palfrey, Thomas, and Howard Rosenthal. 1985. “Voter participation and Strategic Uncertainty.” American Political Science Review 79:62-78.

Simon, Herbert. 1955. “A Behavioral Model of Rational Choice.” *Quarterly Journal of Economics*. 69:99-118.

MODELING THE CO-EVOLUTION OF STATES AND NATIONS

L.-E. CEDERMAN, Harvard University*

ABSTRACT**

It can be hypothesized that wherever states manage to assimilate their peripheries prior to nationalism, nationalist transformations proceed comparatively peacefully. In cases where the cultural penetration is weak, however, the political and cultural maps clash, thus producing tensions that drive national secession, unification, and irredentism. To trace these processes, I propose an agent-based model that embeds nationalist mobilization in a dynamic state system. My preliminary findings confirm the main hypothesis: conflict was found to vary negatively with state-framed cultural centralization. Yet some of the high-assimilation cases feature extreme levels of conflict due to nationalist unification's undermining effect on the balance of power.

INTRODUCTION

This paper employs agent-based modeling as way to cut through the complex interplay between power and culture in world politics. As opposed to standard modeling strategies, which usually treat states and nations as fixed units, the computational framework allows me to endogenize both actor types as distinct, though co-evolving, units. My primary goal is to investigate the effect of specific causal mechanisms on macro-outcomes. More precisely, I investigate the hypothesis that relates nationalist violence to a mismatch between cultural and political borders: wherever cultural penetration manages to centralize political rule prior to nationalism, the transition to the nationalist era is likely to proceed relatively smoothly. By contrast, where states lack the infrastructural capacity to standardize culture, violent transitions involving secession and unification become more likely.

Although still subject to further sensitivity analysis, my preliminary findings confirm the hypothesis: as the strength of state-led cultural assimilation increases, the transition to nationalist politics usually becomes smoother and less violent. In situations characterized by little coupling of cultural evolution and state-formation, however, the geopolitical and cultural maps clash, thus triggering conflictual outcomes as state borders and national communities adjust to the underlying ethnic landscape. But this relationship is not deterministic. In some cases, strong state-framed assimilation removes only partially smaller-scale ethnic cleavages. In other cases, it provokes unification that undermines the preexisting balance of power, thus paving the way for potentially catastrophic turmoil.

---

* *Corresponding author address:* Lars-Erik Cederman, John M. Olin Institute for Strategic Studies, Harvard University, 1737 Cambridge Street, 411, Cambridge, MA 02138; phone: 617-495-2616; fax: 617-495-1384; e-mail: cederman@cfia.harvard.edu (Department of Government, Harvard University, from July 1, 2001).

** An earlier, more extensive version of this paper was prepared for presentation at the Annual Convention of the American Political Science Association, Washington DC, August 31-September 3, 2000. The current version owes much to the excellent computational advice of Laszlo Gulyas and Nick Collier. I am also very grateful for the generous support of the Robert Schuman Centre for Advanced Studies, Florence.

I rely on a new implementation of my previous computational research that focused primarily on state-formation (Cederman 1997). On the theoretical side, the extended model goes well beyond the earlier version by endogenizing both states and nations. At the heart of the new model, there is a geopolitical engine that generates emergent state systems through competitive pressures, including internal and external combat, resulting in conquest and secession. Within this environment, a multidimensional ethnic "landscape" is embedded. Based on cultural identities, nations form and take on behavioral importance for the development of the system's structure. These consequences include national secession from multinational states, national voluntary unification into culturally homogenous polities, and irredentism, involving both processes at the same time.

Methodologically, the new geocultural model also represents significant progress over earlier modeling solutions. In contrast to the latter, which was programmed in Pascal, the current framework is based on RePast, a comprehensive, Java-based simulation package developed at the University of Chicago.

The paper is organized in the three main sections. The first section provides an overview of the model. Then follows a section presenting a sample run with weak assimilation. The third section presents a contrasting run with strong state-formation. Finally, a concluding section discusses the heuristic insights gained from this exercise. It should be emphasized that this paper offers only a flavor of the model. A full technical description together with replication results can be found in Cederman (2000).

OVERVIEW OF THE MODEL

While this study focuses on studying the consequences of nationalism, it is first necessary to create both a cultural landscape and a geopolitical environment in which national identities can develop. Thus, before modeling the repercussions of nationalist transformations, I start by modeling the ethnic backdrop of state-formation as in pre-modern Europe.[1] I assume that the system, while far from perfect as a model of the pre-modern environment, initially consists of a large number of independent actors, which is subsequently followed by a phase of state formation. This process produces a balance-of-power equilibrium that prevents a limited number of states from conquering each other. During this pre-nationalist phase, state-led cultural assimilation takes place, but nationalist activity does not occur until nationalist mobilization is triggered. Assuming that a geopolitical equilibrium has been reached, the era of nationalism then begins. At this point, culture starts to matter for political mobilization. Given suitable conditions, nations will start to appear. What follows is a co-evolutionary process that links nation-building with changes in the state system.

Before turning to the model's behavior, it is useful to summarize the three main phases of the model graphically together with the interaction effects operating in each period (see Figure 1). All simulation runs start with a pre-modern period, which sees the creation of the initial grid of states. This is also when the age of migrations creates cultural landscape. After this

---

[1] To avoid conceptual confusion, it is important to keep apart the definitions of states and nations. For the present analytical purposes, it is sufficient to rely on Max Weber's (1946) classical distinction between states as sovereign organizations exercising legitimate control over a bounded territory and nations as cultural communities that strive to possess their own state.

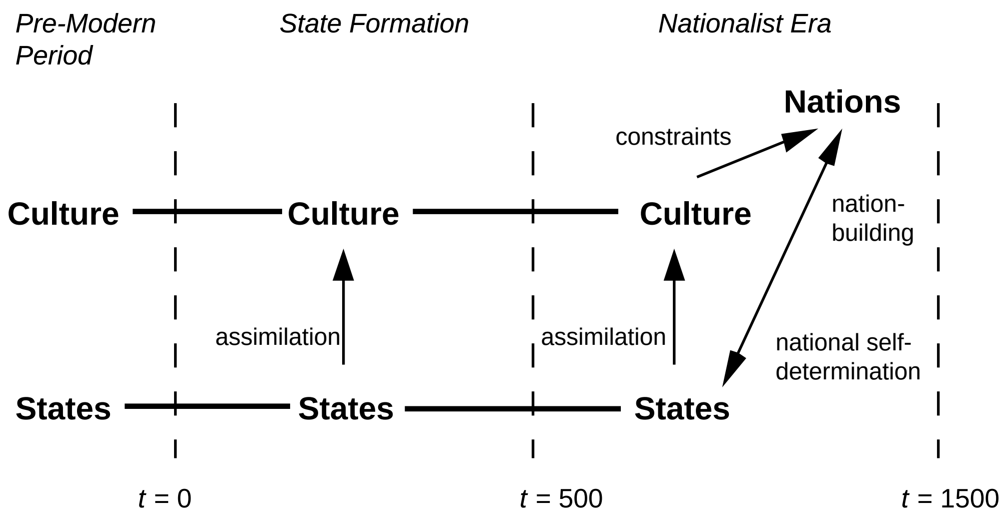

FIGURE 1 The Three-Phase Logic of the Model

stage, the era of state-formation unfolds. Most of the action at this point concerns interactions among states. To the extent that assimilation takes place, it affects the cultural identity of the states (and provinces). In the opposite causal direction, there is no impact at all, since culture is assumed to matter only once nationalism appears after time step 500.

From this point, the causal picture becomes more involved. Under nationalism, capitals still assimilate the culture of their provinces, but, in addition, a third type of entity emerges, namely nations. Their emergence can be attributed to political decisions on the part of states and provinces. Within the constraints of the cultural landscape, they decide either to create new nations or to join already existing national communities. Unlike the early modern period, culture starts to matter for political action, but only indirectly, through national identities. National self-determinations implies that national affiliations profoundly affect the behavior of the states. As illustrated by the sample runs, culturally "unhappy" national minorities resort to national secession, split nations come together, and states take irredentist action. This bi-directional link between states and nations constitutes the co-evolutionary nexus that is the main focus of our attention.

The present framework is based on an extended and stylized version of my previous "emergent polarity" model, developed in Cederman (1997, Chap. 4). As in that model, the world initially consists of a grid of unitary states (for earlier frameworks of this type, see Bremer and Mihalka 1977; Cusack and Stoll 1990). To give room for cultural evolution, however, I use a 15×15 grid for all runs rather than a smaller grid size (see Figure 2). The figure reveals the square arrangement of the initial system's 225 statelets. The black boundaries mark the states' borders and the square dots represent "capitals."

From the very beginning, all actors acquire cultural identities and retain these whether they are sovereign or not. The shading in Figure 1 denotes the cultural differences among the local sites. While the darker squares correspond to the cultural border areas characterized by great differences, the brighter ones refer to "plains" in the cultural landscape. Each actor has a

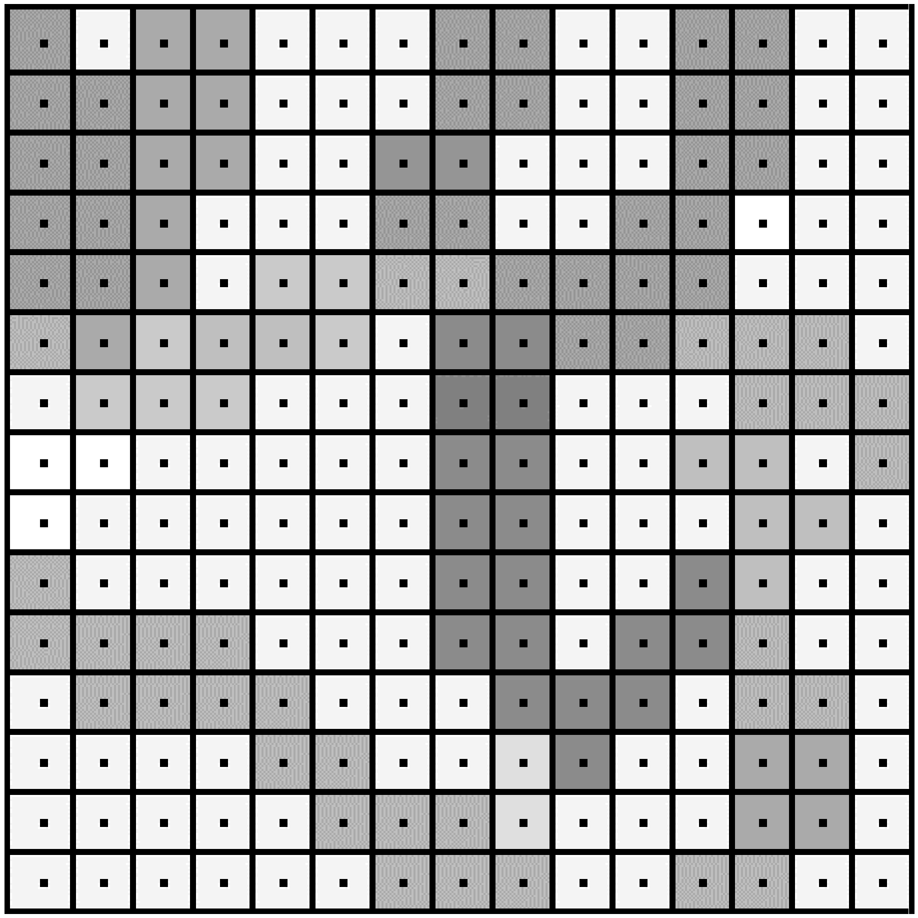

FIGURE 2 The Initial State of the Pre-Nationalist World

string of eight symbols that is assigned in a simulated "age of migrations" prior to the first time period (cf. Axelrod 1997, Chap. 7). This setup phase creates a synthetic cultural map that can be arbitrarily tuned.

Two parameters drive this process: the number of "ethnies" that participated in the cultural settling of the landscape and the "smoothness" defining the cultural similarity within each tribe's area of population. In our example, ten randomly located ethnic groups fill the whole grid with a trait-by-trait similarity of 0.90. This process guarantees that there will be both dialectal nuances as one moves from province to province as well as more abrupt ethnic cleavages.

Once unleashed, these actors will start interacting locally with their immediate territorial neighbors. Interaction can be of two types: either the actors peacefully exist side by side or they engage in combat. Conflict is initiated according to a simple decision rule, which depends on the local power balance between any two states. Roughly speaking, the actors play a "grim trigger" strategy with another. This means that they normally reciprocate whatever their neighbors did to them on their respective fronts. But whenever the power balance in their favor exceeds a preset threshold that governs the offense dominance of the system, they launch unprovoked attacks. The results reported in this paper are based on a fairly sharp threshold of 2.2 with some added noise.[2] The latter feature makes the decisions less than fully deterministic. In essence, this rule implies that all states consider attacking adjacent states once they get a little more than twice as powerful as any of them.

Once conflict erupts between two states, neither of them will resume cooperation until the battle is over. Combat ends with the victory of either party or, in some cases, with a stalemate. As in the decision rule for unprovoked attacks, the criterion for victory is a stochastic one with the same threshold at 2.2, though much more noise is involved here than in the decision function. In other words, the actors start to consider offensive behavior once their superiority becomes slightly more than a factor of two. Normally, the actor can count on winning, but because of the Clausewitzian “fog of war,” this may take a few iterations during which unpredictable things may happen. Most importantly, the attacking state may weaken its other fronts so much that opportunistic neighbors launch new offensives against the original aggressor.

Victorious parties gain the local territory fought over in any specific battle. If conquest destroys the supply lines from the capital under attack to its provinces, those parts that have been cut off regain independence. Should the capital itself be invaded, the entire state collapses. Thus, in offense-dominated systems, conquering states can grow quite quickly. Figure 3 illustrates the situation at time step 57 for the particular system depicted in Figure 1. At this point, the system’s polarity has been drastically reduced. Larger states with emergent boundaries have started to appear, though there are still a few smaller territories left waiting to be absorbed.

The figure also illustrates that action sometimes goes beyond interstate warfare, because in addition to such “horizontal” exchanges, the model features “vertical” two-level action. Graphically, rebellions of this type are marked with crosses, whereas small lines sticking into the neighboring sites denote interstate attacks. Even after their losing sovereignty, all provinces retain some resources and may thus rebel against the capital. This is likely to happen if the capital gets involved in too many other fronts, as illustrated by the state slightly northwest of the grid center. This actor is in deep trouble since it is under both internal and external attack.

The internal fronts are treated very much as any interstate combat theater: the same rules of deterrence and combat apply, with the exception that capitals have no reason to trigger conflicts with their peripheries. Should a rebellion occur, however, they always respond. While victorious rebellions are rewarded with secession, failures have no other consequences than the province remaining inside the state (see Cederman 1997, Chap. 5).

Shortly after the current time-step in the sample system, the neighboring actors around the collapsed state carve up the post-imperial remainders, and the system settles with seven surviving states (see Figure 4). In the following, these states will be referred to by letters ranging from A to G.

Despite the noise built into the decision rules, this is a rather robust equilibrium because even after 1500 time periods, it still remains intact. This is unsurprising given that the current model relies on a very simple territorial rule of resource updating. In fact, the resource

[2] This value was chosen since it generates stable geopolitical with moderate polarity equilibria quickly.

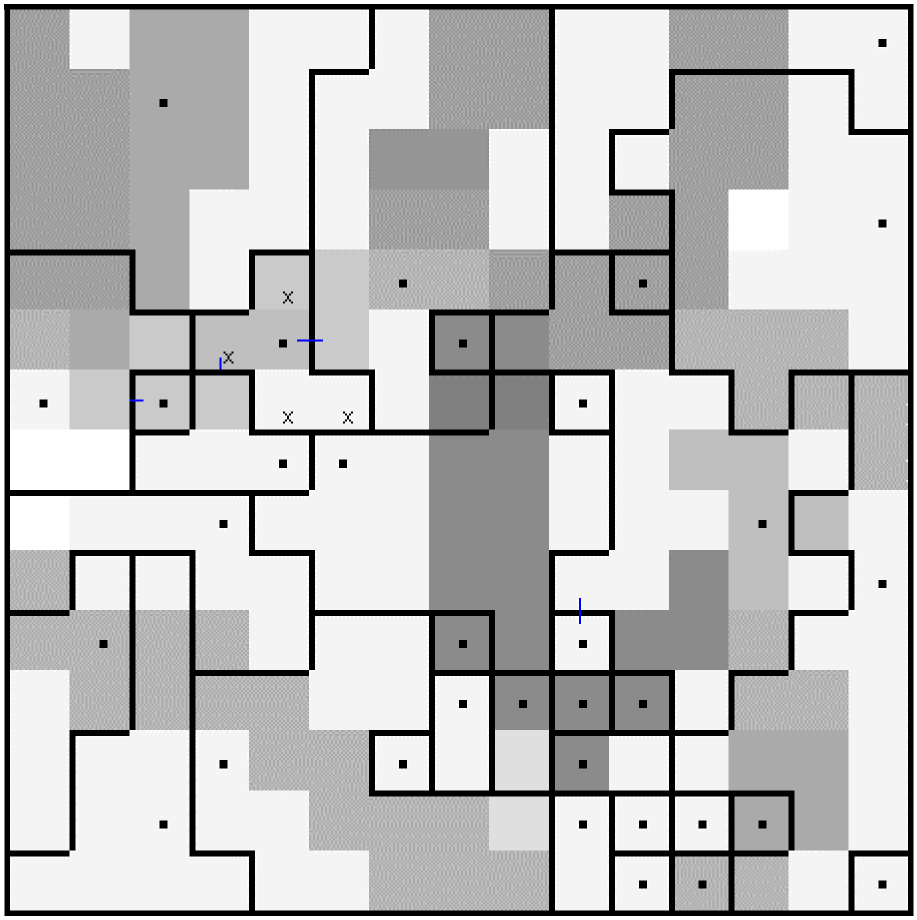

FIGURE 3 Secession Attempts and Interstate Combat at Time 57

"metabolism" of the system depends directly on the extent of the states' territorial holdings.[3] Each square in the grid yields a fixed amount of resources. A distance-dependent "tax rate" yielding, at most, 70% defines the power balance between the capitals and their conquered provinces.

THE NATIONALIST PHASE WITH WEAK NATION-BUILDING

Once having created a stable geopolitical environment without ruling out the possibility of changes, we are ready to introduce nationalism. Rather than treating nationalism as a constant "law," I model the phenomenon as a macro-historical process that exogenously hits an area at a specific point in history. To simplify things, I assume that this point occurs after 500 steps. The pre-national stage of the model was calibrated to generate a reasonably stable multipolar equilibrium well before this point.

---

[3] This mechanisms represents a simplification of the resource scheme used in Cederman (1997, Chaps. 4, 5). In that framework, a stochastic "harvest" was allocated to each state in each round. Rather than including accumulated resources, the current system lets a state's resource level be a direct function of the territory it controls.

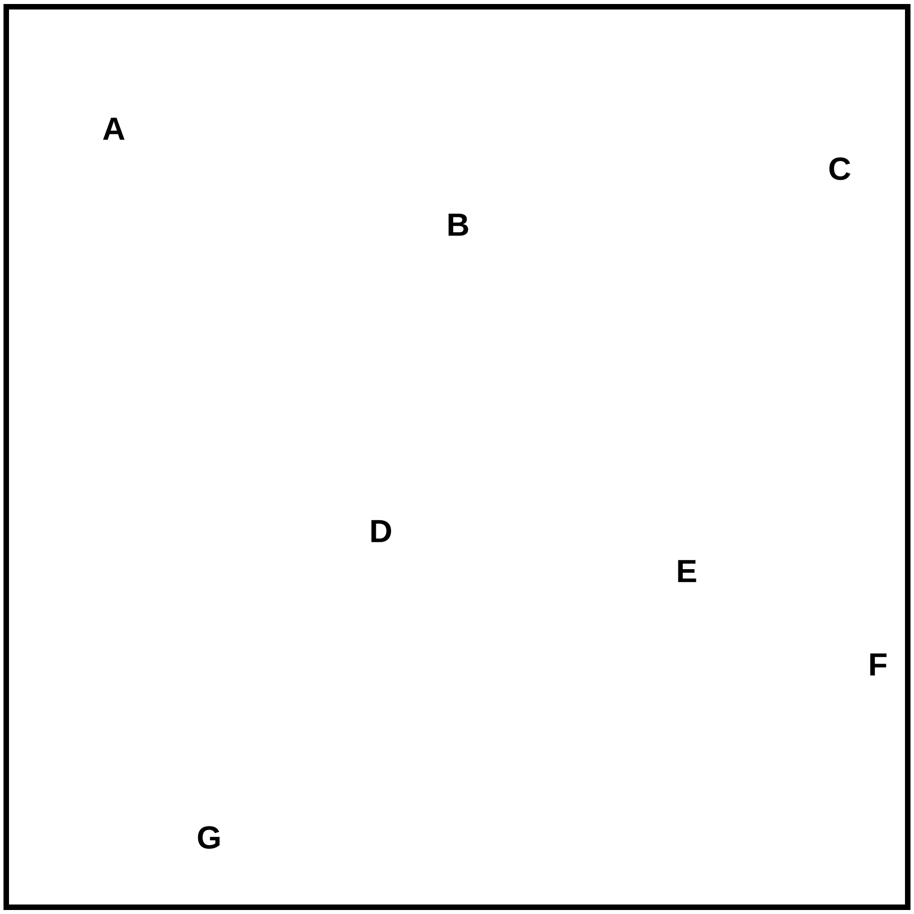

FIGURE 4 The Geopolitical Sample System in Equilibrium at Time 500

Nations are represented as distinct groups to which any state or province can belong if the nation's identity fits the actor's own culture. While both capitals and provinces are eligible members of nations, the probability of launching a nationalist movement depends crucially on the geopolitical status of the territory in question. Thanks to their resources, capitals have a much higher likelihood of founding their own nations, but provinces may sometimes create nationalist platforms in opposition to their respective capitals.

Figure 5 illustrates national mobilization in our sample system. At time period 536, precisely 36 steps after the onset of nationalism, three nations have formed, each one marked with a number surrounded by gray (green) boundaries. Of the seven capitals, three have their own nations. The most interesting development relates to nation 1, which started to develop within state D but soon spilled over into the territory of state A. This was to be expected given the shape of the cultural landscape. Capital D is located in a cultural basin of high similarity that intersects the borders with states A, B, and G. The settlements of 1-nationals in state A are far away from the capital, so it is not surprising that that they started a nationalist secession campaign (see the crosses).

The settlements of nation 1 inside state A violate the most basic principle of nationalism: namely that each nation should possess its own state, or, for short, there should be national

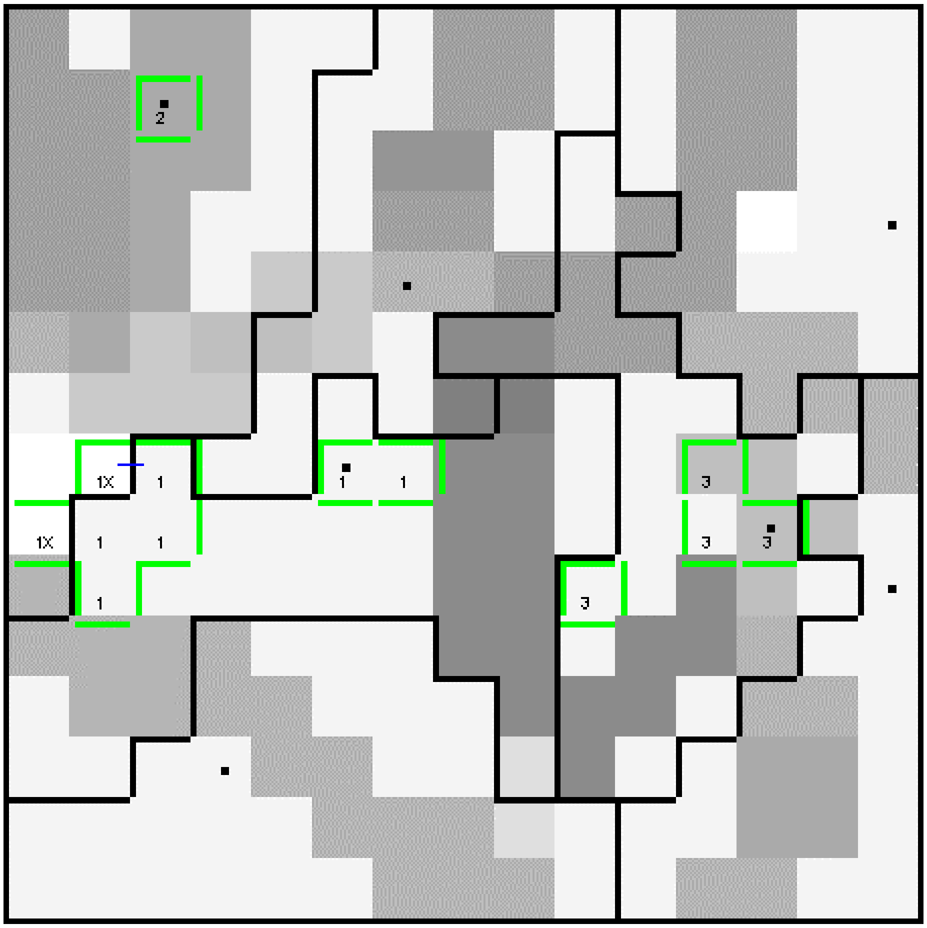

FIGURE 5 Irredentist Action in the Era of Nationalism at Time 536

self-determination. The discrepancy can be rectified both internally and externally. As pointed out by Myron Weiner (1971), these two phenomena often appear together. Whereas the former possibility corresponds to nationalist secession, the latter relates to irredentism. The behavioral rules of the nationalism extension turn all mobilized capitals and provinces into nationalist actors without entirely depriving them of geopolitical "appetite."

In our sample history, the co-nationals' secession campaign obliges state D, which is also a member of nation 1, to take irredentist action in order to help its kin across the border. The resulting interstate warfare, which involves the forces of capital A against those of state D, compounds the secessionist civil war, pitting the provincial nationalists (marked by crosses) against capital A.

In general, nationalist capitals seek to liberate their nationalist kin in other states if these populations do not enjoy "home rule." Provinces that belong to a nation modify their strategy such that they try to jointly break out of "foreign rule." This calculation implies that the national communities' decision making can be coordinated and their resources pooled within the national community. Postulating a more uncertain and risk-taking mode of decision making characterizing nationalist politics, I have set the superiority threshold at 1.5 for both internal and external action with considerably more uncertainty than in the non-nationalist case. Without abandoning entirely power considerations, states and provinces acting on behalf of their nations are thus more trigger happy than when acting on their own behalf. Moreover, there is an automatic obligation to come to the rescue of any kin group fighting a third party. On the whole, these rules are more prone to drag nationalist actors into armed struggle than the purely geopolitical strategy.

