Agent 2003

Conference on

Challenges in Social Simulation

presented by

Argonne National Laboratory

The University of Chicago

Gleacher Center

The University of Chicago
Chicago, Illinois

October 2-4, 2003

ARGONNE NATIONAL LABORATORY • UNIVERSITY OF CHICAGO •

Crescat scientia vita excolatur

ISBN 0-9679168-4-4

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/

Argonne National Laboratory

Argonne National Laboratory, with facilities in the states of Illinois and Idaho, is owned by the United States Government, and operated by the University of Chicago under the provisions of a contract with the Department of Energy.

Agent 2003

Conference on

Challenges in Social Simulation

presented by

Argonne National Laboratory

The University of Chicago

Gleacher Center

The University of Chicago
Chicago, Illinois

October 2-4, 2003

CONTENTS

Approaches to Validation

FRIDAY, OCTOBER 3, 2003

Welcome

Invited Speaker: Steven Bankes

Computational Organization Theory (Parallel Session I)

SATURDAY, OCTOBER 4, 2003

Invited Speaker: Lars-Erik Cederman

Social and Cultural Dynamics (Morning Parallel Session I)

Network Complexities (Morning Parallel Session II)

Political Processes (Afternoon Parallel Session I)

Ecological Interactions (Afternoon Parallel Session II)

FOREWORD

Welcome to the *Proceedings* of the fourth in a series of agent simulation conferences cosponsored by Argonne National Laboratory and The University of Chicago. Agent 2003 is the second conference in which three Special Interest Groups from the North American Association for Computational Social and Organizational Science (NAACSOS) have been involved in planning the program—Computational Social Theory; Simulation Applications; and Methods, Toolkits and Techniques.

The theme of *Agent 2003*, *Challenges in Social Simulation*, is especially relevant, as there seems to be no shortage of such challenges. Agent simulation has been applied with increasing frequency to social domains for several decades, and its promise is clear and increasingly visible. Like any nascent scientific methodology, however, it faces a number of problems or issues that must be addressed in order to progress. These challenges include:

- Validating models relative to the social settings they are designed to represent;
- Developing agents and interactions simple enough to understand but sufficiently complex to do justice to the social processes of interest;
- Bridging the gap between empirically spare artificial societies and naturally occurring social phenomena;
- Building multi-level models that span processes across domains;
- Promoting a dialog among theoretical, qualitative, and empirical social scientists and area experts, on the one hand, and mathematical and computational modelers and engineers, on the other;
- Using that dialog to facilitate substantive progress in the social sciences; and
- Fulfilling the aspirations of users in business, government, and other application areas, while recognizing and addressing the preceding challenges.

Although this list hardly exhausts the challenges the field faces, it does identify topics addressed throughout the presentations of *Agent 2003*.

Agent 2003 is part of a much larger process in which new methods and techniques are applied to difficult social issues. Among the resources that give us the prospect of success is the innovative and transdisciplinary research community being built.

We believe that *Agent 2003* contributes to further progress in computational modeling of social processes, and we hope that you find these *Proceedings* to be stimulating and rewarding. As the horizons of this transdiscipline continue to emerge and converge, we hope to provide similar forums that will promote development of agent simulation modeling in the years to come.

Charles Macal, Director
Michael North, Deputy Director
David Sallach, Associate Director

Center for Complex Adaptive Agent Systems Simulation (CAS^2)
Decision and Information Sciences Division
Argonne National Laboratory

ACKNOWLEDGMENTS

We acknowledge the support of many people in organizing *Agent 2003*. Kathy Ruffatto handled registration, administration, and logistics. Bob Baker managed the conference Web site. Argonne's Information and Publishing Division prepared the program booklet and the proceedings for publication. Margaret Clemmons served as project manager and editor, with assistance from Michele Nelson (graphic design) and the Document Processing Center (word processing support).

ORGANIZING COMMITTEE

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

Thursday, October 2, 2003

Simulation Methods

AIDS TRANSMISSION IN SUB-SAHARAN AFRICA: ISSUES IN MODELING AND METHODS

P. HEUVELINE,* The University of Chicago
D.L. SALLACH, The University of Chicago
T. HOWE, The University of Chicago

ABSTRACT

Hybrid strategies offer advantages over the dominant modeling strategies for agent-based social simulation, which rely on highly simplified assumptions or on empirical patterns that could result in overfitting to particular settings. Hybrid strategies seek to create models that incorporate the advantages of these approaches while incorporating rules of agent behavior that more closely represent complex social dynamics. The present study documents the types of social complexities that make a hybrid strategy desirable with respect to the pattern of AIDS transmission in sub-Saharan Africa, which is structured, complex, and largely hidden. Existing methods cannot capture the underlying dynamics, while emerging modeling techniques depend heavily on assumptions. The goal of the research discussed here is to provide a test bed for the development of prospective hybrid models. The strategy was twofold. First, we constructed a simple model that incorporates generic representations of the sources of variation in the HIV infection patterns. Second, the richness of the model was enhanced to better represent the social complexities from which AIDS emerges, while avoiding the risk of overfitting in producing the resulting hybrid models.

Keywords: hybrid modeling strategies, multilayer interactions, AIDS/HIV, sub-Saharan Africa

INTRODUCTION

AIDS is not only a devastating epidemic, but also one that is challenging to model effectively. The modeling difficulty arises from the intensely interactive nature of its transmission, an interactivity that gives rise to structured empirical patterns, but structures that may shift based upon changes at one or more levels of interaction. The most rigorous types of demographic models (Heuveline 2001) may be unable to capture the nature of these multilayer interactions and their co-evolution. In various parts of the world, the AIDS epidemic is spread by drug usage patterns, heterosexual and homosexual patterns of sexual relations, and infected birth. Each of these sources of HIV infection are likely to be influenced by differing patterns of interaction that are, at best, difficult to represent within a mathematical model.

In Africa, the primary source of AIDS transmission is heterosexual intercourse. However, even considering only this source of HIV infection, social factors rapidly multiply. For example, AIDS is sub-Saharan Africa is recognized to be related to migration patterns (Hunt 1989; Chirwa 1997; Hampshire 2002). Migration patterns may, in turn, be related to seasonal variation and/or

---

* *Corresponding author address:* Patrick Heuveline, Department of Sociology, The University of Chicago, Chicago, IL 60637; e-mail: pheuveli@midway.uchicago.edu.

economic conditions. Migration patterns may also be related to cultural factors, such as coming-of-age rituals. As a result, the patterns from country to country and region to region, may differ significantly.

Socially generated complexity takes other forms as well. The norm governing casual affairs may vary by whether an actor is (1) married or not (as well as when and how particular marriages occur — see, for example, Todd, Billari and Billari 2003) and (2) in a home village or migratory camp (or city). Further, available social networks in each locale determine operative constraints and opportunities. Such social networks may, in turn, be influenced by cultural group proximity and activity, with the result that, while infection rates are significantly shaped by migratory dynamics, these patterns are themselves socially mediated in complex ways.

The strategy of the present study is two-fold. First, we have constructed a simple model that incorporates generic representations of economic variation, migration, social networking and other selected sources of variation in the HIV infection patterns. Where such factors are notional, we incorporate applicable probability distributions, along with a capability of exploring interactive effects through relevant parameter sweeps. Second, we gradually enhance the richness of the model to more fully represent the social complexities from which the AIDS epidemic emerges while, at the same time, avoiding the risk of overfitting in the production of the resulting hybrid models. Ultimately, it is anticipated that the hybrid models may be used in exploring more complex issues, such as the secondary effects of AIDS infection patterns on the structuring of families (Heuveline, Timberlake and Furstenberg 2001; Wachter 2002).

AGENT-BASED MODELING

Over the last several decades, agent simulation has emerged as a novel methodology in the social sciences, one that integrates theory and empirical research, drawing premises and assumptions from the former, and generating aggregate patterns that can be compared with the latter. Although it takes somewhat different forms in the several disciplines, as a method it holds the promise of integrating the insights of multiple types of specialization into unified models.

The early work of Schelling (1978), Maynard Smith (1982) and Axelrod (1984)[1] provided a first wave of exemplars demonstrating the potential of a new approach to social simulation research.[2] Schelling's enormously influential model, which was essentially a thought experiment carried out on a checkerboard, was perhaps the closest to demographic concerns. With minimal technical resources, Schelling demonstrated that segregation at the aggregate level

---

1 For a critique of Axelrod's work from a game theoretic perspective, see Binmore (1998).

2 The waves or generations of agent simulation exemplars identified here are drawn primarily from two areas of social simulation: complex adaptive systems and evolutionary game theory. Parallel developments were occurring in distributed artificial intelligence (AI, or multiagent systems, see Weiss 1999), demography (microsimulation, see Wachter, Blackwell and Hammel 1997), ecological modeling (individual-based simulation, see DeAngelis and Gross 1992), and computational organization theory (see Carley and Prietula 1994). Development in each of these areas followed a different pattern. Early and continuing contributions in distributed AI, for example, were primarily in a variety of technical and problem-solving domains (see Bond and Gasser 1988); only later did multiagent insights begin to be applied in the area of social simulation (cf., Castelfranchi and Werner 1994). It is of inherent interest how the same computational capabilities give rise to similar innovations in various specialized areas of research.

was possible without bias in the micro-level population. In general, these studies illustrate how relatively simple models can provide insights into complex issues, including those with potential policy implications.

A subsequent generation of agent simulation research, including Epstein and Axtell (1996), Axelrod (1997) and Young (1998), provides a second wave of exemplars. They respectively illustrate, *inter alia:* (1) how agent simulation can be applied to an range of interactive social processes, (2) the diversity of social topics that can be addressed using simulation based on simple agents, and (3) the emergence of social institutions and structure from the interaction of agent strategies.

Based upon such foundations, more specialized types of research began to emerge, for example in economics (Sargent 1993), ecology (DeAngelis and Gross 1992) and international relations (Cederman 1997). It is not surprising then that demography, with its tradition of microsimulation, came to apply agent-based techniques as well (Billari and Fürnkranz-Prskawetz 2003). Whether addressing migration, the evolution of the family dynamics, or important historical transitions, agent-based computational demography (sometimes abbreviated as ABCD) provides the means for more deeply probing the complexities out of which demographic processes arise.

MODEL CONSTRUCTION

As indicated, one of the strengths of agent simulation is its ability to model complex interactions. This potential, which provides a focal point for the expression of theoretical generalizations, is also what enables the capability of modeling the complex cultural and social structures through which the AIDS epidemic is transmitted. To fully realize the potential of agent modeling, it is necessary to design relevant mechanisms and also to structure the interactions among such mechanisms. This is fundamentally a theoretical exercise, an activity that draws upon existing theory and by which further theoretical insights can be refined.

The present research project involves the design of four categories of mechanisms: (1) work-related migration, (2) networking and interaction, (3) disease and mortality, and (4) marriage and divorce. Each type of mechanism can be seen as contributing to the larger pattern of AIDS transmission in the South African region, which we selected for being currently the region of highest prevalence. In this paper, we discuss the construction of a basic version of the model, and its gradual elaboration. This basic model has a full architecture in the sense that the four mechanisms are represented, but they are initially represented by aggregate statistical distributions only, as they would be in any other type of micro-simulation. These aggregate parameters can be thought as “place holders” in order to establish the architecture of the full model, but will be gradually replaced by modeling the rules of behavior and interactions between agents that determine the observed distribution. It is only when this is fully implemented that the full potential of an agent-based simulation will be realized. At this time, only the marriage and divorce module has been so implemented.

THE BASIC MODEL

Work-Related Migration. In Southern Africa, migration plays a role in the spread of the AIDS infection (Hunt 1989; Chirwa 1997; Hampshire 2002). Young men migrate to urban areas and/or work camps where the HIV/AIDS rate and the risk of infection are significantly higher than in the villages and rural areas. Specific parameters will vary from population to population, depending *inter alia* on topology, population distribution, and cultural patterns. However, the generalized effect of migration creates a two-tier structure to the spread of the disease to which an HIV/AIDS model must attend. The sources of variation can then be explored by conducting a sensitivity analysis of relevant parameter ranges.

Our baseline migration model is driven by an exogenously determined unemployment rate. At present, the structure of its distribution, which is relative to quasi-discrete bands, is artificially defined. Subsequent refinement can substitute a theoretical or empirical economic base, but the present goal is simply to capture the two-tier structure.

As is typical in employment-driven migration processes, we assume that the propensity to migrate is highest among young adult males. The frequency of migration for specific agents gradually declines as, over time, migrating workers age and marry. Seasonal effects also influence the rate of migration. In the model, the rate of migration return is determined by season, and current migration duration of the agent. The entire migration process can be visualized using Geographical Information System (GIS) capabilities.

Networking and Interaction. Potential sexual partners are found within affinity networks of various types. In the basic model described here, affinity networks operate according to the following rules: (1) new acquaintances (and therefore prospective sexual partners) are introduced by mutual friends, (2) friendships without further contact decay over time, and (3) there is an upper limit on total friendships.

The number of sexual partners for a given agent is reduced by village residence, increased by migration, and is influenced by marital status (i.e., after agents marry, the number of sexual partners in a given time period is reduced). In the current model, frequency of sexual intercourse is based on an empirical distribution shaped by the values of relevant parameters. As discussed before, as the model evolves, any particular component may be refined or replaced.

Disease and Mortality. In the basic model, we have two mortality schedules depending on HIV status. In other words, when an individual becomes infected, she leaves the original age-at-death distribution and her age-at-death follows a second distribution corresponding to her reduced survival chances.

The infectivity of infected agents is also duration dependent, that is, it depends on the length of time between the time of infection and the time of a subsequent sexual contact. As suggested by epidemiological studies, infectivity is assumed to be highest during the first two weeks after the infected agent has been exposed to the virus and lowest immediately thereafter. Subsequently, there is a gradual increase correlated with the length of the agent's infection.

Marriage Formation and Dissolution. In our earliest models, marriage and divorce rates were based on a probability distribution summarizing empirical patterns. During young adulthood (ages 20–29) both marriage and divorce rates are relatively high. Subsequently, both marriage and divorce rates drop to levels that are roughly equivalent. There are aspects of the African cultural context (e.g., polygamy) that are not yet captured in this model.

MODULE SUBSTITUTION

Marriage is the first example in which one of the underlying mechanisms has been refined by replacing the basic mechanism with one that it is more sophisticated and intuitive. Specifically, what might be called the Basic Plus model draws upon and extends the marriage formation model of Todd and Billari (2003). In this model, each agent has a base quality, aspiration level, and courtship duration. Each is assigned randomly from a normal distribution. From about age 13 on, each agent surveys their friends in search of a friend of the opposite sex whose quality level exceeds their aspiration level. When one is found, an offer of courtship is extended. If the potential partner agrees, using the same criteria, a dating relationship is formed.

During courtship agents continue to look for a better relationship with friends of higher quality. Each agent also has a waiting threshold. If they do not participate in a courtship for longer than that threshold, their aspiration level is reduced. If agents date someone whose quality is higher than their aspiration level, the latter is adjusted upward. Alternatively, if an offer of a relationship is rejected, the agent's aspiration level declines as well. Ultimately, if the relationship lasts longer than the courtship duration parameter of both agents, they get married.

This marriage formation model better captures the serial and contingent nature of relationship formation than simple assignment based upon probability distribution does. In the present study, the mechanism has been further adapted by making aspiration levels more concrete and multi-dimensional, specifically using aspiration levels for age and wealth of prospective partners. As model development continues, it as anticipated that cultural criteria will be incorporated as well. This process serves as an example of how a mechanism based demographic model can extended to model more socially and culturally specific processes.

GENERATION AND ASSESSMENT OF RESULTS

The current study has completed its design and development phases and is presently moving into exploratory analysis. This phase initially focuses on how the five family structure variables presented in Table 1 evolve over time.

These endogenous variables will be affected by the parameters governing each of the mechanisms of the model, including the rate of migration, presence/strength of a right of passage, transmission rate (suggestive of agent condom use), proportion affected by network embedding and marriage/divorce rates. The interaction of these effects provide a central focus of ongoing research activities.

CONCLUSION

During the relatively brief history of agent simulation, two modeling strategies have come to dominate. The first, in which simulations are based upon *artificial society* models,

TABLE 1 Family Structure Variables

Male and female prevalence (% HIV positive in the population)
Ratio of male to female currently infected with HIV/AIDS
Proportion of children who are orphans (maternal, paternal or both)
Proportion of adults age 50+ with at least one orphaned grandchild whose deceased parent was their own child
Proportion of population by age category (0–15, 15–25, 25–50, 50+)

involves designs based on highly simplified assumptions. These computational models are used to clarify key relationships without attempting to fully capture the empirical complexities that arise in natural settings. The strength of this approach resides in its transparency and accessibility, while its weakness is the gap between the simple model and the complex structured reality to which it is (designed to be) applicable.

A second strategy attempts to achieve verisimilitude by drawing upon empirical patterns as a means of capturing complex social dynamics. While this approach may appear more representative in applied and/or policy settings, it runs the risk of being overfitted to a particular setting.

Between these poles, hybrid strategies seek to create models that incorporate advantages from each approach. The present study documents the types of social complexities that make a hybrid strategy desirable. More specifically, the pattern of HIV/AIDS transmission in sub-Saharan Africa is structured, complex, and largely hidden. Prevailing methods cannot capture the underlying dynamics, while emerging techniques are heavily dependent upon underlying assumptions. In addition to the goal of reducing the human cost of this pernicious epidemic, the present modeling strategy provides a salient test bed for the development of hybrid methods.

REFERENCES

Axelrod, R., 1984, *The Evolution of Cooperation,* New York: Basic Books.

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

Billari, F.C., and A. Fürnkranz-Prskawetz (Eds.), 2003, *Agent-based Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behaviour,* Berlin: Physica-Verlag.

Binmore, K., 1998, "The complexity of cooperation: Agent-based Models of Competition and cooperation," *Journal of Artificial Societies and Social Simulation,* **1**(1) [http://jasss.soc.surrey.ac.uk/JASSS/1/1/review1.html].

Bond, A.H., and L. Gasser, 1988, *Readings in Distributed Artificial Intelligence*, San Mateo, CA: Morgan Kauffman.

Carley, K.M., and M.J. Prietula (Eds.), 1994, *Computational Organization Theory,* Hillsdale, NJ: Lawrence Erlbaum.

Castelfranchi, C., and E. Werner (Eds.), 1994, *Artificial Social Systems: Fourth European workshop on Modelling Autonomous Agents,* Berlin: Springer-Verlag.

Cederman, L.-E., 1997, *Emergent Actors in World Politics*, Princeton, NJ: Princeton University Press.

Chirwa, W.C., 1997, "Migrant Labour, Sexual Networking and Multi-partnered Sex in Malawi," *Health Transition Review,* **7**(Supp. 3):5–15.