Interstate combat between A and D continues for several rounds, the whole time combined with rebellious activity on the part of the nation-1 settlements in state A. But because the balance is fairly even between the two camps, there are no territorial changes for the time being. At time 568, however, fighting diffuses to state B as consequence of nation 1's mobilization spreading to the province within B immediately west of capital D. This development triggers immediate irredentist action on the part of D (see Figure 6). As this province secedes, it first becomes independent, but thanks to the unification mechanism, it quickly joins state D. Completely unprepared to fight, state B starts to crumble. In time period 572, there is a wave of rebellions in the northern part, reflecting the longer distance from the capital than in other parts of the territory. Moreover, state B's territorial neighbors A and C profit from its weakness.

Nevertheless, while losing some territory, B manages to retain its sovereignty. At time 605, it is not the existence of this state, but of state D, that is threatened (see Figure 7). As a result of continuous irredentist fighting on its western front, it has just lost a nation-1 province to

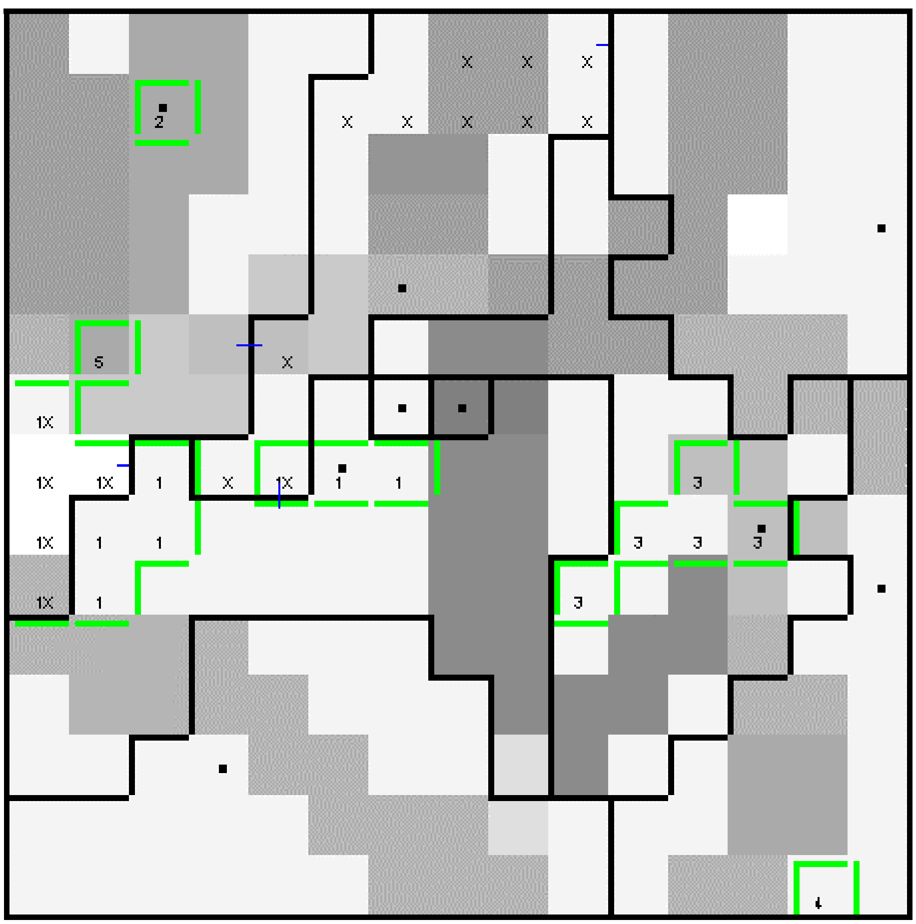

FIGURE 6 State B under Attack at Time 572

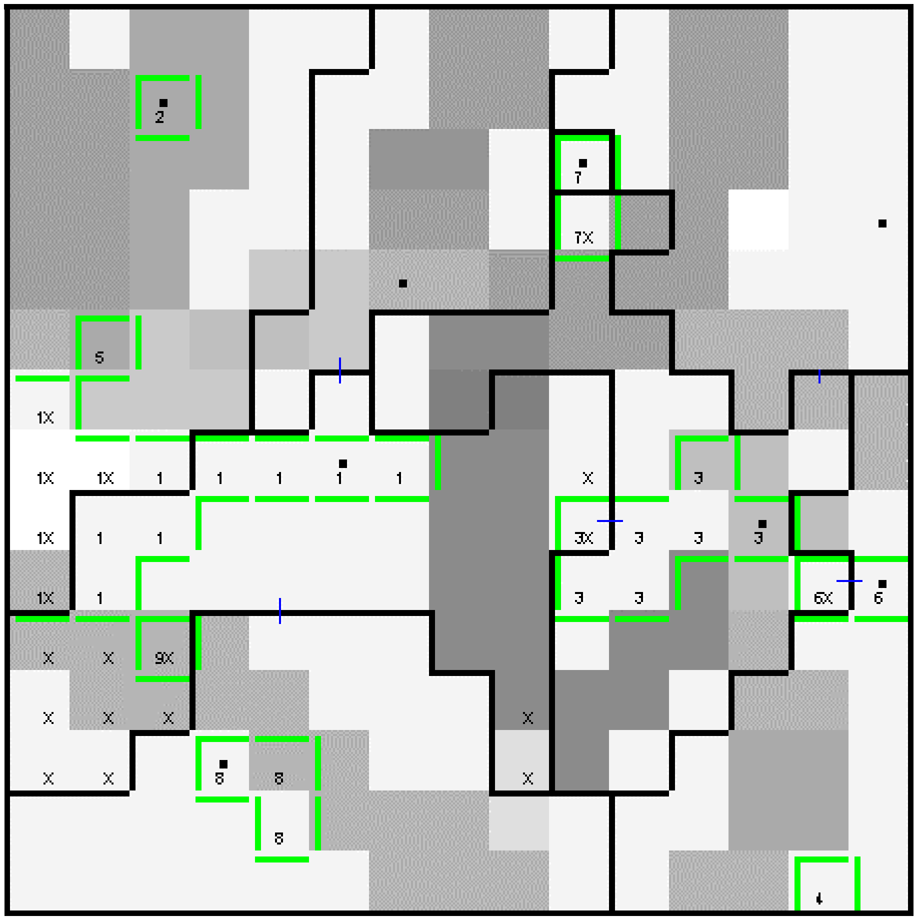

FIGURE 7 State D under Attack on Several Fronts at Time 605

state A. All the other sovereign neighbors have started wars against D, having diluted its power on too many fronts. In fact, on its eastern front, state E is attempting to redeem a newly mobilized group of nation-3 activists.[4] Finally, in the southwestern and southeastern parts of the country, there are regional rebellions.

Where does all this activity lead? Figure 8 portrays the long-run equilibrium after 1,500 iterations. This situation features four surviving states: A, C, E, and G. Unsurprisingly, state D never recovered from the "feeding frenzy" by its internal and external enemies. As a result of this development, nation 1 loses its sovereign political representation and becomes a national minority within state A. Note that some of the sites within this community have been converted to nation 2, the state-carrying identity of A. When state B finally collapsed, the capital of A was engaged in interstate conflict that provoked a rebellion on the part of the nation-1 minority. But once state B disappeared, state A turned all its power against the insurgents and crushed them. Under normal circumstances, national communities can never be destroyed, but in the event of a central power's subjugation of a nationalist secession attempt, the core sometimes

[4] State E needs to watch out, however, because state C has initiated a campaign against it, and state F is acting on an irredentist claim to a province inhabited by nation 6.

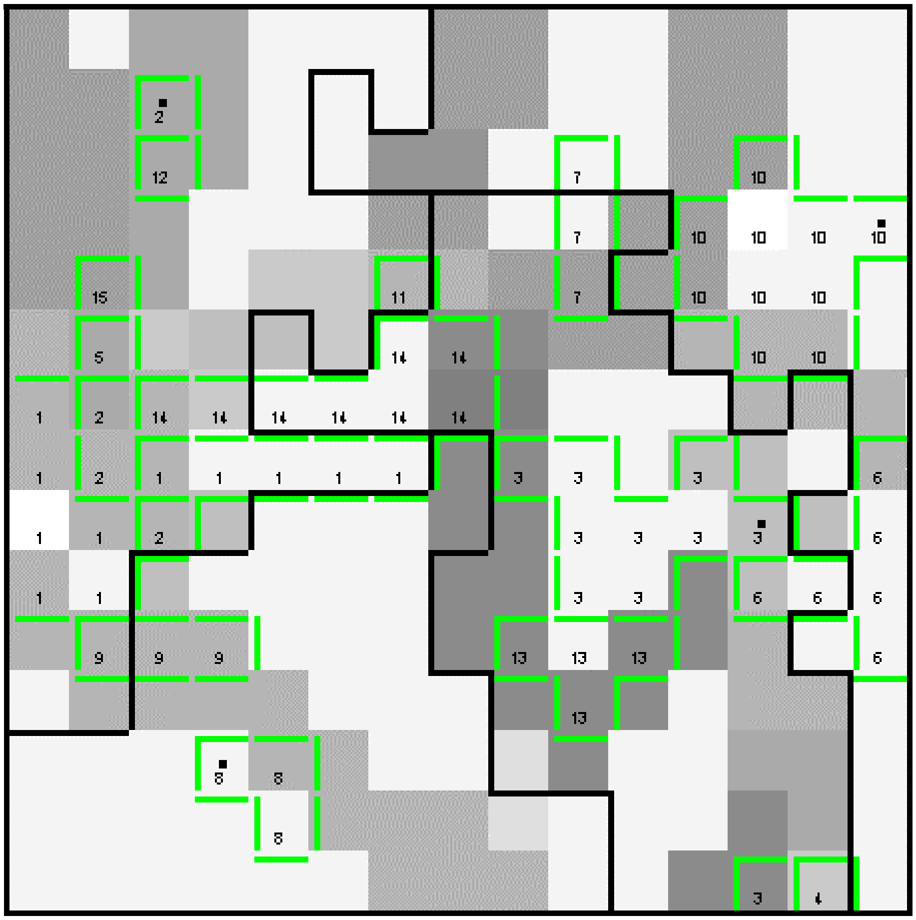

FIGURE 8 The Sample System in Equilibrium at Time 1500

gets the chance to "ethnically cleanse" the province through migration from one of the co-nationals residing within its territory. Such a transition also shifts the national identity to that of the capital. This is exactly what happened in our example. Indeed, the cultural plain that was previously monopolized by nation 1 has become fragmented by the nation-2 colonies.

Other nations have faired equally badly. The once state-controlling nation 6 has been split between states C and E after unsuccessful irredentist challenges to E's power. A few national communities are "unhistorical," since they formed too late to gain a state on their own (see nations 7, 13, and 14; each nation's number reflects its temporal appearance). The lucky nations are those that grew up around the capitals. There are two major basins of national activity (see nations 3 and 10). The cultural conditions in states A and G are less favorable for state-framed national mobilization. Capital A is surrounded by a "rugged" cultural region, preventing it from spreading the state's national identity 2 other than in the cleansed areas. In the case of G, it does not help to have the capital located in such a basin, because too specific a national identity prevents nationalism from catching on beyond the most immediate area.

AN ALTERNATIVE NATIONALIST HISTORY WITH STRONG STATE-LED ASSIMILATION

The assumption that the capitals have no influence over the cultural landscape within their states does not capture the conditions of nation-building in all areas. True, in some cases, as in Eastern Europe or the Third World, states have left a rather modest imprint on the cultural landscape (Schieder 1991; Gellner 1997). Even in the weakest cases of nation-building, however, one would expect some cultural spill-over beyond what the cultural background offers. In some prominent cases, states have indeed had a very strong impact on culture through centralized assimilation, even if it takes time for assimilation to penetrate large territories (E. Weber 1976). In early-modern Western Europe, such processes proceeded mostly unintentionally, as a side-effect of administrative standardization, military service, or commerce (Mann 1992). The model includes a parameter controlling the states' capacity to assimilate the culture of their provinces both before and after the onset of nationalism. As indicated in the introduction, the cultural assimilation parameter plays a central role as independent variable in this paper, and it is subject to controlled manipulation in the replications reported on below.

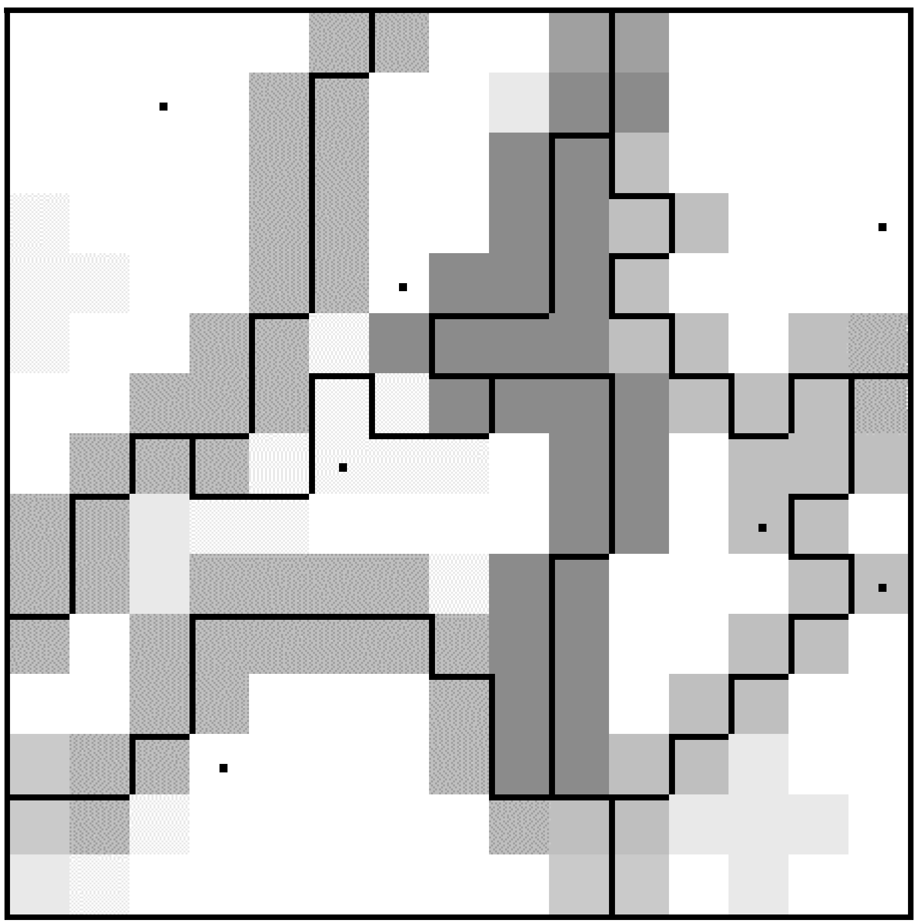

FIGURE 9 The Assimilatory System at Time 500

What would happen if we were to replay the "tape of history" with a high level of state-framed cultural penetration? As already stated, one of the main purposes of this paper is to investigate the geopolitical repercussions of different levels of cultural assimilation. So far, the grid has resembled eastern, rather than western, Europe, or perhaps even the Third World. Agent-based modeling makes it possible to rewrite history massively and to study the systemic effect of complex historical processes such as nationalism.

Starting with an identical system as in Figure 2, the states are now allowed to standardize culture quickly from the very beginning. Figure 9 shows the result of this changed specification at time 500. Compare this new setting to that of Figure 4 (as well as Figures 5 through 8), in which the ethnic landscape coincides with the initial cultural map. In the present case, clear boundaries coinciding with state borders have started to emerge. (Recall that whereas darker shades correspond to cultural border sites, the brighter ones denote areas with similar culture.)

Once nationalist mobilization is triggered, it will follow very different lines from those characterizing the culturally decentralized system (see Figure 10). In the counterfactual sample run, the activity starts in capital E and spreads quickly throughout its provinces at time 527, shortly followed by a similar process in state D. Thanks to the coincidence of political and cultural boundaries, national mobilization campaigns remain safely inside the states' borders and are therefore less likely to cause geopolitical havoc.

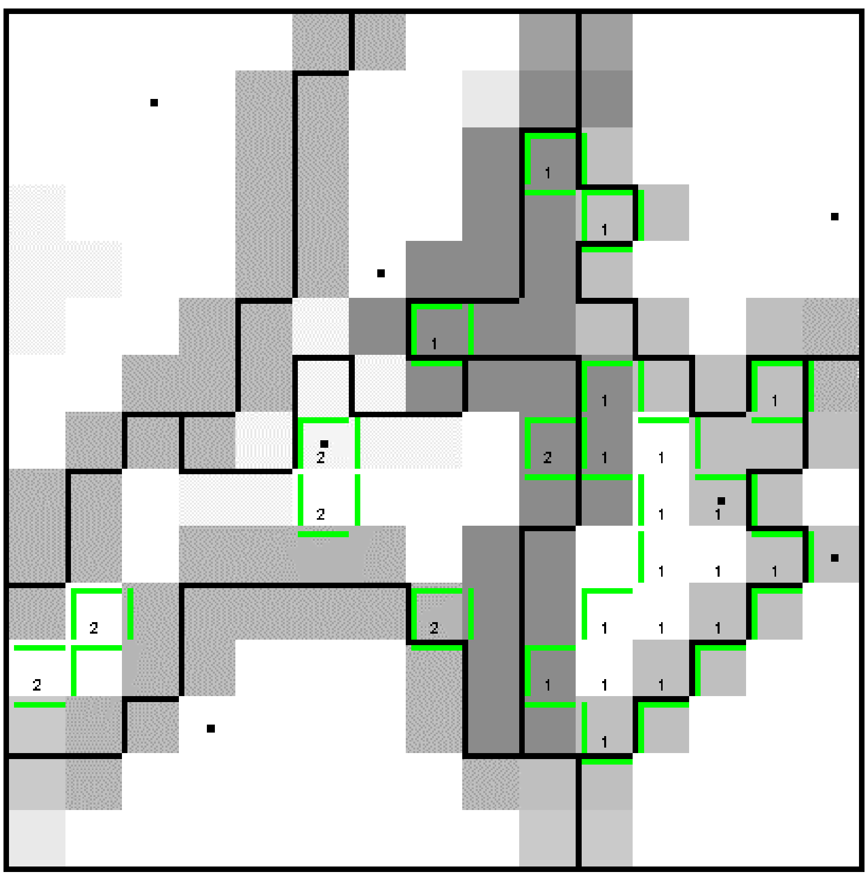

FIGURE 10 State-Framed Nation-Building Progresses Quickly at Time 527

Yet, even with high assimilation speeds, the balance of power can be undermined even in culturally centralized systems. The most obvious threat is unification in cases where the cultural border between two states is insufficient to stop border-transgressing mobilization. A less obvious source of instability derives from the differentiated resource extraction rule. It is assumed that taxation of co-national provinces proceeds without distance discounting.[5] Thus, uneven timing of nationalist mobilization, as, for example, Napoleon's pioneering use of *la levée en masse*, could upset the balance.

All the same, in our sample run, no such disturbances occur. Focusing on time period 1500, Figure 11 illustrates how the entire system has undergone a nationalist revolution without this affecting the geopolitical map. As opposed to the outcome of the culturally fragmented run reported on in Figures 5 through 8, states B, D, and F are saved by state-led assimilation and the lack of external irredentist temptations or irredenta located inside their territories. Moreover, as an emergent result, the fit between states and nations is perfect. All nations own their own states, and the states are ideal nation-states without any minorities.

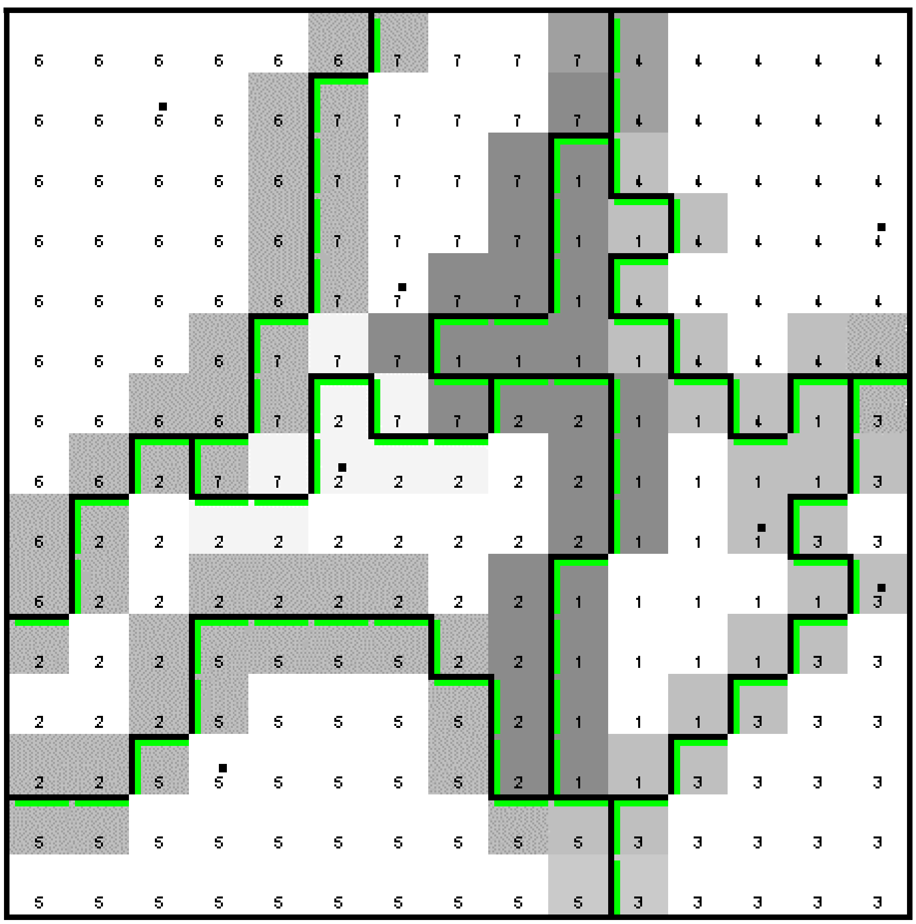

FIGURE 11 State-Framed Nation-Building Produces Perfect Nation-States at Time 1500

[5] The higher level of loyalty through the printed word or modern mass media justifies the absence of a distance gradient (cf. Anderson 1991).

CONCLUSION

Notwithstanding lack of comprehensive sensitivity analysis,[6] this study has generated a number of insights. First of all, it highlights the value added of a truly systemic perspective that views nationalist conflict as a side-effect of fundamental sociological transformation. Whereas most analysts have contented themselves with treating states and nations as given entities while focusing on the behavioral interactions between them, the computational framework allows us to trace the macro effects of nationalism, including structural transformations of both cultural and political boundaries. While far from the only factor influencing such processes, state-led assimilation has played the central role as the main determinant of this geocultural process.

States' cultural penetration, which includes ostensibly "soft" factors such as education and language policy, explains why states and nations coincide in some historical cases, whereas in others they do not. Endogenizing such a process puts the analysis on a more secure, constructivist basis than has been proposed by IR scholars.

Agent-based modeling helps us put the systemic effects of geopolitics and nationalism into perspective. In analogy with Schelling's (1978) famous segregation model that traces residential clustering of two initially interspersed ethnic groups, it is not too fanciful to view the entire international system as a sorting mechanism, where national self-determination acts as an institutionalized motivational rule to reduce the overall "frustration" of the system (i.e., the deviation from the nationalist idea of one-nation-one-state). At the level of mechanisms, however, there are important differences. Whereas migration drives convergence in Schelling's framework and the more general class of Tiebout models (Kollman, Miller, and Page 1997), systemic adjustments in world politics happen through state-transforming events, such as secession, unification, and irredentist conquest, and the nation-altering processes of assimilation and nation-building.

The present geocultural model also clarifies some nontrivial aspects of constructivist nationalism theorizing. Rather than assuming an essentialist one-to-one correspondence between cultural cores and national identities from the outset, the separation of cultural landscapes from nationalist mobilization facilitates analysis of the conditions under which such emergent outcomes become likely. As we have seen, in addition to preexisting ethnic conditions, state-led assimilation plays a key role in the generation of national identities. Yet, there is no reason to expect cultural identities to be easily "malleable," especially once nationalist mobilization has taken off. This is ultimately an empirical issue that has to be determined in particular cases. Nor is there any need to postulate that nationalism is fundamentally an irrational force. While I have assumed nationalist politics to be characterized by higher levels of uncertainty and risk-taking behavior, all the simulations in this study were based on a power-sensitive decision rule. Yet perhaps the most important contribution relates to the systemic context of nationalism. So far, most of the specialized literature on nationalism has adopted a one-country focus while ignoring geopolitical interaction effects. By providing a fundamentally co-evolutionary design that problematizes states and nations as distinct entities, the current model places constructivist theories of national mobilization within an ecological context of state interactions.

[6] Cederman (2000) reports on systematic replication results from this model.

REFERENCES

Anderson, Benedict. 1991. *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. London: Verso.

Axelrod, Robert. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton: Princeton University Press.

Bremer, Stuart A. and Michael Mihalka. 1977. “Machiavelli in Machina: Or Politics Among Hexagons.” In *Problems of World Modeling*, ed. Karl W. Deutsch. Boston: Ballinger.

Cederman, Lars-Erik. 1997. *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*. Princeton: Princeton University Press.

Cederman, Lars-Erik. 2000. “The Co-Evolution of States and Nations: Modeling Nationalist Conflict as a Side-Effect of Macro-Historical Change” Paper presented at the Annual Convention of the American Political Science Association, Washington DC.

Cusack, Thomas R. and Richard Stoll. 1990. *Exploring Realpolitik: Probing International Relations Theory with Computer Simulation*. Boulder: Lynnie Rienner.

Gellner, Ernest. 1997. *Nationalism*. New York: NYU Press.

Kollman, Ken, John H. Miller, and Scott E. Page. 1997. “Political Institutions and Sorting in a Tiebout Model.” *American Economic Review*.

Mann, Michael. 1992. “The Emergence of Modern European Nationalism.” In *Transition to Modernity: Essays on Power, Wealth and Belief*, ed. John A. Hall and Ian Jarvie. Cambridge: Cambridge University Press.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: W. W. Norton.

Schieder, Theodor. 1991. *Nationalismus und Nationalstaat: Studien zum nationalen Problemen im modernen Europa*. Göttingen: Vandenhoeck und Ruprecht.

Weber, Eugen. 1976. *Peasants into Frenchmen: The Modernization of Rural France, 1870-1914*. Stanford: Stanford University Press.

Weber, Max. 1946. *From Max Weber: Essays in Sociology*. New York: Oxford University Press.

Weiner, Myron. 1971. “The Macedonian Syndrome: An Historical Model of International Relations and Political Development.” *World Politics* 23: 665-683.