DeAngelis, D.L., and L.J. Gross (Eds.), 1992, *Individual-Based Models and Approaches in Ecology,* New York: Chapman and Hall.

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

Hampshire, K., 2002, "Fulani on the Move: Seasonal Economic Migration in the Sahel as a Social Process," *Journal of Development Studies* **38**(5):15–36.

Heuveline, P., 2001, "HIV and Population Dynamics: A General Model and Maximum Likelihood Standards for Eastern Africa," Paper presented at the General Conference of the International Union for the Scientific Study of Population, Salvador, Brazil.

Heuveline, P., J.M. Timberlake, and F.F. Furstenberg, Jr., 2001, "An International Comparison of Children's Experiences of Family Structures," Paper presented at the General Conference of the International Union for the Scientific Study of Population, Salvador, Brazil.

Hunt, C.W., 1989, "Migrant Labor and Sexually Transmitted Disease: AIDS in Africa," *Journal of Health and Social Behavior,* **30**(4):353–373.

Maynard Smith, J., 1982. *Evolution and the Theory of Games*, New York: Cambridge University Press.

Sargent, T.J., 1993, *Bounded Rationality in Macroeconomics,* New York: Oxford University Press.

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

Todd, P.M., and F.C. Billari, 2003, "Population-wide Marriage Patterns Produced by Individual Mate-Search Heuristics," pp. 117–137 in F.C. Billari and A. Fürnkranz-Prskawetz (Eds.). *Agent-Based Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behavior,* New York: Physica-Verlag.

Wachter, K.W., 2002, “AIDS and the Elderly of Thailand: Projecting Familial Impacts,” *Demography*, **39**(1):25–41.

Wachter, K.W., D. Blackwell, and E.A. Hammel, 1997, “Testing the Validity of Kinship Microsimulation,” *Journal of Mathematical and Computer Modeling*, **26**:89–104.

Weiss, G. (Ed.), 1999, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence,* Cambridge, MA: MIT Press.

Young, H.P., 1998, *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions,* Princeton, NJ: Princeton University Press.

MODEL ROBUSTNESS VERSUS PARAMETER EVOLUTION: ASSORTATIVE INTERACTION WITHIN A BARGAINING GAME

MARK H. GOADRICH,* Computer Sciences Department
University of Wisconsin–Madison

ABSTRACT

Agent-based models that explore aspects of social behavior invariably contain multiple parameters, such as population size, heterogeneous makeup, and spatial distribution. A common way to validate a model is to ensure robustness; that is, the model must produce consistent results independent of the initial parameter settings. When information can be learned about the prior probability of some parameter settings, however, robustness requirements on these parameters should be relaxed. The focus instead should be on the results produced from using these more likely settings. Brian Skyrms investigates a two-player noncooperative one-shot bargaining game called "Divide the Cake." Placed in an evolutionary setting, where players' claims are genetically hardwired and pairings are made at random, only 67% of initial population distributions result in all players using the "fair" strategy. Skyrms introduces correlation among players and shows that it precipitates the evolution of fairness from 100% of initial populations. Critics argue, however, that his exploration of correlation is lacking; other correlation models yield much worse performances. This paper examines the evolution of these nonrandom correlations, known as assortative interactions, through two separate agent-based models — a social network and a Schelling segregation model. The experiments show convergence to the fair strategy occurs approximately 90% of the time. This paper concludes that evolving the assortative interactions between players to find likely correlations, as opposed to guaranteeing model robustness, leads to a much more realistic picture of a model's behavior.

Keywords: Model robustness, assortative interaction, social networks, evolutionary game theory, agent-based models

INTRODUCTION

One of the common dimensions used to classify agent-based models is the degree of complexity. Models can be abstract, such as an iterated Prisoner's Dilemma (Axelrod, 1984), or realistic, such as an attempt to "investigate where prehistoric people of the American Southwest would have situated their households based on both the natural and social environments in which they lived" (*Village Project*, 2003). Abstract models usually have broad applicability and are pursued to explain the general mechanisms underlying a particular process, but these simple models are criticized for not capturing the complex details of the real world.

As we move toward realistic models, however, the size and scope of what is being simulated explode. As a result, our agents might have to cope with heterogeneous thresholds and

---

* *Corresponding author address:* Mark H. Goadrich, University of Wisconsin–Madison, 6785 Medical Sciences, 1300 University Avenue, Madison, WI 53706; e-mail: richm@cs.wisc.edu.

diverse landscapes, among other complications. Each new aspect brings into the simulation new parameters that must be tested, as these models are open to being overly sensitive to any one choice of parameters. Ultimately, our models should be "robust" and produce consistent behavior independent of our parameter settings, but how reasonable is this goal? Testing for robustness implies that all parameter settings are equally likely, yet this is not always the case. What if prior information were known about the parameter likelihood, a situation that brings into question the strict pass or fail test for robustness?

To demonstrate this situation, this paper explores a simple bargaining game made popular by Brian Skyrms. In his book, *Evolution of the Social Contract*, Skyrms explores the use of evolutionary game theory to explain our concepts of fairness (Skyrms, 1996). As in other fields, Skyrms hopes this direction will help to explain human social behavior when theories that rely on rational deliberation are lacking. His initial abstract model quickly becomes complicated when he introduces correlation among his agents, and the model is no longer robust when it takes into account these new parameters. The following two sections briefly summarize the current literature on this topic.

Divide the Cake

Skyrms' first example involves dividing a chocolate cake between two players, C_1 and C_2. Each player demands a certain amount of cake; when the total cake demand is less than or equal to the whole cake, each player receives her demand. Should the total demand exceed one, however, the cake is discarded, and the players leave empty handed. Our natural inclination when presented with this game is to divide the cake evenly — one-half for C_1 and one-half for C_2. But why do we consider this split fair? Skyrms points out that an infinite number of polymorphic solutions, or Nash equilibriums (e.g., C_1 demands 30% and C_2 demands 70%), exist. Rational deliberation does not help us distinguish between the "fair" solution and the polymorphic splits. This distinction opens the door to other explanations of fairness, namely, that evolution may have a hand in deciding our social behavior.

An evolutionary model is constructed by creating a finite population of players, each with preset and constant cake demands. This scenario assumes the use of the D'Arms et al. (1998) finite population and discrete simulation rather than the Skyrms continuous equations. Later, we explore larger numbers of strategies, but to simplify the analysis, we start with three:

- S_1: Always demand one-third of the cake (*modest*)
- S_2: Always demand one-half of the cake (*fair*)
- S_3: Always demand two-thirds of the cake (g*reedy*)

Individual cake games are conducted by independently and uniformly drawing C_1 and C_2 without replacement until all players are exhausted. A player's fitness score is the portion of the cake, if any, received in a game. The next generation of players is determined by the relative success or failure of each strategy for this game in combination with the current population distribution, a selection process known as the *replicator dynamics* (Weibull, 1995). This iterative process is continued until convergence of the population reaches a steady distribution. Skyrms states that the percent of initial population distributions, which evolve to a population where all players demand 50% of the cake, is 74%, not exactly the degree of success that we might expect.[1]

Skyrms solves the problem by introducing positive correlation among the strategies, or nonrandom mating of like-minded players. His players are given the ability to determine self- versus non-self-relationships among opposing players: greedy most likely plays with greedy, fair with fair, and modest with modest. This assumption breaks the polymorphic barrier, and Skyrms reports that only minimal correlation is necessary to cause widespread outbreaks of fairness quickly reaching 100%.

We can incorporate this correlation into our model by allowing the first player to influence the choice of an opponent. The initial finite-size population is still created according to a random population distribution, and C_1 is selected randomly from the current population distribution, $P(S_i)$. A player's preference for other strategies can be as defined by a function *pref*(i, j), the preference of a player using strategy i for a player using strategy j. Table 1 shows the correlation matrix when using Skyrms' assumptions of nonrandom mating. The selection of C_2 is governed by the following formula:

$$P(C_2 = S_j \mid C_1 = S_i) = \alpha < pref(S_i, S_j) * P(S_j) > ,$$

where α is the normalization constant. If the total demand of $C_1 + C_2$ is less than 1, a record is made of a successful game for each player's strategy. This process continues to sample without replacement until all players are exhausted. The average fitness for a strategy is calculated on the basis of successful games, and the players are redistributed accordingly for the next round. To assist in evaluating correlations, we define the *strength* of a correlation matrix in terms of the scale between preferred and nonpreferred strategies. For example, Table 1 is based on strength 8 because since fair is eight times more likely to choose fair over either greedy or modest.

TABLE 1 Skyrms' positive correlation matrix

Strategy *i*	pref(*i*, Modest)	pref(*i*, Fair)	pref(*i*, Greedy)
Modest	0.8	0.1	0.1
Fair	0.1	0.8	0.1
Greedy	0.1	0.1	0.8

Anti-Correlation Rebuttal

Skyrms' positive correlation is part of a broad class of correlations known as assortative interactions. Assortative interaction is usually discussed in the context of choosing a mate for reproduction as opposed to random mating strategies; in general it describes the tendency for individuals to choose their associates. In Divide the Cake, C_1 is still randomly selected, but the selection of C_2 is now weighted by the preferences of C_1.

1 Skyrms' result of 62% (presented in his book) is calculated for a population with 10 possible cake divisions, which are explored later. In the simulations of the three divisions described here, this number was 74%.

D'Arms (1996) quickly replies with questions about the assumption of positive correlation. He proposes that a model is robust if the result is virtually independent of the starting parameters. Skyrms' positive correlation makes the model robust with respect to initial population distributions, but correlation is now a parameter and should be examined with the same scrutiny. Finding one particular correlation that works is not a very robust argument.

D'Arms, et al. (1998) expand this claim into a model that allows for both correlation and anti-correlation as shown in Table 2. A greedy player using anti-correlation should wish to face anyone but another greedy in competition for cake. Fair still uses positive correlation and prefers fair players, and modest is happy playing against all three strategies. Unfortunately, anti-correlation enlarges the basin of attraction for a greedy/modest polymorphism to 54%, and their results hold across many strengths. D'Arms, et al. conclude that Skyrms' model is not robust with respect to variations in correlations.

TABLE 2 Anti-correlation matrix from D'Arms, et al.

Strategy *i*	pref(*i*, Modest)	pref(*i*, Fair)	pref(*i*, Greedy)
Modest	0.33	0.33	0.33
Fair	0.1	0.8	0.1
Greedy	0.47	0.47	0.06

MODELS OF CORRELATION

Both Skyrms and D'Arms, et al. use a scatter-shot approach to find reasonable correlation assumptions. While D'Arms, et al. succeed in their goal of providing a counter-example to Skyrms, the discussion should not end here. What other models of correlation are possible; how do they influence the evolution of fairness; and, more important, are some more likely than others?

In a separate critique of Skyrms' model, Barrett, et al. (1999) describe what they believe is the most natural correlation matrix (shown in Table 3): players choose associates with a mind toward their own utility. Modest players still freely associate with all players equally, but fair players prefer fair and modest opponents, while greedy players exclusively prefer modest opponents.[2] In general, players seek opponents who will not tip their combined demand over one. Using the preference matrix from Table 3, the experiments discussed in this paper show that 90% of initial populations evolve to all fair. If fair players constitute at least 8% of the initial population, this evolution is guaranteed. As the strength of this correlation increases, fairness approaches 100%.

Why is there such a benefit for fair players? Since fair is content with either fair or modest opponents, it steals some of the necessary modest players from the greedy players. The

2 Barrett, et al. suspect the resulting fairness model will evolve similarly to D'Arms, et al. with a broad basin for polymorphism.

TABLE 3 Utility preference correlation matrix

Strategy *i*	pref(*i*, Modest)	pref(*i*, Fair)	pref(*i*, Greedy)
Modest	0.33	0.33	0.33
Fair	0.47	0.47	0.06
Greedy	0.8	0.1	0.1

greedy strategy is never able to act with full power and, therefore, is at an evolutionary disadvantage. This strategy is opposite that of the D'Arms anti-correlation, where greedy players were stealing fair players and disrupting the average utility of fair. But for realistic anti-correlation, some greedy players must be willing to sacrifice themselves for the good of the strategy. This option is unlikely considering that they demand two-thirds of the cake.

Another possible correlation is created when players search for opponents who are seeking an equal portion of cake, a correlation suggested by Ernst (2001). Table 4 shows this "efficiency" correlation at strength 8. Ernst considers competition between groups of players, as opposed to a single population, and finds that efficient populations fare better than those that leave cake behind. How this situation could arise within a population is not exactly clear, since there is currently no benefit to consuming all of the offered cake. Nevertheless, this correlation is possible, and it exhibits behavior similar to that of utility preference.

Table 5 lists each correlation matrix discussed and shows the effect on fairness evolving as the strength of correlation increases. With all correlations except anti-correlation, greater

TABLE 4 Efficiency preference correlation matrix

Strategy *i*	pref(*i*, Modest)	pref(*i*, Fair)	pref(*i*, Greedy)
Modest	0.1	0.1	0.8
Fair	0.1	0.8	0.1
Greedy	0.8	0.1	0.1

TABLE 5 Effects of various assortative interactions on evolution of fairness

Strength	Positive (%)	Anti-Correlation (%)	Utility (%)	Efficiency (%)
0	74	74	74	74
2	98	63	77	67
4	100	59	83	70
8	100	56	90	79
16	100	56	95	87

strength brings about a greater evolution of fairness. While neither is as successful as positive correlation, utility preference is the closest. The use of a utility preference correlation matrix would very beneficial to the evolution of fairness.

Rather than relying solely on robustness as the criterion for success, it is also important to discriminate between correlations to find those that could arise naturally from player interactions. To properly understand the relationship between assortative interactions and the evolution of fairness, we must consider the evolution of correlations. The remainder of this paper explores two models used for discovering such correlations. First, players consciously construct a social network to help them learn what types of players will benefit their own claim. Second, players unconsciously employ a Schelling segregation model on a two-dimensional lattice; players randomly select new locations, without looking at their opponents' strategies, when their current utility falls below a threshold.

SOCIAL NETWORK MODEL

Learning preferences among players is not as hard as it may seem. In fact, nature provides evidence that these interactions exist. Sober and Wilson (1998) cite an experiment that examined the interactions of guppies in the context of altruism:

> A separate experiment allowed three guppies to inspect predators in an aquarium divided by transparent panels into three lanes. The guppies were placed in an apparatus that allowed the fish that occupied the center lane to indicate a preference for one of the two side fish by swimming over to join it as a companion. The side fish that moved closer to the predator was consistently chosen as a future associate (p. 140).

If simple-minded guppies can learn preferences that increase their utility, Sober and Wilson contend, how much more likely is it that humans with all our faculties can do the same? A simple way to learn the preferences for our agents is to randomly pair them with opponents and then record whether a game is successful or not. We assume each strategy is assigned a tag, which can be recognized by other players, and records are kept based on those tags, not on individual players. The resulting correlation matrix is consistent with the utility preference matrix from Table 3, but this process seems too easy.

Skyrms and Pemantle (2000) suggest a more complex mechanism to dynamically learn a social network between game players. To make things more interesting, the number of strategies is now 9, from 0.1 to 0.9. We redefine greedy and modest players as those demanding more or less than one-half the cake, respectively. Players begin with a uniform preference for all other players. Each player is given 1,000 rounds to play Divide the Cake, choosing opponents according to her preference vector. Whenever a game is successful, the player initiating the visit updates her preference vector to increase the chance of revisiting this cooperative strategy. To add noise, unsuccessful games are recorded favorably 20% of the time.

While this model might be expected to evolve like the simple model, there are key differences. Tables 6 and 7 show two resulting correlation matrices. Although all of the

TABLE 6 Sample results of dynamic social network model: resulting fairness 99%

Strategy *i*	pref(*i,0.1)*	pref(*i,0.2)*	pref(*i,0.3)*	pref(*i,0.4)*	pref(*i,0.5)*	pref(*i,0.6)*	pref(*i,0.7)*	pref(*i,0.8)*	pref(*i,0.9)*
0.1	0.08	0.17	0.02	0.12	0.06	0.21	0.14	0.13	0.07
0.2	0.12	0.00	0.07	0.07	0.15	0.24	0.08	0.27	0.00
0.3	0.06	0.03	0.02	0.44	0.41	0.03	0.01	0.01	0.00
0.4	0.55	0.07	0.11	0.07	0.13	0.04	0.00	0.01	0.01
0.5	0.19	0.16	0.18	0.25	0.21	0.00	0.00	0.01	0.01
0.6	0.39	0.03	0.41	0.15	0.00	0.00	0.01	0.00	0.01
0.7	0.01	0.84	0.13	0.00	0.01	0.01	0.00	0.00	0.00
0.8	0.58	0.40	0.00	0.00	0.00	0.00	0.01	0.00	0.00
0.9	0.94	0.00	0.01	0.00	0.01	0.01	0.00	0.01	0.00

TABLE 7 Sample results of dynamic social network model: resulting fairness 58%

Strategy *i*	pref(*i,0.1)*	pref(*i,0.2)*	pref(*i,0.3)*	pref(*i,0.4)*	pref(*i,0.5)*	pref(*i,0.6)*	pref(*i,0.7)*	pref(*i,0.8)*	pref(*i,0.9)*
0.1	0.08	0.03	0.25	0.03	0.15	0.16	0.01	0.18	0.11
0.2	0.02	0.01	0.06	0.44	0.32	0.06	0.08	0.00	0.00
0.3	0.02	0.21	0.01	0.23	0.04	0.10	0.40	0.00	0.00
0.4	0.17	0.25	0.10	0.02	0.05	0.41	0.01	0.00	0.00
0.5	0.23	0.43	0.22	0.09	0.01	0.01	0.00	0.00	0.00
0.6	0.06	0.15	0.62	0.14	0.00	0.00	0.00	0.00	0.01
0.7	0.25	0.07	0.66	0.01	0.00	0.00	0.00	0.00	0.00
0.8	0.63	0.34	0.00	0.00	0.00	0.00	0.01	0.01	0.00
0.9	0.97	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.00

preference is concentrated on opponents that provide a positive outcome, this preference is no longer uniform; random choices of initial opponents cause the revisiting and reinforcement of certain players rather than other equally acceptable players.

For experiments described in this paper, the dynamic social network model was used to create 800 networks. Each correlation matrix was then tested for the resulting fairness, with 1,000 randomly selected initial populations. Figure 1 shows a histogram for the distribution of fairness percentages. The mean fairness was 89.3%, but the median score was 94.6%, showing evidence of a distribution skewed heavily toward fairness evolving. In social nets where fairness did not dominate, evidence of a tight network is shown between other demands. Table 7 shows the close preferences of 0.3, 0.4, and 0.6, as well as the fair strategy of 0.5 preferring opponents of 0.2. Exactly why these correlation matrices do not evolve fairness is still being investigated.