TRANSMISSION OF CULTURAL TRAITS BY EMULATION: A VECTOR VOTING MODEL

R.G. REYNOLDS, Wayne State University*
R. WHALLON, University of Michigan
S. GOODHALL, Compuware

ABSTRACT

The goal of this research is to assess the impact of Culture on decision-making behavior. Specifically we are concerned whether the emergence of human culture provided humans with an adaptive advantage over non-human primate counterparts in terms of hunter-gathering capabilities. Reynolds has proposed several mathematical models of hunter-gatherer and primate decision making based upon differences in human and primate cultural traits [1], these were labeled the cultural and vector voting algorithm models respectively. In this paper an agent-based implementation of the vector voting model using Swarm is presented. Learning takes place there by emulation. The performance of this model in a variety of landscapes is compared with that of a random walk model. The results suggest that the vector voting model can produce a variety of emergent patterns that can be considered adaptive and that reflect human foraging patterns as well.

INTRODUCTION

The goal of this research is to assess the impact of Culture on decision-making behavior. Specifically we are concerned whether the emergence of human culture provided humans with an adaptive advantage over non-human primate counterparts in terms of hunter-gathering capabilities. Reynolds has proposed several mathematical models of hunter-gatherer and primate decision making based upon differences in human and primate cultural traits [1], these were labeled the vector voting and cultural algorithm models respectively. In the vector voting model each individual's vote was based upon their own knowledge and knowledge was not shared between individuals. The decision made by the group was a consensus based upon the weights and opinions of the members and was based upon patterns of interaction seen among primate groups. In the cultural algorithm model the individuals knowledge was pooled and used by a central decision maker to produce a decision.

The basic context in which these two decision-theoretic models were compared was a two-dimensional cellular space divided into R discrete sub-regions or cells each of unit area. The task facing the model groups was to compute the answer to various spatial predicates or queries about the region based upon the agents current knowledge. The models were analyzed theoretically and it was shown that the ability to form a collective intelligence through the pooling of knowledge had some distinct advantages.

* *Corresponding author address*: Robert G. Reynolds, 432 State St., Wayne State University, Detroit, MI 48202; e-mail: 220057@wayne.edu.

In particular, predicates such as the best direction within a region in which to forage, the one containing the most resources, was limited by the maximum area over which each individual had knowledge. On the other hand, pooling of that knowledge theoretically allowed a group to make these decisions over the entire region.

However, even in a social system where knowledge is not directly pooled, learning can take place. For example, Tomasello and Call [2] state that there are many similarities between how humans and primates understand their social worlds. Each has its own cultural system. In an extensive survey of primate cognition studies they conclude that, “all primates live in basically that same type of social world, in which they individually recognize conspecifics and appreciate both the vertical (dominance) and horizontal (affiliative) relationships that hold between group members. They also have the ability to predict the behavior of conspecifics in many situations based upon combinations of cues and insights, and in some cases to affect the behavior of groupmates via various social and communicative strategies.”

They state that the basic difference between primate and human cultures is that in the latter the “intersubjectivity of human linguistic symbols — and their perspective nature as one offshoot of this intersubjectivity — means that linguistic symbols do not represent the world directly, in the manner of perceptual or sensory-motor representations, but rather they are used by people to induce, to construe, certain perceptual conceptual situations — to attend to them — in one way rather than another.

Thus, learning in a primate social system relies heavily on **emulative learning**. An individual watches another perform an action and observes the state changes that result. Thus, learning in this context is directly associated with sensory motor activities relative to objects in the environment. While humans can acquire knowledge in this way as well, they are able to support the **imitative learning** of concepts. With imitative learning “an individual understands others as intentional agents, like the self, that have a perspective on the world that can be followed into, directed, and shared.”

The idea is that even when a group makes a decision based upon the knowledge of each individual without pooling, the physical results of that decision can be observed by everyone and learning can take place in an emulative fashion. The question of interest in this paper is what additional behaviors emerge from a group that uses the vector voting approach along with an emulative learning process.

Here we propose an agent-based implementation of the vector voting model using the Swarm simulation environment. First, we describe the vector voting model and the emulative learning method that serves as the basis for primate cultural transmission here. Next, we describe the results of running this model in a variety of environments with a variety of social configurations, and then we summarize our conclusions.

IMPLEMENTING CULTURAL TRANSMISSION IN THE VECTOR VOTING MODEL

The vector voting model represented how primate groups made consensus-based decisions. One of those decision was the direction in which to move during the day. It was shown in [3] that it was not possible for a vector voting model to always select the direction with the maximal amount of resources within a region R. However, theoretical limits aside, how does the vector voting model perform when coupled with emulative learning, the type of learning that is frequently observed in primate cultures?

In emulative learning the observer makes a cognitive connection between what action is performed and the state changes it produces. For example, a primate can observe another rolling over a log and exposing a number of insects. That action can be viewed as producing a state change which can be stored in memory. In our model, the result of a directional decision produces a trajectory through the landscape. As a result of that path each individual has an opportunity to get fed. In each cell that the group enters the resources there are divided among group members by various strategies such as priority or fixed order access or equal sharing of resources.

At the end of the day each individual can store a memory, not of the decision, but of the result. An individual does this by associating a visible landmark with the degree to which they were fed (satisfaction scale) the day they saw that landmark. A memory can have more than one landmark attached to it and different individuals in the group can associate different landmarks with the memory. Each individual has a maximum number of memories that it can store, and a memory is forgotten after a certain number of days (memdepth) unless it is used again.

Emulative learning using memories associated with the icons encountered relative to a group's previous decision can now be used to impact future decisions. Here, our region R has a number of landmarks (numLandmarks) which are distributed through the space. The cellular space is of size N by M and a cell can have at most one landmark assigned to it. Each day after the group decides on a direction to move based upon the memories associated with the landmarks currently visible from their location (visibility).

Each individual effectively pools the satisfaction scores for the memories associated with the visible landmarks in each direction. Here the scores are represented by a preference scale from –5 to +5 where the + direction represents satisfaction and – dissatisfaction. The direction with the highest score is the direction of choice for an individual. That choice is weighted by their status in the group. Each member then moves in the direction of their choice and is observed by the others. The group moves in the direction which achieves the highest consensus.

Group size can change based upon the extent to which individuals are fed. Individuals who have not gotten sufficient resources over a given period die and a group is removed when all of its individuals have died. On the other hand, if the group has been able to feed all of its members over a given period it can add a new individual up to a maximum group size. When it reaches that size it can fission into two new groups.

EXPERIMENTAL RESULTS

The vector voting model augmented with emulative learning was applied to a variety of different environments, each with a different resource distribution pattern. Figure 1 gives an example of a patchwork environment for a given set of 10 runs for each of two different configurations. A table with the set of parameters for each of the competing configurations is given in Table 1. The goals of these experiments are threefold:

1. To compare the model to a baseline random walk through the environment.

2. To observe any emergent patterns of foraging behavior that correspond to those exhibited in primate and hunter-gatherer groups.

3. To observe the relative survivability of groups in terms of various cognitive and social parameters, e.g., the number of memories an individual can have.

While a number of experiments have been conducted, we will summarize some of the results here.

When compared to a random walk the vector voting model invariably plateaued out at a number of surviving individuals that was below the carrying capacity of the environment but substantially above that of the random walk model which converged to a zero population. That exact location of the equilibrium point is a function of a number of model parameters.

Figure 2 gives the number of surviving individuals for the run associated with the environment in Figure 1 and compares the vector voting model to the random walk. As exhibited there, not being able to use knowledge in making a best direction choice will ultimately cause the system to collapse.

Surviving group also exhibited certain patterns of foraging behavior observed in primate and human groups. Specifically, each group began to forage in a area associated with positive landmarks in a cyclic fashion. The size of the territory reflected the spacing between. Figure 3 shows the distances between groups for 10 runs of the vector voting model versus that of the random walk. Notice that the distribution of distances between groups is much more focused than that for the random walk.

It is also interesting to note that those surviving groups for random walk are larger than the vector voting group. This is due to the fact that more individuals foraging can compensate for the lack of specific information. This is shown in Figure 4.

A final observation can be made in terms of the memories. Substantially fewer memories were needed to produce the behaviors above for vector voting groups than were used. This can be seen in Figure 5. Overall, random search produced more bad memories than memory based search with the vector voting model. Notice that in the surviving groups both positive and negative memories were associated with the territorial and cyclic foraging behavior of surviving bands.

CONCLUSIONS

In this paper we augmented the vector voting model of consensus-based decision-making in primate groups with emulative learning. Emulative learning was the basic form for cultural transmission used here. Extensive runs of this model suggest that the consensus based approach based upon icon memories was sufficient to produce regional stability below the carrying capacity of the region, territorial behavior, and a cyclical foraging patterns. It should be mentioned however that the territories could shift over time as different groups disappeared from or were added to the environment. As such, it was flexible to changes in group numbers within the region.

The next step in our project is to implement the Cultural Algorithm paradigm which allows the individuals to pool their memories. Theoretically, this sharing of information will allow the group to make regional decisions about certain predicates more reliably than in the vector voting model. In particular, the pooling of information via a belief space will allow the group to decide on predicates that are not tied to particular locations (such as norms) within the environment. These are called position invariant predicates. Being able to do this makes imitative learning, a form of learning largely unique to the human species, possible [4]. Both activities are influenced by the structure of the language used to perform the pooling and within which the learned information can be articulated. This will be the subject of a subsequent paper.

ACKNOWLEDGMENTS

Dr. Reynolds' research is supported by NSF grant IIS-9907257.

REFERENCES

[1] Reynolds, R.G., (1978) "On Modeling the Evolution of Hunter-Gatherer Decision-Making Systems," *Geographical Analysis*, Vol. X:31-46.

[2] Tomasello, M. (1999) "The Human Adaptation for Culture," *Annual Review of Anthropology*, 28:509-529.

[3] Reynolds, R.G., and Zeigler, B.P. (1979) "A Formal Model for the Operation of Consensus-Based Hunting-Gathering Bands," in C. Rebfrew, K. Cooke (Eds.) *Transformations: Mathematical Approaches to Culture Change*, New York, Academic Press, pp. 405-418.

[4] Pedone, R., and Conte, R. (2000) "The Simmel Effect: Imitation and Avoidance in Social Hierarchies," paper to be presented at Workshop on Cultural Transmission, Sienna, Italy.

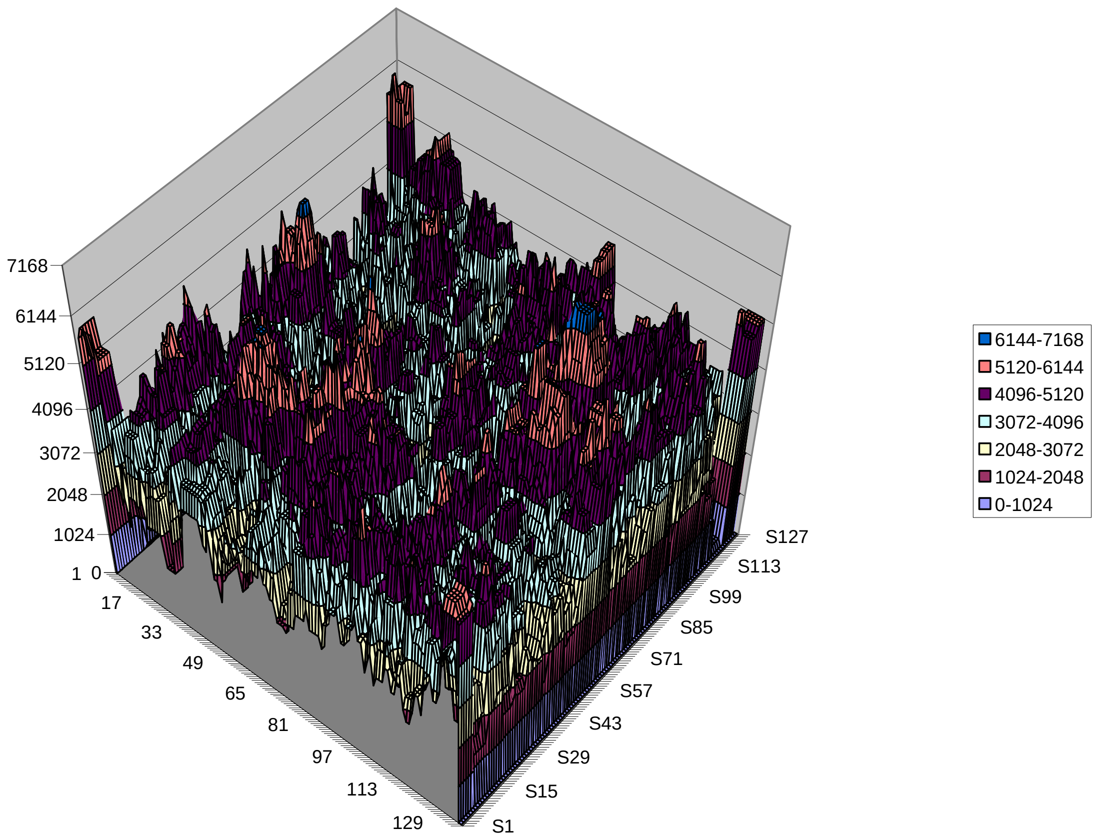

FIGURE 1 A patchwork environment, representing the yields for each cell in the example region

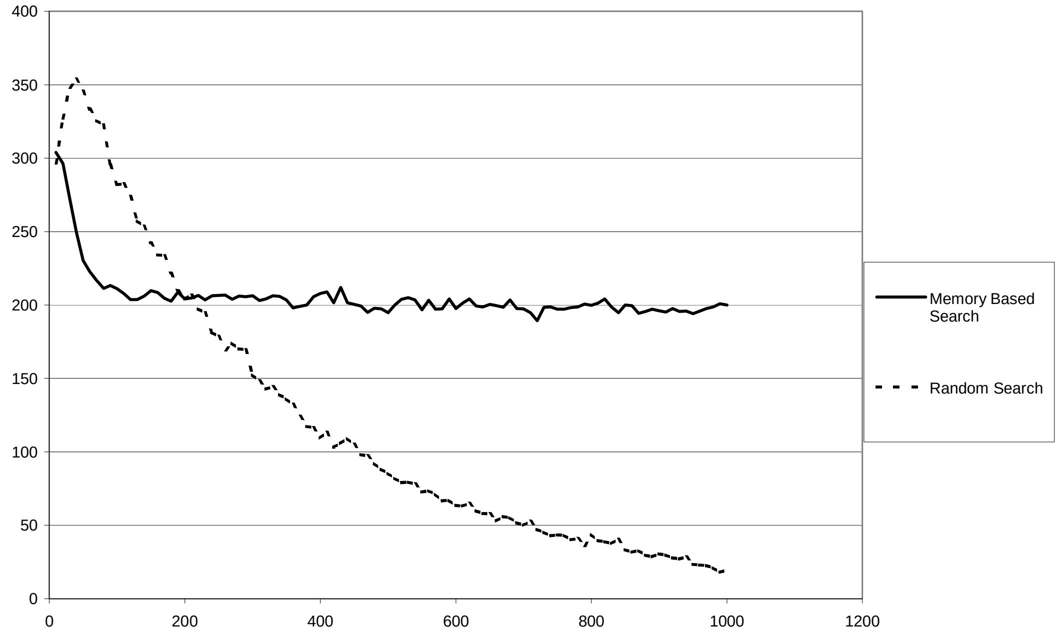

FIGURE 2 The number of surviving individuals in the vector voting model versus the random walk model, summarized over 10 runs each

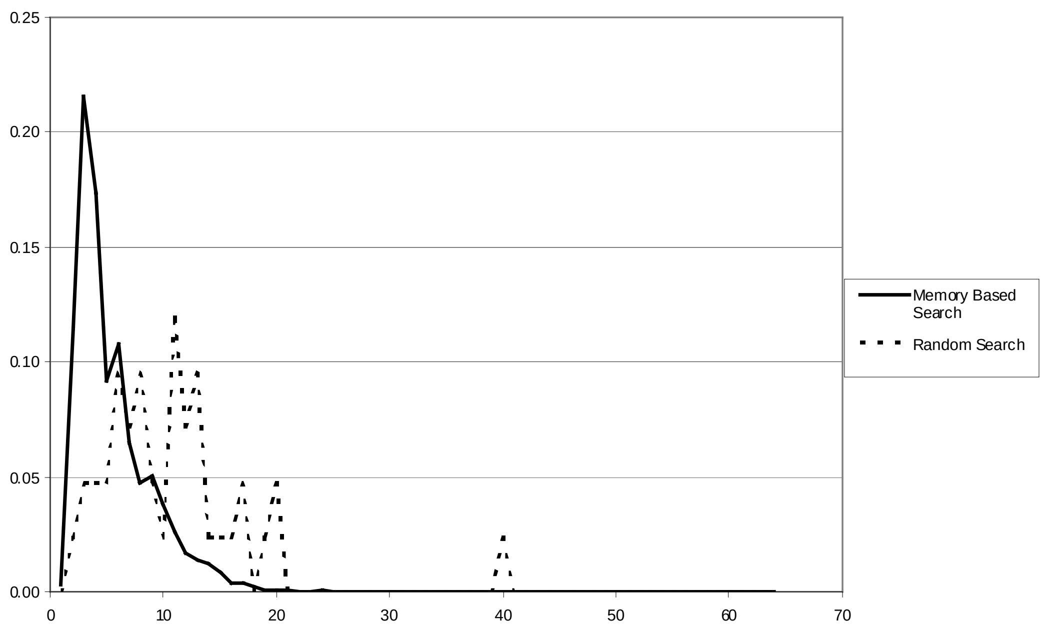

FIGURE 3 Distance between groups at the end of the simulation in the vector voting and random walk models

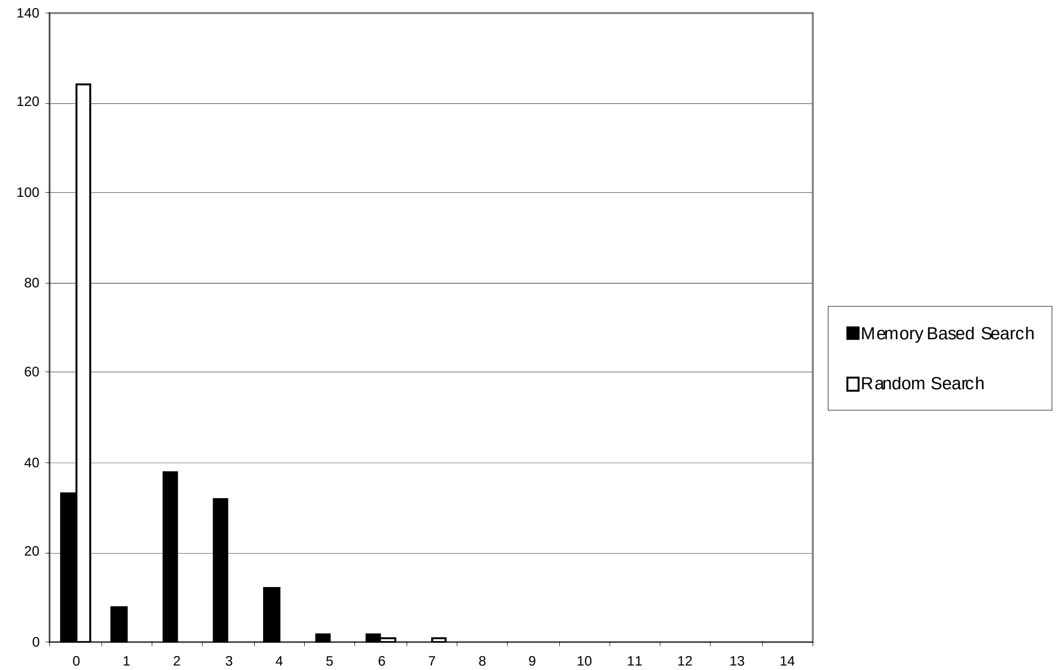

FIGURE 4 Size of the surviving groups at the end of 1000 time steps for the vector voting and random walk models

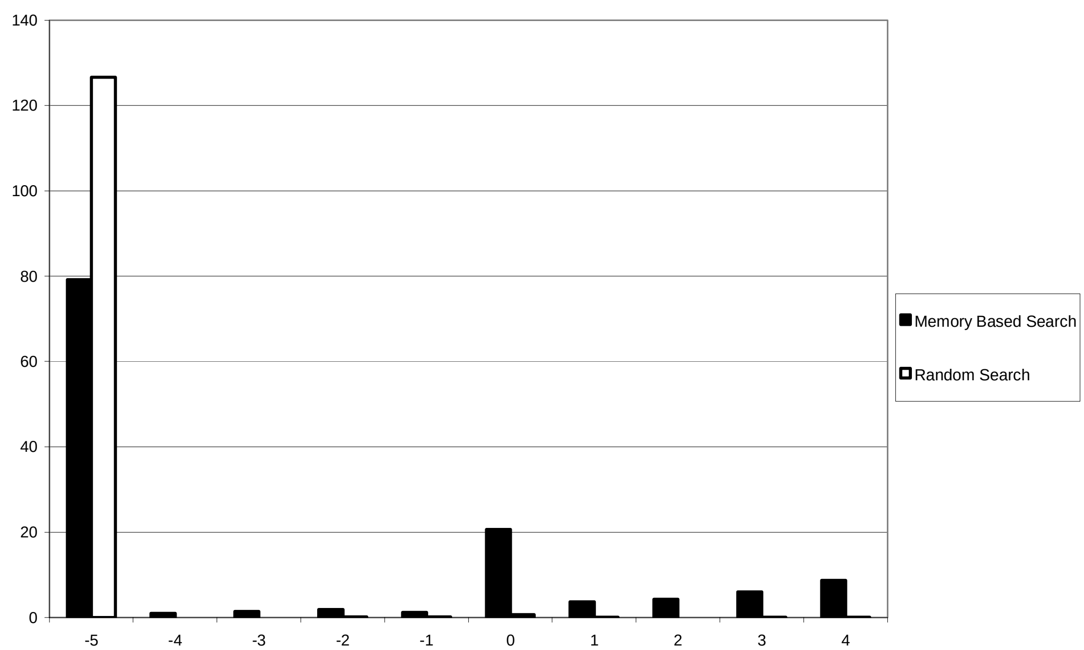

FIGURE 5 The number of memories of different types that are held by groups at the end of 1000 time steps in the simulation for the vector voting and random walk models

TABLE 1 Parameters That Can Be Set for the Vector Voting Simulation System (reflecting two specific configurations selected for testing)

ParmName	Parameter Description	Parm Value1	Parm Value Description1	Parm Value2	Parm Value Description2	Mismatch
bandSize	Maximum size of each band (number of members)	12		12		
bandSplit	Size at which band can split	0		0		
bandStart	Size of each band at the beginning of the run (number of members)	3		3		
baseMDReq	Base minimum daily food requirement for an adult	1500		1500		
bounceMode	Behavior mode when a band reaches the edge of the space or collides with another band	1	make new decision on bounce	1	make new decision on bounce	
CAAcceptance Interval	Cultural Algorithm Acceptance Interval (days): this is the interval at which the CA component accesses the data in the model space	10		10		
CAInfluence Interval	Cultural Algorithm Influence Interval (days): this is the interval at which the CA component provides feedback to the model space	100		100		
CAMode	Selects the CA algorithm to be used	1	basic landmark history	1	basic landmark history	
caption	Trial caption	Memory Based Search		Random Search		YES
consumeMode	Mode of consumption during foraging	0	fixed order	0	fixed order	
daysAhead	Number of days of consumption that an individual can stockpile to recover from shortfalls or to get ahead	+2.0000000000000000e+00		+2.0000000000000000e+00		
daysToDie	Number of days without MDR before an individual dies	7		7		

TABLE 1 (Cont.)

ParmName	Parameter Description	Parm Value1	Parm Value Description 1	Parm Value2	Parm Value Description2	Mismatch
decayDays	Number of days for a dead band to decay and disappear from the display. Until this interval expires, the square occupied by the band is blocked to other bands.	7		7		
diagBand	A band number to trigger diagnostic outputs that are limited to a selected band	0		0		
diagnosticDest	Destination for diagnostic outputs 1 = > console 2 => log 3 => both	0		0		
diagnosticInterval	Diagnostic interval in days: interval at which diagnostic outputs will be written	10		10		
diagnosticMask	Diagnostic mask: a bitmap selecting various classes of diagnostic outputs	0		0		
flipMode	Select a yes no parameter to be flipped for each trial	0	No Parameter Flipping	0	No Parameter Flipping	
foodDistFactor	First computational factor for food distribution. Used differently by each mode. Not always used.	+5.0000000000000000e-01		+5.0000000000000000e-01		
foodDistFactor2	Second computational factor for food distribution. Used differently by each mode. Not always used.	+5.0000000000000000e-01		+5.0000000000000000e-01		
foodDistMode	Food distribution function	6	patchwork around landmarks	6	patchwork around landmarks	

TABLE 1 (Cont.)

ParmName	Parameter Description	Parm Value1	Parm Value Description 1	Parm Value2	Parm Value Description2	Mismatch
foodDistRange	Food distribution range in squares. Usage varies by algorithm. Currently used only for mode 6, patchwork around Landmarks, to control the size of the patches.	12		12		
foodMaxValue	Maximum quantity of food that can exist in a single cell	16384		16384		
forageMax	Number of squares to forage in Mode 1	0		0		
forageMode	Forage mode 0 = en masse while traveling 1 = individual after traveling (unimplemented)	0		0		
lmDistMode	Landmark distribution mode	0	random (seeded corners)	0	random (seeded corners)	
lmPerMemory	Controls how many landmarks are attached to a memory when it is formed	2		2		
logDataInterval	Interval to write to data log (days)	0		0		
memDepth	Memory depth in days, number of days after a memory is formed before it is forgotten	64		64		
memMode	Memory formation mode	0	use n closest landmarks in any direction	0	use n closest landmarks in any direction	
numBands	Number of bands to generate	128		128		
numLandmarks	Number of landmarks to generate	128		128		
quitAfter	Terminate the mode after this many steps. Results in program failure. Superseded by batch mode operation.	0		0		

TABLE 1 (Cont.)

ParmName	Parameter Description	Parm Value1	Parm Value Description 1	Parm Value2	Parm Value Description2	Mismatch
randomDays	Number of days of random movement before using memory (during this interval, bands do not starve)	0		0		
regenRate	Rate of food regeneration per day (≤ 1.0)	+5.0000000000000003e-02		+5.0000000000000003e-02		
reincarnateDays	Number of days before a dead band is reincarnated	9999		9999		
reproduceDays	Number of days where the band is fully fed before it can add a member	10		10		
searchMode	Search mode	1	memory based vector voting	0	random	YES
seedProb	Probability of food in a cell (random distribution). Also controls the minimum amount of food. Not currently used in other modes.	+1.0000000000000000e+00		+1.0000000000000000e+00		
starveMode	Mode for determining starvation	1	deficit > MDR * days to die	1	deficit > MDR * days to die	
stepsPerDay	Maximum number of steps to move in one day	4		4		
stopEvery	Interval (days) at which the model will pause	100		100		
trialName	File name prefix for the trial (yyyymmdd_hhmmss). This is not specified externally but is generated and exported at run time.	20000729_160227		20000729_151713		YES
visibility	Default landmark visibility in squares: for how many squares is a landmark visible	32		32		

TABLE 1 (Cont.)