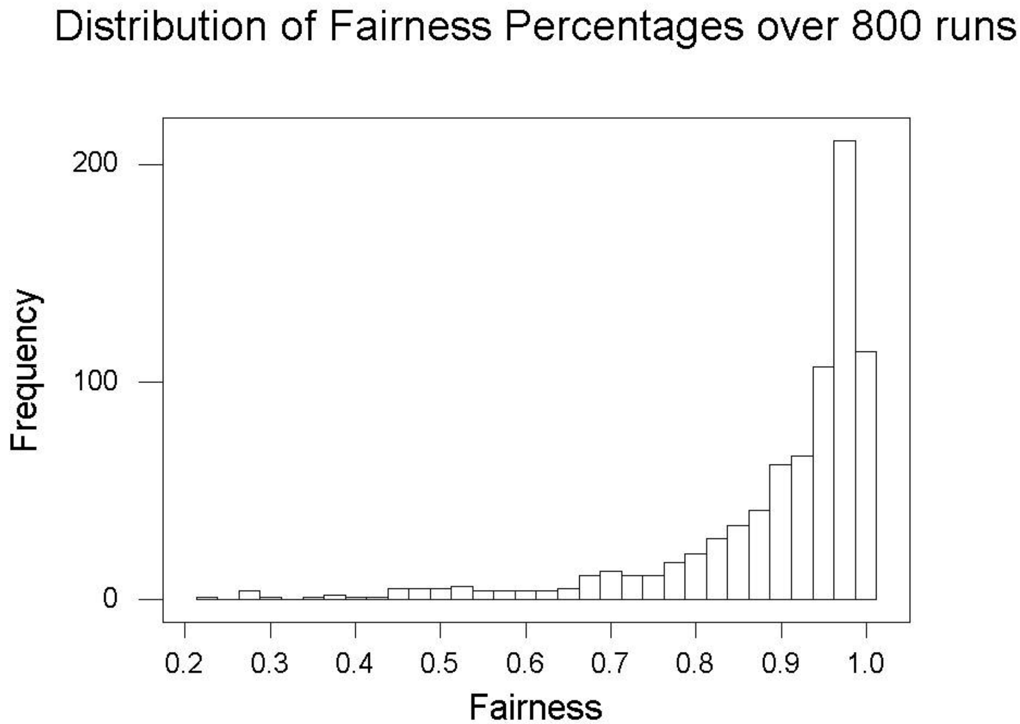

FIGURE 1 Resulting distribution of fairness evolution using 800 learned social networks

SCHELLING SEGREGATION MODEL

Evolving assortative interaction matrices can also be approached from the perspective of a spatial model. Again, we tried to make a minimal number of assumptions, which are reasonable and relatively benign. Under this model, we removed the previous assumption that players can distinguish between other players based on strategy. First, players are spatially distributed. Second, players are allowed to change their location if they deem it unsuitable. Finally, a player's goal is to maximize utility — a common assumption in game theory. Skyrms and Alexander (1999) have explored spatial models of Divide the Cake, but they only allowed players to change strategies.

These assumptions can be readily modeled in a common framework borrowed from economics. Schelling's famous segregation model demonstrates that minor preferences of satisfaction within your neighborhood can have striking results for the overall distribution of individuals (Schelling, 1978). He specifies a simple game to be played with pennies and dimes on a chessboard. First, place about 45 dimes and pennies randomly on the board.[3] The neighborhood of a coin is defined as the eight surrounding squares, with both the horizontal and vertical edges wrapping around as in a torus. Second, assign certain preferences to both dimes and pennies; for instance, dimes prefer neighborhoods with at least one-third dimes, and pennies are only happy when surrounded by at least one-half pennies. Third, determine who is unhappy

3 Similar segregation behavior should evolve independent of the initial distribution of dimes to pennies.

in the initial board and move them to a new random location.[4] Finally, repeat this process until either all the coins have reached stability or oscillations develop. The overall behavior of the game gravitates toward patterns of segregation, even though both dimes and pennies would be satisfied under certain layouts of integration.

Divide the Cake naturally fits into this framework. To continue our simple model of three strategies, we now have three types of players — one for each strategy. A player is defined as unhappy in her neighborhood as follows:

$$\text{unhappy}(C_i) = \begin{Bmatrix} \text{true}: & \bar{u}(C_i, N_j) < t \times \text{demand}\ (C_i) \\ \text{false}: & \text{otherwise} \end{Bmatrix},$$

where

$\bar{u}(C_i, N_j)$ = average utility C_i receives against all her neighbors N_i,

demand (C_i) = demand of the player, and

t = parameter in the range [0,1] indicating the threshold a player has for receiving no cake.[5]

In other words, players are looking for neighborhoods to maximize their total possible gain. Repeatedly moving unhappy players and examining the resulting neighborhoods exposes the preferences for each strategy.

Simulations of the Schelling model were tested for population sizes from 1,000 to 5,000 by using random samples of population distributions. Players were allowed to assort for 20 time steps before evolving into the next generation based on their current fitness levels. Unsettled populations were terminated after 100 generations and recorded as a failure to evolve fairness. Our new parameters for this more complex model are the size of the board and the tolerance t at which a person is unhappy. Figure 2 shows a sample run for an initial board size of 31×31, 9 player categories, with a distribution of 50 players per strategy, and a tolerance value 0.75. Table 8 reports the evolved correlation matrix after 20 time steps. This matrix was calculated by counting the neighboring strategies for each individual and then normalizing to one.

When looking at the preference of happy players, this run of the Schelling simulation appears to evolve a utility preference matrix similar to that shown in Table 3. Since not all players are happy, the utility preference matrix is an asymptote. When all players are considered, greedy individuals display an overall preference to choose themselves because of the unavailability of suitable modest players. Figure 3 shows the change in fitness scores due to assortment of the players.

---

4 Schelling recommends starting at the upper left corner and proceeding row by row. He claims the order of movement is unimportant; however, this procedure leads to waves of unhappy players moving down the board.

5 Here, 0 means the player is happy no matter how much cake she receives, and 1 means the player must fully receive her demand to be happy.

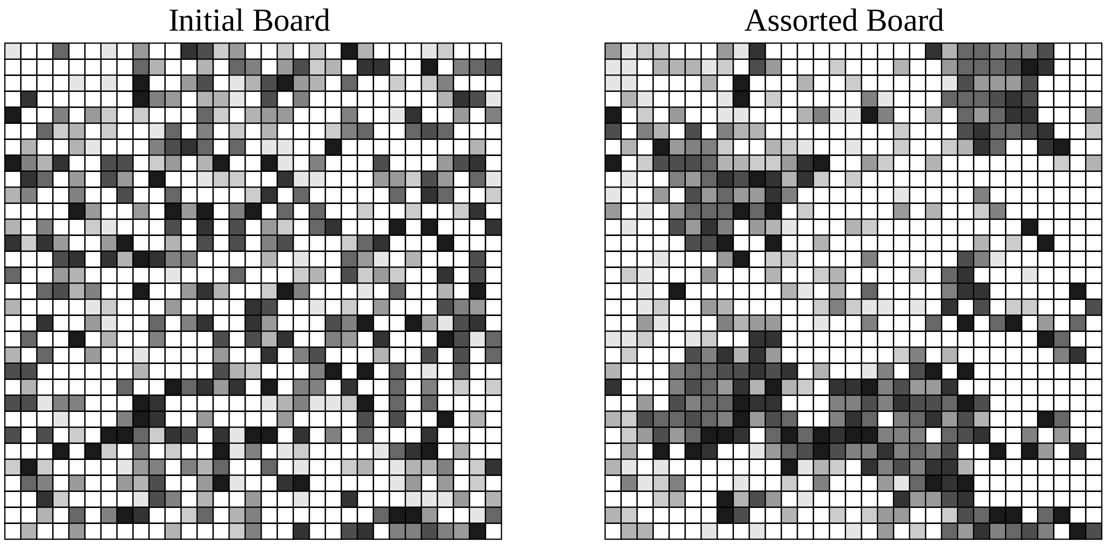

(shades range on the gradient from 0.1 = black to 0.9 = light grey)

FIGURE 2 Rearrangement of players based on Schelling's model

TABLE 8 Sample results of Schelling's spatial model: resulting fairness 89%

Strategy *i*	pref(*i,0.1)*	pref(*i,0.2)*	pref(*i,0.3)*	pref(*i,0.4)*	pref(*i,0.5)*	pref(*i,0.6)*	pref(*i,0.7)*	pref(*i,0.8)*	pref(*i,0.9)*
0.1	0.10	0.17	0.14	0.12	0.13	0.12	0.07	0.06	0.07
0.2	0.13	0.18	0.16	0.13	0.13	0.12	0.08	0.03	0.02
0.3	0.10	0.16	0.12	0.18	0.13	0.18	0.06	0.04	0.02
0.4	0.09	0.12	0.17	0.26	0.13	0.17	0.02	0.01	0.03
0.5	0.11	0.14	0.15	0.15	0.23	0.02	0.05	0.07	0.07
0.6	0.11	0.14	0.21	0.21	0.02	0.12	0.05	0.07	0.07
0.7	0.10	0.14	0.11	0.04	0.07	0.07	0.14	0.22	0.11
0.8	0.09	0.06	0.07	0.02	0.11	0.11	0.23	0.13	0.18
0.9	0.09	0.03	0.04	0.05	0.11	0.11	0.11	0.17	0.28

Modest players rarely move from their initial random locations. The only reason they would be unhappy is if they were lonely and had no neighbors; otherwise, they would be content to play against anyone. Also, fair players settle down and find groups much more easily than do greedy players. This fact is true irrelevant of the initial distribution of players, which could result because fair players can find neighborhoods of either fair or modest players, whereas greedy players must find near-exclusive modest neighborhoods complementary to their own demand to be satisfied. As more and more greedy players surround modest players, the average utility for each greedy player falters and places her on the move again.

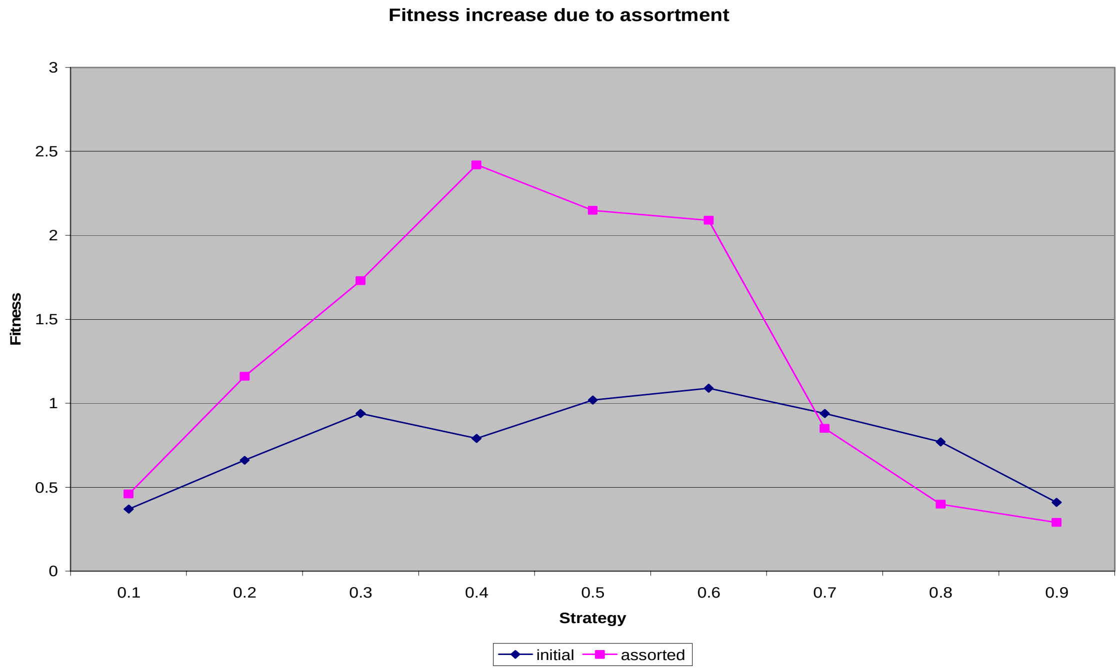

FIGURE 3 Change in fitness due to Schelling's assortment

Settings for the board size parameter were tested for up to three times what is necessary to fit all the players. Differences in the resulting evolution of fairness were minimal; however, extra space made it easier for players to find favorable opponents. With more elbow room, greedy players can surround modest players while still avoiding each other. This ability shifts the correlation matrix closer to the efficiency correlation of Table 4.

Variations of the tolerance threshold produced more interesting results. Figure 4 (on the following page) shows the average fairness evolution when tolerance was varied from 0 to 1. Values from 0.6 to 0.85 result in close to 90% fairness, while higher values, such that players are only happy with receiving their demand, show a return to polymorphic solutions over fair evolution.

CONCLUSIONS

Skyrms shows that a certain model of correlation effectively promotes the evolution of fairness across all initial populations. But once he introduces correlation, he is open to criticisms from D'Arms, et al. that other correlation schemes produce opposite results. We feel that the examination of alternate correlation systems should also proceed in an evolutionary environment to bring out those correlations that could naturally emerge from player interactions. While certainly not robust with respect to alternative correlations, the approach of learning our probable parameter values gives a much more accurate picture of the model.

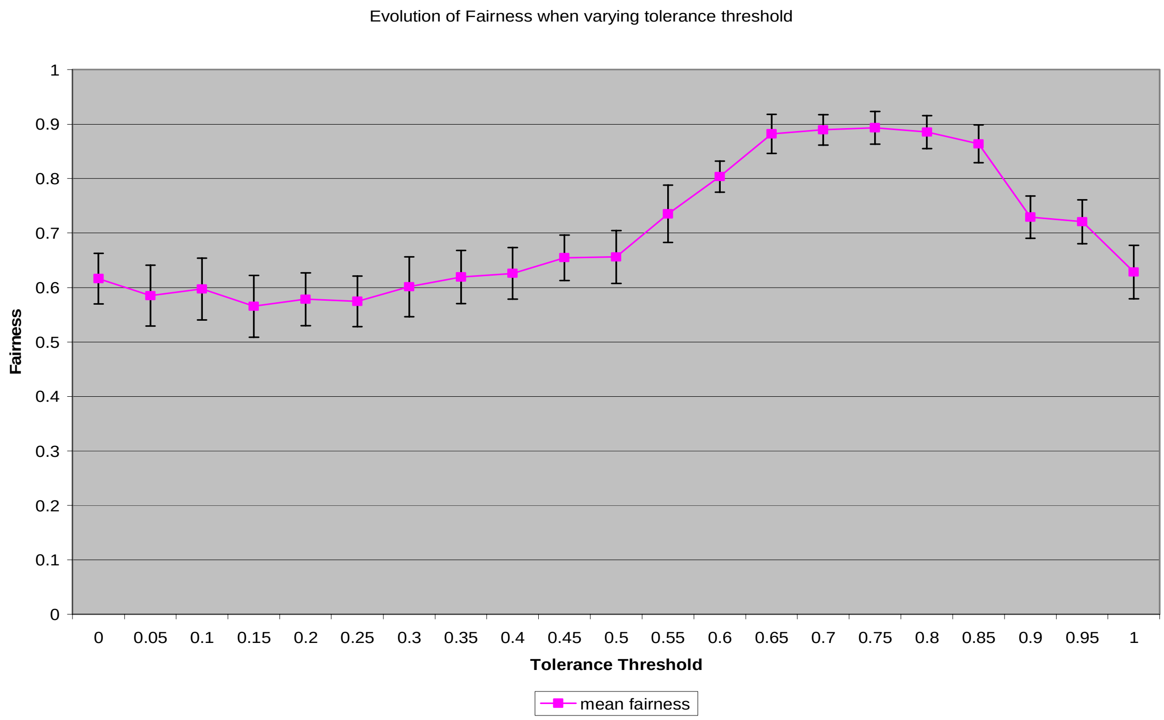

FIGURE 4 Variation of tolerance threshold parameter from 0 to 1 (error bars for one standard deviation)

The use of a social network model for player types rather than actual players could be seen as overly simplistic; a more complete model would have each player learn a distinct preference vector for every other player. Also, Skyrms and Pemantle (2000) discuss other variations on their dynamic social network formation, such as reciprocal visiting and decay in memory, which should be investigated in the context of the bargaining game.

The results from the spatial Schelling model are very promising. To reinforce the claims made in this paper, a number of extensions should be made to the model. First, the space of possible tolerance values should be examined. With this new parameter, we should examine ways of letting each player learn her own tolerance, as the implications of heterogeneous tolerance by strategy and by player could have very drastic implications and need to be explored. Second, at this time, unhappy players are randomly relocated to a new location; better relocation packages for displaced players should be explored, such that a player could select the best from *n* randomly chosen new locations. In addition, the cost of obtaining preferences as discussed in D'Arms, et al. has been totally ignored. A cost could be assessed per player based on how many times they must move to be happy.

Each model was not entirely successful in showing a complete evolution of fairness; however, these results are significantly different than when using a random correlation and bear further investigation. The approach shown here can be readily incorporated into other agent-based models, allowing us to delve deeper into those relevant areas of the model. Although it requires an additional step to tune the model parameters, the benefits can be drastic. We should focus our attention on essential areas rather than quibble over irrelevant parameter values.

ACKNOWLEDGMENTS

The author would like to thank Elliott Sober and Brian Skyrms for their helpful suggestions on an earlier draft of this paper. Part of this research was performed under the support of National Library of Medicine training grant 1T15LM007359-01.

REFERENCES

Axelrod, R., 1984, *The Evolution of Cooperation*, New York City, NY: Basic Books.

Barrett, M., E. Eells, B. Fitelson, and E. Sober, 1999, "Models and Reality: A Review of Brian Skyrms' 'Evolution and the Social Contract'," *Philosophy and Phenomenological Research,* **59**:237–242.

D'Arms, J., 1996, "Sex, Fairness, and the Theory of Games," *The Journal of Philosophy* **93**:615–627.

D'Arms, J., R. Batterman, and K. Gôrny, 1998, "Game Theoretic Explanations and the Evolution of Justice," *Philosophy of Science*, **65**:76–102.

Ernst, Z. 2001, "Explaining the Social Contract," *British Journal for the Philosophy of Science*, **52**(1):1–24.

Schelling, T., 1978, *Micromotives and Macrobehavior*, New York City: NY: W.W. Norton & Company.

Skyrms, B., 1996, *Evolution of the Social Contract*, Cambridge, U.K.: Cambridge University Press.

Skyrms, B., 1999, "Reply to Critics," *Philosophy and Phenomenological Research* **59**:243–254.

Skyrms, B., and J. Alexander, 1999, "Bargaining with Neighbors: Is Justice Contagious?" *The Journal of Philosophy* **96**:588–598.

Skyrms, B., and R. Pemantle, 2000, "A Dynamic Model of Social Network Formation," in *Proceedings of the National Academy of Sciences of the United States of America*, pp. 9340–9346.