ParmName	Parameter Description	Parm Value1	Parm Value Description 1	Parm Value2	Parm Value Description2	Mismatch
vision	Default vision in squares: how many squares can an individual see a landmark	32		32		
worldXSize	X dimension of the model space	128		128		
worldYSize	Y dimension of the model space	128		128		
zoomFactor	Swarm zoom factor (0 = automatic) 1-	0		0		

DISCUSSION: POLITICS

J. PADGETT, University of Chicago, Moderator

[Presentation by Lustick]

Michael North: I think what you presented here is very interesting. What I would like to see in future models is environments that change from very rational to irrational, because I perceive that as happening in society right now in many cases — where the environment itself is changing partway through the evolution of history.

Ian Lustick: First we wanted to find out a little bit about the world without that [variability]. One of the areas that we are working in is globalization, and in that context we will change the parameters of the global state and then watch what happens in the local state under varying conditions, so that would be a [way of looking at what you suggest].

[Presentation by Bendor]

Claudio Cioffi-Revilla: I found your presentation very stimulating because you've taken an important step in the direction of an alternative to the rational choice paradigm. It seems to me that there is a fundamental flaw in the rational choice paradigm. In this instance, it assumes that people vote because they believe that they're going to determine the election. I doubt there is any voter in the world who imagines that he will individually be able to determine the election. If that's the test that the rational choice theory offers, it's misleading about the way people behave and what they believe they can or cannot do.

It's probably true that most people vote because they think they *should* vote, because they feel an obligation to vote. The notion of obligation is important, because, like the notion of aspiration, which you're using, it looks in an entirely different direction for the motivation for voting. Regarding obligation, there's some very interesting literature on deontic logic. Did you look in that direction and decide not to pursue it? There are aspects of deontic logic that are very formalized.

Jonathan Bendor: That's a very good question. We have thought about that. Right now, the only way the model can represent obligation is in a fairly crude way by allowing the cost of voting to go negative to reflect a sense of obligation, a civic duty to vote. And let me just briefly report to you what happens. Suppose people have a negative cost of voting because they feel guilty if they don't vote. What do you think will happen? From a game-theoretic point of view, if people have negative costs of voting, then voting is a dominant strategy — you can't lose by voting. So you should see complete turnout. But we don't get complete turnout. Why? Because the aspiration levels rise when you have negative costs of voting. You can still care about the collective, electoral outcome, you can still be disappointed if you turn out and lose; therefore the model cannot stabilize at a propensity of one. I do think the question of obligation is a very important direction of research.

Maurits van der Veen: I have two related questions, one theoretical, one empirical. The first is, most people don't vote in 1,000 elections over their lifetime. So I would be interested to see how quickly your model gets to the equilibrium outcome of about 50% turnout. And second, the model gives some suggestion that as people get older and presumably have voted in more elections, their turnout should essentially approach 50%, and that should be empirically testable.

Bendor: It is testable, but note that turnout may approach 50% by declining from above that value, rather than rising to it. As a matter of fact, a colleague of ours who studies Eastern Europe suggested to us that voting participation rates after the fall of communism began to decline after a while, and we have a result in the paper that explains that phenomenon. When you start with very high aspirations that the government can't deliver, participation rates will tend to fall.

About the rate of time to convergence: we make no real-time claims in the paper, because, as in the linguistics paper that was given yesterday [Satterfield], we have no reason to believe that the speed-of-adjustment parameters reflect, or are calibrated to, the real world. We simply don't know.

van der Veen: But every period is one vote experience.

Bendor: Yes, but how long it takes to converge to 50% depends critically on the speed-of-learning parameter, on the alphas and the betas.

van der Veen: Right. Well, can you tweak the speed of learning so that the [times] become realistic?

Bendor: Yes. First of all, the starting case of "all shirk" is not realistic. That's just a thought experiment, as it were, to show that you cannot make the system stay at very low levels of participation — even if you start it there, you would see a breakout of participation.

One commentator on an earlier version of the paper, someone who studies elections, advised us that rate of convergence is not an issue in stable democracies. Turnout rates change very slowly, and the new generation is likely to "inherit" propensity rates that are quite close to their parents' — so we don't see a breakout of participation. Empirically, that would be a very unusual situation. We were simply demonstrating a theoretical property.

Michael Heaney: I'm Michael Heaney, University of Chicago. I wonder if you've thought about putting a couple of other features into the model: one to reflect the heterogeneity of people's psychic involvements in the election, and another that would allow for variations in the social pressure to vote. I see these two factors as interacting. Some people are always going to vote, because they are psychically involved, they're interested; other people don't care; and there is a large degree of variation in between. And there's the way people are connected to social networks. So people like myself who are always talking about politics are putting social pressure onto others. We'd expect that if a person had low psychic involvement combined with little social pressure, they'd be unlikely to vote. I think that including these variables in the model would be a mechanism to explain why some people vote and some don't. Your model seems to assume that everyone is psychically involved in the election.

Bendor: Right now we can distinguish between people in the two factions. We can have heterogeneity and exogenous parameters within a faction. It's time-consuming, but it's relatively straightforward to reprogram the simulation. I think the network idea is definitely a worthwhile thing to do, and will yield estimable predictions.

[Presentation by Cederman]

Miles Parker: Miles Parker, Brookings. I think this is a very nice, rich model. I'm wondering if you've looked at the flip side of the nationalism issue, say, national bifurcation or new identities arising within nations. I'm thinking of the pre-Civil War South, for instance.

Lars-Erik Cederman: Actually, the current system doesn't allow for splits of nations as opposed to states. Once a national community has been formed, the nation can change only if it's completely broken up by a state capital waging a war against the periphery. But what you're saying is entirely right. Our model is a simplification. It's entirely possible to include rules that would allow both for mergers of split nations — smaller nations, say, that may be close to each other in culture — as well as for breakups of existing nations.

Parker: I have another question. It may just be happenstance, but I thought I noticed an episodic or punctuated flavor to the periods when a lot of changes happened. Have you noticed that?

Cederman: Very much so. This model is very close to the [Stephen Jay] Gould perspective of punctuated equilibria. This is what makes it so hard to study nationalism with micro-level theory, because within small timeframes nothing may happen. Much like what happens before an earthquake, there is a tendency for tensions to build up over time. And because of the threshold-type logic that I have built into this system for mobilization and conquest, you get periods of extensive change — a bit like in the Schelling model again. One family moving to another neighborhood may cause a chain of changes. So, absolutely, this episodic tendency is built into the system.

Lustick: I have a question about research strategy. Every time you get an intuition from another kind of theory, say mobilized populations, or secessionist theories, then the model becomes more complex. And when something doesn't happen the way you expected, it's much more difficult to determine why. After my presentation, I got a great suggestion to vary the environment. You can now vary the environment during a run. The number of opportunities to insert things that are not only plausible but that we also know are relevant is outrunning our ability to explore the space that we have created. So how do you balance it?

Cederman: There's no simple answer, but I would say first of all that you have to make a choice from the beginning, because it is a tradeoff between tractability and realism. I've been driven very much by my own intuitive or theoretical knowledge of nationalism that I've drawn mostly from the qualitative literature. My model is an attempt to get things down in a more formalized, "cleaner" fashion than in the qualitative literature.

We're not done — this is just a starting point. But let me emphasize that when you decide how complex a model should be, you always have to be concerned about the underlying theoretical assumptions. There's always a minimum level of simplicity below which you cannot go without violating the most fundamental assumptions of the qualitative literature. But our model is far from the very simple rational choice models that have become influential in the literature. I'm not saying that those models are not useful; I'm just saying that our model offers another perspective.

I believe there is a need to do more robustness checking of this model. I should emphasize that the goal of a model should not be to sweep the entire parameter space, because that's impossible. This is a very complex model. A full parameter sweep may also be a Holy Grail that we're not interested in. If you can get reasonably close to stylized scenarios from history and you can say something meaningful about possibilities, if you can get a feel for things you didn't expect, then you can go back to the historical record and look with this new intuition in mind. For instance, I didn't expect national unification to be as destructive as it was. So this is a qualitative and more heuristic use of simulations. I think that people have misunderstood the value of simulation in this regard — they have confused simulation with the hubris that existed in the older literature.

[Presentation by Reynolds]

David Sallach: Was the environment structured in regions? In other words, could the agent say it is dry in the west or cold in the north? I'd also like to ask about the extent of generalization or classification in the learning process.

Robert Reynolds: We can put in any function that we want to. For example, in one, the middle is awful, but if you go along the edges, it's good.

Sallach: So they were learning regional patterns?

Reynolds: They were learning regional patterns. For these examples, the regional patterns are pretty obvious to us. The patchwork environment that I gave you actually has regional patterns in it, but I can't see them. In fact, the resources are organized probabilistically around the landmarks, and we can adjust how much a focus a particular landmark is. So we can make landmarks in the north more focused for resources than ones in the south. They are learning those types of patterns, but it's very hard individually to see them.

Sallach: So it's iconic patterns as well.

Reynolds: Exactly.

Environmental Processes

THE PROCESS-INTERFACE-TOPOLOGY MODEL: OVERLOOKED ISSUES IN MODELING SOCIAL SYSTEMS

H.V.D. PARUNAK, Environmental Research Institute of Michigan*

ABSTRACT**

The Process-Interface-Topology (PIT) approach to modeling social systems considers the Processes executed by individual agents and by their environment, the Interfaces between participants, and the overall Topology of their interconnection. The PIT perspective focuses attention on important details that the conventional bipartite discussion of individual agents and agent organization sometimes overlooks. This paper identifies some of these details and argues that effective modeling of social systems must include integrated, disciplined analysis of all three aspects.

INTRODUCTION

Discussions of the design and engineering of multi-agent systems commonly focus on two aspects: the design of the individual agents, and design of their interactions. For instance, the Gaia methodology (Wooldridge, Jennings et al. 2000) analyzes a system in terms of roles and interactions, and produces a design around the services that individual agents perform and the acquaintances of each agent. The binary distinction characterizes our own ontology of agent applications (Parunak 1996; Parunak 1998), and is reflected in Jennings' distinction between the knowledge level and social level in agent systems (Jennings 1994). Holarchic (University of Hannover 2000) and compositional (Brazier, Jonkers et al. 1998) approaches emphasize that more layers can exist than just "agent" and "system," and discuss interactions as the mechanism by which entities at one level form an entity at a higher level.

This ontology is fine as far as it goes, but overlooks some key issues.

(1) A real social system is made up not only of agents, but also of an environment in which those agents exist and through which they interact. This environment must be taken into account in analyzing, modeling, and engineering social systems. The case for this position has been articulated for some years in the embodied cognitive science community, most recently as the principle that one must design an agent's ecological niche along with the agent itself (Pfeifer and Scheier 1999).

(2) Design of agent interactions is typically restricted to symbolic communication protocols. Other forms of interaction, often relying on nonsymbolic physical actions, can also be important.

* *Corresponding author address:* H. Van Dyke Parunak; ERIM, P.O. Box 134001, Ann Arbor, MI 48113-4001; e-mail: vparunak@erim.org.

(3) The behavior of the community emerges from the interactions of individual agents with one another and with the environment in a way that is not immediately obvious from the definitions of the individual components. Recent results in complexity science (Watts 1999) suggest that the topology of agent interaction is critical to the nature of this emergent behavior.

To accommodate these concerns, we have begun to analyze systems that we wish to model using a tripartite model, focusing on the *Processes* that take place in the system (whether in individual agents or in the environment), the *Interactions* among agents and between agents and the environment, and the *Topology* of those interactions. The successive sections of this paper discuss each of these areas, with special emphasis on details that one might be likely to overlook from the traditional bipartite perspective.

PROCESSES

A process P_i may be defined formally as

$$P_i = \langle V_i, R_i \rangle \quad \mathbf{(1)}$$

where

V_i is a set of variables whose assignments change over time, and

$R_i \subset V_i^+ \times t \rightarrow V_i^+$ is a set of rules governing how those changes take place over time.

$\boldsymbol{P} = \bigcup_i \{P_i\}$ is the set of all processes.

Processes come in many flavors. The variables may be numerical or symbolic. Time may be continuous or discontinuous. The rules may be expressed computationally (in terms of rewrite rules) or as a system of differential equations. In classical AI, processes usually concern symbolic variables manipulated computationally in discontinuous time, while physics focuses on numerical variables whose values over continuous time are defined by differential equations, but some impressive examples apply physics-like processes to cognitive problems (Port and vanGelder 1995; Schöner, Dose et al. 1995).

We have elsewhere argued for the definition of an agent as a bounded process (Parunak 1997). Most current discussion on the architecture of individual agents focuses on defining that process in terms of computational mechanisms such as rule-based reasoning, multi-threaded object architectures, or even the integration of differential equations. Clearly, anyone who constructs an agent-based model of a social system must define the behavior of the individual agents. By drawing the analyst's attention to "processes" rather than "behaviors," we wish to emphasize that the system's behavior depends not only on the behaviors of individual agents ("bounded processes"), but also on processes that are not "bounded," but form part of the environment in which agents interact.

Social norms and governmental regulations are an important example of such processes in a social system, as are acts of God such as weather, earthquakes, and the natural deterioration of transportation facilities over time. In modeling a social system, sometimes it makes sense to instantiate an agent representing the "environment" to handle such influences. But we must first recognize that the influences exist, and that centralizing them in an "environment agent" is an implementation compromise. Agents that correspond to real bounded entities in the domain are qualitatively distinct from agents that serve to encapsulate diffuse, distributed influences, and we must be aware of the possibilities of unexpected dynamics in their interaction. Two examples of the atypical nature of environmental agents come to mind:

(1) "Environmental" agents may not respond to actions taken by "real" agents as predictably as other "real" agents (Ferber and Müller 1996). Error-correcting protocols and automated retransmission mechanisms can guarantee that a message sent by one agent reaches another, thus changing the state of the recipient (at least with respect to its mailbox). But a robotic agent might push on a boulder all day long and never make any difference to the environment.

(2) The environment is usually spatially distributed, and influences are often bounded in space and time. Thus an agent-based model of the environment may need to consist of a network of agents rather than a single environment agent. Experimental evidence (Parunak, Brueckner et al. 2000) shows that lumping spatially distributed processes into a single agent can distort the result of the overall model.

Control theoreticians are accustomed to making this distinction when they analyze the separate and combined effects of the dynamics of the controller and the dynamics of the plant (the entity being controlled). We recently encountered this effect in a project providing routing control for military aircraft, reported in detail elsewhere (Parunak, Brueckner et al. 2000). Briefly, we constructed a pheromone-based mechanism for guiding friendly aircraft through a space populated with targets that they should attack and threats that they should avoid, and then analyzed their performance as we varied the composition of friendly and adversarial forces. The outcome showed strong nonlinearities. As we studied these irregularities, we hypothesized that they were largely due not to our control mechanisms, but to the rules that governed the outcome of combats between the two forces. We were able to verify this hypothesis by constructing an abstract model that embodied only the dynamics of the environment, and comparing the performance landscape it produced with that resulting from the behavior of agents embedded in that environment.

Thus analysis of a social system to be modeled using agents must include a comprehensive review of the processes that are observed in the domain. Conformably with conventional approaches to agent-based modeling, many of these processes will be localized within bounded entities. In fact, it is reasonable to argue that only domains in which most processes are so localized are really good candidates for agent-based modeling. However, it is rare to find a domain in which all processes are bounded, and development of a faithful model requires that the analyst pay special attention to unbounded environmental processes as well.

INTERFACES

One important distinction between a monolithic software system (a "single-agent system") and a multi-agent system is the fact that agents (and more generally, processes) interact. The interface between two processes specifies how changes in one process's state variables affect the evolution of the other processes' variables. Formally, an interface I_j among a set of processes is itself a process that includes the union of the other processes, as well as additional rules R_I specifying the coupling across the original processes:

$$I_j = \left\langle \bigcup_i V_i, R_I \cup \bigcup_i R_i \right\rangle \quad \textbf{(2)}$$

where

i is an index ranging over the processes in the interfaced set,

$R_I \subset \left(\bigcup_i V_i \right)^+ \times t \rightarrow \left(\bigcup_i V_i \right)^+$ (the "interface rules") is a set of rules spanning the variables of different processes and governing how those changes take place over time, and

$\boldsymbol{I} = \bigcup_j \{I_j\}$ is the set of all interfaces I_j.

In conventional MAS's, interfaces are restricted to inter-agent protocols. Such protocols are an essential component of the interfaces in modeling a real social system, but others must be considered as well. For example, in a market economy, protocols may support the negotiation over the nature and prices of services and goods to be exchanged. With the advent of electronic cash even the payment for those services and goods can be embedded in a protocols, and some services (notably information services) can be delivered through a protocol as well. But if a purchaser is buying a lawn-mowing service or a new car, at some point a physical transaction must take place that falls outside the neat definition of an information protocol. If this transaction is not represented, the overall model will be defective.

We experienced this problem in our early work on YAMS (Yet Another Manufacturing System), a control system for a flexible manufacturing system that distributed tasks across multiple manufacturing workstations (Parunak 1987). We carefully designed the protocol through which the workstations negotiated for allocation of tasks, taking care to construct a formal proof showing that the protocol would not deadlock. Then we implemented the agents, installed the system on physical machinery, and turned it on. Before very long, the system deadlocked! Our proof was not in error. However, the formal analysis covered only the movement of electronic messages among the agents representing the workstations. In the real world, physical parts also moved among the workstations, and their movement (and the information they conveyed by moving) was not included in our analysis.

Agent-based systems for information services can sometimes equate interfaces with protocols. Agent-based systems that model social systems cannot. Social systems engage people in interactions with one another and with the real world that include a variety of physical as well as informational influences, and our models must incorporate these aspects. For implementation reasons, we may need to model (say) the shipment of a car as a special message in an electronic protocol, just as we may need to model an environmental process as an agent. As with processes, so with interfaces, we need to recognize that such an implementation is a compromise, and may obscure important distinctions in the system we are trying to model.

TOPOLOGY

Interfaces induce a graph-like structure over the set of processes:

$$T = \langle \boldsymbol{P}, \boldsymbol{E} \rangle \qquad \textbf{(3)}$$

where

$\boldsymbol{E} = \{E_1, \ldots, E_m\}$

$E_j \subset \Pi \boldsymbol{P}$ is a multi-edge connecting the processes in I_j.

It is commonplace in designing agent-based systems to construct a diagram more or less isomorphic to UML's acquaintance diagram, indicating which agents interact with which other agents. Just as agents do some but not all of the work of processes and protocols do some but not all of the work of interfaces, so acquaintance diagrams satisfy only part of the need for attention to topology. At least two defects encourage us to generalize the concept.

The first defect reflects once more the rude intrusion of the environment into situated agent-based systems, such as those that are intended to model social systems. To a first approximation, the topology of a system in which agents both influence and sense a shared environment is a star with the environment at the hub and agents at the ends of spokes, augmented by direct agent-to-agent communications. A serious weakness of this first approximation is its inattention to spatial and temporal locality. Typically, an agent can influence and sense only a limited area in its environment, and influences in the environment propagate at a finite speed, frequently decreasing in intensity as they propagate and as they age. Insect pheromones are a canonical example of environmentally-based interaction. One insect can sense the pheromones deposited by another only if it comes close enough to the original deposit, within a time window determined by the evaporation rate of the original deposit. Thus the environment links agents whose trajectories in space-time come close enough together. If agents can move, the topology of the system is a function of time. Synthetic pheromone infrastructures (Brueckner 2000; Parunak 2000) suggest one way to implement such a topology. The environment is modeled as a network of places, among which agents move. A place provides a number of services to agents that occupy it: agents can deposit pheromones on a place, and sense pheromone strengths at the place and its immediate neighbors. In addition, the place evaporates pheromones over time, and propagates them to its neighbors. Places themselves are naturally represented as (non-mobile) agents, subject to the caveats in Section 2. A static graph is clearly inappropriate to describe the topology of such a system. At a minimum, one might capture the topology by presenting one acquaintance diagram showing the interconnectivity of the places, a second representing the connectivity of an arbitrary place with one or more arbitrary agents that are resident on it at a given time, and a representation of the constraints on agent movement and the propagation and evaporation of influences that agents can exert on places.

A second defect of the acquaintance diagram as a representation of the topology of an agent-based system is its inattention to the dynamical implications of interconnectivity. Recent work (Watts 1999) shows that processes interacting through graphical structures can behave very differently depending on graph-theoretic features of those structures, features such as characteristic path length and clustering coefficient. Analysts and designers of multi-agent systems and models should recognize this work by paying attention to the potential for topology-dependent dynamics in their systems. Agent populations should be tested for their behavior on a range of topologies, including random, regular, and intermediate structures, to determine how the overall behavior may vary with connectivity.

SUMMARY

Classical analysis of agent-based systems is bipartite, focusing on individual agents and their community relations. The Process-Interface-Topology (PIT) model looks at issues that are orthogonal to this classical division (Table 1), and encourages us to identify details that might otherwise be missed, such as the role of an active environment, the possibility of indirect agent interaction through a shared environment, the implication of agent mobility on locality of interaction, and the influence of system topology on overall behavior.

TABLE 1 Comparing the Models

		Tripartite PIT Model		
		Processes	**Interfaces**	**Topology**
Classical Bipartite Model	**Individual Agents**	Agents are bounded processes.	Agent I/O may include sensors and actuators as well as digital communications.	Agent mobility can cause topology to change over time.
	Community	The environment, which spans multiple agents, may also support processes that impact the behavior of the system as a whole.	Agents interact both directly (classical protocols) and indirectly (through interactions with a shared environment).	The dynamics that emerge from processes on a graph can vary nontrivially based on details of the topology.

When the time comes to implement an agent-based model or system, the software engineer must still work in terms of individual agents and community mechanisms. In model construction, the physical environment itself must be represented computationally. We have identified several work-arounds that can accommodate these constraints, as long as care is taken to avoid certain pitfalls. Specifically,

- Environmental processes can be instantiated as agents. These agents may need to respond to actions from other agents differently than ordinary agents would. If the environment has spatial extent, an adequate representation will typically consist of a network of agents rather than a single environment agent.

- Physical interactions may be modeled as digital protocols. These protocols must model the inconclusive nature of physical actions (in contrast with the deterministic nature of digital communications).

- Acquaintance graphs are a starting point for capturing topology, but should be constructed with attention to the implications of environmentally-based interactions and agent mobility for variation over time, and should be analyzed for graph-theoretic characteristics that may impact the emergent dynamics of interacting agent and environmental processes.

REFERENCES

Brazier, F. M. T., C. M. Jonkers, et al., 1998. *Principles of Compositional Multi-Agent System Development*. Proceedings of the 15th IFIP World Computer Congress, WCC'98, Conference on Information Technology and Knowledge Systems, IT&KNOWS'98, Chapman and Hall.

Brueckner, S., 2000. Return from the Ant: Synthetic Ecosystems for Manufacturing Control. *Computer Science*. Berlin, Germany, Humboldt University Berlin.

Ferber, J. and J.-P. Müller, 1996. *Influences and Reactions: a Model of Situated Multiagent Systems*. Second International Conference on Multi-Agent Systems (ICMAS-96).

Jennings, N. R., 1994. *Cooperation in Industrial Multi-Agent Systems*.

Parunak, H. V., S. Brueckner, et al., 2000. *Distinguishing Environmental and Agent Dynamics: A Case Study in Abstraction and Alternative Modeling Technologies*. Engineering Societies in the Agents' World (ESAW'00), Berlin, Germany.

Parunak, H. V. D., 1987. Manufacturing Experience with the Contract Net. *Distributed Artificial Intelligence*. M. N. Huhns. London, Pitman: 285-310.

Parunak, H. V. D., 1996. Applications of Distributed Artificial Intelligence in Industry. *Foundations of Distributed Artificial Intelligence*. G. M. P. O'Hare and N. R. Jennings. New York, John Wiley.

Parunak, H. V. D., 1997. "'Go to the Ant': Engineering Principles from Natural Agent Systems." *Annals of Operations Research* **75**: 69-101.

Parunak, H. V. D., 1998. What can Agents do in Industry, and Why? An Overview of Industrially-Oriented R&D at CEC. *Cooperative Information Agents II*. M. Klusch and G. Weiss. Berlin, Springer. **1435**: 1-18.

Parunak, H. V. D., 2000. Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization. H. V. D. Parunak, J. Sauter and R. S. Matthews, ERIM CEC.

Pfeifer, R. and C. Scheier, 1999. *Understanding Intelligence*. Cambridge, MA, MIT Press.

Port, R. F. and T. vanGelder, 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA, MIT Press.

Schöner, G., M. Dose, et al., 1995. "Dynamics of behavior: Theory and applications for autonomous robot architectures." *Robotics and Autonomous Systems* **16**: 213-245.

University of Hannover, 2000. Holonic Manufacturing Systems. *HMS, IMS*, University of Hannover.

Watts, D., 1999. *Small Worlds : The Dynamics of Networks Between Order and Randomness*. Princeton, NJ, Princeton University Press.

Wooldridge, M., N. R. Jennings, et al., 2000. "The Gaia Methodology for Agent-Oriented Analysis and Design." *International Journal of Autonomous Agents and Multi-Agent Systems* **3** (Forthcoming).

AN OBJECT-ORIENTED, AGENT-BASED, SPATIALLY EXPLICIT ENVIRONMENTAL MODEL: A DISCUSSION OF THE APPROACH TO IMPLEMENTING THE SYSTEM

C.C. REWERTS, Engineering Research and Development Center*
P.J. SYDELKO, Argonne National Laboratory
J.H. CHRISTIANSEN, Argonne National Laboratory
J.E. DOLPH, Argonne National Laboratory
A.-M. SHAPIRO, Engineering Research and Development Center
T.N. TAXON, JR., Argonne National Laboratory

ABSTRACT

In this paper we discuss the implementation of an object-oriented, agent-based, spatially explicit simulation of the red-cockaded woodpecker (*Picoides borealis*) (RCW). The goal of this particular modeling effort is to provide a tool for Army land managers who must balance the mission for military training with the protection of the RCW. The goal of the authors is to explore, develop, and use tools for ecological modeling in a wider frame of reference, with this particular effort as a component of the overall ambition. The purpose of this paper is to describe our design goals, why we chose the tools we used, and how we employed them to accomplish the task.

INTRODUCTION

An environmental modeler makes many choices when addressing a problem. The real world, the living inhabitants, and their processes are abstracted to fit the model and modeling environment. To complicate matters, all aspects of modeling technology are in a constant state of change. As technology changes, the modeler is faced with a number of moving targets when choosing the tools, data, and methods from which to build a model of the environment. The environmental processes in the real world are a symphony of interactions. As modelers abstract the real world, they will, for practical reasons, need to focus on a limited number of players and interactions. Meanwhile, other modelers may have focused on other aspects of the real world in their models. Thus, for example, could a wildlife model, a habitat model, and a model of human encroachment be connected to work together?

This paper explains the criteria we used when designing an environmental model that eventually will be multifaceted and applied to assist in land management decisions. We wanted the model to represent entities and processes as they are understood in the real world. We began with an agent-based population model of a rare woodpecker, with plans to include aspects of the environment such as habitat quality and human activities. As we expanded the model, we wanted the option of including models developed by other people, even if the other models were written in a different programming language or ran on a different computer on the Internet. We knew that

* *Corresponding author address*: Chris C. Rewerts, U.S. Army Construction Engineering Research Laboratory, Engineering Research and Development Center, P.O. Box 9005, Champaign, IL, 61826-9005, e-mail: Chris.C.Rewerts@erdc.usace.army.mil.

if we integrated other models or more of our own submodels, we wanted to do it in such a way that the various models and processes could be updated through time without requiring a complete overhaul of our system. We wanted a spatially explicit model, to take advantage of geographic information system (GIS) data and functionality. There are several potential users of our model, each of whom will have various data formats or goals for the model; consequently we wanted ubiquitous flexibility. Finally, we wanted a method of modeling such that the effort invested in answering one specific problem could be applied as much as was practical to other, unforeseen problems.

BACKGROUND

The red-cockaded woodpecker (*Picoides borealis*) (RCW) is a resident of the old-growth longleaf pine woodlands of the southeastern United States. The RCW has been listed as a federally endangered species since 1970. The decline in RCW populations has been attributed to loss of habitat and alteration of habitat, primarily due to suppression of fire (USFWS 1985, Walters 1990, Jackson 1994).

United States law places specific requirements on federal land-owning agencies, such as the Department of the Army, for management of federally listed endangered species. The Army cannot harass, harm, kill, or disrupt the natural behaviors of listed species, including the RCW. The Army must take proactive steps to enhance the numbers of RCWs on its lands. Any significant change to activities or construction projects on Army lands must be coordinated with the regulatory agency, the U.S. Fish and Wildlife Service. Because the RCW is known to exist on many Army installations throughout the southeastern United States, its management has been coordinated among many locations to ensure uniform compliance with the law. Formal procedures and decisions are necessary that can benefit from long-term projections of population numbers and responses of the species to human activities. We anticipate that our simulation model will assist the Army in meeting these legal and policy requirements.

The RCW has highly specific habitat requirements. It is the only woodpecker known to excavate cavities for nesting and roosting in living pine trees (Ligon et al. 1986). Furthermore, these pine trees are typically infected with heart-rot fungus (*Phellinus pini*) and tend to be the older trees in the forest (Hooper 1988, Hooper et al. 1991). Older and infected trees can be more easily exploited for cavity excavation than younger trees (Hooper 1988). RCWs depend on their cavities for roosting, nesting, and rearing their young (Ligon 1970). Limited availability of appropriate cavity trees has been thought to be a major factor in the bird’s decline (Hooper 1988).

The RCW is also unique in that it is a cooperative breeder. Some of the young males remain on their natal territory for up to seven years as nonbreeding helpers (Ligon 1970, Lennartz et al. 1987, Walters et al. 1988). If the breeding male within the natal territory dies, helpers often inherit the natal territory and become breeders. Helpers may also disperse to a nearby territory to become breeders. Other young males, and nearly all females, disperse during their first year to search for a breeding vacancy (Walters et al. 1992). Because of the RCW’s relatively short dispersal distances, the spatial distribution of clusters of cavity trees appears to play an important role in the population dynamics of the species (Engstrom and Mikusinski 1998).

Army land managers must provide high-quality military training while supporting endangered species management and other natural resources objectives. As part of these efforts, environmental data collection and field research have occurred over many years at the larger Army installations. In recent years, the Army has invested in the development of simulation models to help assess the impact of human activities on natural resources. To date, scientific data and computer technologies are often applied to one objective or one facet of the mission at a time. As interdisciplinary approaches and integrated management become more common, decision support systems must be able to combine the various aspects of the ecosystem with the land use mission and land management activities. The Army has launched a large initiative called Land Management Systems (LMS) to promote the integration of technological support capabilities and to solve the varied technical hurdles to increased integration of dynamic simulation models. We hope that our work also helps to meet this Army need.

This paper describes the expandable object-oriented modeling framework we used to implement an agent-based, spatially explicit population model for the RCW. The framework chosen, called the Dynamic Information Architecture System (DIAS), meets our numerous design criteria. It provides dynamic interaction between the population model and other future models and applications. We will project RCW populations dynamically (through space and time) and then expand the model to include other natural processes and various land use and land management influences that are acting within the ecosystem. An earlier application of DIAS produced the Object-Oriented Integrated Dynamic Analysis and Modeling System (OO-IDLAMS). OO-IDLAMS was developed to demonstrate the advantages of an object-oriented architecture approach to integrated natural resources decision support (Sydelko et al. 1998).

APPROACH

We implemented a dynamic RCW population model developed separately by Lechter et al. (1998) within DIAS as the first step in providing a flexible, robust simulation tool for the Army. The Agent-Based Model (IBM) approach was chosen for implementation over a more traditional aggregate population modeling approach, because of its ability to (1) describe the population traits with distributions rather than mean values, (2) represent agent performance and local interactions, and (3) provide a mechanistic rather than a descriptive approach to modeling (DeAngelis and Rose 1992). Furthermore, the model described by Letcher et al. (1998) is spatially explicit, accounting for the importance of spatial distribution to RCW population dynamics.

Currently, the model is being validated using Fort Benning, Georgia, as a case study. Plans are currently underway to incorporate a forest growth model, a RCW foraging and nesting habitat model, and military activity projections within DIAS to further provide modeling and simulation to support decisions impacting the RCW at Fort Benning.

Implementation of RCW Population Model

The dynamic behaviors (or processes) described in the Letcher et al. (1998) model were broken out into five distinct “submodels” for implementation in the DIAS RCW application: breeding, competition, dispersal, fledgling role change, and mortality. These Letcher et al. submodels address specific behaviors (DIAS “Aspects”) of specific classes of objects (DIAS “Entities”).

We utilized Lechter's model by capturing individual birds, their population groups, and their territories as object entities, and by assigning separate, specific properties and behaviors to each entity.

The Entity object classes developed for the RCW application (RCWIndividual, RCWTerritory, RCWPopulation) contain the attributes that describe the state of the environment throughout a simulation. These Entities also contain links to the processes (the Letcher et al. [1998] submodel behavior) via their Aspects.

Design Criteria for RCW Model

The following sections describe characteristics of DIAS that met our design criteria for this application.

Modularity

Creating a modular, compartmentalized structure allows us to more closely mimic the associations and relationships that exist in the real world and provides a level of modeling flexibility that is well beyond what we could achieve using traditional modeling practices (GIS-based or procedural). By defining the object entities and articulating how they interact through process code, we provide a foundation upon which we can logically articulate and include other relationships and connections in the future. While the application development is very specific (in this case, RCW population dynamics), the application component development is generic. This avoids the pitfalls of "monolithic" models where development proceeds as a series of submodel additions, one on top of another, resulting in a highly unwieldy and inflexible model structure that is time-consuming and costly to expand.

Furthermore, because the RCW Entities and their associated behaviors are developed within a framework that already houses a diverse array of environmental and non-environmental objects (Hydrologic, Atmospheric, and Vegetative), making the connection to other models and processes is easier. We do not have to build everything from the beginning — we can utilize existing objects from the DIAS library and add attributes and behaviors as new applications dictate.

It is important to note that this modular approach allows the user to readily plug in alternative RCW behaviors (submodels) without time spent recoding or reworking the existing application. One would simply substitute a new algorithm for one of the behaviors (write a new method for the Process object), but the connections between the Entities and the events that trigger behaviors remain the same. Thus, altering models or substituting a new model is not disruptive to the overall system, and it is therefore more efficient and more cost-effective to develop applications in this type of framework.

Systems that serve environmental managers need to include opportunities for updating and changing information as knowledge is advanced through related research and monitoring. Overall, DIAS provides an efficient framework within which to bring together disparate data and software for integrated resource planning. It is flexible and robust enough to readily assimilate new models into an existing application.

Code/Object Reusability

Entity class objects need only be designed and built once. Any future application or model that requires use of RCWs can utilize these existing objects as is, or by simply adding any additional attributes or process linkages that may be required for a new application. In this way, objects continue to mature, but are not recoded, as is often the case with more conventional model development. This obviously points to the need to thoroughly design Entity objects up front and to make certain that they are generic enough to be utilized by diverse domains. The reusability of objects and code can and does represent significant cost savings. Oftentimes, applications can be more quickly developed in the DIAS framework because of the benefits of object/code reusability.

Expandability

Our intention was to design the RCW model application so that future implementations could predict RCW populations by incorporating various land use and land management influences acting in the ecosystem. Again, because the approach taken was modular, and the behaviors are "contained" and distributed to the appropriate Entities, we were able to build a system that will not require recoding to incorporate management impacts. The RCW state — whether Individual, Territory, or Population — will be affected by changes in the environment, regardless of what is producing the environmental change. These cause-effect linkages are built into the system. For instance, a change in vegetation whether produced by fire, management practice or disease, will produce a resulting change in RCW habitat that could, in turn, affect the bird. A multitude of influences can impact the natural environment. These influences change the state of the environmental objects first (directly), and then this effect is propagated down through the birds (or any other object in the simulation) as appropriate. In this sense, linkages are not "hard-coded," but occur naturally, reflecting our current scientific understanding of interrelationships. The birds register to receive events that are "of interest" to them. If an environmental parameter changes, regardless of what changed it, and this parameter is of interest to the bird, the bird will receive notice of the change and will respond accordingly.

Adding in management influences to the existing application will simply be a matter of articulating the management practice/procedure and coding a model that reflects the practice (most likely a CourseOfAction type [COA] model). As management affects the state of the environment, these changes will impact the RCW.

Expressiveness of DIAS COA Objects

We used the FACET (Framework for Addressing Cooperative Extended Transactions) object suite to code the breeding behavior as a COA model. FACET objects are used to represent complex interactions between objects, or "agents," in a simulation. For this application, the agents are individual RCW birds. The breeding behavior COA is basically a flowchart of agent actions, and acts in the same manner as a Process object to implement behaviors of Entities. While this particular COA is not very complex, it illustrates the ability of DIAS to handle social process models that involve cooperative behavior between agents.

FACET COA type models will be critical for the development of land use and land management plans to assess the interplay of human impact on the RCW environment.

System Dynamics and Feedback Mechanisms

Traditional GIS-based systems are static in nature and do not lend themselves very well to dynamic, inter-process modeling. In general, most models and applications are designed to operate independently, even though effective decisions call for assessing several components of the ecosystem simultaneously, in terms of their relationship to each other as well as how they affect management decisions. DIAS provides a framework for developing applications that address inter-process dynamics in a highly realistic way. DIAS allows us to articulate the dynamics of an ecosystem much more closely to the way in which we understand them to operate in the real world. At the same time, it does not impose a single worldview (one discipline's view) on the development of the application.

DIAS has an event-driven simulation environment. Events in a simulation invoke the behaviors of the Entities at the appropriate time, following the dynamics of the system, as we have articulated them. We specify what objects will be in the simulation (the "playing pieces"). These objects carry their state as attributes. At any time during a simulation, we can pause and assess the state of the environment by evaluating these attribute values. Any process operating in the simulation can potentially modify the state of each object. These processes can be associated with internal or external models, natural/physical models, or management plans. In this way, the DIAS inter-process flows are realistic and reflect the dynamics of the system as we understand them.

CONCLUSION

We have described using DIAS, an object-oriented, event-driven architecture, as an enabling technology to construct a flexible framework for simulating RCW populations, and eventually, habitat and management factors that will affect the populations. Dunning et al. (1994) point out that spatially explicit population models are not a panacea for predicting the locations of individual animals with a high degree of accuracy. Rather, simulation efforts are building a bridge to link ecological research with applied fields such as wildlife management and conservation biology (Turner et al. 1995). Modeling exercises can serve to collect knowledge, data, and theory about ecological processes to improve understanding and communication; serve to screen potential management actions; and help to identify gaps in knowledge or data needed to assess the impacts of human activities

REFERENCES

DeAngellis, D.L. and K.A. Rose. 1992. Which individual-based approach is most appropriate for a given problem? In DeAngelis, D.L. and Gross, L.J. (Eds.). *Individual-based Modeling and Approaches in Ecology: Populations, Communities, and Ecosystems*, Chapman and Hall, New York, New York USA, pp. 67-87.

Dunning, J.B., D.J Stewart, B.J. Danielson, B.R. Noon, T.L. Root, R.L. Root, R.H. Lamberson, and E.E. Stevens. 1995. Spatially explicit population models: Current forms and future uses. Ecological Applications 5(1): 3-11.

Engstrom, R.T and G. Mikusinski. 1998. Ecological Neighborhoods in Red-Cockaded Woodpecker Populations. Auk 115(2): 473-478.

Holling, C. S., editor. 1978. *Adaptive Environmental Assessment and Management,* John Wiley, New York, New York, USA.

Hooper R.G. 1988. Longleaf pines used for cavities by red-cockaded woodpeckers. J. Wildl. Manage. 52: 392-398.

Hooper, R.G., M.R. Lennartz, and H.D. Muse. 1991. Heart rot and cavity tree selection by red-cockaded woodpeckers. J. Wildl. Manage. 55: 323-327.

Jackson, J.A. 1994. Red-cockaded woodpecker (*Picoides borealis*). The birds of North America, number 85. The American Ornithologists' Union, Washington, D.C., USA, and The Academy of Natural Sciences, Philadelphia, Pennsylvania, USA.

Lennartz, M.R., R.G. Hooper, and R.F. Harlow. 1987. Sociality and cooperative breeding in red-cockaded woodpeckers. Behav. Ecol. Sociobiol. 20: 77-88.

Ligon, J.D. 1970. Behavior and breeding biology of the red-cockaded woodpecker. Auk 87: 255-278.

Ligon, J.D., P.B. Stacey, R.N. Conner, C.E. Bock, and C.S. Adkisson. 1986. Report of the American Ornithologists' Union Committee for the Conservation of the Red-cockaded Woodpecker. Auk 103: 848-855.

Letcher, B.H., J. A. Priddy, J.R. Walters, and L.B. Crowder. 1998. An individual-based, spatially-explicit simulation model of the population dynamics of the endangered red-cockaded woodpecker, *Picoides borealis*. Biological Conservation 86: 1-14.

Turner, M.G., G.J. Arthaud, R.T. Engstrom, S.J. Hejl, J. Liu, S. Loeb, and K. McKelvey. 1995. Usefulness of spatially explicit population models in land management. Ecological Applications 5: 12-16.

Walters, C. J. 1986. *Adaptive Management of Renewable Resources*, McMillan, New York, New York, USA.

Walters, J.R., P.D. Doerr, and J.H. Carter III. 1988. The cooperative breeding system of the red-cockaded woodpecker. Ethology 78: 275-305.

Walters, J.R. 1990. Red-cockaded woodpeckers: a "primitive" cooperative breeder. Pages 67-102 in P.B. Stacey and W.D. Koenig (Eds), *Cooperative Breeding in Birds: Long-Term Studies of Ecology and Behavior*, Cambridge University Press, Cambridge, United Kingdom.

Walters, J.R., P.D. Doerr, and J.H. Carter III. 1992. Delayed dispersal and reproduction as a life-history tactic in cooperative breeders: fitness calculations from RCWs. American Naturalist 139: 623-643.

DISCUSSION: ENVIRONMENTAL PROCESSES

D. SALLACH, University of Chicago, Moderator

Robert Reynolds [to Chris Rewerts]: In your habitat model, are different aspects of the habitat considered agents as well? For example, your birds are part of a food chain: they have predators and they prey on others. And the Army itself could be viewed as an agent, allowing the model to predict the impact of various maneuvers on the bird population.

Chris Rewerts: Yes, although the predator-prey issue is not a major one in this case, if you maintain the habitat properly. For instance, if the birds are in big, open areas, then, yes, kestrels come in and get them. Another predator is a snake that climbs up the tree and gets into the hole. The birds take care of that by making the sap run, which is an irritant to the snakes. Another problem is other creatures who try to use the holes. The site staff manages that problem by putting restrictor plates on the holes so only the birds can fit through — they're not afraid to use hardware out there.

The modeling of training is something that I thought I would be doing soon. As it turns out, though, the field studies done so far show that the impacts of the noise, maneuvers, and smoke on the birds is minimal. The Army has management guidelines — 1996 is the latest iteration of those — that say how it can train around the nesting sites. Units are required to stay 200 meters away, I believe, from a nesting site during maneuvers.

Modeling Army training is a bit problematic, because, for security reasons, they don't want to tell us much about what they do. So we look at Army training the same way we look at any other animal: we take measurements then extrapolate upon the landscape. Another project I worked on did just that — it looked at impact probabilities for different types of training events. We could say, for example, "In this footprint of area they're making this type of noise and they're using smoke and obscurance," and we could characterize the time and the spatial location of the training attributes. So, yes, we can get to that point, but so far we've prioritized in different directions based on what we've been learning about the effects of training.

Pamela Sydelko: I'm Pam Sydelko from Argonne. I've been working with Chris on this project. One of the interesting things is that there are many different objectives for managing the land. We run into this issue with the military, but it also applies to any land-managing agency. One of the things the Army has a real struggle with, at Fort Stewart or Fort Benning, for instance, is that primarily they are just managing for this bird. In some installations, the Army may be managing for a bird but also has a water quality or erosion problem. And there are many stakeholders that have to be considered.

I'm interested in how we can represent those different stakeholders as agents, to feed into our modeling suites that look at long-term satisfaction with land management decisions and the tradeoffs involved. There's only so much money you can spend in an installation to manage all of these things. If you make erosion better, you might not have enough money to maintain habitat for an endangered species. Stakeholders drive many of these decisions. And the stakeholders are as dynamic as the environment itself. For example, we're putting dynamics, such as a drought, into the environment as we do this modeling. But what if all of a sudden the red-cockaded woodpecker is not the issue anymore, but an endangered tortoise instead. Your management plan wasn't concerned with soil compaction during training because that's not an issue for the bird, but it's important to the tortoise. It would be helpful to have some agent-based modeling ideas about this type of dynamic — and this would apply for national park and forest management, as well — because these changes in policy and focus happen all the time.

Rewerts: That's a very good point. It sounds as though we're focusing on one species to exclusion. But really our goal is to move toward ecosystem management, even though that science is still wide open.

Christopher Langton: I agree that that's a very important point. We tend to believe that all we have to do is get the science right, and then we can go to the policymakers and say, "Here's the truth," and they'll say, "Fine." But it's not that easy. The policymakers have to consider the special interest groups. This has come up in some models for the Columbia River, where people are trying to get salmon to come back into the river. A solution may be known, but implementing the solution would create conflicts with the special interests. And so it's incredibly important to consider, and probably even put into our models, all of the special interest groups that are working with the policymakers on these issues.

Benjamin Schoepfle: I'm Benjamin Schoepfle from Argonne. Pam's comments are very important. I've done similar work on military bases in the South helping to locate corridors for tanks in relation to a species of woodpecker. One issue is that many of the base commanders have short tenures of about two years, so a new commander may not know the history, especially the finer nuances of these very spatially-tuned policies. I think a context of history is important for modeling: it's important not only to look toward the future, but also to replicate the past. Modeling can give a historical record, and record-keeping is extremely important because many managers are on the job for such a short time.

Randal Picker: I'm Randy Picker of the University of Chicago Law School. To follow up, it's not obvious to me that that an agent-based framework is the right way to capture interactions between policymakers — who are the people I spend a lot of my time with — and the kinds of models you are building. The agent-based framework strikes me as particularly good for dealing with population situations, but I'm not sure it's the right tool when you've got principally five actors who interact with each other, maybe on a repeat basis, in a very small setting.

Rewerts: There's a couple of different ways of looking at this, and I'm going to pick the easy one. One reason to model is for consensus-building. All models are based on assumptions. As you put these things together and see how they operate, you're testing your theories of what you think is actually going on out there in the ecosystem. If everybody agrees that you've got a valid theory being played out in this simulation, you can make policy decisions based upon it. If you're not using a tool like this to help make decisions, you're just flying by the seat of your pants. But you're talking about whether or not to put various decision-makers in as agents?

Picker: Exactly. It's not a question of whether your model does a good job of capturing what's going to happen to woodpeckers. I suspect it does, and it seems to me the Army's doing exactly the right thing in looking to you. But my question is how to get at how some other group is going to react to that, whether the Sierra Club is going to decide, "Woodpeckers aren't important; it's snail darters that we care about." I don't think an agent-based model is necessarily the best way to get at that.

Sydelko: Actually, I think one of the things that we're talking about is not so much the tail-end arguments of the policymakers looking at the decisions made by the red-cockaded woodpecker model and deciding whether it does what they expected. What would be nice is more of a front-end engine, a suite of models, that actually balances these many factors and concerns in ecosystem management. For example, we have vegetation models that show training impacting vegetation, the woodpecker reacting to the vegetation changes, and erosion occurring because of the vegetation changes.

It's a challenge to make a management plan that 20 years out is going to satisfy all conditions. Many of these conditions are generated not by land managers, but by other parties. It would be helpful to have a series of scenarios based on various players' perspectives and what they value. We had a workshop with a group of land managers and military people to try to determine what they value most about the land. Is it protecting endangered species or controlling erosion? Is training most important? We wanted to see whether we could generate a finite set, perhaps even using genetic algorithm tools, of good possible scenarios, run them through the model, and determine how well we meet the criteria. This might give us an idea of where to start. There's a lot of choices to be made about land management, not only about *what* to do but *where* to do it. If we had some idea of where we could use the agents and consensus-building to drive the scenario-building, that would be worth looking at.

Energy and Infrastructures

BUILDING ELECTRICITY MARKET PARTICIPANT STRATEGIES WITH ADAPTIVE AGENTS

C. RICHTER, Alstom ESCA Corporation, Bellevue, Washington*

ABSTRACT

Strategies that reduce risk and increase profit are important to electricity market participants. Financial instruments that hedge against uncertainty and manage risk are important components of such strategies. One such vehicle is the fixed transmission right (FTR) on a transmission path. The FTR holder is entitled to the transmission congestion rents collected on that path. However, if congestion occurs in an unanticipated direction, the holder of an FTR is obligated to pay. An interesting compromise derivative is the FTR option. For a premium, the FTR option gives its holder the right to the congestion rents if they are positive, without the obligation to pay if they are negative. This paper builds on previous work [13] and discusses how adaptive agents can be used to develop strategies that help the market participant value and utilize FTR options as part of a comprehensive market strategy.

INTRODUCTION

Electricity markets are being re-regulated to promote competition [3,4,5]. Entry barriers to new market participants are being removed; formerly monopolistic utilities have been forced to separate their generation, transmission, and distribution into independent companies known as GENCOs, TRANSCOs, and DISTCOs, respectively. New market entities are emerging (e.g., energy service companies (ESCOs) that purchase wholesale electricity and repackage it for resale to the end-consumer). Most of these market participants are profit-based and benefit from well-designed bidding strategies. The strategies must each be designed for a specific set of market rules, which vary from region to region.

One promising technique for developing strategies involves intelligent or adaptive agents. In addition to the promise these techniques hold for developing profitable strategies for market participants, the same techniques can be used by market regulators and operators to test the market structure, find loopholes that might allow gaming, and improve the stability of the market. The technique begins with building a simple model of the market participants. Part of the rules making up the agent must be adaptable. A training, learning, or evolutionary process (e.g., a genetic algorithm, weight updating) presents the agents with various sets of inputs and rewards or punishes them based on their performance, as judged by a variety of metrics.

Electricity flows according to the laws of physics over a network that is subject to congestion. Real-world market operation, pricing, and product offerings in many markets are designed to reduce the impact of transmission congestion. Simulations using agents to model market participant behavior or develop strategies for a competitive electricity market should include a model of the electricity network

---

* *Corresponding author address*: Charles W. Richter, Jr., Alstom ESCA Corporation, 11120 NE 33rd Place, Bellevue, WA 98004; e-mail: charles.richter@alstom.esca.com.

if they are to produce results of interest to many of the companies seeking to make profit in the new world of electricity.

ELECTRICITY NETWORKS AND COMPETITIVE MARKETS

In some areas, organized exchanges centralize the trading of electricity products. Some areas calculate a region-wide price for electric energy, while others develop locational marginal prices (LMPs) for each node in the network [6]. Some exchanges remove contract default risk by becoming an intermediate partner to all trades. Centralized exchanges facilitating physical (intended for delivery) power trading must interface with an independent system operator (ISO), responsible for secure electricity network operations. Other regions rely on decentralized bilateral trading. Bidding and operational strategies for bilateral trading may vary widely from those used for centrally allocated markets.

Electricity markets should promote economic efficiency and secure power system operation. This means balancing the simplifying assumptions required for creating a liquid market with the reality of power system operation. Power flows through a network according to Kirchoff's laws, making it difficult to predict how a particular transaction might impact the network without considering all other transactions. Additional complexity comes with the market treatment of supportive services (e.g., reactive power and transmission) related to the physical flow of electricity, which may vary widely with the market implementation.

Real-time electricity prices can be largely impacted by transmission congestion. Transmission congestion can prevent the transport of electricity over the network, thereby causing the price of electricity at the demanding end to rise, and the price at the sending end to fall. Congestion costs are often measured by the differences in sending and receiving end LMPs. This amount is collected from those transporting electricity across the congested path, thus providing an incentive to shift generation to the other side of the congested transmission path.

To reduce transmission congestion uncertainty, interested participants can purchase fixed transmission rights (FTRs) from an injection point to a delivery point. The FTR cannot guarantee physical delivery, but can make the holder financially whole by entitling the holder to congestion rents collected along that path in that direction. If network congestion requires a supplier to pay large collection rents to deliver the product, the FTR allows that supplier to offset the charges. Unfortunately, if other transactions on the network cause the power to flow in the opposite direction (counter-flows), the FTR holder is obligated to pay the negative congestion rent.

To further reduce the transmission congestion risk, many market participants have expressed an interest in FTR options. Options on FTRs allow the holder to collect positive congestion rents, but do not obligate them to pay negative congestion rents. Several markets regions are on the verge of adopting FTR options. Several more markets are considering the sale of flow gate rights, which share the nonobligatory nature of FTR options. The FTR option should be considered a valuable part of market strategies.

FIXED/FINANCIAL TRANSMISSION RIGHTS

To illustrate how FTRs are utilized, an example with and without an FTR is presented. Losses are ignored in this example. Suppose that generator *G* is able to generate electricity for a cost of \$10/MWh. *G* wants to lock-in a price for generation and has bilaterally contracted to supply 100 MW for an hour to load *L* for \$12.50/MWh. *G* (Node A) and *L* (Node B) are connected through a meshed network as shown in Figure 1. If the network is not congested and LMPs at nodes A and B are \$10/MWh, then *G* generates the 100 MWs, and *L* consumes 100 MWs. *G* collects \$250 more from the bilateral contract than it would have collected from selling to the spot market. *L* paid \$250 more than it would have paid from the spot market. Even though in this scenario *G* profited \$250 and *L* "lost" \$250, both *G* and *L* are happy because they locked in their price ahead of time, thus reducing their exposure to the uncertainty of the energy spot price.