Sober, E., and D.S. Wilson, 1998, *Unto Others*, Cambridge, MA: Harvard University Press.

Village Project, The, 2003, Washington State University, July 28. Available at http://www.wsu.edu/~village/.

Weibull, J.W., 1995, *Evolutionary Game Theory*, Cambridge, MA: MIT Press.

TOWARD SIMULATION-BASED, PROBLEM-SOLVING ENVIRONMENTS FOR CONFLICT MANAGEMENT IN COMPUTATIONAL SOCIAL SCIENCE

LEVENT YILMAZ,* Auburn University
TUNCER Y. ÖREN, University of Ottawa

ABSTRACT

With the increasing power and utility of computational tools and infrastructures in performing social science research, the need to move forward to simulation-based advanced problem-solving environments (PSEs) and computational social science laboratories is evident. The field of computational social conflict modeling and analysis is growing rapidly. PSEs, such as those suggested in this paper, offer a new way of performing simulation-based social science research. To this end, this paper focuses on supporting computational social scientists in conflict modeling. PSEs are integrated computer systems that provide computational facilities necessary to solve a target class of problems efficiently. By definition, PSEs extend the program-compile-execute cycle of model development and simulation to high-level, problem-solving activities. While existing simulation-based methods suggest a program-compile-execute cycle, this paper emphasizes the significance of a simulation modeling environment that integrates model building, simulation management, collaboration, intelligent distributed simulation, and sophisticated analysis tools. This paper also discusses fundamental features of social conflict modeling and analysis PSEs and argues the limitations of the existing simulation conceptual frameworks in modeling realistic conflict scenarios. Existing problems in PSE technology are discussed, and several recommendations are provided to address these limitations.

Keywords: Conflict management, simulation-based problem-solving environments, social simulation, social agents

INTRODUCTION

The social science research community is focusing more than ever on simulation-based computational models. The capability of modeling and simulating sophisticated social phenomena and understanding the implications of mechanisms based on abstractions of reality facilitates reasoning about complex social systems. As the emphasis in social science computation shifts from low-level simulation programming and execution to high-level, problem-solving environments (PSEs) to specify models and scenarios and to test hypotheses, PSEs will become even more important for performing social science research; this movement follows the trends in engineering and the physical sciences.

Among the social phenomena worth studying are conflicts because they affect the quality of life everywhere. Conflicts have occurred frequently throughout human history; national and international conflicts are ubiquitous (Balencie and de La Grange, 1999). As common

* *Corresponding author address:* Levent Ýilmaz, Department of Computer Science and Software Engineering Samuel Ginn College of Engineering, Auburn University, Auburn, AL 36849; e-mail: yilmaz@eng.auburn.edu.

occurrences during the 20th century (Grant, 1992), conflicts are at least as worthy of study as the Cold War (Arquilla and Ronfeldt, 1997; Khalilzad and Lesser, 1998). Even more important, perhaps, is the study of conflict management (i.e., conflict avoidance and conflict resolution). For example, on the basis of behavioral science's prospect theory, Davis and Arquilla (1991) assert that "possible opponents are likely to become increasingly and unreasonably risk-accepting as they become emotionally more dissatisfied with [the] current situation and trends." Game theory has been applied to social problems (Shubik, 1964). Evidence exists, however, that classical game theory fails in cases where opponents have different value systems. Schelling's (1980) pioneering work of analytical game theory recommends identification and consideration of focal points, which are the perceived mutual expectations, obsessions, sensitivities, appreciation, and the like for conflict resolution in search of win-win conditions.

Conflict systems are complex social systems. Some modeling approaches available for resolving conflicts are based on, for example, different types of game theories (e.g., sequential, differential, evolutionary, and hyper games). Several other approaches, such as bounded rationality, deterrence theory, and crisis destabilization, are also used for their solutions. Some novel simulation modeling formalisms, which are not in competition with already proven theories and approaches, may be useful for the proper formulations and resolutions of conflicts. Waldrop (1993) and Kaufmann (1996) have investigated the subject of complexity. There are examples in conceiving complexity in elegant ways. For example, fractals can be used to generate a complex system based on simple initial knowledge (Barnsley, et al., 1988). A catastrophic manifold can represent interesting, and sometimes contradictory and counterintuitive, patterns of behavior (Casti, 1979). Cybernetics has been considered as a source of paradigm for simulating complex systems, including social systems (Knight, et al., 1971; Ören, 1978). For a bibliography on contemporary sociocybernetics studies, see Geyer and van der Zouven (1998). Recently, computational social science initiatives have emerged to facilitate systemic and intelligent study of societies; yet, computational studies of conflict management are not as pervasive as economics and other social phenomena. Furthermore, unlike researchers in the life and physical sciences, social scientists who study conflict management are not yet equipped with state-of-the-art, domain-specific computational laboratories. To this end, the goal of the research reported in this paper is to develop a *problem-solving environment* for computational social scientists to *rapidly* compose multi-level, multi-faceted artificial societies to facilitate experimentation with sophisticated intervention and conflict negotiation mechanisms.

Why Is This Problem Important?

The way we perceive reality affects our actions. Ideally, we need appropriate paradigms and modeling methodologies to perceive, conceive, and foresee conflicting situations to avoid them and, if they are inevitable, to resolve them (Ören, 2001). Regardless of their type and origin, conflicts are parts of social systems; like other social phenomena, they are difficult to model. Social systems are sometimes labeled in the literature as "soft" or "ill-defined" systems, where the usefulness of traditional mathematical representations is questioned (Spriet and Vansteenkiste, 1982, p. 42). In a major effort, Davis (1986) used the structure of war gaming and included artificial intelligence models (rule-based systems) to represent national and international leaders and commanders. Zeigler (1990) used these works as an example of the more general approach of variable structure agent-based simulation. Many studies have been conducted on a special type of conflict, namely, war gaming. In war gaming, military decision makers (i.e., commanders at different levels) can obtain “war experience” in peacetime by using gaming simulations, also called constructive simulation, in defense applications. Today, war-gaming studies use computers extensively, although such studies predate computers. For example, in a bibliography on professional war gaming, early studies date back to the second half of the 1880s (Riley and Young, 1957). Two types of war games exist: one for professionals and one for hobbyists. In war gaming, it is much easier to model equipment than it is to model humans. Recently, studies have been performed to remedy the situation (Pew and Mavor, 1998). It is argued that conflict avoidance and conflict resolution deserve levels of effort similar to war gaming. Like war-gaming experience, military and civilian decision makers can enhance their conflict management skills through conflict management simulation studies.

What Is Required?

The premise of the outlined research is to develop appropriate modeling paradigms, such as multi-aspect and multi-stage modeling formalisms, and associated enhanced simulation formalisms (i.e., multi-simulation) to simulate conflict avoidance and conflict resolution. The suggested modeling and associated simulation formalisms would allow simultaneous experimentation with different — even contradictory — aspects of reality. Results of the experiments with multi-stage models can be displayed at the same time by taking advantage of the possibilities offered by virtual and augmented realities. Such modeling and simulation formalisms might also be useful in modeling other social phenomena and hence useful for sociocybernetics studies.

TOP-DOWN CONCEPT OF THE COMPUTATIONAL ENVIRONMENT

Simulation-based, Problem-solving Environment for Computational Conflict Analysis and Management

As integrated computational environments, PSEs allow users to access relevant knowledge and software tools to solve problems. Software tools are used to specify problems; to check consistency and completeness of the specifications; to transform the problem specifications into executable computer programs; and to run these programs to generate, analyze, document, display, and store the results and other relevant aspects of the problems. The PSE would mentor and advise users on several aspects of knowledge about conflicts and conflict management. The simulation ability of the PSE would enable users to test the effect of decisions on the conflict management process. The simulation system would be a discrete event-driven system. Discrete event abstractions represent dynamic systems through discretely occurring events that can be triggered on the basis of conditions that occur outside and inside the system model. In discrete event systems, control over time can be expressed explicitly and flexibly, along with its essential constraints on complex adaptive system behavior and structure. Also, the capability to efficiently represent loosely coupled distributed semiautonomous processes through either synchronous or asynchronous communication provides insight into the behavior of the system as well as the interactions among its components. Moreover, because it can run with trace data, discrete event abstraction facilitates parallel experimentation with real-world data.

Figure 1 illustrates the conventional usage of simulation-based problem environments. The simulation is often initialized with domain-specific configuration parameters and uses the underlying hard-wired assumptions and strategies modeled at design time. The extension of this basic problem-solving mode is discussed in the following section to argue for potential extensions to deal with realistic conflict scenarios.

MODES OF USING AN ADVANCED PSE

Three types of system usage are envisaged:

- *A systematic source of knowledge about conflict and conflict management*. Advanced PSE can be used as a separate service and/or within the following two types of simulation studies:

- *A conventional stand-alone simulation system*, where simulation is helpful for analysis, education, training, and research.

- *An embedded simulation system* (integrated with the real system). The simulation system can support the operation of the real system by allowing parallel experimentation while the real system is running. The simulation system provides predictive displays for decision making as well as calibration of the knowledge embedded in the PSE while monitoring the predictions of the simulation system and the occurrences of the real phenomena. Embedded simulation systems are well known, especially in training associated with equipment operation (embedded simulation system [ESS], Simulation Training and Instrumentation Command [STRICOM]); however, they are also applicable to decision systems (Beer 1975).

Role of Personality and Cultural Knowledge

The PSE would include *the five-factor personality traits* knowledge of human behavior in order to take into account knowledge about the personality traits in human behavior simulation (Ghasem-Aghaee and Ören, 2003a). Knowledge about *cultural* backgrounds of the participants and their *value systems* has to be included in the PSE knowledge base because they are often sources of disagreements as well as essential elements in consensus building (Laszlo, et al., 1977; Huntington, 1996; Lewis, 1999).

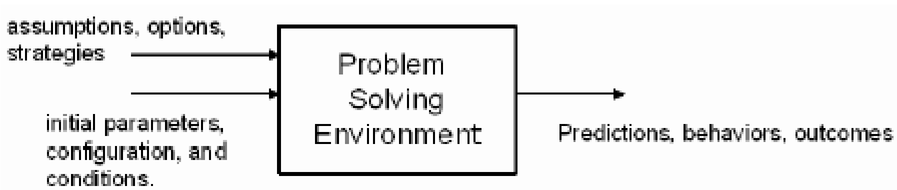

FIGURE 1 Traditional problem-solving practice

Figure 2 illustrates the extension of the basic model using the knowledge that incorporates the social, cultural, and psychological context. One strategy to encode this knowledge into simulation models is to use the computational intelligence methods as discussed below.

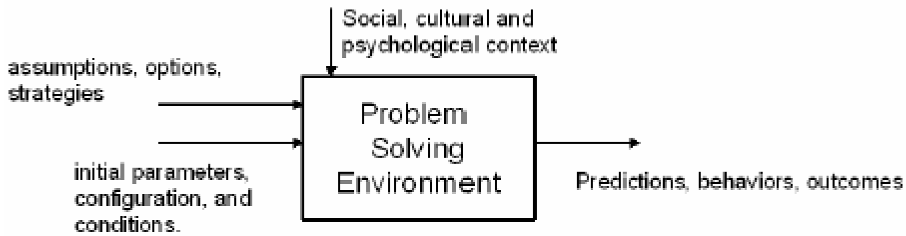

FIGURE 2 Incorporating social, cultural, and psychological context

Modeling Intelligent Entities, Artificial Intelligence, and Soft Computing

"Soft computing differs from conventional (hard) computing in that, unlike hard computing, it is tolerant of imprecision, uncertainty and partial truth. In effect, the role model for soft computing is the human mind. The guiding principle of soft computing is: Exploit the tolerance for imprecision, uncertainty and partial truth to achieve tractability, robustness and low solution cost" (Zadeh, 1975).

Software Agents and Agent Simulation

Software agents are entities that function continuously and autonomously in a particular environment, often inhabited by other agents and processes (Shoham, 1993). These agents possess some of the following characteristics to a certain level of degree:

1. *Reactivity* (selectively sense and act),

2. *Autonomy* (i.e., goal directness, proactive and self-started behavior),

3. *Collaboration* (i.e., work with other agents and entities to achieve a common goal),

4. *Knowledge-level communication ability* (i.e., communicate with other entities in a language like speech-act, higher level than symbol-level, program-to-program protocols),

5. *Inferential capability* (i.e., act on abstract task specifications, using models of self, situation, and/or other agents),

6. *Temporal continuity* (i.e., show persistence of state and personality),

7. *Personality* (i.e., manifest attributes of a believable agent),

8. *Adaptability* (i.e., learn and improve with experience), and

9. *Mobility* (i.e., migrate from one host to another in a self-directed way).

The envisaged PSE would utilize and extend the multi-agent simulation (MAS) paradigm by the novel concepts briefly discussed in the following sections. Software agents constitute the fundamental components of MAS. The MAS paradigm brings a radically new solution to the very concept of modeling and simulation in social sciences by offering the possibility of directly representing individuals, their behavior, and their interactions. The MAS paradigm makes it possible to model complex situations and synthetic worlds whose overall structures emerge from interactions between individuals, that is, to cause structures on the macro level to emerge from models on the micro level, thus breaking the level of barrier in classical modeling (Ghasem-Aghaee and Ören, 2003b).

Fuzzy Agents and Systems

By their vary nature, digital simulations consider only quantitative parameters and seem powerless when faced with multitudes of qualitative data collected by researchers in the field. Fuzzy set theory is a mathematical apparatus for the formal representation, processing, and utilization of data and information characterized by nonprobabilistic uncertainty and vagueness. The extension of discrete-event simulation agents of the envisaged PSE with this theory can allow the creation of agent behavioral models that reason on imprecisely and ambiguously defined terms, relations, and mechanisms of approximate inference, which are typical of human reasoning (Ghasem-Aghaee and Ören, 2003a).

Holonic Systems, Cooperation, and Holonic Agents

Holonic agent simulations can allow exploration of the effects of curtailing autonomy of the holons (or some subsystems) to optimize the performance of the entire system to provide a basis for negotiations (Ghasem-Aghaee and Ören, 2003b). Furthermore, between competition and full cooperation, there is an important possibility, namely, cooperation in some areas, but competition in other areas (i.e., co-opetition). Methodologies have to be developed to model and explore co-opetition.

Novel Simulation Paradigms for Conflict Modeling and Analysis

Multi-models and Multi-aspect Models

A multi-model is a modular model where only one model module is active at a certain time. Each model module is an alternate model. A multi-model provides a conceptually clean way of representing system entities. With multi-models, similar to any conventional simulation study, only one aspect of reality can be simulated at a given time. The concept is applicable to continuous, discrete, and memoryless models, as well as to other modeling formalisms, such as discrete-event systems, rule-based models and software agents, including intelligent agents and mobile agents (Ören, 1987, 1991, 2001). Two special cases of multi-models are metamorphic models and multi-aspect models.

A metamorphosis can be represented by a metamorphic model which, in turn, can be represented as a special case of a multi-model. For example, alternate models can represent egg, larva, pupa, and butterfly; alternate models can be selected under well-defined conditions. In this case, however, there is a predefined sequence for the alternate models; that is, transitions from alternate models would be rather limited.

A multi-aspect model is another special case of a multi-model where the condition of having only one alternate model active at a given time is relaxed. An example usage might be representation of solid, fluid, and vapor phases of the same mass of material (e.g., ice, water, and vapor) and the transitions from one phase to another. In the example, alternate models representing both water and vapor can exist concurrently with a mass transfer from one to another alternate model. The direction of the transfer of an entity — in the example, water, or vapor — depends on whether energy is given to or taken from the multi-model. Similarly, multi-models can be used to represent turmoiled and law-abiding groups that can co-exist with transitions from one group to another based on the emerging/created/engineered conditions.

Multi-stage Models and Multi-simulation

In a multi-stage modeling formalism, several aspects of reality can be formulated by sets of component models. Normally, all the multi-stage models are not known *a priori*. For example, only the initial model M1 may be known. In this case, one can attempt to model alternative models to prepare for contingencies. Multi-stage model formalism can allow multi-simulation. A multi-simulation can allow the experimentation with several (even contradictory) aspects of reality simultaneously as shown in Figure 3. When some previously unforeseen conditions arise (i.e., under emerging conditions), one can add emerging successor models to existing models to explore behavior of alternative system models. Multi-simulation may be the simulation paradigm to experiment with Schrödinger's cat, which can be alive and dead at the same time (Marshall and Zohar, 1997). In non-quantum theoretic realm, it is argued that ability to experiment with several — even contradictory — aspects of reality may bring new vistas in conflict management. Multi-stage models facilitate exploratory analysis. But exploratory analysis becomes

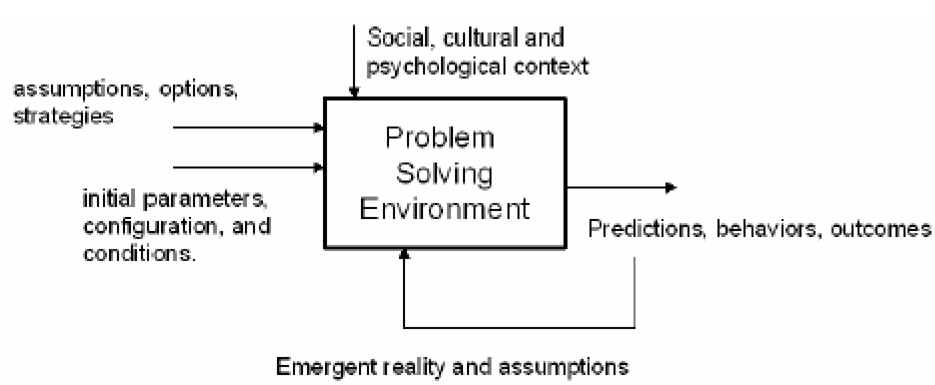

FIGURE 3 Dealing with emergent reality with multi-stage modeling

computationally difficult to manage as the level of model detail increases. Multi-resolution modeling with metamodels (Davis, 2000) is suggested to deal with this issue.

Substantive Theories Explaining Social Phenomena

The knowledge about social, cultural, and psychological context needs to be formulated under consistent and powerful conceptual frameworks to facilitate explanation and reasoning. Our long-term goal is to develop a full-fledged formalization of context that can be used, among many others, in conflict modeling and analysis. Situation theory is a mathematical theory of information that can be used to capture abstract situations that designate real-world counterparts.

In situation calculus, the world is conceived as a tree of situations, starting at an initial situation, S_0, and evolving to a new situation through the performance of actions by the opponents in conflict. The state of the world is expressed in terms of relations and functions that are true or false or have a particular value in situation *s*. The major contribution of situation theory in a PSE would be deductive plan synthesis to plan sequences of real-world actions and preference options over a search space. To this end, situation theory offers a powerful framework that might be useful in realizing the exploratory modeling and simulation concept introduced by Davis (2000). Figure 4 illustrates the inclusion of context modeling through a substantive theory, such as situation theory.