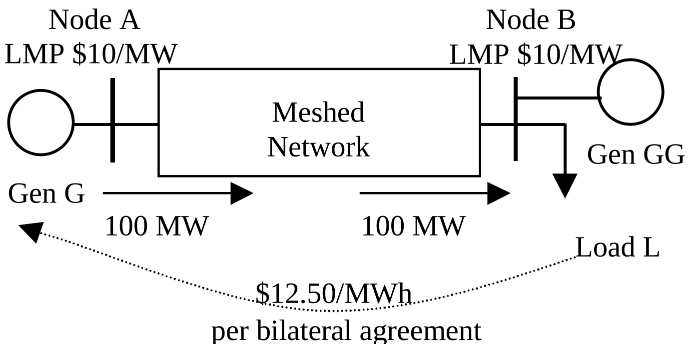

FIGURE 1 Simplified network, without congestion

However, if the network between Node A and Node B is congested, as shown in Figure 2, the LMPs at both of the nodes will be different. *G* may not be able to physically deliver its generation, and may have to purchase generation at Node B to meet its obligation to *L*. Because generation at Node B could be quite expensive, *G* is now exposed to considerable risk.

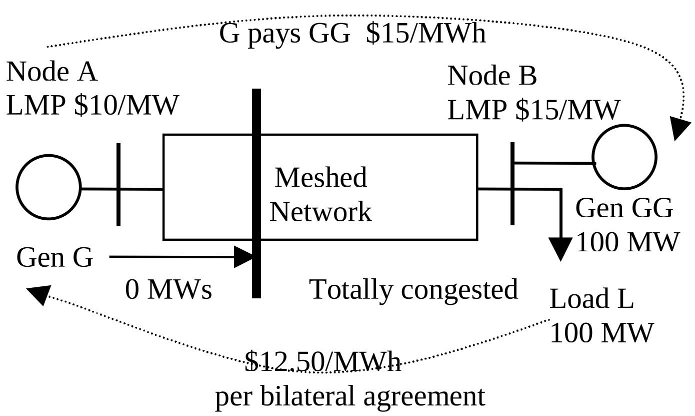

FIGURE 2 Simplified network with congestion (decentralized dispatch)

An example will help illustrate this exposure to price uncertainty. If the path from A to B was completely congested, as shown in Figure 2, and LMP at B was \$15/MWh, then *G* would be forced to purchase power at a higher LMP at B (e.g., \$15/MWh) and deliver it to *L*, who would pay only \$12.50/MWh. This means that *G* would be losing \$2.50/MWh. If LMP at B was lower than at A, then congestion is in the opposite direction and *G* will generate from its own unit. Note that in a less congested scenario, *G* may be able to provide some power from its own generator and would purchase the remainder from *GG*. Table 1 shows *G*'s profit in a decentralized environment for the bilateral transaction with *L* under scenarios with varying LMPs at Node B. In the table, *G*'s profit is calculated as:

$$profit = revenue - total\ costs\ ,$$

where

revenue = *L*'s payment to *G*,
total cost = (*G*'s gen cost) + (*G*'s payment to *GG*),
G's gen cost = (\$10/MWh) × (*G*'s gen), and
G's payment to GG = (LMP@B) × (*GG* MWs).

TABLE 1 *G*'s Profit for a 1-h 100-MW Bilateral Contract under Four Scenarios

LMP A (\$/MWh)	LMP B (\$/MWh)	*L* Pays *G* (\$)	*G* gens (MW)	*G*'s gen (\$)	*G* pays *GG* (\$)	*G*'s total Cost (\$)	*G*'s "profit" (\$)
10	20	1,250	0	0	2,000	2,000	-750
10	15	1,250	0	0	1,500	1,500	-250
10	10	1,250	100	1,000	0	1,000	250
10	5	1,250	100	1,000	0	1,000	250

In a centrally dispatched market, each generator is dispatched by the ISO, which is essentially running a power flow that considers physical network limits, bids, and offers. The LMPs are a result of the central dispatch and allocation process. The ISO generally would not consider any bilateral agreements between *G* and *L* in the centralized dispatch. The ISO would collect and distribute the money associated with the centrally allocated transaction. Table 2 demonstrates results for the same size transaction and the same nodal price scenarios when coordinated by the ISO.

TABLE 2 Settlement for *G* to *L*, 1-h 100-MW Transaction under ISO

LMP A (\$/MWh)	LMP B (\$/MWh)	*L* Pays ISO (\$)	*G* Gens (MW)	ISO pays *G* (\$)	ISO pays *GG* (\$)	*G* pays *L* (\$)	*G*'s cost (\$)	*G*'s profit (\$)
10	20	2,000	0	0	2,000	750	750	-750
10	15	1,500	0	0	1,500	250	250	-250
10	10	1,000	100	1,000	0	-250	1,250	250
10	5	500	100	500	0	-500	1,000	250

Table 3 demonstrates generation and payments allocation in the presence of congestion-induced transmission limits. The ISO is clearing (or not clearing) offers submitted by generator *G* at Node A (100 MW offer @ $10/MWh) and by generator *GG* at Node B (100 MW @ $15/MWh). Load L requires 100 MW in this example. The table illustrates generation levels and settlement under different levels of congestion-limited transmission. When LMP A is greater than LMP B, *G* is not responsible for any congestion payments (CP's). When the LMPs at all nodes are the same, there is no congestion in the network. However, when LMP A is less than LMP B, *G* is liable for congestion payments (implicitly defined in the previous table) [(LMP B − LMP A) × MW] as *G* has not hedged against this congestion.

TABLE 3 Effects of Congestion on *G*'s Profit without FTR

Cong Limit A-B (MW)	Supply Cleared @ *G* (MW)	LMP @ A ($/MWh)	Supply Cleared @ *GG* (MW)	LMP @ B ($/MWh)	*L* Pays ISO ($)	ISO Pays *GG* ($)	ISO Pays *G* ($)	*G* pays ISO CP ($)	*G*'s Profit ($)
100	100	10	0	10	1,000	0	1,000	0	250.0
75	75	10	25	15	1,500	375	750	375	125.0
50	50	10	50	15	1,500	750	500	250	0.0

To hedge against paying the uncertain LMP at B under congestion, *G* can chose to purchase an FTR of 100 MW (in this example, the price is assumed to be $1/MW) to ensure that it is hedged. *G*'s profit (including the cost of purchasing FTR and the profit of selling electricity at $2.50 more than generation costs) is shown in Table 4. LMP A is kept the same, while LMP B is reduced from $15 to $5. The congestion limit is assumed to be 100 MW.

TABLE 4 *G*'s Profit with an FTR

LMP @ A ($/MWh)	LMP @ B ($/MWh)	FTR (MW)	FTR ($/MWh)	*G* Pays ISO CP ($)	ISO Pays *G* CR ($)	*G*'s Profit ($)
10	15	100	1.0	500	500	150
10	10	100	1.0	0	0	150
10	5	100	1.0	500	0	-350

When LMP B is greater than LMP A, *G* has a profit of $150, as it collects $500 in congestion rent and has paid $100 to purchase the FTR, and has a profit of $250 for the bilateral contract with *L*. When there is no congestion, then it also stands to make a profit of $150. However, when LMP B is greater than LMP A, then *G* stands to lose $350 because of the counter flow in the opposite direction of the purchased FTR. Hence as the price differential increases, and the LMP A is larger, then *G* will start to make large losses although it has purchased an FTR and hedged against congestion. *G* is always financially liable for any congestion charges because of counter flow in the opposite direction of the FTR purchased. FTRs come with their own risk. If power flows in an unanticipated direction (counter-flows), the FTR holder is obligated to pay the negative congestion rent.

FIXED TRANSMISSION RIGHTS OPTION

A fixed transmission rights option (FTRO) gives its holder the right to the congestion rents when power flow is in the indicated direction, without the obligation to pay when power flows in the wrong direction. A party "writes" an FTRO in return for the premium, or FTR option purchase price, and is then obligated to collect/pay congestion rents when the flow is opposite the direction anticipated.

Figure 3 shows the profitability of an FTRO to the congestion rents from point A to B. The purchaser has paid a premium (e.g., $1) to the FTRO writer, thereby giving the purchaser the right to positive congestion rents (e.g., $0). The purchaser has reduced risk by limiting losses to the premium. The right side of Figure 3 shows what happens from the FTRO writer's point of view. The FTRO writer receives the premium for assuming the risk and is obligated to pay the congestion rents when they are negative.

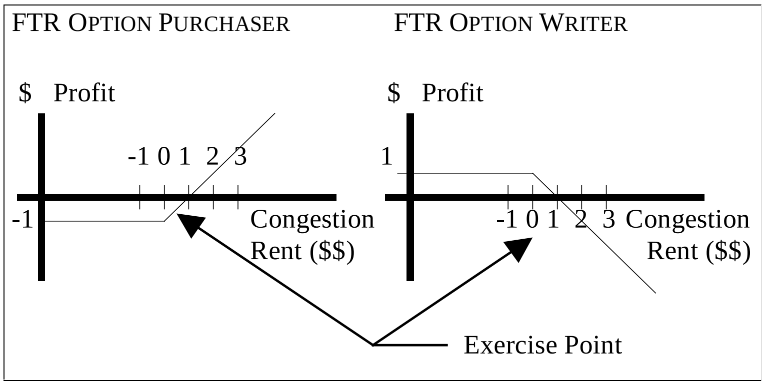

FIGURE 3 FTRO profitability

The FTRO price should reflect the value of the FTRO to the potential holders. The worth of an FTRO may vary from trader to trader because of risk preferences, makeup of portfolios (collection of assets and contracts), etc., but is largely dependent on the uncertainty in the underlying asset. Large uncertainties in the direction of flow on the FTR path and in the amount of congestion rent to be collected translate to high premiums (FTRO prices).

The method used to value options must consider the idiosyncrasies of electricity. Many permutations of Black-Scholes have emerged. Another approach that can help identify an upper bound on the value of the FTRO is to determine expected monetary value (EMV) and value at risk (VAR) under scenarios of interest (see Figure 3).

FTR VERSUS FTRO

The FTRO is an additional tool for managing risk. Here the only FTR example is augmented with FTRO for comparison. Revisiting the simple system depicted in Figure 1, Table 5 provides an example of hedging with FTRO. The FTRO protects *G* from congestion charges because of counter flows. The profit of *G* if it purchases an FTRO from Node A to Node B is shown in Table 5. The congestion limit is assumed to be 100 MW.

From Table 5, when LMP A is greater than LMP B, then *G*'s profit is $100, which is the profit of selling electricity to *L* minus the cost of purchasing the FTRO. This is also the case when there is no congestion, and when LMP A is greater than LMP B. Hence with an FTRO, *G* is not responsible for any counter flow from Node B to Node A.

TABLE 5 *G*'s Profit with an FTRO

LMP A ($)	LMP B ($)	FTRO (MW)	FTRO ($/MWh)	*G* Pays ISO CP ($)	*G* Gets from ISO CR ($)	*G*'s Profit ($)
10	15	100	1.5	500	500	100
10	10	100	1.5	0	0	100
10	5	100	1.5	0	0	100

G may decide to use both an FTR and an FTRO from Node A to Node B to hedge its financial risks. The question for the market participant becomes not which instrument to use, but rather how much of each should be used to achieve the desired risk to premium ratio.

AGENT-BASED TECHNIQUES

Developing market strategies that work well in many market scenarios can be a difficult task. Theoretical macroeconomic assumptions do not always hold (e.g., rational behavior) in a microeconomic setting. One way to develop rules is to combine experimental economics with adaptive agent techniques. The agent can be (by some people's definitions) autonomous, intelligent, and adaptive (meaning that it's behavior changes over time based on the input-output successes and failures that the agent has encountered in the past). Computerized agents (consisting of evolving or adapting rules that process state information) can represent market participants. Markets are simulated in which the adaptive agents are allowed to buy and sell and test strategies. Even without a complicated model for each individual agent, interesting behaviors can be observed when many of these agents are allowed to react in an environment.

The agents (or the strategies of which they consist) can be developed to model various problems that market participants face in the evolving deregulated market place. For example, agents may learn a set of rules for producing bids or offers for forward electricity. The agents may evolve rules for hedging against electricity network congestion by valuing and bidding on or offering products like FTRs or FTROs. The rules can be combined to include responses for either problem, depending upon which inputs are presented to the agent. In most cases, agents produce a response or an action when presented with a set of inputs or state information. Possible actions for an agent representing an agent playing in the deregulated electricity market place might be as shown in Table 6. Some of the many possible inputs or environment variables that agents may wish to observe as part of their strategy (depending on the experiment) are as follows:

- Forecasted price,
- Forecasted demand,
- Lowest and highest bid and offer observed during previous negotiations,
- Average market price,
- Measure of market power,
- Measure of market depth,
- Scheduled outages,
- Forecasted network congestion,
- Fuel costs, and
- Competitor's historical actions.

TABLE 6 An Example of an Agent Market Participant Strategy Action

Take Action	Yes/No	Time Last Performed
Request load forecast?		
Request price forecast?		
Update pricing model?		
Perform unit commitment?		
Run trade-mix/risk optimizer?		
Negotiate fuel contracts?		
Initiate cent. market purchase?		
Accept cent. market purchase?		
Initiate cent. market sale?		
Accept cent. market sale?		
		With whom?
Initiate bilateral energy purchase?		
Accept bilateral energy purchase?		
Initiate bilateral energy sale?		
Accept bilateral energy sale?		
Initiate FTR purchase?		
Accept FTR purchase?		
Initiate FTR sale?		
Accept FTR sale?		
Etc.		

GP-Automata

One way to process the state information to develop the input-output rules is through GP-Automata. GP-Automata combine finite state automata with genetic programs (GPs). They were first described as such by Ashlock [11] and were used by Ashlock and Richter [12]. The typical finite state automaton specifies an action and "next state" transition for a given input or inputs. With only one or two binary inputs to work with, it can be fairly simple to develop a finite state diagram to cover the possible input/output relations. When the number of inputs is large, the task is much harder. The number of transitions needed to cover all possible combinations of inputs grows exponentially (e.g., 10 inputs each having 5 possible values would require 5^{10} transitions). This is where genetic programming comes in. The GP-trees are bandwidth compressors. GP-Automata uses them to select which inputs to consider and to perform computations involving these inputs. Table 7 gives an example of a GP-Automaton.

TABLE 7 A Four-State GP-Automaton

	IF ODD		IF EVEN		
State	Action	Next State	Action	Next State	GP (Decider)
1	14.5	1	U	1	lte (mul(10, abs (hbb))
2	*	1	37	3	ite(max(10, asb), hbb, lbb)
3	12	2	5	1	avg (5, abb)
4	U	3	*	2	47
Initial Action		24	Initial State	2	

Reading the rule encoded in the GP-Automaton in Table 7 is fairly simple. The automaton begins by bidding the number in the “initial action” field. Following the initial action, the “initial state” indicates which state is used next (in this case, 2). The GP-Automaton in the table has four states. Coupled with each of these states is a GP-tree termed a decider. When executed, the decider returns a value between 0 and 100. On the basis of that returned value, one of the following two things will happen: (1) if that value is even after truncation, the action listed under “IF EVEN” is taken and we move to the next state listed under “IF EVEN”; (2) if the returned value is odd after truncation, then we use the action and next state listed under “IF ODD.” The “action” is the number listed in the action field of the automaton, with two exceptions. The first exception is the “U,” which indicates that the value returned by the decider should be taken directly as the action. The second exception is a “ *,” which indicates that further computation is necessary, and, hence, the GP-Automaton refrains from acting immediately. Instead, it immediately moves to the next state. This gives rise to the possibility of complex (multistate) computation as well as infinite loops. To prevent infinite loops, one can specify a maximum number of “ * ”s to be honored, after which a valid action is selected at random.

The GP-Automata population evolves as in any GA. After selecting parents, as described in Part 4, offspring are produced using crossover and mutation. Crossover for the GP-Automata involves selecting (with a uniform probability) a crossover point ranging from zero to the number of states. We then copy parent1’s states from zero to the crossover point to child1, and parent2’s states to child2. Following the crossover point, child1 gets parent2’s state information and child2 gets parent1’s state information (including the associated decider). Before replacing less fit members of the population, each child is subjected to various types of mutation.

SUMMARY

Electricity derivative instruments are an essential part of strategies that enable market participants to maximize/increase profit and manage risk in the deregulated electricity industry. Accurate modeling of the transmission network and the effects of congestion are becoming increasingly important as the transmission systems become stressed by the added transactions under competition. A financial hedge against congestion is the FTR; for a premium, the FTR’s downside risk can be eliminated through options. This paper presents the FTRO in an easy to understand manner. It suggests that an agent-based framework is a useful tool in developing comprehensive market participant strategies, and for populating simulations with behaviors that mimic realistic market behavior. The proposed framework can be applied to forward as well as future markets. Among the challenges in operational implementation of the presented model are predicting spot market prices and demand, measuring and quantifying other state information for the agents' use, and designing easy-to-understand adaptive data structures that can encapsulate good market rules.

REFERENCES AND BIBLIOGRAPHY

[1] Hull, J.C. *Options, Futures, and Other Derivatives*, USA: Prentice Hall, 1997.

[2] Marshall, J. *Futures and Option Contracting: Theory and Practice*, South Western Publishing, 1989.

[3] Kumar, J., and G. Sheblé, "Framework for Energy Brokerage System with Reserve Margin and Transmission Losses," 1996 IEEE/PES Winter Meeting, 96 WM 190-9 PWRS.

[4] Sheblé, G., "Electric Energy in a Fully Evolved Marketplace," Presented at the 1994 North American Power Symposium, Kansas State University, KS, 1994.

[5] Sheblé, G., "Priced Based Operation in an Auction Market Structure," Presented at the 1996 IEEE/PES Winter Meeting, Baltimore, MD, 1996.

[6] Caramanis, M.C., R.E. Bohn, and F.C. Schweppe, "Optimal Spot Pricing: Practice and Theory," IEEE Transactions on Power Apparatus and Systems, Vol. PAS-101, No. 9, 1982, pp. 3234-3244.

[7] Howard, R.A. "Decision Analysis: Practice and Promise," *Management Science*, 1988:679-675.

[8] MIDAS, "Multi Objective Integrated Decision Analysis System," *EPRI Journal*, Vol. 12, 1987:57-58.

[9] Kumar, J. and G.B. Sheble, "A Framework for Transaction Selection Using Decision Analysis Based upon Risk and Cost of Insurance," Proc. of the 29th North American Power Symposium, Kansas State University, KS, 1994:548-557.

[10] Kumar, J., and G. Sheble, "A Decision Analysis Approach to Transaction Selection Problem in a Competitive Market," *Electric Power Systems Research Journal*, 1996.

[11] Ashlock, D., "GP-Automata for Dividing the Dollar," Mathematics Department, Iowa State University, Ames, IA, 1995.

[12] Ashlock, D., and C. Richter, "The Effect of Splitting Populations on Bidding Strategies," GP-97, Proceedings of the Second Annual Conference, pp. 27-34, San Francisco, CA, Morgan Kaufmann Publishers.

[13] Richter, C., A. Jayantilal, and J. Kumar, "Improving Market Participant Strategies with FTR Options," forthcoming in Proceedings of 2001 PICA Conference, Sydney, Australia.

AN ABS INVESTIGATION OF GENERATOR MARKET POWER IN THE ENGLAND AND WALES ELECTRICITY MARKET USING AN EXCEL/VBA PLATFORM

J. BOWER and D.W. BUNN, London Business School*

EXTENDED ABSTRACT

In November 2000, the UK electricity market regulator (OFGEM) proposes to replace the existing trading arrangements in the England & Wales Electricity Pool ("the Pool") with a bilateral market. This action has been taken in response to persistently high and volatile prices that are symptomatic of the extent to which generating firms have been able to exercise market power. In addition to changing the trading arrangements OFGEM has also forced the industry to restructure and as a result the largest incumbent generators divested plant to new entrant firms from the United States during 1999.

Under these New Electricity Trading Arrangements (NETA) consumers will contract directly with generators for supplies of physical power rather than a centralised market place. Self despatch of generating plant is envisaged as the main mechanism of delivery, and an optional Balancing (spot) Market will be used to maintain system security in which generators and consumers make firm bids and offers for increments and decrements of power, in real-time, for each half hour of the day. In contrast to the current Pool, in which all generators get paid the bid price of the marginal generating plant in each period (SMP) plus a capacity payment for making plant available to the system, generators would only be paid their own bid price. Underlying the proposed reform is the strong belief that the wholesale electricity market should operate more like other competitive commodity markets and that paying generators SMP only serves to increase the potential for gaming and exploitation of market power, in particular, that of the mid-merit generators. The rationale is simple; pay-as-bid pricing should reduce market power and, hence, wholesale electricity market prices will fall as a result.

Since the beginning of 2000 it is clear that Pool prices have been significantly lower than in previous years. OFGEM has publicly stated that it believes this fall in prices is due to the market's expectation of the imminent introduction of NETA in November 2000. However, this conclusion ignores the potential impact that industry restructuring during 1999 may also have had on market prices. To examine whether the recent price falls are due to proposed changes in the trading arrangements, industry restructuring, or both, we have built an agent-based simulation (ABS) model of the Pool and NETA. In this modelling environment, each generator is represented as an autonomous adaptive agent that submits a separate daily bid price, from each of its plants, at which it is prepared to supply electricity to the market. Equipped with a rudimentary learning ability, each agent is capable of developing its own bidding strategies in response to changes in trading arrangements, and the bidding strategies of other agents. Despite the simplicity of the agents' behavioural rules, complex bidding strategies emerge spontaneously. These are both consistent with behaviour seen in the real electricity markets and with economic theory. Additionally, as the model relies on the micro-simulation of pricing decisions for

* *Corresponding author address*: John Bower, London Business School, Sussex Place, Regent's Park, London NW1 4SA, United Kingdom; e-mail: jbower@london.edu.

individual power plants, it has been possible to investigate the impact of plant divestment as an alternative means of reducing generator market power.

The ABS model is a discrete event simulation platform that replicates the daily bidding activity, market clearing, and settlement processes in the Pool and NETA. Generating firms (agents) are represented by data arrays containing plant capacity, operating costs, plant availability, and bidding algorithms held on Excel spreadsheets that are manipulated with Visual Basic commands. As Excel is primarily designed to handle large arrays of data and formulas it is possible to simulate four years worth of trading activity, representing over 2 million separate bids, in approximately 30 minutes. Excel appears to offer a higher level of performance, at least for this specific purpose, than could be achieved using alternative agent-based tools such as Swarm or commercial simulation packages such as Mathematica or Powersim. Scaling or amending the ABS model, for example to increase the number of agents or represent an entirely different national different market, only requires the use of ‘cut and paste’ commands rather than reprogramming. Built in graphical interfaces, analytical and charting tools, data export facilities, and easy portability offer additional advantages.

The simulated results show that, far from increasing competition, the NETA proposal would actually magnify market power by allowing generating firms to segment the market on a half hour-by-half hour basis. In contrast, when recent forced divestments, and other changes to industry structure are included in the model there is a significant reduction in the ability of firms to coordinate their actions, resulting in a significant fall in simulated Pool prices. Our conclusion is that, adopting the NETA reforms proposed by OFGEM could therefore have a detrimental effect and that the industry structure reforms already carried out should be sufficient to curtail the market power of generators.

REFERENCES

Bower, J., Bunn, D. W. (2001) “Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales Electricity market,” *Journal of Economic Dynamics and Control*, Vol. 24, No. 3/4, pp. 561-592.

Bower J., Wattendrup, C., Bunn, D. W. (2000) “Implications of German Electricity industry consolidation on competition: An agent-based analysis,” *London Business School Working Paper* (forthcoming in *Energy Policy*, 2001)

Bower, J., Bunn, D. W. (2000) “Model-Based Comparisons of Pool and Bilateral Markets for Electricity,” *The Energy Journal*, Vol. 21, No. 3, pp. 1-29.

ACKNOWLEDGMENT

This research was partly funded by ESRC Postgraduate Training Award No. R00429724400 and a London Business School Financial Award.

An ABS investigation of generator market power in the England & Wales Electricity market using an Excel/VBA platform

Agent2000, Chicago

5-7 October, 2000

John Bower & Derek Bunn

Overview

- Introduction and summary
- An agent-based model of strategic rivalry
- Results
- Conclusion

Introduction and summary

UK electricity prices went up after deregulation as generators exercised market power….

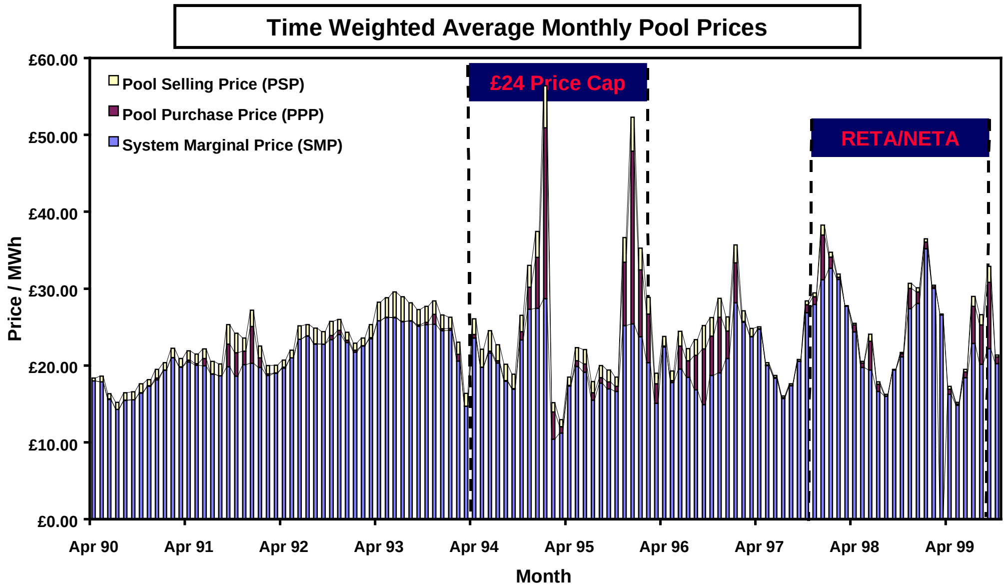

Introduction and summary

…. and the 1998 RETA (Review of Electricity Trading Arrangements) produced NETA....

Proposed Solutions to Market Power

Market reform: Uniform Price (Pay SMP) Pool day-ahead auction replaced with Discriminatory (Pay Bid) auction in Bilateral Model

Plant Divestment: National Power and PowerGen each to divest 4,000 MW (~25%) of their plant

Gas moratorium: Ban on new gas fired generating plant until Pool reform and divestment has taken effect

Introduction and summary

…. (New Electricity Trading Arrangements) to replace the Pool with the Bilateral Model….

Reform Objectives

Reform Proposals

1. **Pay-as-Bid (WYBIWYG)**
2. **Firm Bidding**
3. **Simple Bidding**
4. **Transparent Bidding**
5. **Demand Side Bidding**
6. **Security of Supply**

1. **Forward / futures / options**
2. **Short-term (24 - 4 hr)**
3. **Balancing (4hr - despatch)**
4. **Settlement Process**

Introduction and summary

…. at a cost of over £500 million but with no hard evidence that it would actually work!

Research Questions

What is the impact of the following factors on wholesale price:

Market mechanism (i.e. wholesale trading arrangements)?

Industry structure (i.e. size and number of generating firms)?

Plant technology (i.e. type and distribution of plant)?

An agent-based model of strategic rivalry

Urgent need to test but few insights from empirical observation, economics, game theory

Operations Research Response

An agent-based simulation approach offered a potential solution:

Discrete event simulation: agents replicate repeated nature of daily trading

Artificial Intelligence: agents build own strategies with reinforcement learning

Behavioural modelling: agent strategies emerge not imposed by modeller

Level of aggregation: micro 'bottom up' not aggregate 'top down'

Experimental method: an economic laboratory with perfect controls

An agent-based model of strategic rivalry

We use a "bottom up" agent-based method to analysing market power in generation ….

Firms and plants operating in England & Wales Pool

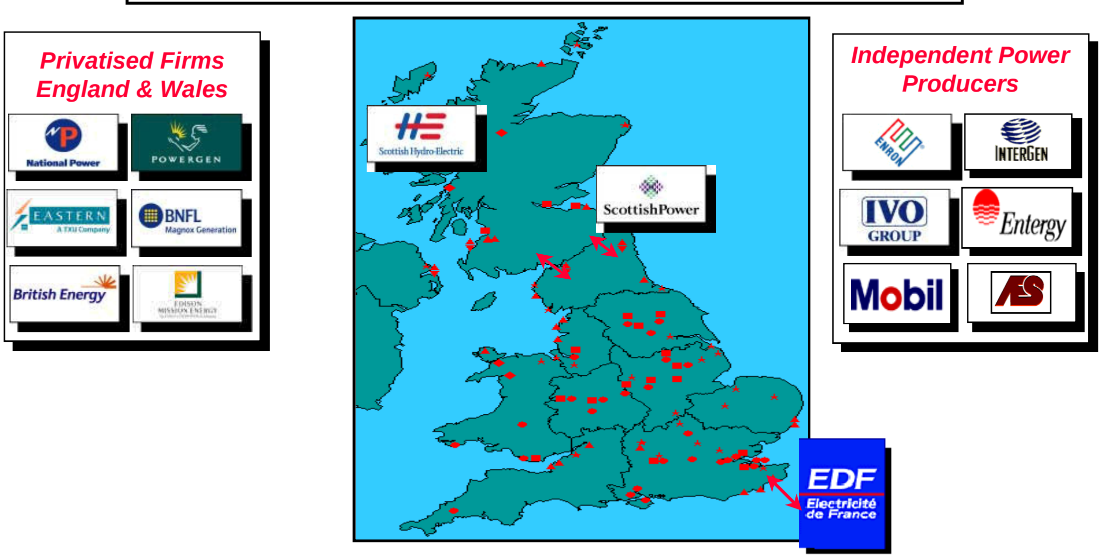

An agent-based model of strategic rivalry

…. which consists of four components and allows us to …..

Model components

- **Economic environment:** A series of interchangeable auction market types through which electricity is traded
- **Agents (Supply):** Each of the generating firms operating in the Pool is individually represented at the level of its plants
- **Agents (Demand) :** Consumers are represented as an aggregate demand curve estimated from empirical data
- **P&L Archive:** Calculator and data base where results of daily trading evaluated and stored

An agent-based model of strategic rivalry

…. capture crucial features of the micro-structure of the supply side of the market

Model Inputs

- **Plant capacity/ownership** (National Grid Company SYS reports)
- **Plant availability** (genset outage rates from private industry sources)
- **Plant efficiency** (heat rates from industry / environmental reports)
- **Fuel costs** (coal, gas, oil prices from Reuters data feeds)
- **Plant operating constraints** (OFGEM / industry reports)

An agent-based model of strategic rivalry

…. but demand side agents are aggregated as assumption is they have no market power

Seasonal Aggregate Demand (Working Days)

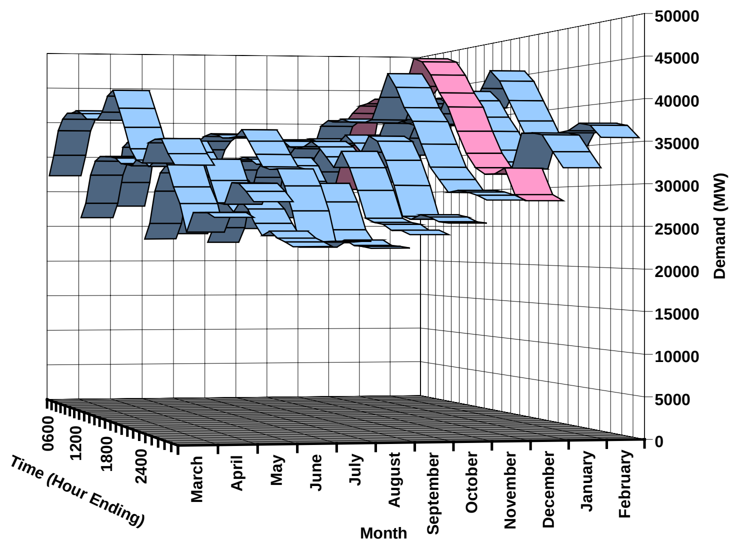

An agent-based model of strategic rivalry

Generating agents submit a bid for each hour of the day, for each plant, for four years ….

ABS Model Daily Bid Cycle

An agent-based model of strategic rivalry

…. so built custom VBA/Excel tool to run this large scale simulation in ● 30 mins on PC

Simulation Bidding Statistics

90 PLANTS
x
4 YEARS
x
248 DAYS
x
24 HOURS

2,140,000 BIDS

An agent-based model of strategic rivalry

Using naive reinforcement learning each agent develops it bidding strategies through time

Agent behavioural rules

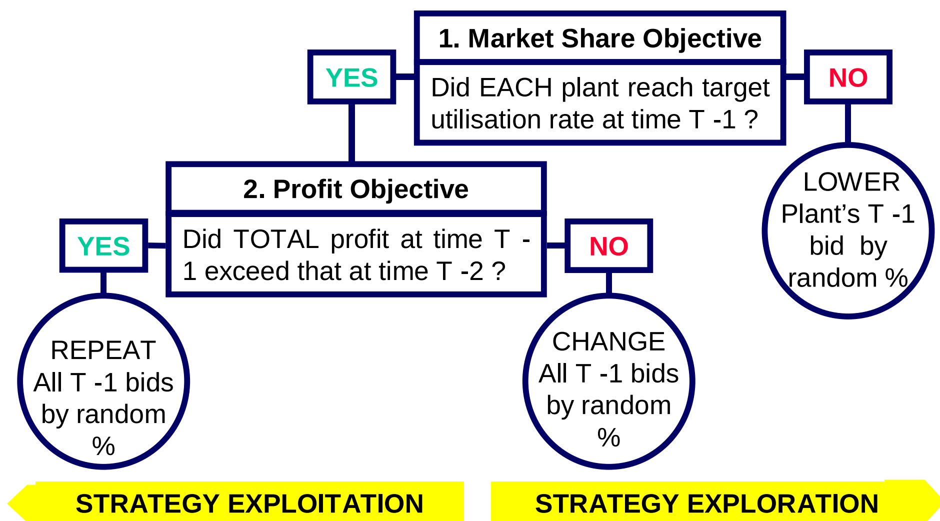

Results

Calibrated model by adjusting Target Utilisation rate and replicate one day of trading….

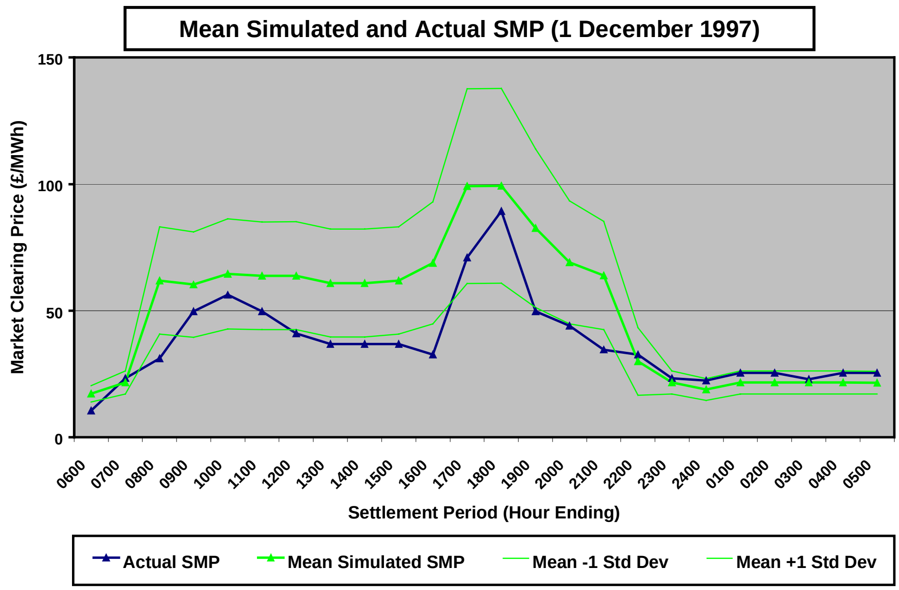

Results

… and then compared prices from simulated trading in the Pool and Bilateral Model

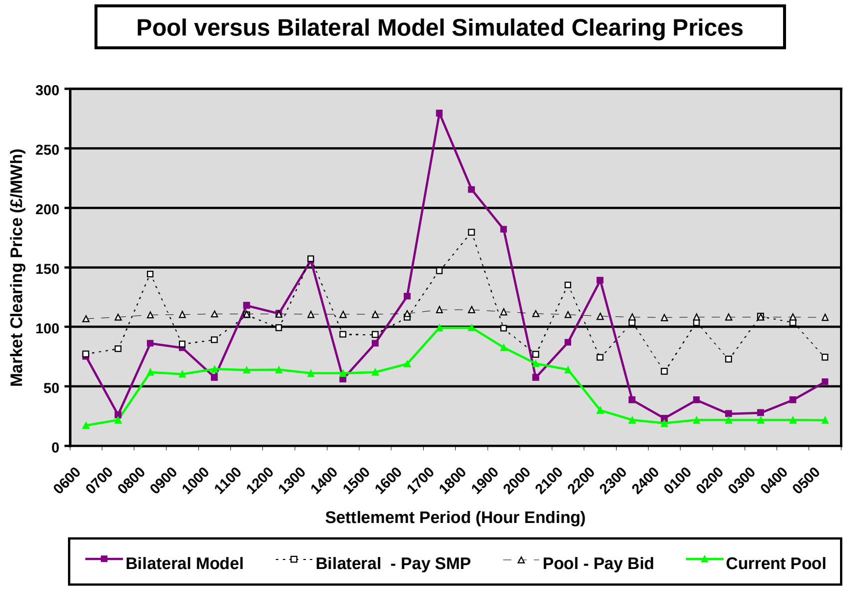

Results

Agent-based simulation has produced results contrary to the regulators expectations

Summary Results

Bilateral Model would make market less competitive than in current Pool

1. Large generators gain an information advantage over small generators
2. Large generators can raise operational risk for small new entrants
3. Large generators segment market more easily and raise price in peak hours

Ofgem believes that the market abuse licence condition is necessary …. and that there is scope for future abuse of wholesale electricity market under both the present electricity trading arrangements and under NETA.

Source: Introduction of the market abuse condition into the licences of certain generators. Ofgem's second submission to the Competition Commission, OFGEM June 2000, http://www. ofgem.gov.uk/public/pub.htm

Conclusion

This is the only model-based analysis of NETA and it is widely used by industry/regulators

Model Application

- Results cited in industry's representations to UK regulator (OFGEM)
- OFGEM forced to respond with own real-time trading experiment
- Results put before the UK DTI Select Committee on Energy
- Requested to present evidence to UK Competition Commission
- Model used to test merger proposals in German electricity market
- Future application to European cross-border electricity market

Conclusion

By combining a number of OR methodologies we addressed a critical economic issue

Impact and Contribution

- **Government Policy:** Regulator now agrees NETA will not control market power and Govt. is putting in place a tighter regulatory framework
- **Auction Theory:** Showed multi-unit, multi-period auction mechanisms are not *Revenue Equivalent* where bidders exercise market power
- **ABS Methodology:** One of the first applications of a large scale ABS model to solve a real problem with a substantial economic impact
- **OR in Practice:** Produced insights into strategic behaviour that have eluded conventional economic and game theoretic analysis

Speaker

John Bower is a Doctoral student at London Business School and is a member of the Energy Markets Group within the Decision Technology Centre. His research interest is in the study of market power and evolution of cross-border commodity trading in electricity. In particular, he is using agent-based simulation approaches to study the emergence of complex strategic behaviour between firms in these new markets.

John's previous career was in the commodity industry and his experience ranges from energy trading, at Marc Rich & Co, to risk management consultancy, with Coopers & Lybrand, advising commodity traders, producers and processors in base metal, precious metal, 'softs' and energy markets. Before joining the PhD programme he was Global Controller Metals/Commodities at Deutsche Morgan Grenfell.

John also has an MBA from London Business School and an MA in Biochemistry from Oxford University.

London Business School
Sussex Place, Regent's Park
London NW1 4SA
United Kingdom

Telephone: +44 (0)20 7834 6736
Facsimile: +44 (0)20 7834 6736
e-mail: jbower@london.edu

AGENT-BASED MODELING OF COMPLEX INFRASTRUCTURES

M. NORTH, Argonne National Laboratory*

ABSTRACT

Complex Adaptive Systems (CAS) can be applied to investigate complex infrastructures and infrastructure interdependencies. The CAS model agents within the Spot Market Agent Research Tool (SMART) and Flexible Agent Simulation Toolkit (FAST) allow investigation of the electric power infrastructure, the natural gas infrastructure and their interdependencies.

KEYWORDS

Complex adaptive systems (CAS), Agent-based modeling (ABM), Electric power system modeling, Natural gas system modeling, Infrastructure interdependency, Swarm, RePast, CAVE, Virtual reality (VR)

INTRODUCTION

Many insights can be gained by viewing energy analysis from a Complex Adaptive Systems (CAS) agent-based modeling perspective. Argonne has taken such a perspective to produce integrated models of the electric power and natural gas markets. The agents within the present Spot Market Agent Research Tool (SMART) and the future Flexible Agent Simulation Toolkit (FAST) allow investigation of the electric power infrastructure, the natural gas infrastructure and their interdependency.

THE PRESENT AND FUTURE

Several tools presently exist:

- SMART Version 2.0 (SMART II) is a Swarm model with an integrated set of agents and interconnections representing the electric power marketing and transmission infrastructure.

- SMART II VR is a virtual reality (VR) interface for SMART II.

* *Corresponding author address*: Michael J. North, Decision and Information Sciences Division, Complex Adaptive Systems Section, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439; e-mail: north@anl.gov.

- SMART II+ is an extension to SMART II that includes an integrated set of agents and interconnections representing the electric power infrastructure, the natural gas infrastructure and connections between them in the form of natural gas-fired electric generators.

FAST is currently under construction. FAST is a complete redesign of SMART II+ that includes improvements in the modeling environment, model detail and representational fidelity.

SMART II

SMART II is a Swarm-based [1] model that uses a set of agents and interconnections to represent electric power systems. SMART II is the Swarm Development Group 2000 Conference (SwarmFest 2000) Best Presentation winner. SMART II itself builds on several other models [2-3]. The SMART II interface is shown in Figure 1. SMART II includes three different kinds of components as follows:

- Generation agents produce electric power.
- Consumer agents use electric power.
- Interconnections represent the transmission grid.

SMART II considers important economic issues such as production costs, investment capital, demand growth for successful consumers, new generation capacity for profitable producers, and bankruptcy for noncompetitive organizations.

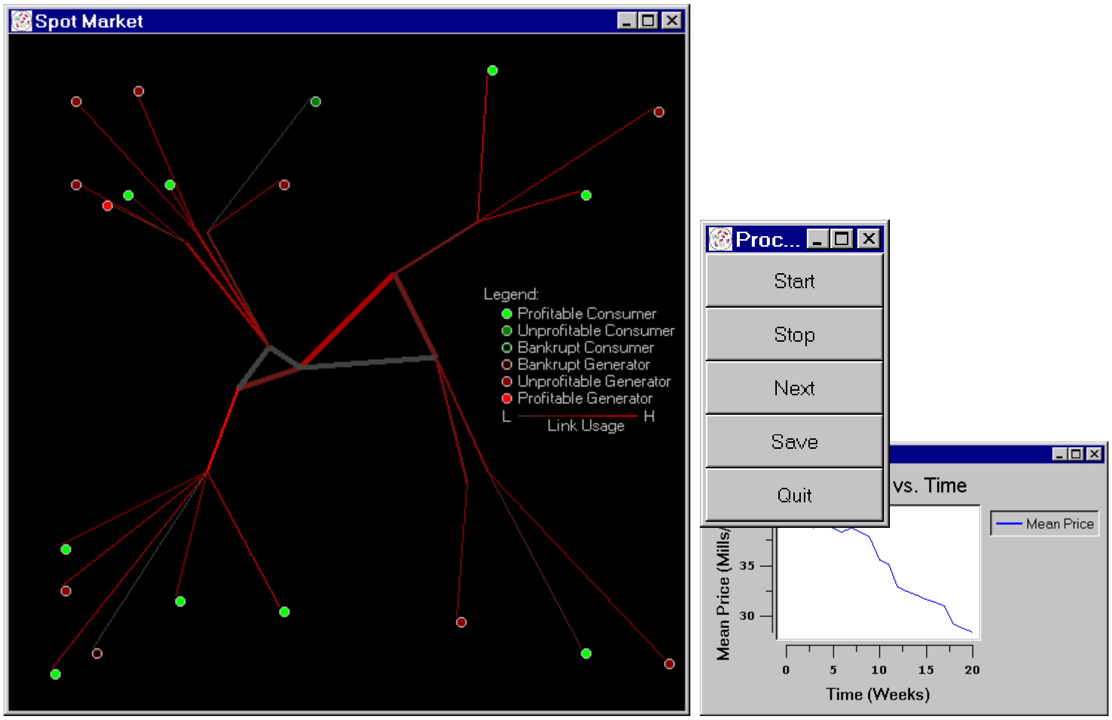

FIGURE 1 The SMART II Interface

SMART II has undergone initial qualitative validation by matching its outputs to the following basic analytic predictions:

- Markets with a single superior producer among a large number of higher cost competitors have been tested.
- Markets with many identical participants have been tested.

Much more work is clearly required to quantitatively validate and calibrate SMART II therefore only limited qualitative insights are currently being derived.

As originally presented at SwarmFest 2000, qualitative insights from SMART II indicate that certain transmission line configurations may encourage price spikes. Soon after SwarmFest 2000, this insight was borne out.

As was specifically noted at SwarmFest 2000, the California electrical grid has a configuration of a type that may cause price spikes. Substantial price spikes of the kind predicted by SMARTII where recently observed in this market.

Further qualitative insights suggest that greater electrical market price stability may be gained by consciously avoiding specific configurations that encourage instabilities. In other words, qualitative insights from SMART II can help us make things better by not making things worse.

SMART II VR

SMART II VR is a prototype agent visualization tool. SMART II VR is intended to explore the use of advanced interactive three-dimensional visualization in agent-based modeling.

SMART II is a CAVE Automatic Virtual Environment (CAVE)-based virtual reality interface for SMART II. The CAVE is a virtual reality library co-developed by the University of Illinois at Chicago and Argonne. From the CAVE User's Guide [4]:

> The CAVE is a projection-based VR system that surrounds the viewer with four screens. The screens are arranged in a cube made up of three rear-projection screens for walls and a down-projection screen for the floor; that is, a projector overhead points to a mirror, which reflects the images onto the floor. A viewer wears stereo shutter glasses and a six-degrees-of-freedom head-tracking device. As the viewer moves inside the CAVE, the correct stereoscopic perspective projections are calculated for each wall. A second sensor and buttons in a wand held by the viewer provide interaction with the virtual environment.

SMART II VR includes an interactive multifunction wand and two rendering modes.

Detail rendering mode focuses on rendering quality. Directional lighting is included. Agents are rendered as lighted spheres. A texture-mapped floor with shadows and first order reflections is included. This mode allows SMART II VR to take advantage of computers with high graphics performance. An example is shown in Figure 2.

Speed rendering mode focuses on rendering time. Agents are rendered as flat shaded cubes. This mode allows SMART II VR to be used on low performance personal computers. An example is shown in Figure 3.

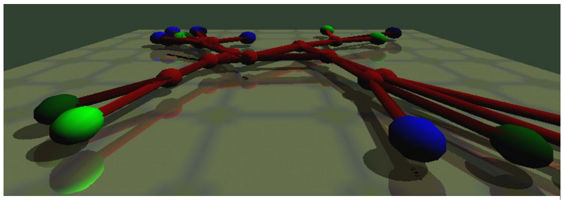

FIGURE 2 SMART II VR Detail Mode

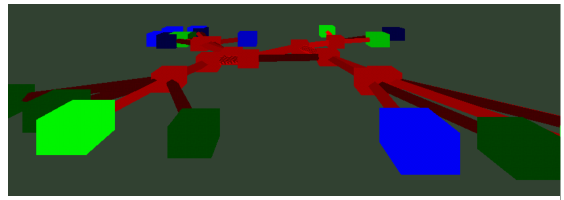

FIGURE 3 SMART II VR Speed Mode

In SMART II VR, generation agents are shown as green spheres or cubes. Spheres are shown in detailed rendering mode and cubes are shown in speed rendering mode. The size of each object represents its total normalized investment capital level. Size can be interactively changed with the CAVE wand. Each object's color intensity represents its hourly profit level.

In SMART II VR, consumer agents are shown as blue spheres or cubes. Spheres are shown in detailed rendering mode and cubes are shown in speed rendering mode. The size of each object represents its total normalized investment capital level. Sizes can be interactively changed with the CAVE wand. Each object's color intensity represents its hourly profit level.

In SMART II VR, interconnections are displayed as red tubes. The size of each tube represents its normalized transmission capacity level. Sizes can be interactively changed with the CAVE wand. Each tube's color intensity represents its hourly utilization level.

SMART II+

SMART II+ is a Swarm-based [1] extension to SMART II. SMART II is the Swarm Development Group 2000 Conference (SwarmFest 2000) Best Presentation winner. SMART II itself builds on several other models [2-3].

SMART II+ includes an integrated set of agents and interconnections representing each of the following:

- The electric power marketing and transmission infrastructure.
- The natural gas marketing and distribution infrastructure.
- The interconnections between the two infrastructures in the form of natural gas fired electric generators.

Both of the infrastructures modeled in SMART II+ include many features:

- Two different kinds of agents, producers and consumers, represent the market participants.

- Interconnections represent transmission or distribution systems with capacities on each line or pipe and complex routing.

- Important economic issues are considered such as investment capital, demand growth for successful consumers, new generation capacity for profitable producers, and bankruptcy for noncompetitive organizations.

- Components can be disabled in real time to simulate failures.

The electric power infrastructure includes the added feature of natural gas fired electric generators. These generators buy fuel from the natural gas market. The resulting electricity is then sold in the electric power market.

SMART II+ PRODUCER AGENTS

SMART II+ producers determine their production level based on the potential profit to be made. Each producer has investment capital that is increased by profits and reduced by losses. If a producer reaches a predetermined level of investment capital it can purchase additional production capacity in the form of new electric generators or new natural gas sources. New producers are similar to their owner and can connect to the distribution network in either the same location or a new one. Producers that run out of investment capital go bankrupt and no longer participate in the market. Producers choose whether or not to sell energy based on either their cost curves or natural gas prices. Standard producers derive their costs and capacities from cost curves with maximum generation limits as shown in Figure 4. Both costs and capacities are exogenous.

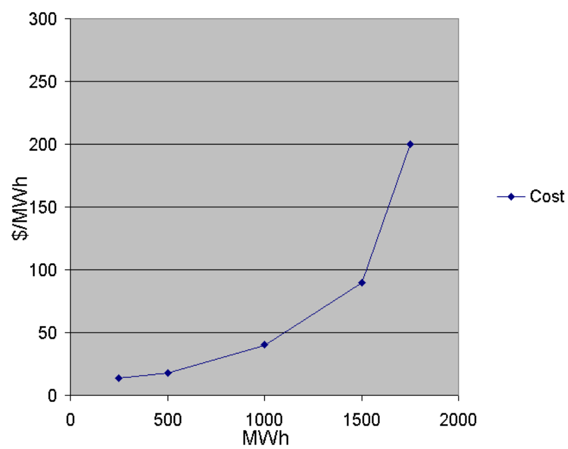

FIGURE 4 An example producer cost curve

Natural gas fired electric generators derive their costs and capacities from the endogenous natural gas market. These generators are consumers in the natural gas marketplace. Their costs are based on the price they pay for natural gas. Their capacities are based on both the amount of natural gas they can purchase and their design limits.

Producer simulation display appearance depends on current profit levels (Figure 5).

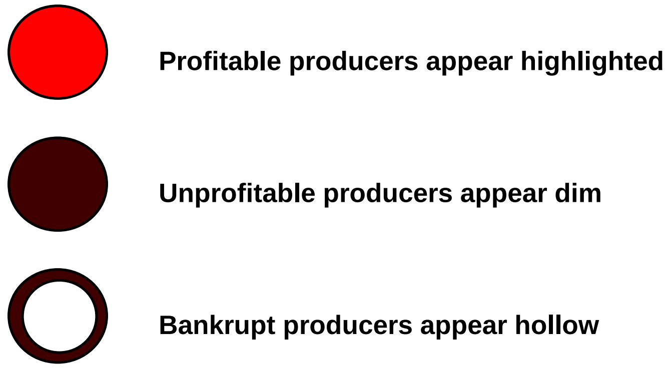

FIGURE 5 Producer Appearance

SMART II+ CONSUMER AGENTS

SMART II+ consumer agents buy energy for their own use. Businesses buy fixed amounts of energy to remain in business. Populations buy fixed amounts of energy to live their lives. Natural gas fired electric generators buy natural gas to produce salable electric power.

Each consumer has investment capital that is increased by profits and reduced by losses. If a consumer reaches a predetermined level of investment capital it can grow in the form of new consumers. Consumers that run out of investment capital go bankrupt and no longer participate in the market.

Consumer simulation display appearance depends on current profit levels (Figure 6).