PROBLEMS AND RECOMMENDATIONS FOR PSE DEVELOPMENT

The current state of computational social science problem solving has limitations. These limitations are due primarily to the lack of mature technologies that would support the identified requirements and features. In this section, we review principal technical problems and suggest recommendations to facilitate further progress in this area.

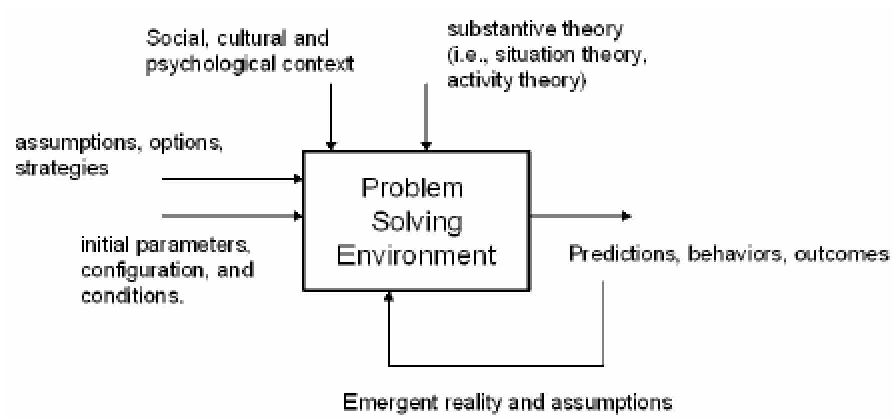

FIGURE 4 Using substantive theory for context modeling

Monolithic Problem-Solving Practice

Most existing problem solvers in the computational social science domain focus on a specific problem. Flexibility is achieved through parameterization of the problem inputs, rather than customization of the model configuration and components at the time of construction. That is, given a slightly different problem, such hard-wired problem solvers cannot be reconfigured to model a different problem with new domain-specific constraints. Furthermore, these problem solvers are stand-alone systems independent of the services or models provided by potential collaborators. The plug-and-play paradigm can increase the flexibility of existing problem solvers. The paradigm suggests taking one algorithm or new implementation and substituting in place of another existing model component without causing conflicts. Specification-directed model generators can help ease model derivation for a wide variety of problems as long as the specification is expressive enough to communicate constraints and requirements of a general problem area.

Lack of Architecture, Technology, and Methodological Support for Scalable Problem Solving

One of the fundamental barriers to problem solving is the lack of environments and methods that can scale to handle realistic artificial societies. Existing distributed simulation infrastructures, such as high-level architecture, immediately degrade in performance as the number of federations that join to the simulation increases. It is well known that it is difficult to develop and manage large complex software systems. As simulation-based PSEs become more and more software intensive, the scientific community that relies on simulation to analyze scientific phenomena is affected by the lack of a “silver bullet” that can deal with software complexity.

Lack of Flexible Model Adaptation and Assembly

It is difficult and unrealistic to have a single model that is useful for many purposes. Such models immediately become overcomplicated and hard to maintain. Hence, developing flexible and adaptable components can facilitate having reconfigurable designs that can be adapted to satisfy the constraints of emerging scenarios. Parameterized modeling of scenarios and components enables not only adaptation, but also composition through parameter instantiation.

Lack of Principled Design Methodologies for Cognitive Modeling of Human Behavior Simulation Components

In a system (or model) without memory, an input can be transformed to an identical output according to the transfer function of the system as many times as the input is applied to the system. In state-determined systems (or systems with memory), a given input may induce different outputs corresponding to the state of the system. Human behavior is not only state determined (i.e., past experience influences the current outcome), but several filters affect the outcomes (decisions, reactions, etc.). For example, personality acts as a filter. Two individuals who may have similar past experiences are expected to react differently on the basis of their personalities (Ghasem-Aghaee and Ören, 2003a,b). Furthermore, mood, cultural background, and value system of an individual (group) also act as filters to affect decisions (or reactions) of the individual (or the group). Development and proper consideration of these filters as well as emotion management knowledge are not to be excluded in conflict management studies (and simulations).

ROLE OF THE PSE IN CONFLICT EDUCATION AND TRAINING

Educating students enrolled in social science disciplines at all levels in emergent next generation paradigms in their own disciplines is an immediate and paramount goal for the continued vitality of the country's technical infrastructure. Recently, this objective has been the subject of intense debate at various National Science Foundation meetings and panels. It has been recognized that the effectiveness of technical education lies not only in improved facilities, but also in the social aspects involving pedagogy, presentation, and dissemination. The chief educational frontier in computational social science thus refers to the design of leading-edge tools, software, and learning modules that use innovative methodologies for transforming the educational experience. Simulation-based PSEs are definitely superior to other types of learning environments. Simulation studies facilitate experimentation with dynamic models of real systems under any conceivable and even extreme conditions and allow generation and observation of knowledge pertinent to the behavior of the model under the experimentation scenarios. This type of rich knowledge about a system could not be represented without using simulation.

Pedagogical Uses of PSEs

The conventional wisdom and approach in teaching the application of computational methods in social sciences (computational science, in general) are to emphasize that it is easier to change the problem to suit the algorithms and models than vice versa. The goal in PSEs and problem formulation methodology discussed here is to select and adapt algorithms to suit the problem at hand. Matching problems to appropriate algorithms is an integral aspect of a scientist's formative training, and its importance in educational circles is widely recognized. The central idea is to promote the use and integration of PSEs into the social science curriculum. An increasing number of educational tools are needed, such as modeling platforms with prebuilt intellective models that are available to students to manipulate and run virtual experiments. To this end, problem formulation, model selection, simulation, and the ability to explain how certain outcomes emerge are the fundamental components of a pedagogical plan. One important aspect that is an immediate implication of the proposed multi-modeling (i.e., multi-stage) approach is the ability to explain why certain recommendations are made. This aspect can be advantageously utilized toward the formalization of decision processes for the student and the development of a learning module that emphasizes a recommendation and simulation-based approach to solving computational social science problems.

CONCLUSIONS

After the conquest of the material world which led to material wealth, it is extremely challenging to start to understand ourselves and to learn how to manage our conflicts. Simulation, which helped us in many ways in the material world, may also be useful for achieving these goals. Diligently, we must focus on exploring the synergy of several related fields by gleaning useful knowledge from humanities, developing appropriate modeling and simulation methodologies and technologies, and taking advantage of ubiquitous computational power. Some possibilities are pointed out in this paper.

ACKNOWLEDGMENT

This paper is based on Ören (2001), and some relevant passages are freely adopted from it.

REFERENCES

Arquilla, J., and D. Ronfeldt, 1997, *A New Epoch — and Spectrum — of Conflict,* Santa Monica, CA: RAND Corporation. Available at http://www.rand.org/publications/MR/MR880/.

Balencie, J.-M., and A. de La Grange, 1999, *Monde rebelles*, *Guerres civiles et violences politiques — Encyclopedie des conflits*, Paris, France: Editions Michalon (in French).

Barnsley, M.F., et al., 1988, *The Science of Fractal Images*, New York City, NY: Springer-Verlag.

Beer, S., 1975, *Platform for Change*, New York City, NY: John Wiley and Sons.

Casti, J.L., 1979, *Connectivity, Complexity, and Catastrophe in Large-scale Systems*, Chichester, U.K.: Wiley-Interscience.

Davis, P., 1986, *Applying Artificial Intelligence Techniques to Strategic-level Gaming and Simulation*, in M.S. Elzas, ed., *Knowledge Based Modelling and Simulation Methodologies*, Amsterdam, the Netherlands: North Holland Publishing Company (also available as RAND P-7120, RAND Corporation, 1986).

Davis P.K., 2000, "Exploratory Analysis Enabled by Multiresolution, Multiperspective Modeling," in J.A. Joines, R.R. Barton, K. Kang, and P.A. Fishwick, eds., *Proceedings of the 2000 Winter Simulation Conference*, RP-925, Santa Monica, CA: RAND Corporation.

Davis, P.K., and J. Arquilla, 1991, *Thinking about Opponent Behavior in Crisis and Conflict: A Generic Model for Analysis and Group Discussion*, RAND/R-4111-JS, Santa Monica, CA: RAND Corporation. Information about Embedded Simulation Systems (ESS) is available at http://www.ets-news. com/emb_sim.htm.

Elzas, M.S., T.I. Ören, and B.P. Zeigler, eds., 1986, *Modelling and Simulation Methodology in the Artificial Intelligence Era*, Amsterdam, the Netherlands: North Holland Publishing Company, pp. 315–338.

Geyer, F., and J. van der Zouven, 1998, *Bibliography on Sociocybernetics*, 3rd Ed., Amsterdam, the Netherlands: SISWO, July. Available at http://www.unizar.es/sociocybernetics/bibliografia.html.

Ghasem-Aghaee, N., and T.I. Ören, 2003a, "Towards Fuzzy Agents with Dynamic Personality for Human Behavior Simulation," in *Proceedings of the 2003 Summer Simulation Multiconference*, Montreal, PQ, Canada, July 20–24, 2003.

Ghasem-Aghaee, N., and T.I. Ören, 2003b, "Soft Systems Methodology and Fuzzy Holonic Agents for the Simulation of Cooperative Behavior" (in press).

Grant, N., 1992, *Illustrated History of 20th Century Conflict*, London, U.K.: Reed Consumer Books.

Huntington, S.P., 1997, *The Clash of Civilizations and the Remaking of World Order*, New York City, NY: Touchstone.

Kaufmann, S., 1996, *At Home in the Universe: The Search for Laws of Self-organization and Complexity*, Oxford, U.K.: Oxford University Press.

Khalilzad, Z., and I.O. Lesser, 1998, *Sources of Conflict in the 21st Century — Regional Futures and U.S. Strategy*, Santa Monica, CA: RAND Corporation.

Knight, D.E., H.W. Curtis, and L.J. Fogel, eds., 1971, "Cybernetics, Simulation, and Conflict Resolution," *Proceedings of the 3rd Annual Symposium of the American Society for Cybernetics*, New York City, NY: Spartan Books.

Laszlo, E., et al., 1977, in *Goals for Mankind*, New York City, NY: E.P. Dutton.

Lewis, R.D., 1999, *When Cultures Collide – Managing Successfully across Cultures*, London, U.K.: Nicholas Brealey.

Marshall, I., and D. Zohar, 1997, *Who's Afraid of Schrödinger's Cat*? New York City, NY: Quill-William Morrow.

Ören, T.I., ed., 1978, *Cybernetics & Modelling & Simulation of Large-scale Systems*, Namur, Belgium: International Association for Cybernetics.

Ören, T.I., 1987, "Model Update: A Model Specification Formalism with a Generalized View of Discontinuity," in *Proceedings of the Summer Computer Simulation Conference*, Montreal, Quebec, Canada, July 27–30, pp. 689–694.

Ören, T.I., 1991, "Dynamic Templates and Semantic Rules for Simulation Advisors and Certifiers," in P.A. Fishwick and R.B. Modjeski, eds., *Knowledge-based Simulation: Methodology and Application*, Berlin: Springer-Verlag, pp. 53–76.

Ören, T.I., 2001, “Towards a Modelling Formalism for Conflict Management,” in H.S. Sarjoughian and F.E. Cellier, eds., *Discrete Event Modeling and Simulation: A Tapestry of Systems and AI-based Theories and Methodologies*, New York City, NY: Springer-Verlag, pp. 93–106.

Pew, R.W., and A.S. Mavor, eds., 1998, *Modeling Human and Organizational Behavior — Application to Military Simulation*, Washington, DC: National Academy Press.

Riley, V.G., and J.P. Young, 1957, *Bibliography on War Gaming*, Operations Research Office, The John Hopkins University, Chevy Chase, MD.

Schelling, T., 1980 (reprinted from 1960), *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.

Shoham, Y., 1993, “Agent Oriented Programming,” *Artificial Intelligence* **60:**51–92.

Shubik, M., ed., 1964, *Game Theory and Related Approaches to Social Behavior*, New York City, NY: Wiley.

Spriet, J.A., and G.C. Vansteenkiste, 1982, *Computer-aided Modelling and Simulation*, London, U.K.: Academic Press.

Waldrop, M.M., 1993, *Complexity: The Emerging Science at the Edge of Order and Chaos*, Carmichael, CA: Touchstone Books.

Zadeh, L.A., 1975, “Outline of a New Approach for the Analysis of Complex Systems and Decision Procedures,” *IEEE Transactions on System Man and Cybernetics*, SMC-3, pp. 28–44.

Zeigler, B.P., 1990, *Object-oriented Simulation with Hierarchical, Modular Models — Intelligent Agents and Endomorphic Systems*, London, U.K.: Academic Press.

DISCUSSION:

SIMULATION METHODS

(Thursday, October 2, 2003, 1:00 to 3:15 p.m.)

Chair: *Michael North, Argonne National Laboratory*
Discussant: *Roger Burkhart, Deere Company*

AIDS Transmission in Sub-Saharan Africa: Issues in Modeling and Methods

Roger Burkhart: We would like to take questions on each session and leave time at the end for a general discussion. We would like to start with questions or comments for David Sallach.

Unidentified Speaker: What do you think an agent-based model would provide you? Would it give epidemiologic models?

David Sallach: Traditional methods of epidemiology do not address this epidemic. This epidemic is rooted in specific forms of interaction. Therefore, I think that one of the things you can achieve is a higher level of accuracy [with agent-based models]. Beyond that, however, there are policy-oriented issues that have to do with lifestyle facts and so forth. By having a much deeper understanding of the social mechanisms involved, the potential for addressing those policy issues increases.

Unidentified Speaker: The model that you described is incredibly complex and the ones you want to do may become even more so. I'm assuming this is done in Repast. Is that correct?

Sallach: Yes.

Unidentified Speaker: Can you explain more about what this model looks like? How are you encoding all this? How many agents do you have running, and how much is happening per cycle, per step for each agent? What's going on?

Sallach: We do not have a tremendously large number of agents — only a couple of hundred agents or so, because we are specifically looking at the migration cycle. In other words, we are not modeling a country. In that respect, the model has some of the characteristics of an artificial society model, in spite of the complexity that we have built into it. The reason is that we are really looking at the multi-tier process by which the infection spreads, say the way that seasonal and age affects the migration cycle, which increases the risk of exposure, which is then brought back into the village. Even when it is brought back into the village, though, it is mediated by the affinity network. It is really a simple model. There is one affinity network, but, of course, the idea is to multiply. The simple model is a friendship network; however, we want a workplace network and a cultural network along similar lines.

So, this is just one process by which young people (by a certain probability that is mediated by about three or four considerations) who do or do not migrate can increase their exposure when they do migrate. At some point, they come back into a relationship with the duration of the time stay. When they come back, if they have been infected, there is an increase in the probability that the disease, or the epidemic, will spread. But the disease will spread in social network-mediated ways. That is a basic description of how the simple model works.

Model Robustness Versus Parameter Evolution: Assortative Interaction within a Bargaining Game

Burkhart: Our next speaker is Mark Goadrich from the University of Wisconsin–Madison.

Mark Goadrich: I am currently a graduate student in computer science at the University of Wisconsin–Madison. Today, I'm going to speak about some issues that have come up in a "philosophy of science" setting, such as evolutionary game theory. Next, I will move into agent-based systems, including the complexity discussed by David Sallach, that is, moving from a simple model into a very complex model. I will discuss some of the problems that will come up, and some of the ways that we might try to approach them.

[Presentation]

Burkhart: Thank you, Mark. We have time to take a couple of questions.

Unidentified Speaker: With regard to your last comment, if we assume that there will always be some level of uncertainty, isn't this a bit of a false dichotomy? There is always going to be a range of values to measure parameters, and you are going to have to worry about robustness inside of that range and even about absent measurements. Your robustness is across a larger range, but the result is that the data restrict the range across which you need to worry about robustness. Will the dimensionality of the problem remain the same regardless of whether you compare the data or not?

Goadrich: I hope that once we start modeling the parameters, as seen with the threshold, there will still be another parameter to model; however, I hope it will be reduced. If you learn a model for one parameter, you can move on to another parameter and try to model that parameter. The number of parameters that you end up with is much less than if you tested across everything.

Unidentified Speaker: Can you address robustness in terms of the initial numbers of the distribution of the correct research within the population? Most theoretical grid-based games are very sensitive as to how many agents are of each type and start the simulations. For example, if you have very few greedy agents, they tend to dominate or win, mostly because they take advantage of the others. So can you say a few words about this type of robustness?

Goadrich: That's the parameter that [I used to keep] robustness in these models. I did test across many different populations, some with only 10 greedy people, 900 fair people, and 10 modest people. That is where I get the average fairness from a model. I still kept that in because cannot be sure what happens. At that point, you want robustness to help you out. It is a situation where you do not know what those parameters should be. In situations where you know what your outcome should be, however, in some ways, you know that we observe this phenomenon of fairness. We want a model to show us fairness. We can still have depth within that.

Unidentified Speaker: When you talked about adjusting the parameters one by one, did you assume that these parameters did not correlate with each other? If you did not, the minute you adjust one and then move on to the other one for the new value of the second parameter, the first one may not be at the appropriate place. As a result, you will have a continuous iterative process that would be very long and time-consuming.

Goadrich: That is a great point. What I was saying does assume independence of your parameters. But if they are not independent, you might have to either model them together or just leave it to robustness.

Burkhart: We have time for one more question. Any additional questions or comments can be given in the discussion.

Unidentified Speaker: I was wondering whether you have ever incorporated altruistic behaviors in this model. If you talk about fairness, one aspect of that characteristic would be agents benefiting from how well the other agents do in their utility function.

Goadrich: That would be a great extension. No, I have not incorporated altruistic behavior in the model. The fitness is only based on that individual model. However, if we do extend it to where the fitness of an individual in the Shelly model is based on their eight-person neighborhood and how well they are doing because they could share, we would definitely want to look at that extension.

Toward Simulation-based, Problem-solving Environments for Conflict Management in Computational Social Science

Burkhart: Our next speaker is Levent Yilmaz in a joint effort with Tuncer Oren. Levent is from Auburn University, and Oren is from University of Ottawa.

Levent Yilmaz: Thank you. This presentation is quite different than those of my colleagues in this session. Their main interest involved finding solutions to interesting problems — unique, normal problems — in terms of conventional methodologies or simulation methodologies. The point of our paper, however, is to look at ways to extend or enhance simulation methodologies to solve certain problems for which agent-based social simulation by itself is not sufficient.

[Presentation]

Burkhart: Thank you Levent. Let’s take a couple of quick questions at this time. Please hold additional questions for the discussion session.

Joanna Bryson: I want to revisit one of your slides midway through your presentation. You talked about needing fuzzy logic, and after that, multiple other logics to represent emotions and further roles.

First, you said the BDI architecture does not have anything like that. I think that statement is true of Woolridge and Jennings and Rowlings, or theory-based people, but if you go back further, when it was actually running on robots, I believe Karen Myers and Kurt Konolige did something called PRS Lite, which updated PRS to use fuzzy rules. That was working on robots.

Yilmaz: They might have; I am familiar with their work. Basically, when I referred to fuzzy logic, I not only meant the framework of the agents but also how to describe personality dynamics in terms of fuzzy variables.

Bryson: I don't understand. I'm agnostic about this. I have never found it necessary to have anything except deterministic control for performing the modeling that I do because the environment has a great deal of uncertainty. From my perspective, when I have seen fuzzy logic or probabilistic talks, it seems that the probabilistic work and the Bayesian stuff are very clean. What is the advantage of the fuzzy logic and multi-modal logic?

Yilmaz: I think probabilistic ideas and fuzzy ideas are different. Probabilistic reasoning means that we are talking about the certain probability of going from one state to another. It is not about greatness or uncertainty. Rather, it is a probability associated with a certain transition from one state to another. And that probability is 0.6%: you go from State A to State B. It does not say anything about notions. For instance, "I am very hungry. I am very, very, very hungry." You see, it is uncertain and vague. It is not probabilistic, and that is the main difference between fuzzy logic and probabilistic reason. So, in my opinion, they are not the same.

Panel Discussion

Burkhart: Mark [Goadrich], please join the panel so that all three speakers [from this session] can respond to questions. Thank you.

This session was not about toolkits specifically. It was, however, titled "Simulation Methods," and as a toolkit developer myself, I thought it was quite interesting. Taken together, the three talks actually help to build a case for some of the models we are attempting to build. Certainly, all of the speakers raise the issue of the simple, of the theoretical, of what David [Sallach] called the "artificial society," but almost the artificially simplified reality of trying to define and discover often-simple mechanisms in isolation. Together, they are concerned with how to actually test alternative mechanisms and also with the results of many different mechanisms working together.

The paper on problem-solving environments used multi-models, multi-simulation, multi-aspect, and multi-stage. It strikes me that the theme of these talks is multi-mechanism models, whether it deals with how we validate or even distinguish them. We heard three different mechanisms for possible fairness games — from the assortative roles to the spatial to the social networks.

Another question also raised by the problem-solving environments is: Have we begun to discover and put together mechanisms to use in more roles than just the theoretical or scientific role as suggested by "problem-solving environment?" Are we switching to the applications of some of these models, and certainly with policy implications and problems such as AIDS?

I would like to open this discussion with some basic questions, one for each speaker, as a way of priming discussion. I invite everyone to participate.

For David [Sallach], the question involves all of the mechanisms. You are the one who really defined “social mechanism” — this stylized causal chain, which I think is very important for identifying the potential building blocks as we start to work with these larger models. The question is: Can we actually validate the individual mechanisms empirically in addition to trying to recognize the results of all of them put together? I would like to hold that question so that I can present all three questions to the speakers up front.

For Mark [Goadrich], the question is regarding the robustness criterion: When I read the paper, I wondered why is robustness the rule? I think we really are trying to create the observed phenomenon of fairness, but when doing the *a priori*, more theoretical game theory models, are such criteria enough? Or are they weaker, thus ending some sort of extra-empirical validation or other source, and ultimately being necessary to distinguish?

Finally, for Levent [Yilmaz], you want to build models that have more mechanisms. I was struck by not only the mechanism that might be there, but also by the ones endogenously created in successive multi-stage models. The model recommender module, the run-time recommender, seemed interesting. Do you think it is enough to anticipate the kinds of mechanisms, but then go to run-time switching? Or is it necessary to either run-time load or run-time generate, or merge the kinds of mechanisms that might be present?

Feel free, everyone, to comment on each other’s talks and also to address this question of what expanded roles we might be looking at for in these complex multi-mechanism models.

Unidentified Speaker: I would say that I’m not very close to addressing the validation issue at this point. The scope of the issue I wanted to address had to do more with the kind of multi-dimensional complexity that you find in a real-world problem and how you get that under control conceptually and in terms of modeling. I also looked at how to begin to integrate it or synthesize it in ways that you can talk about — a reduction of the parameter space, which I think necessarily means a higher level of abstraction. Of course, I agree that at some point we have to turn to creative ways for validating models. But if you look at the factors that I considered, you could see that even if those were aligned very well from Malawi, all you would have to do is shift to another country, where you have a different mix of factors. Those factors might be migration factors or the presence or absence of IV drug use or any number of things, where if you could model it for Malawi, it would not necessarily be applicable to the other.

So, from my perspective, at least for that type of problem, looking for a great deal of rigor and validation is premature. What I am looking more for kind of abstractions that will help us to structure the problem first.

Unidentified Speaker: Yes, that’s an interesting intermediate point as we go from the theory construction models to the explanatory models for specific problems. Actually, I think these mechanisms can help to build the vocabulary, identify the building blocks.

I think the question has to do with the multiple mechanisms that are being tested and with checking the robustness of the parameter. Is that a strong enough criterion, or do you really need other ways of distinguishing, or ultimately answering, which mechanisms are the appropriate ones to explain fairness?

Unidentified Speaker: In the concept of fairness, much more would come into play. This model is very restrictive, and you only get one demand. There are many more models that ask: "Do I get to change my policy? Do I get to change my demand? Do I get to change, not just where I live, but do I get to change my demand based on people around me?" In our situation where we had such limited information, you need robustness to validate across numerous things. As you get more specific and as you add more information, however, the robustness seems to fade away, because more is understood about what's happening. You also know more about the actual situation, and you can limit yourself to looking in-between certain areas.

Unidentified Speaker: The final question concerned whether the mechanisms might be dynamic to the point where they could emerge or be introduced after the model starts.

Yilmaz: This is the reason I like to come to conferences. That is an excellent question, which I had not considered before. If you try and anticipate what type of situation will emerge, how do you come up with some counter approach to be able to recommend that?

The best answer to that question would be to incorporate learning mechanisms. That is, under certain conditions, if an actor is doing a good job in terms of the tactics, strategies, and outcomes, you might reinforce that particular conflict management procedure captured by a model and incorporate it as often as possible in the future.

But that is a very difficult problem. You do not have a table that describes under what conditions a particular model would be able to help an actor in a conflict situation, so that makes the problem more difficult to handle.

You asked an excellent question. You bring an excellent question to Dr. Sallach about how do you modularly valvate. There are different social mechanisms in your model; it is very complex. How did he compose them into different social process mechanisms, valvate them separately, and bring them together modularly in a tractable manner? This also is an excellent question, and there are certain studies in social engineering and modeling that deal with modular composition validation of different models that might help.

Burkhart: At this time, I would like to open the discussion so that you can ask questions of any the speakers on these things or other topics that you think the papers raise.

Greg Madey: My comment on the presentation has to do with fairness. You ignored need. Suppose someone was hungry and someone else was not hungry, or someone was 200 pounds and someone else was 100 pounds. In those scenarios, your equilibrium point would be different.

Unidentified Speaker: Definitely. Many assumptions went into the games that we talked about. But if there were a need, we would want to have a model that produced the appropriate behavior, given that need. I think we would still go through some of the same mechanisms and follow the same procedure to try to validate a model of fairness, given who needs what.

Bryson: I also have a question for the same person, Mark [Goadrich]. However, the other panelists may want to talk about this, too. My comment actually goes back to the question that Roger [Burkhart] asked. To some extent, if you show that your model is very robust for a particular parameter, you are showing that that parameter does not matter, at least within your model. You have come up with an agent that can learn to deal with the fact that it may have individual variation.

So, what you really want is to identify what variables you cannot do that for, because that's just a critical attribute that could actually start explaining some of the data. It may be that sometimes when you are not robust to a variable, that it is extremely important information. That might have confused me, and until you asked me that question, I had not figured out what was bothering me.

Goadrich: That seems to make sense. If you are exploring and trying to identify the correct model for a situation, you might want to focus on parameters that have some give-and-take. But in cases where you know what's going on, you want your model to apply in many different situations. That was the focus of this work. If you are trying to discover things, however, you want to look at the interesting parameters, the ones you are talking about.

My model is a good explanation for fairness. How do I go about doing that? One method that was proposed was to test across all different parameters. Perhaps that is not the situation, though, because many models fail when you do that. If you focus on the likely parameters, or what the situation is now, how does it change things?

Jesse Voss: First, I want to say … try to interpret your present modeling in terms of that call, and second, would you agree that a conclusion that Idee's discussion of Body 1, Body 2, and the dimension of technology he uses in his recent book would be appropriate tools for visualizing what you are trying to do? If I am correct, you are trying to create a new scientific visualization to describe the AIDS problem in Africa.

Unidentified Speaker: I'm sorry. I missed part of your second question. Could you repeat that part, that is, before the new visualization?

Voss: Don Idee has come out with another book in which he describes Body 1 in the introduction. The Body 1 concept involves our experiences of the body that come through virtue of thought; our biological echo system in the Body 2 concept is the dimension that's possibly culturally constrained. According to Idee, in one culture, a person would be aroused by looking at a … and in another culture by another body part, so those body responses would be another type of thing. The third dimension is that of technology … or language or those types of things that connect. I was wondering if it would be appropriate to interpret your present work as a …

Unidentified Speaker: So, does the first part address the hermeneutic methodology, design methodology, and so forth?

I think that the place that that would come in is less, for example, on the simple model, less on the basic migration and things like that, and more on the things that we are trying to move toward, for example, the cultural effect, where different priorities, different interpretations, and different meanings are possible. These meanings might evolve endogenously, so I would not say that we have used a tremendous amount of that methodology in the first phase, but I think it is present in the horizon.

As for the second part, it is not directly influenced by the Idee work that you cite, but I do think that something is present here that we all face, because we are in a very qualitative pre-Newtonian, social science mode, and we are trying to account for the tremendous range of diversity and flux of social settings. Therefore, one way to look at it is moving toward abstraction and determining the appropriate level of abstraction. Another way of looking at it, though, is to move in more situated frameworks. And that is what I see: a body-centric framework is one form that “situatedness” takes, just like temporal situatedness, spatial situatedness, and so forth. If you push that down to the point, we arrive at the question of how we relate to our physiology — the role of hunger, the role of thirst, the role of sexual attraction, and so forth. We are physically grounded beings.

But that doesn’t answer the question. We know that. In many different fields, this is becoming more and more important, but it does not answer the important modeling question, which is, “Yes, okay, we’re physiologically grounded people, and we need to take that into account, but what’s the right modeling level to take that into account?”

On a personal level, I am most interested in the interaction between emotion and cognition, the way in which emotion drives cognition; the way in which the cognitive framework will adapt itself to emotion. Sometimes, though, there is a reciprocal effect, and a new fact can actually begin to shift the emotional balance. But that’s just my focus. Other people may be more interested in the neurological side of it. A great deal of very interesting research is going on in that realm, and it seems that Idee’s approach is a little different.

Unidentified Speaker: That was another interesting suggestion posited by of a couple of the papers. That is, that the situatedness may actually help provide the context as we build these multiple mechanisms, both for the conditions in which they apply and the parameterization by which we correlate them. If we are going to do run-time recommenders of these or possibly some of the structures, we will need to provide these multi-mechanism models.

Toolkits and Techniques

MASON: A JAVA MULTI-AGENT SIMULATION LIBRARY

S. LUKE,* Department of Computer Science
G.C. BALAN, Department of Computer Science
L. PANAIT, Department of Computer Science
C. CIOFFI-REVILLA,* Center for Social Complexity
S. PAUS, Department of Computer Science
George Mason University

ABSTRACT

Agent-based modeling (ABM) has transformed social science research by allowing researchers to replicate or generate the emergence of empirically complex social phenomena from a set of relatively simple agent-based rules at the micro level. Swarm, Repast, Ascape, and others currently provide simulation environments for ABM social science research. Since the development of Swarm, arguably the first widely used ABM simulator employed in the social sciences, subsequent simulators have sought to enhance available simulation tools and computational capabilities by providing additional functionalities and formal modeling facilities. Our system, called MASON (Multi-Agent Simulator of Neighborhoods), follows in a similar tradition that seeks to enhance the power and diversity of the available scientific toolkit in computational social science. MASON provides a core of facilities useful not only to social science but also to other ABM fields, such as artificial intelligence and robotics. This flexibility can foster useful "cross-pollination" between such diverse disciplines. Furthermore, MASON's additional facilities will become increasingly important as social complexity simulation matures and grows into new approaches. The new MASON simulation library is illustrated with a replication of HeatBugs, and a demonstration of MASON is applied to two challenging case studies: ant-like foragers and micro-aerial agents. Other applications are also being developed. The HeatBugs replication and the two new applications give an idea of MASON's potential for computational social science and artificial societies.

Keywords: MASON, agent-based modeling, multi-agent social simulation, ant foraging, aerial-vehicle flight

1 INTRODUCTION

Agent-based modeling (ABM) in the social sciences is a productive and innovative frontier for understanding complex social systems (Berry, et al., 2002). Object-oriented programming from computer science allows social scientists to model social phenomena directly in terms of social entities and their interactions in ways that are inaccessible through either statistical or mathematical modeling in closed form (Axtell and Epstein, 1996; Axelrod, 1997;

---

* *Corresponding authors' addresses*: Sean Luke, Department of Computer Science, Evolutionary Computation Laboratory, George Mason University, 4400 University Drive MSN 4A5, Fairfax, VA 22030; e-mail: sean@cs.gmu.edu. Claudio Cioffi-Revilla, Center for Social Complexity, George Mason University, 4400 University Drive MSN 3F4, Fairfax, VA 22030; e-mail: ccioffi@gmu.edu.

Gilbert and Troitzsch, 1999). The multi-agent simulation environments developed in recent years are designed to meet the needs of a particular discipline:

- TeamBots (Balch, 1998) and Player/Stage (Gerkey, et al., 2003) emphasize robotics.

- StarLogo (Massachusetts Institute of Technology, 2002) is geared toward education.

- breve (Klein, 2002) specializes in physical modeling and artificial life.

- Repast (University of Chicago, 2003), Ascape (Brookings Institution, 2003), and Swarm (Swarm Development Group, 2003) have traditionally emphasized social complexity scenarios with discrete or network-based environments.

Social science ABM applications-based environments in this final category are well documented in earlier proceedings of this conference (Macal and Sallach, 2000; Sallach and Wolsko, 2001) and have contributed substantial new knowledge in numerous domains of the social sciences, including anthropology (hunter-gatherer societies and prehistory), economics (finance), sociology (organizations and collective behavior), political science (government and conflict), and linguistics (emergence of language) — to name only a few examples.

We present MASON, a new Multi-Agent Simulator of Neighborhoods developed at George Mason University as a joint collaborative project between the Department of Computer Science's Evolutionary Computation Laboratory and the Center for Social Complexity. MASON seeks to continue the tradition of improvements and innovations initiated by Swarm. Because it is a more general system, however, MASON can also support core simulation computations outside the human and social domain in a strict sense. More specifically, MASON is a general-purpose, single-process, discrete-event simulation library intended to support diverse multi-agent models across the social and other sciences, artificial intelligence, and robotics, ranging from three-dimensional continuous models, to social complexity networks, to discretized foraging algorithms based on evolutionary computation. MASON is of special interest to the social sciences and social insect algorithm community because one of its primary design goals is to support very large numbers of agents efficiently. As such, MASON is faster than scripted systems such as StarLogo or breve, while still remaining portable and producing guaranteed replicable results. Another MASON design goal is to make it easy to build a wide variety of multi-agent simulation environments (for example, to test machine learning and artificial intelligence algorithms or to cross-implement for validation purposes), rather than provide a domain-specific framework.

This paper is organized as follows. Section 2 describes the new MASON environment in greater detail, including our motivation for creating MASON, and its main features and modules. Section 3 argues for MASON's applicability to social complexity simulation, including a comparison with Repast and a simple case study replication of HeatBugs (a common Swarm-inspired ABM widely familiar to computational social scientists). Section 4 presents two additional case studies of MASON applied to areas somewhat outside of the computational social science realm, but which point in directions of interest to the field in the future. Section 5 provides a brief summary.

2 MASON

2.1 Why MASON? History and Justification

MASON originated as a small library for a very wide range of multi-agent simulation needs, from robotics to game agents to social and physical models. The impetus for further development of MASON stemmed from our needs as the original architects of the system (Luke, Balan, and Panait). As computer scientists, we specialize in artificial intelligence, machine learning, and multi-agent behaviors. We needed a system in which to apply these methods to a wide variety of multi-agent problems. Previously, various robotics and social agent simulators were used for this purpose (notably TeamBots). Domain-specific simulators tend to be complex, however, and can lead to unexpected bugs if modified for use in domains for which they are not designed.

Our approach provides the intersection of features needed for most multi-agent problem domains, rather than the union of them, and makes it as easy as possible for the designer to add increased domain functionality. This “additive” approach to simulation development is less prone to problems than the “subtractive” method of modifying an existing domain-specific simulation environment. As such, MASON is intentionally simple, but highly flexible.

Machine learning methods, optimization, and other techniques are also expensive, requiring a large number of simulation runs to achieve good results. Thus, we needed a system that ran efficiently on back-end machines (such as Beowulf clusters), while the results were visualized, often in the middle of a run, on a front-end workstation. Because simulations might take a long time, we further needed built-in checkpointing to disk so we could stop a simulation at any point and restart it later.

Finally, our needs tended toward parallelism in the form of many simultaneous simulation runs, rather than one large simulation spread across multiple machines. Thus, MASON is a single-process library intended to run on one machine at a time.

While MASON was not conceived originally for the social agents community, we believe it will prove a useful tool for social agent simulation designers, especially as computational social science matures and grows into new approaches that require functionalities such as those implemented by the MASON environment. MASON’s basic functionality has considerable overlap with Ascape and Repast, partially to facilitate new applications as well as replications of earlier models in Swarm, Repast or Ascape; indeed, we think that developers accustomed to these simulators will find MASON’s architecture strikingly familiar. Finally, our motivation also includes the need to replicate simulation results as an essential strategy in advancing computationally based claims (Cioffi-Revilla, 2002), similar to the role of replication in empirical studies (Altman, et al., 2001).

2.2 Features

MASON was conceived as a core library around which one might build a domain-specific custom simulation library, rather than as a full-fledged simulation environment. Custom simulation library “flavors” might include robotics simulation library tools, graphics and physical modeling tools, or interactive simulator environments. MASON also provides enough simulation tools that it is quite usable as a basic “vanilla” flavor library in and of itself; indeed, the applications described later in this paper use plain MASON without any particular simulator flavor wrapped around it.

To achieve the flavors concept, MASON is highly modular, with an explicit layered architecture: inner layers have no ties to outer layers whatsoever, and outer layers may be completely removed. In some cases, outer layers can be removed or added to the simulation dynamically during a simulation run. We envision at least five layers: a set of basic *utilities*, the *core model library*, provided *visualization toolkits*, additional custom simulation layers (flavors), and the simulation applications using the library. These layers are shown in Figure 1.