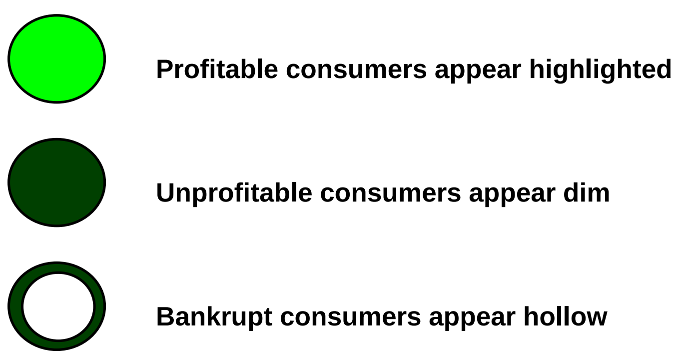

FIGURE 6 Consumer Appearance

Investment capital represents several things. For industrial users it is their total financial capital. For individuals it is the employment and personal opportunities that keep them in an area or encourage them to leave.

SMART II+ INTERCONNECTIONS

Interconnections represent transmission lines or distribution pipes each with an individual capacity limit. Individual capacity limits vary by interconnection type. Central transmission lines or main distribution pipes have high capacity limits and are drawn with thick marks. Outlying transmission lines or secondary distribution pipes have moderate capacity limits and are drawn with medium marks. Feeder lines or pipes have low capacity limits and are drawn with thin marks. Interconnection color represents contents and usage as shown in Figure 7.

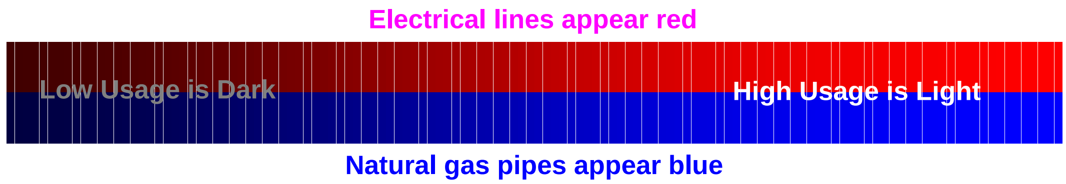

FIGURE 7 Interconnection Appearance

SMART II+ MARKET INDICATORS

The key SMART II+ market indicators are market prices, unserved energy and natural gas fired electrical generator market share. All key SMART II+ indicators are represented by graphs updated in real time.

Market price is the per unit purchase price of the given energy resource. Electric power prices are given in tenths of a cent per kilowatt-hour (Mills/kWh). Natural gas prices are given in dollars per thousand cubic feet ($1,000 cubic feet). The SMART II+ price graphs are shown in Figure 8.

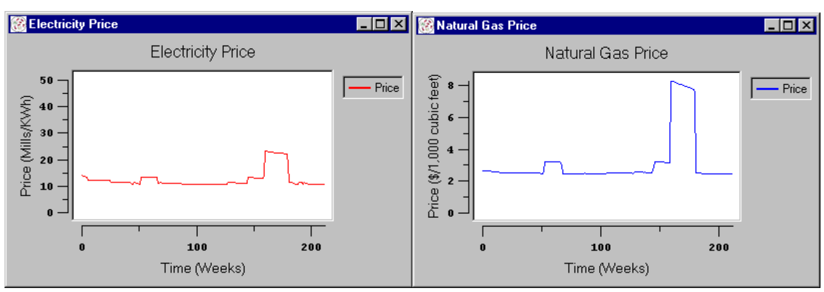

FIGURE 8 Price Graphs

Unserved energy (UE) is the energy demand that was not met by the market. UE represents a form of market failure. UE is given as a percentage of total energy demand. The SMART II+ UE graph is shown in Figure 9.

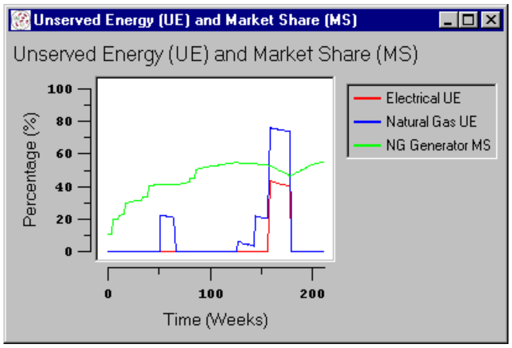

FIGURE 9 UE and NG Generator MS Graph

Natural gas fired electric generator market share (NG Generator MS) is a measure of the electric generation capacity that is supplied by natural gas units. NG Generator MS is key to infrastructure interdependency. NG Generator MS is given as a percentage of total capacity. The SMART II+ NG Generator MS graph is also shown in Figure 9.

SMART II+ NETWORK DISPLAY

The geographical SMART II+ display is based on an equivalenced network. An example notional SMART II+ network is shown in Figure 10.

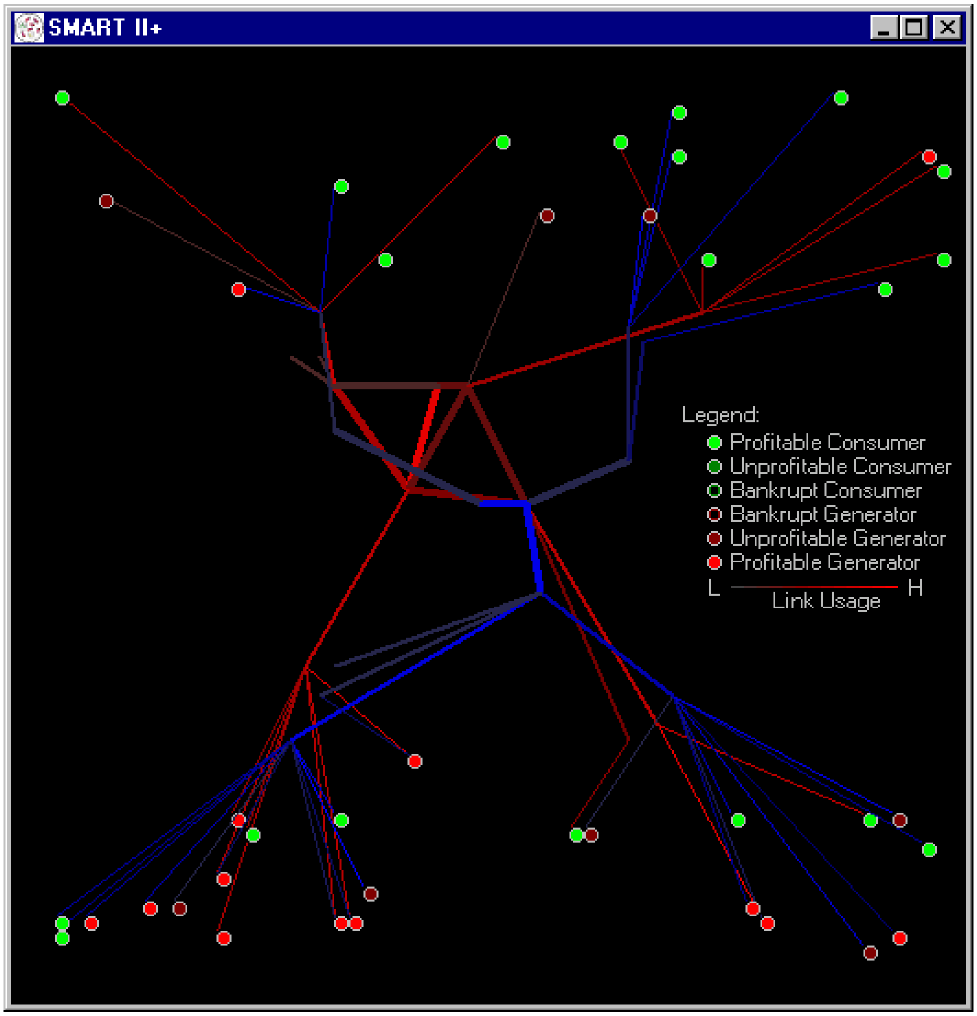

FIGURE 10 Example Notional SMART II+ Network

SMART II+ VALIDATION AND CALIBRATION

As with SMART II, SMART II+ has undergone initial qualitative validation by matching its outputs to basic analytic predictions:

- Markets with a single superior producer among a large number of higher cost competitors have been tested.

- Markets with many identical participants have been tested.

SMART II+ has undergone initial qualitative calibration by comparing the model's natural gas-fired electric generator market share trends to those found in real systems.

Much more work is clearly required to quantitatively validate and quantitatively calibrate SMART II+. Therefore, only limited qualitative insights are currently being derived.

SMART II+ INSIGHTS

As originally presented to our research sponsors in May 2000, preliminary insights from SMART II+ indicate that:

- Rising natural gas-fired electrical generator market share radically increases market interdependence.

- Increasing market interdependence can pit the electric power and natural gas markets against one another during simultaneous failures since both markets are fighting for the same underlying resource, natural gas.

This interdependency insight was borne out in the aftermath of the recent El Paso natural gas pipeline explosion.

What is the state of the world today? Nationwide natural gas-fired electrical generator market share is roughly 15% to 20%. Nationwide natural gas-fired electrical generator market share is expected to radically increase over the next five years. J.P. Morgan analysts predict that there is expected to be a 31% increase in generation capacity [5]. These analysts predict that roughly 95% of new electrical generation capacity will come from natural gas-fired units [5]. An example is the midwestern region dominated by Commonwealth Edison (ComEd).

ComEd presently gets less than 10% of its current 20,000 MW generation capacity from natural gas-fired generators. Permits are being issued for the construction of 8,000 MW of new capacity, over 95% of which will be natural gas-fired.

The interdependency between the electric power and natural gas markets implies that when natural gas-fired electrical generator market share becomes high enough a single energy resource, “virtual natural gas,” is being traded in both markets. Viewing energy systems from the perspective of virtual natural gas suggests that future electrical system capacity expansion planning should explicitly feature the natural gas distribution infrastructure as a key component. Power Systems Engineers should note that electrical models might be substantially incomplete without explicitly including the natural gas infrastructure. Highly distributed electrical generation plans including local load servicing schemes may especially benefit from this view since they rely heavily on the existence of other energy sources such as natural gas.

FAST

FAST is an integrated infrastructure model based on SMART II+. FAST includes many of the features of SMART II+ along with improvements in modeling infrastructure, detail and fidelity. FAST is currently under construction. FAST has three components:

- FAST:Run is the runtime infrastructure that will be merged with RePast [6].

- FAST:E is the electric power system model.

- FAST:G is the natural gas system model.

FAST:Run is designed to be a lightweight large-scale system with the following major features:

- FAST:Run is written entirely in Java.
- FAST:Run is fully distributed.
- FAST:Run has a multithreaded scheduler that focuses on maximizing parallel execution.

The underlying design paradigm of FAST is that of a time continuum ranging from decades to seconds:

- On the scale of decades the focus is long term human decisions constrained by economics.
- On the scale of years the focus is short-term human economic decisions constrained by economics.
- On the scale of months, days and hours the focus is short-term human economic decisions constrained by economics and physical laws.
- On the scale of minutes or less the focus is on physical laws that govern energy distribution systems.

Modeling over the full range of time scales is necessary to understand the complex infrastructure interdependency between the electric power and natural gas markets.

FAST includes a large number of different agents to model the full range of time scales. The focus of agent rules in FAST varies to match the time continuum. Over longer time scales human economic decisions are emphasized. Over shorter time scales physical laws dominate.

Many FAST agents are relatively “thick” compared to typical agents. FAST agents are highly specialized to perform diverse tasks ranging from acting as Independent System Operators to being transmission lines. To support specialization, FAST agents include large numbers of highly specific rules.

The FAST system and its component agents will be subjected to rigorous quantitative validation and calibration.

CONCLUSION

Developing the initial capability to create CAS models requires substantial organizational investment. Once this initial investment has been made tools can be created that allow many insights.

ACKNOWLEDGEMENT

This work is sponsored by the U.S. Department of Energy under contract number W-31-109-ENG-38.

REFERENCES

1. Burkhart, R., M. Askenazi, and N. Minar. *Swarm Release Documentation*. Available as http://www.santafe.edu/projects/swarm//swarmdocs/set/set.html.

2. VanKuiken, J.C., T.D. Veselka, K.A. Guziel, D.W. Blodgett, S. Hamilton, J.A. Kavicky, V.S. Koritarov, M.J. North, A.A. Novickas, K.R. Paprockas, E.C. Portante, and D.L. Willing, *APEX User's Guide (Argonne Production, Expansion, and Exchange Model for Electrical Systems) Version 3.0*, ANL/DIS/TM-21, Argonne National Laboratory, Argonne, Ill., 485 pages (November 1994).

3. Veselka, T.D., E.C. Portante, V.S. Koritarov, S. Hamilton, J.C. VanKuiken, K.R. Paprockas, M.J. North, J.A. Kavicky, K.A. Guziel, L.A. Poch, S. Folga, M.M. Tompkins, and A.A. Novickas, *Impacts of Western Area Power Administration's Power Marketing Alternatives on Electric Utility Systems*, ANL/DIS/TM-10, Argonne National Laboratory, Argonne, Ill., 216 pages (January 1994).

4. Pape, D., C. Cruz-Neira, M. Czernuszenko *CAVE User's Guide*. Available as http://www.evl.uic.edu/pape/CAVE/prog/CAVEGuide.html (1997).

5. Anderson, J.A., "What's Electrifying Natural Gas," *Business Week*, (September 19, 2000).

6. Collier, N. and D. Sallach. *RePast*. Available as http://repast.sourceforge.net/.

DISCUSSION: ENERGY AND INFRASTRUCTURES

R. CIRILLO, Argonne National Laboratory, Moderator

[Presentation by Bower]

Tony Andrews: If they bid zero, are you saying that they are not bidding?

John Bower: No, when they bid zero they're not charging for their product, they're giving it away for nothing. The trading rule is they get paid the marginal plant bid. They want to make sure they run, though, which is why they're bidding zero.

Gale Boyd: Gale Boyd from Argonne National Laboratory. Have you done any simulations looking at changes in market structure? If so, do bilateral market prices still come out higher? In other words, does the more competitive supply side affect the bilateral solution?

Bower: Yes. At the limit, if you have a very competitive market structure, say with 50 firms, it really doesn't matter what kind of market you have; you're going to end up with marginal pricing. If you have a monopoly, it also doesn't matter what kind of trading arrangements you have, you're going to get the monopoly price. In between, when you have a duopoly or an oligopoly of firms, the types of trading arrangements *do* make a difference. As for changing the market structure, we have simulated this in some detail using our model, and the way we have changed the structure has been sufficient, we think, to curtail the market price. And in fact we've seen the market price begin to fall. So we hope it's worked.

Boyd: I have another comment. I found a great deal of similarity between your presentation and yesterday's paper by Jing Yang ["Price Efficiency and Risk Sharing in Two Inter-Dealer Markets"]. Her paper talked about dealer-to-dealer as a bilateral trade, the broker as essentially a pool, and some of the circumstances where there's a mixture of these elements. But you dealt with these elements as either/or. Have you considered a mixed market, where people can make both bilateral trades and trades in the pool?

Bower: We have thought about modeling some of the questions that Jing Yang talked about, because at one point there was a proposal to have both an on-the-telephone kind of bilateral market and an on-screen market, and we wondered whether that would make any difference. In fact, the regulators canned the idea because they said they didn't want to prescribe how people should trade. So they're not going to have a screen-based market. But what we found was really crucial here is the information available to different agents. Big agents, big firms, gain an informational advantage in the bilateral market over small firms, and that allows them to drive the price up. So information — what's on the screen, what isn't on the screen — it's absolutely crucial.

Prakash Thimmapuram: I'm Prakash Thimmapuram from Argonne. You said they can withdraw capacity by the generator. At what time can they withdraw the capacity?

Bower: They can make their bid in the pool, and they can withdraw their capacity one second before it's dispatched.

Thimmapuram: Is there any penalty for withdrawing the capacity?

Bower: No.

Thimmapuram: Another question: You show that the payer's bid is higher than the marginal price. Can you explain why this is so?

Bower: Under marginal price, the large, low-cost generators always bid zero — so they're always undercutting the coal plants, which can't bid zero. That's in the pool. But in a bilateral market, everyone has to bid around the marginal price. Occasionally, a low-cost operator will overbid a coal plant, and that means the competition against the coal plants is reduced. Another reason is that bidding is now by the hour rather than the whole day, and that means the large generators can segment the market. They can bid much higher prices for peak hours than they bid for the base-load periods. In fact, base-load prices are virtually the same in the bilateral market and the pool. As you know, segmenting the market is a great way of discriminating against certain customers. There's not much elasticity in this market.

Randal Picker: What is the source of the information advantage for the big traders in the bilateral market? Is it that they observe more trades?

Bower: Yes. If you've got 10 plants you're putting in a lot more bids than the one-plant operator. Also, if you're a small, marginal-cost operator, you're really quite risk averse because you've got large amounts of debt. You want to run your plant all the time. Often we find the small generators are having to shave their bids dramatically. So really the small generators are nowhere near the marginal price some of the time, and they haven't many plants, so they just don't have the same capacity to learn.

Richard Cirillo: The issue was raised yesterday about whether agent-based simulation techniques, at least in some domains, were ready to be used to help shape and formulate policy. What's your experience with the agent-based simulation technique in this domain — is it ready for prime time?

Bower: In fact, this model is being used by the government in the U.K. already, by the Competition Commission, and it's being widely cited by the industry. The regulators dislike the result so much that they set up their own experiment with real people. They had 50 students trade the new market. One of the students found a way of getting the price up, but they only allowed the students to trade for 12 days, so, in my opinion, the students had no chance to learn. Their results showed no difference between the pool trading and the bilateral trading. So they just carried on and implemented the marketplace.

[Presentation by North]

Picker: What happens to the generating capacity if the agents are destroyed?

Michael North: The generating capacity is reused. It's available for some other potential entrepreneur, or even an existing large company, to take it over. Given that this is a very competitive market, it's usually acquired very quickly.

Picker: As you know, FERC [Federal Energy Regulatory Commission] has proposed moving from ISOs [independent system operators] to regional transmission grid operators. How easy is it for you to capture that kind of institutional change in your agents?

North: It depends on the amount of change you want to make. It's actually relatively easy to replace agents in the model. Building the model of the agent that you're working with, though, can be a very difficult task. It's not really an infrastructure issue; it's a matter of thoroughly identifying and characterizing a different agent. We hope it's going to be easier with the new infrastructure because one of the things we're focusing on is changing market roles. It's just one of our interests.

Blake LeBaron: Blake LeBaron, Brandeis University. You've got a couple of complicated infrastructure investment problems underneath, both in transmission capacity and generating capacity. How do your agents do forecasting? This is a difficult problem, especially the transmission one. I'm curious about how the grid forms based on forecasts.

North: Infrastructure investment *is* a difficult problem, and I don't think there is any one ideal solution. Our agents use a relatively simple learning process, very similar to the other learning that you've seen presented today. They look at the past, they keep a record of what they did, and they find out what the market did. They don't know about one another's bids or involvement. They know about costs; they know how much it costs to create one thing or change another. They use a relatively simple multivariable linear model, essentially, to calculate what they're going to do next.

We think that it's a weakness in the model, that right now we have relatively thin agents, and we're working on changing that. We've talked about embedding neural network systems in the new "thick" agents, for instance. We've also been considering genetic algorithms, so we'll probably use a combination of those.

Andrews: I'm Tony Andrews, with the Navy. Many generators now are looking at futures and options to minimize their risk. How might you be able to incorporate that? I think the real test is going to be whether it will drive the price down, as they anticipate.

North: It's an open question. People like to think that those things drive prices down, but it's not really guaranteed. In the SMART models, we don't look at those issues, because we're dealing with the spot market, which is essentially short term. For FAST, though, we've created data structures for looking at different packages or bids that people can make, not only for the next hour, but also for the next year. So that's the next step.

Bower: Where does the market price come from in your model, and how do the agents influence it?

North: We have basically an ISO who's taking a series of bids into the marketplace. The agents look at what they've bid in the past, what the market price was, and whether their bids were accepted, and they use that information to adjust their current bidding. Then the ISO basically does the merit ordering to allocate people.

We deal with failures that are user-imposed in our model. When I select something from the screen and make it unavailable, for example, that could be a unit that was allocated but didn't actually run. But beyond that, we're not having explicit failures outside of user input.

[Presentation by Richter]

Cirillo: When you run the simulation and the agents are adapting to the change in the market price, have you experimented with the rate of adaption or the amount of adaption that the agents are allowed to utilize, and does that have a significant effect on the results?

Charles Richter: I've experimented with that some, yes. And it does have a significant effect.

Hong Lei: I'm Hong Lei from Warwick Business School. This is a question I've been forming over these two days about the overall topic of agent-based simulation. I believe it is, of course, useful to learn what these models do; however, it seems they are presented as closed systems. I would like to learn more about this useful methodology by knowing how the models are organized.

Cirillo: I believe the question as stated is what can we do to reveal more of the structure and inner workings of the models so that others can learn how they were put together. Anyone want to comment on that?

North: Michael North from Argonne National Laboratory. I think open source is a good step in that direction. We can't open up a source code for everything all the time — there are proprietary and competitive concerns — but when we can, I think it's the most helpful teaching tool of all. And obviously, publications where you describe a model's organization have a value, too.

Jonathan Bendor: Jon Bendor, Stanford University. Two quick comments on the issue of making the models comprehensible. I think that's a very important issue. First, we could all try to keep our models simple. Not only will the models be more comprehensible to other people, they'll be more comprehensible to us. And going back to the earlier discussion, when things go wrong, as they always will, it will be easier to fix them. So start simple and work incrementally. Resist the temptation — and here I'm seconding Ian Lustick's point enthusiastically — resist the temptation to throw in everything you know about the real world that pertains to your question.

Secondly — and this is related to the model that Diermeier and Ting and I worked on, and why we think that deduction and computation complement each other so well — what one can do, both to gain insight into your own model as well as to communicate it, is deductive work on small aspects of the model. It's probably utopian to try and solve any of these complex models completely analytically; that's why we're simulating in the first place. But what you can do is to freeze an aspect of the model and then solve little parts of it. You're triangulating via deduction and gaining insight into the properties and characteristics of the model. At the same time, you can generalize the unfrozen aspects of the model and say, "Well, this doesn't just hold for Bush-Mosteller, it holds for a whole class of adaptive rules."

So I think a useful strategy is start simply, work incrementally, and marry deduction to computation.

Lars-Erik Cederman: Another simplification would be to have standardized models to use as starting points. This was one of the goals in the Swarm project. Then you wouldn't have to start from scratch in explaining the inner workings of your components. For instance, game theorists talk about "prisoner's dilemma," and you don't have to redescribe the prisoner's dilemma, because everybody knows exactly what is meant by that. I think we haven’t got very far along these lines because there is very little overlap. Apart from using the same basic package that relieves us from reinventing the wheel when it comes to showing graphs and setting up schedules, we have much more work to do in trying to build up clusters of modeling components or libraries. I hope that when we get together in maybe five years we will have made serious progress along these lines.

Closing Panel

CLOSING PANEL

C. MACAL, Argonne National Laboratory
J. PADGETT, University of Chicago
R. PICKER, University of Chicago

Charles Macal: At this point in the program we will wrap up with closing comments. We also invite the audience to bring up any points for discussion. It's certainly not meant to be anything more than a brief wrap-up, possibly reviewing main points of the workshop or offering general comments, perhaps raising some questions that somehow didn't get answered during the workshop. John, do you have some comments?

John Padgett: It's well known in this area that the comparative advantages of the agent-based simulation approach fall on two dimensions. The first is heterogeneity: the degree to which you care about heterogeneity of agents is the degree to which this approach looks appealing. The second is topological issues: the advantages depend on the degree to which there is a structure to the interaction rather than an averaging sort of behavior. So topology and heterogeneity are the standard arguments I hear for why this approach is useful.

Where do the models that we see here fit on those two standard dimensions? First, on the topology issue, it's interesting what Peyton Young said, that an agent-based model is instantly recognizable, qualitatively. I think it's amazing, the degree to which that statement fundamentally comes down to two-dimensional spatial topologies. You see a two-dimensional graph with some nearest-neighbor topology and entities that move around on the graph and change colors. And that's how you recognize an agent-based model.

Now, regarding a more general topology, everybody in the field understands that it's very useful. It's particularly useful when you're talking about cases of real space, and we've had some of those, such as the ecological and housing segregation models. There are many domains for which topology is terribly important, and these tend to be the domains in which you most often see success in these models — when you have some real spatial framework that makes this two-dimensional fit really work.

Miles Parker: I guess I tend to think of the 2D lattice as just a feature, a way in which the field is developing, and a natural representation that people use. So in that sense it's merely a historical accident. There's nothing that particularly reifies the 2D lattice, and there's nothing that says the 2D lattice is any worse in general than other kinds of graph configurations for these kinds of models. I'm sure that there are arguments that could be made; Rob Axtell has a recent paper on different topologies of interactions, so now a lot of work is beginning to develop in that area. I do think it's a natural place to go, but we've certainly done work in other kinds of structures, as well. I see it as simply another development that makes some of the analysis, and certainly some of the methodology, a bit simpler.

Padgett: I understand that as well. I'm just trying to point out that this historical accident has potentially self-limiting qualities that might blind us to broader application of these tools.

Christopher Langton: Yes, but as one of the speakers said, architecture matters. And because we have to study many of these systems as wholes, it's often the case that there are multiple architectures at play here. In ecology, for instance, there may be something you would like to model as a smooth flow for hydrology, but at the same time there are discrete animals moving around that you might be able to model on a lattice — and those animals might be humans with cell phones talking to each other over a telecommunications network. I think it's critical not to try to force all of the processes into a single architecture.

Padgett: I agree completely that in fact architecture does matter, but I'm trying to point out that for whatever historical reason, we are self-limiting ourselves away from that idea, which is the promise of the whole approach.

But returning to the issues of topology and heterogeneity, it's tempting to say that since both are benefits of agent modeling, then more of both is better. And that runs us right into what I'll call the "Bendor critique," that you get ever more complicated, ever more unmanageable models that way. And so my exhortation to greater richness, while I deeply believe it, is very dangerous, for the reasons that Jon has mentioned.

And I think the solution really comes down to what Lars-Erik [Cederman] said earlier. We have to identify in this huge space of possible agents a few areas that our models can actually come to grips with and say something new about. It's not enough to just populate the model world with all sorts of disparate things, which then go off centrifugally in 50,000 different directions. You have to concentrate on some sort of puzzle, which brings us back to the question: What is the theoretical framework? It's good to be focused in on the powers of these methods, but you do have to ultimately concern yourself with the theoretical framework, such as evolution, within which you are seeking to illustrate phenomena.

Two candidates have been expressed here: the evolution framework and the learning framework. We didn't have as many evolution papers as I would have liked, but certainly Blake [LeBaron's] paper fits very nicely in that framework. The learning framework is also very much lurking in the background. And I say the task at hand is to be a bit more self-conscious, not just about the tools, but about evolution and learning. What are the puzzles? What are the dynamical questions? Without that, we're just doomed to go off in all these different directions and not really congeal as a field.