Two additional MASON design goals are portability and guaranteed replicability. Replicability means that for a given initial setting, the system should produce identical results regardless of the platform on which it is running, and whether or not it is being visualized. Replicability and portability are crucial features of a high-quality scientific simulation system because they guarantee the ability to disseminate simulation results not only in publication form, but also in repeatable code form. To meet these goals, MASON is written totally in Java.

Java’s serialization facilities and MASON’s complete divorcing of model from visualization permit the model to easily perform checkpointing; at any time, the model can be serialized to the disk and reloaded. As shown in Figure 2, models can be checkpointed and loaded with or without visualization. In addition, serialized data can be reused on any Java platform. For example, one can freely checkpoint a model from a back-end Intel platform running Linux, then load and visualize its current running state on Mac OS X.

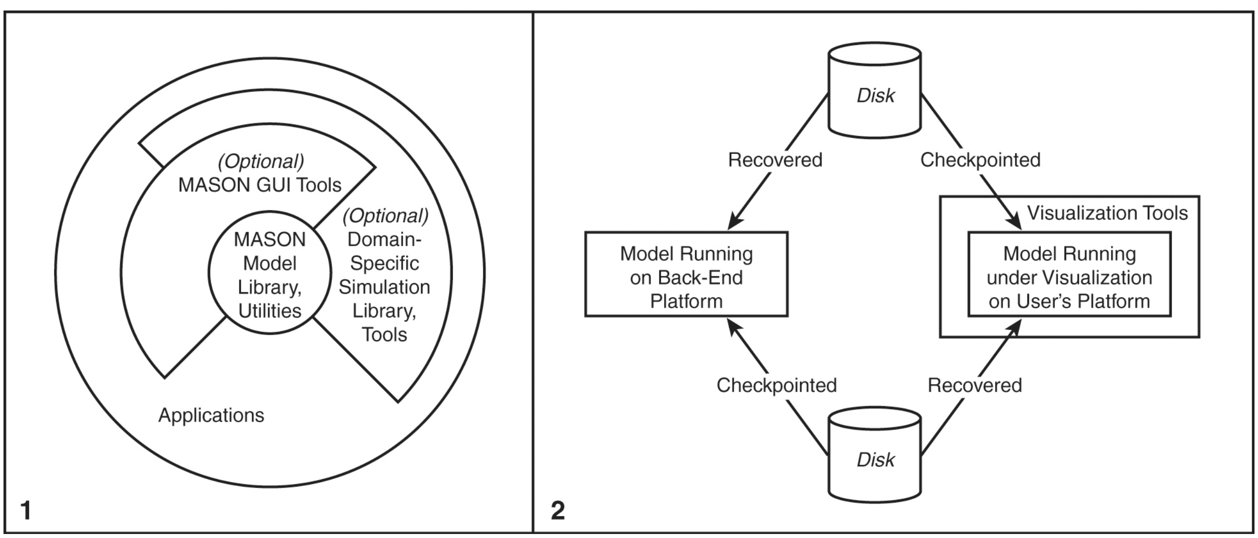

FIGURES 1 and 2 (1) MASON layers and (2) checkpointing architecture

Despite its Java roots, MASON is also intended to be fast, particularly when running without visualization. The core model library encourages direct manipulation of model data, is designed to avoid thread synchronization wherever possible, has carefully tuned visualization facilities, and is built on top of a set of utility classes optimized for modern Java virtual machines.[1] Although MASON is a single-process, discrete-event library, it still permits multi-threaded execution in certain circumstances, primarily to parallelize expensive operations in a given simulation.

2.3 Model and Utilities Layers

MASON's model layer, shown in Figure 3, consists of two parts: *fields* and a discrete-event *schedule*. Fields store arbitrary objects and relate them to locations in some spatial neighborhood. Objects are free to belong to multiple fields or, in some cases, to the same field multiple times. The schedule represents time and permits agents to perform actions in the future. A basic simulation model typically consists of one or more fields, a schedule, and user-defined auxiliary objects. There is some discrepancy in the use of the term *agents* between social sciences and computer sciences. We refer to agents as entities that can manipulate the world in some way: they are brains rather than bodies. Agents are very often *embodied* — physically located in fields along with other objects — but are not required to be so.

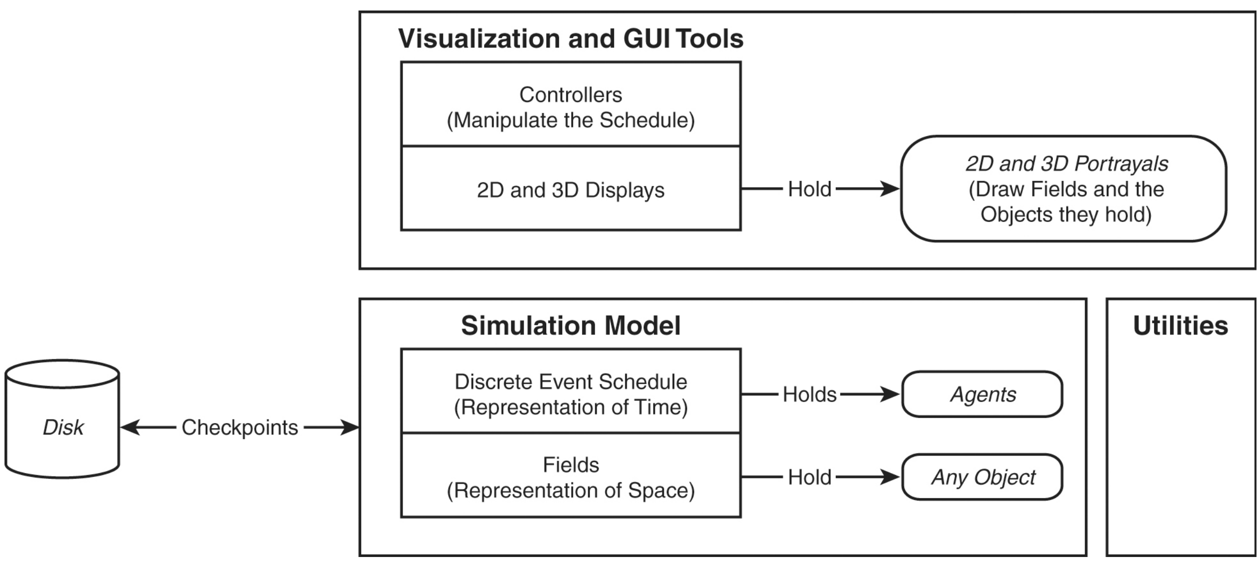

FIGURE 3 MASON utilities, model, and visualization layers

1 One efficiency optimization issue not settled yet is whether to use Java-standard multi-dimensional arrays or to use so-called "linearized" array classes (such as used in Repast). MASON has been implemented with both of them for testing purposes. In tight-loop microbenchmarks, linearized arrays are somewhat faster; but in full MASON simulation applications, Java arrays appear to be significantly faster. This is likely due to a loss in cache and basic-block optimization in real applications as opposed to simple microbenchmarks. We are still investigating this issue.

The model layer comes with fields providing the following spatial relationships, but other fields can be created easily:

- Bounded and toroidal discrete grids in 2D and in 3D for integers, doubles, and arbitrary objects (one integer/double/object per grid location)
- Bounded and toroidal hexagonal grids in 2D for integers, doubles, and arbitrary objects (one integer/double/object per grid location)
- Efficient sparse bounded, unbounded, and toroidal discrete grids in 2D and 3D (mapping zero or more objects to a given grid location)
- Efficient sparse bounded, unbounded, and toroidal continuous space in 2D and 3D (mapping zero or more objects to a real-valued location in space)
- Binary-directed graphs or networks (a set of objects plus an arbitrary binary relation)

When combined with certain classes of the utilities layer, models can run by themselves. They can be launched from the command line with no visualization or graphical user interface (GUI) code attached.

The utilities layer consists of Java classes free of simulation-specific function. Such classes include *bags* (highly optimized Java collection subclasses designed to permit direct access to integer, double, and object array data), immutable 2D and 3D vectors, and a highly efficient implementation of the Mersenne Twister random number generator.

2.4 Visualization Layer

As noted earlier, MASON simulations can operate either with or without a GUI and switch between the two modes in the middle of a simulation run. To achieve this, the model layer is kept completely separate from the visualization layer. When operated without a GUI, the model layer runs in the main Java thread as an ordinary Java application. When run with a GUI, the model layer is kept essentially in its own “sandbox;” it runs in its own thread, with no relationship to the GUI and can be swapped in and out at any time. Besides the checkpointing advantages described earlier, another important and desirable benefit of MASON’s separation of model from visualization is that the same model objects may be visualized in radically different ways at the same time (in both 2D and 3D, for example). The visualization layer, and its relationship to the model layer, is shown in Figure 3.

To perform the feat of separation, the GUI manages its own separate auxiliary schedule tied to the underlying schedule, queuing visualization agents that update the GUI displays. The schedule and auxiliary schedule are stepped through a controller in charge of running the simulation. The GUI does not display or manipulate the model directly, but through *portrayals* that act as proxies for the objects and fields in the model layer. Objects in the model proper may act as their own portrayals but do not have to.

The portrayal architecture is divided into various *simple portrayals* and *field portrayals*. Simple portrayals are stored in a field portrayal and used to portray various objects in the field portrayal's underlying field. Field portrayals are, in turn, attached to a display, which provides a GUI environment for them to draw and manipulate their fields and field objects. Portrayals can also provide auxiliary objects known as *inspectors* (approximately equivalent to "probes" in Repast and Swarm) that permit the examination and manipulation of basic model data.

MASON provides displays and portrayals for both 2D and 3D space and can display all of its provided fields in 2D and 3D, including displaying certain 2D fields in 3D. Two-dimensional portrayals are displayed by using the Abstract Windowing Toolkit and Java2D graphics primitives. Three-dimensional portrayals are displayed by using the Java3D scene graph library. Examples of these portrayals are shown in Figure 4 in Section 3.2.

3 APPLICABILITY TO SOCIAL COMPLEXITY ENVIRONMENTS

MASON was designed with an eye toward social agent models, which may be of value to social science researchers. MASON shares many core features with social agent simulators, such as Swarm, Ascape, and Repast. This section specifies the primary differences between MASON and Repast, followed by a simple example in which MASON is used to simulate the well-known HeatBugs model.

3.1 Comparison with Repast

We provide a brief enumeration of most of the differences between the facilities provided by MASON and those of Repast; the latter evolved from Swarm to model situated social agents.

3.1.1 Differences

- One of the key differences between MASON and Repast is that MASON provides a full division between model and visualization. As a result, MASON can either separate or join the two at any time and easily provide cross-platform checkpointing. In addition, MASON objects and fields can be portrayed in radically different ways at the same time, and visualization methods can be changed even during an expensive simulation run.

- MASON has facilities for 3D models and other visualization capabilities that remain largely unexplored in the social science realm, but that are potentially insightful for social science ABM simulations.

- In our experience, MASON generally has faster models and visualization than Repast, especially on Mac OS X; it also has more memory-efficient sparse and continuous fields. MASON's model data structures have computational complexity advantages.

- MASON has a clean, unified approach for handling network and continuous-field visualization.

- Repast provides many facilities, notably, a geographic information system, Excel import/export, charts and graphs, and SimBuilder and related tools. Because of its design philosophy, MASON does not include these facilities. We believe they are better provided as separate packages rather than bundled. Furthermore, many of these tools can be trivially ported to MASON.

3.1.2 Differences in Flux[2]

- Repast uses linearized array classes for multi-dimensional arrays. MASON currently has facilities for both linearized arrays and true Java arrays but may reduce to using one or the other.

- Repast's schedule uses doubles, whereas MASON's schedule presently uses longs.

- Repast allows objects to be selected and moved by the mouse.

- Repast provides deep inspection of objects; MASON's inspection is at present shallow.

3.2 Replicating HeatBugs

HeatBugs is arguably the best known ABM simulation introduced by Swarm and is also a standard demonstration application in Repast. It contains basic features common to many social agent simulations, for example, a discrete environment defining neighborhood relationships among agents, residual effects (heat) of agents, and interactions among them. The ability to replicate models like HeatBugs, Sugarscape, Conway's Game of Life (or other cellular automata), and Schelling's segregation model in a new computational ABM environment should be as essential as the ability to implement regression, factor analysis, ANOVA, and similar basic data facilities in a statistical analysis environment.

Indeed, a 100×100 toroidal world, 100-agent HeatBugs model was MASON's very first application. In addition to this classic HeatBugs model, we implemented several other HeatBugs examples. Figure 4 includes partial screenshots of two of them. Figure 4b shows HeatBugs on a hexagonal grid (fittingly called "HexaBugs" in Repast). Figure 4f shows 2D HeatBugs visualized in 3D space, where vertical scale indicates temperature; HeatBugs on the same square are also shown stacked vertically. Whereas the original HeatBugs model is based on a 2D grid of interacting square cells (connected by Moore or von Neumann neighborhoods), HexaBugs is more relevant in some areas of computational social science where hexagonal cells are more natural (e.g., computational political science, especially international relations) and four-corner situations are rare or nonexistent (Cioffi-Revilla and Gotts, 2003).

---

2 These features will probably be changed in the final version of MASON.

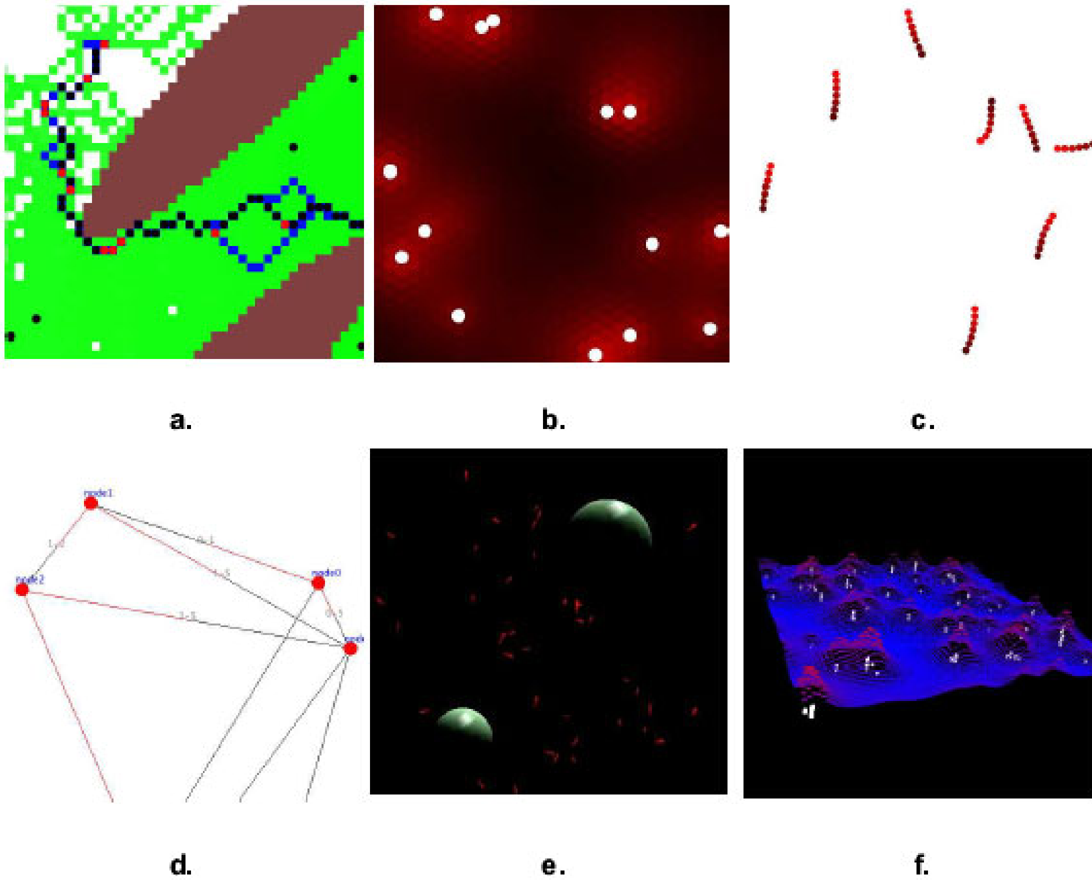

FIGURE 4 Sample field portrayals (applications in parentheses): a. discrete 2D grids (ant foraging), b. hexagonal 2D grids (hexagonal HeatBugs), c. continuous 2D space (“Woims:” flocking worms), d. networks in 2D (a network test), e. continuous 3D space (3D “Woims”), and f. discrete 2D grids in 3D space (HeatBugs)

4 CASE STUDIES

Although it has been in use for only six months, MASON has already been used in a variety of research and educational contexts. We are also conducting tests to port Repast, Swarm, and Ascape models to MASON, by modelers not immediately familiar with MASON. These ports include a model of warfare among countries, a model of land use in a geographic region, and a model of the spread of anthrax in the human body.

This section describes the implementation and results of two research projects that used the MASON simulation library. In the first case, MASON was used to discover new ant colony foraging and optimization algorithms. In the second case, MASON was applied to the development of evolved micro-aerial vehicle flight behaviors. These are not computational social science models per se, but they are relevant enough to prove illuminating. The first case uses a model that is similar to the discrete ABM models presently used, but it is applied to an automated learning method, which demonstrates the automated application of large numbers of simulations in parallel. The second case uses a continuous 2D domain environment and interaction, which points to a future area for ABM research. Neither of these more advanced applications is currently implemented in Swarm, Repast, or Ascape; both take advantage of features specific to MASON. In both cases, experiments were conducted running MASON on the command-line in several back-end machines; progress was analyzed by attaching the simulators to visualization tools on a front-end workstation. In addition, the second case involved a continuous, scalable field that is both memory- and time-efficient (both *O*[#agents], rather than *O*[spatial area]).

The projects described in Sections 4.1 and 4.2 both have an evolutionary computation (EC) component. To save repetition, we provide a brief explanation of evolutionary computation. EC is a family of stochastic search and optimization techniques for “hard” problems for which there is no known procedural optimization or solution-discovery method. EC is of special interest to certain multi-agent fields because it is *agent oriented*: it operates not by modifying a single candidate solution, but by testing a “population” of such solutions all at one time. Such candidate solutions are known as “individuals,” and each individual’s assessed quality is known as its “fitness.”

The general EC algorithm is as follows. First, an initial population of randomly generated individuals is created and each individual’s fitness is assessed. Second, a new population of individuals (the next generation) is assembled through an iterative process of stochastically selecting individuals (tending to select the ones who are most fit), copying them, breeding the copies (mixing and matching individuals’ components and mutating them), and then placing the results into the next generation. The new generation replaces the old generation; its individuals’ fitnesses are in turn assessed, and the cycle continues. EC ends when a sufficiently fit individual is discovered, or when resources (notably time) expire. The most famous example of EC is the genetic algorithm (Holland, 1975), but other versions exist as well. We discuss genetic programming (Koza, 1992) as one alternative EC method below.

4.1 Ant Foraging

Ant foraging models attempt to explain how ant colonies discover food sources and then communicate those discoveries to other ants by leaving pheromone trails, the proverbial “bread crumbs” to mark the way. This area has become popular not just in biology but curiously, in artificial intelligence and machine learning because pheromone-based communication has proved to be an effective abstract notion for new optimization algorithms (known collectively as *ant colony optimization*) and cooperative robotics.

Previous ant foraging models have relied to some degree on *a priori* knowledge of the environment, in the form of explicit gradients generated by the nest, by hard-coding the nest location in an easily discoverable place, or by imbuing the ants with the knowledge of the nest direction. In contrast, the case study presented here solves ant foraging problems by using two pheromones — one applied when leaving the nest and one applied when returning to the nest. The resulting algorithm is orthogonal, simple, and biologically plausible, yet ants are able to establish increasingly efficient trails from the nest to the food even in the presence of obstacles.

Ants are sensitive to one of the two pheromones at any given time; the sensitivity depends on whether they are foraging or carrying food. While foraging, an ant stochastically moves in the direction of increasing food pheromone concentration and deposits some amount of nest pheromone. If there is already more nest pheromone than the desired level, the ant deposits nothing. Otherwise, the ant “tops off” the pheromone value in the area to the desired level. As the ant wanders from the nest, its desired level of nest pheromone drops. This decrease in deposited pheromone establishes an effective gradient. When the ant carries food, the movement and pheromone-laying procedures use the pheromones opposite those used during foraging.

The model assumes a maximum number of ants per location in space. At each time step, an ant moves to its best choice among nonfull, nonobstacle locations; the decision is made stochastically with probabilities correlated to the amounts of pheromones in the nearby locations. Ants move in random order. Ants live for 500 time steps; a new ant is born at the nest at each time step unless the total number of ants is at its limit. Pheromones both evaporate and diffuse in the environment.

Figure 4a shows a partial screenshot of a small portion of the ant colony foraging environment. The ants have laid down a path from the nest to the food and back again. Part of the ground is colored with pheromones. The large oval regions are obstacles. The MASON implementation was done with two discrete grids of doubles (two pheromone values); discrete grids of obstacles, food sources, and ant nests; and a sparse discrete grid holding the ants proper. Each ant is also an agent (and so is scheduled at each time step to move itself). Additional agents are responsible for the evaporation and diffusion of pheromones in the environment and for creating new ants when necessary.

In addition to successfully designing hard-coded ant foraging behaviors, we also experimented with allowing the computer to optimize those behaviors. For this purpose, we connected MASON to the ECJ (evolutionary computation in Java) system (Luke, 2000); ECJ handled the main evolutionary loop. An individual took the form of a set of ant behaviors applied to every ant in the colony. To evaluate an individual, ECJ spawned a MASON simulation with the specified ant behaviors. The simulation was run for several hundred time steps. At the end of the simulation, the amount of food foraged indicated the individual’s fitness.

To evolve ant behaviors, we used *genetic programming* (Koza, 1992). In genetic programming, individuals are actually computer programs in the form of one or more parse trees. We do not describe parse trees here, except to explain that breeding consisted of swapping subtrees among individuals. Our EC individuals (the behaviors) consisted of two such genetic programming trees. The execution of one tree returned the amount of pheromone to deposit, and the execution of the other tree yielded the direction to move. The same behavior was used for both foraging and food-carrying states, but the pheromones specified in the behaviors (food pheromone vs. nest pheromone) were different for each state.

A first experiment scaled the number of ants (50, 50, 500), the number of simulation time steps (501, 1001, 2501), and the world size (10×10, 33×33, 100×100). In each case, the EC populations converged rapidly to simple but reasonably high-performing ant foraging behaviors. Increasing the world size led to longer convergence times (from a mere two generations in the 10×10 case to ten generations on average in the 100×100 case). Interestingly, these behaviors differed from one another in meaningful ways. When the three highest-performing behaviors were compared, the results showed that more difficult domains led to the discovery of more robust foraging strategies. Additional details on this work can be found in Panait and Luke (2003a,b).

4.2 Micro-aerial Vehicle Simulation

An unmanned aerial vehicle (UAV) is a flying device, often an airplane, operated by remote control usually for military functions (such as surveillance, reconnaissance, or attack). The UAV most familiar to the general public is the Predator, a flying drone by General Atomics. The Predator has flown surveillance missions over Afghanistan and Iraq. Large UAVs such as the Predator are expensive to produce; moreover, even though they have no on-board pilot, UAVs require a large team of controllers on the ground to fly the vehicle. One recent thrust in UAVs has been the micro-aerial vehicle (MAV), a tiny (less than 1 meter), inexpensive UAV primarily intended for surveillance. Because they are inexpensive, MAVs are often designed to fly in "swarms" of up to hundreds of vehicles. Such swarms mean that a unique human controller cannot feasibly be allocated for each MAV. Instead, it is hoped that an entire MAV swarm can be controlled by a small team of controllers. To achieve this, MAVs must be *semiautonomous*; they receive high-level commands from human controllers, but most of the work is achieved by the MAVs themselves. Figure 5 shows an MAV simulation in MASON.

The University of Central Florida (UCF) and George Mason University recently worked on a joint project with the Defense Advanced Research Projects Agency (DARPA) demonstrating the feasibility of having swarms of MAVs learn behaviors via simulated evolution. The research system that was developed combined an evolutionary computation system and a library of dominance hierarchies developed at UCF with a MASON simulation environment. The system was recently completed; published results are forthcoming.

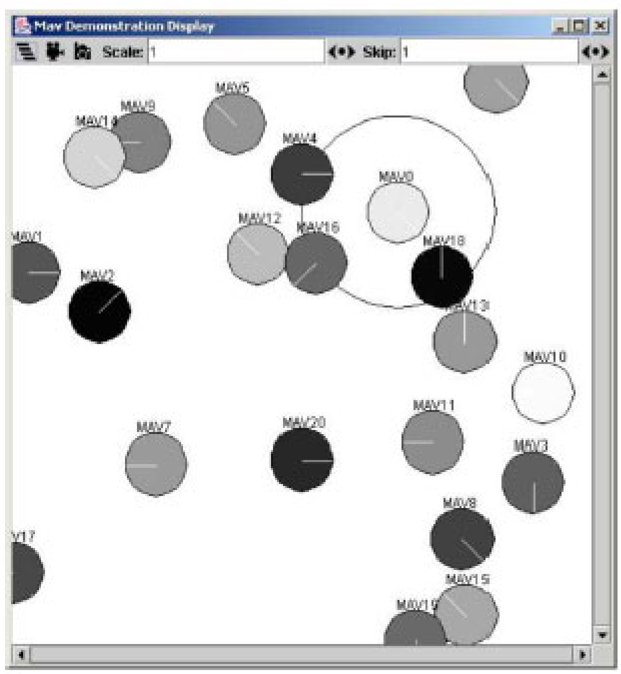

FIGURE 5 Micro-aerial vehicle simulation in MASON (Vehicles [circles] appear much larger than they actually are. Gray values indicate level of dominance, and lines indicate orientation of each vehicle.)

The MASON simulation held MAVs in a continuous 2D neighborhood along with *regions*, colored shapes “painted” on the ground, over which the MAVs would fly. Our goal was to develop MAV behaviors that caused them to fly over specific colored regions as much and as often as possible while avoiding collisions with one another. If MAVs collided, they were removed from the simulation.

Our MAV behaviors consisted of sets of basic *sensor values*→*action* rules: a rule might say, for example, that if the MAV were directly above the appropriate region color, and there was another nearby MAV to the upper left, then the MAV should turn right. An additional “sensor” available to the rules was the dominance of the MAV relative to its neighbors; nearby MAVs established dominance hierarchies among themselves by using a method developed by Tomlinson (2002).

After some number of time steps, the swarm quality of the MAV was assessed by adding up the total time that each MAV was located over a region of interest. We applied an evolutionary computation system to learn MAV behaviors that, when used by an MAV swarm in the simulator, produced the highest quality assessments possible.

The system was constructed in MASON as follows. Each MAV was an embodied MASON agent and was stored in a continuous 2D field. Regions were stored in a second continuous 2D field. Each MAV held a ring of eight “sonar sensors” (rays emanating in eight directions from the MAV). At each time step, each MAV called the provided dominance library to update its dominance values based on the relative values of nearby MAVs. It then determined the distance to the closest MAV that intersected each sonar ray. These eight distance values, plus the value indicating the color of the region presently below the MAV, plus the current dominance value of the MAV, formed the MAV’s 10 sensor values. The MAV then determined which rule in its rule set most closely matched its current sensor values and performed that action. An action consisted of 1 of 8 directions in which the MAV could turn. After turning in that direction, the MAV moved forward some distance. If it then collided with other MAVs, they were all eliminated from the simulation.

Once again, MASON was used as a subsidiary process to an evolutionary computation system, this time one devised by Prof. Annie Wu at the UCF. The individuals (the MAV behaviors) were represented in a genetic algorithm as vectors of numbers indicating the direction to fly given various sensor values. An individual’s fitness was assessed by creating an MAV simulation in MASON, plugging the behavior into the MAVs in the model, running the model for some *N* time steps, then assessing the total time that MAVs stayed over appropriate target regions.

5 SUMMARY

Agent-based modeling has already begun to transform social science research — “the third way of doing science” (Axelrod, 1997) — by allowing researchers to replicate or generate the emergence of empirically complex social phenomena from a set of relatively simple agent-based rules at the micro level. One of the keys to this transformation has been object-oriented modeling (Gulyás, 2002), which moves beyond models in closed form (Taber and Timpone, 1996). Swarm, Repast, Ascape, and other simulation environments already provide numerous capabilities for ABM social science research. Since the development of Swarm, arguably the first widely utilized ABM simulator employed in the social sciences, subsequent simulators have sought to enhance available simulation tools and computational capabilities by providing additional functionalities and formal modeling facilities. In this paper, we present MASON (Multi-Agent Simulator of Neighborhoods), which follows in a similar tradition and seeks to enhance the power and diversity of the available scientific toolkit in computational social science, artificial intelligence, and other multi-agent areas. We argue that besides its immediate use in conducting social complexity simulations, MASON provides a general framework to serve as a core for a wide range of multi-agent needs, many of which will become increasingly important as social complexity simulation matures into new approaches. We illustrate the new MASON simulation library with a replication of HeatBugs and a demonstration of two challenging MASON applications as case studies: ant-like foragers and micro-aerial vehicles. Other applications are also being developed to demonstrate and enhance MASON's features. The HeatBugs replication and the two new applications provide an idea of MASON's potential for computational social science and artificial societies.

6 ACKNOWLEDGMENTS

Initial development of the MASON project was made possible through joint funding by the Evolutionary Computation Laboratory and the Center for Social Complexity of George Mason University. The micro-aerial vehicle simulation work was funded in part by DARPA via a U.S. Department of the Interior award, grant number NBCH1030012.

7 REFERENCES

Altman, M., L. Andreev, M. Diggory, G. King, S. Verba, D.L. Kiskis, and M. Krot, 2001, "A Digital Library for the Dissemination and Replication of Quantitative Social Science Research: The Virtual Data Center," *Social Science Computer Review* **19**(4).

Axelrod, R., 1997, "Advancing the Art of Simulation in the Social Sciences," *Complexity* **3**(2):193–199.

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

Balch, T., 1998, *Behavioral Diversity in Learning Robot Teams*, Ph.D. Dissertation, College of Computing, Georgia Institute of Technology. For information on TeamBots, see http://www.teambots.org.

Berry, B.L., L.D. Kiel, and E. Elliott, eds., 2002, "Adaptive Agents, Intelligence and Emergent Human Organization: Capturing Complexity through Agent-based Modeling," Vol. 99, *Proceedings of the National Academy of Sciences*, National Academy of Sciences.

Brookings Institution, 2003, Ascape Home Page. For information on Ascape, see http://brook.edu/es/dynamics/models/ascape/default.htm.

Cioffi-Revilla, C., 2002, "Invariance and Universality in Social Agent-based Simulations," *Proceedings of the National Academy of Sciences of the USA* **99** (Supp. 3): 7314–7316.

Cioffi-Revilla, C., and N.M. Gotts, 2003, "Comparative Analysis of Agent-based Social Simulations: GeoSim and FEARLUS Models," *Journal of Artificial Societies and Social Systems* **6**(4).

Gerkey, B., R.T. Vaughan, and A. Howard, 2003, "The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems," *Proceedings of the 11th International Conference on Advanced Robotics*, pp. 317–323. For information on Player/Stage, see http://www.playerstage.sourceforge.net.

Gilbert, N., and K. Troitzsch, eds., 1999, *Simulation for the Social Scientist*, Buckingham and Philadelphia: Open University Press.

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

Holland, J., 1975, *Adaptation in Natural and Artificial Systems*, Cambridge, MA: MIT Press.

Klein, J., 2002, "breve: a 3D Simulation Environment for the Simulation of Decentralized Systems and Artificial Life," *Proceedings of Artificial Life VIII, the 8th International Conference on the Simulation and Synthesis of Living Systems*," Cambridge, MA: MIT Press. For information on breve, see http://www.spiderland.org.

Koza, J., 1992, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Cambridge, MA: MIT Press.

Luke, S., 2000, *Issues in Scaling Genetic Programming: Breeding Strategies, Tree Generation, and Code Bloat*, Ph.D. Dissertation, Department of Computer Science, University of Maryland, College Park. See http://cs.gmu.edu/~eclab.

Macal, C.M., and D. Sallach, eds., 2000, *Proceedings of the Workshop on Agent Simulation: Applications, Models, and Tools*, October 15–16, 1999, Chicago, IL: University of Chicago and Argonne National Laboratory.

Massachusetts Institute of Technology, 2002, StarLogo Home Page. For information on StarLogo, see http://education.mit.edu/starlogo.

Panait, L., and S. Luke, 2003a, "Ant Foraging Revisited," in T. Balch and C. Anderson, eds., *Proceedings of the 2nd International Workshop on the Mathematics and Algorithms of Social Insects*, December 15–17, 2003, Atlanta, GA: Georgia Institute of Technology.

Panait, L., and S. Luke, 2003b, "Evolving Foraging Behaviors, in T. Balch and C. Anderson, eds., *Proceedings of the 2nd International Workshop on the Mathematics and Algorithms of Social Insects*, December 15–17, 2003, Atlanta, GA: Georgia Institute of Technology.

Sallach, D., and T. Wolsko, eds., 2001, *Proceedings of the Workshop on Simulation of Social Agents: Architectures and Institutions*, October 6–7, 2000, Chicago, IL: University of Chicago and Argonne National Laboratory.

Swarm Development Group, 2003, Swarm Home Page. For information on Swarm, see http://www.swarm.org.

Taber, C.S., and R.J. Timpone, 1996, “Computational Modeling,” in M.S. Lewis-Beck, ed., *Series: Quantitative Applications in the Social Sciences*, Vol. 07-113, Newbury Park, CA: Sage Publications.

Tomlinson, W., 2002, *Synthetic Social Relationships for Computational Entities*, Ph.D. Dissertation, Media Arts & Sciences, Massachusetts Institute of Technology, Cambridge, MA. Available at http://web.media.mit.edu/~badger/pubs.html.

University of Chicago, 2003, Repast Home Page. For information on Repast, see http://repast.sourceforge.net/.

SIMULATION AND DISTRIBUTED ARCHITECTURE OF MULTI-AGENT–BASED BEHAVIORAL ECONOMIC LANDSCAPE (MABEL) MODEL WITHIN SWARM

Z. LEI,* Michigan State University
B.C. PIJANOWSKI, and K.T. ALEXANDRIDIS, Purdue University

ABSTRACT

The Multi-Agent–based Behavioral Economic Landscape (MABEL) model introduces a distributed modeling architecture framework that supports the simulation of land-use changes over time and space over large regions. The model is based on the Swarm modeling software package, which is supported by a unique client-server framework with multiple interfaces built around a geographic information system (GIS), statistical analysis and database (SPSS), and Bayesian network software. The model architecture supports an integrated simulation environment with remote data retrieval, distributed and parallel scenario simulations, centralized decision-making algorithms, graphic displays for both client and server model components, and analysis capabilities. On the client side of MABEL, computational agents represent Bayesian relations among geographic, environmental, human, and socio-economic variables, with respect to land-use changes occurring across landscapes. A multi-agent simulation environment is created within Swarm, which simulates the buying, selling, and keeping of land by different types of agents. Agents are allowed to participate in an abstract market model. The characteristics of the server side of MABEL include (1) remote data retrieval via multiple interfaces with GIS software (ArcGIS and Arcview) and statistical database software (SPSS), and (2) coordinated agent decision making that allows for decision requests of agents from clients to be made to centralized Bayesian network agent profiles located on the MABEL server.

Keywords: Multi-agent systems, MABEL, client-server framework, Swarm, Bayesian networks, land-use change

INTRODUCTION

Agent-based modeling (ABM) is a form of artificial intelligence simulation in which autonomous agents interact, communicate, evolve, learn, and make complex decisions within a real-time simulation framework (Holland, 1975). Multi-agent systems present a bottom-up approach to modeling individuals' artificial intelligence (Kohler and Gumerman, 2000). Multi-agent intelligent systems are constructed to represent and simulate problem-solving situations, where collaborative and conflict behaviors can co-occur as they do in real human and natural environments of our daily life (Murch and Johnson, 1999). Recently, ABM approaches have been applied to simulate land-use changes (Alexandridis, et al., 2003; Alexandridis and Pijanowski, 2002; Parker, et al., 2001).

---

* *Corresponding author address:* Zhen Lei, Department of Computer Science & Engineering, Michigan State University, East Lansing, MI 48824; e-mail: leizhen@cse.msu.edu.

Developing large-scale, multi-agent–based simulations that exist in a dynamic spatial context presents several technical challenges. First, simulations with large numbers of agents require high-end computational capabilities. Second, many of the current ABM environments lack robust modeling tools found in other software packages, such as a geographic information system (GIS), which allows researchers to manage and analyze spatial data. Integrating agent-based applications, such as Swarm, with other software applications presents many technical challenges to modelers who want to seamlessly integrate various software tools into one model. Finally, building models that can operate on several computers simultaneously requires the introduction of computer networking technologies that provide for the foundation of a distributed modeling environment.

The purpose of this paper is to introduce the distributed modeling architecture (DMA) of the Multi-agent–based Behavioral Economic Landscape (MABEL) model. The DMA is based on a client-server architecture that separates the specific simulation's scenarios in Swarm from agents' general decision-making and policy rule models stored in Microsoft Bayesian network databases. It bridges these two parts with an efficient decision-making and model results delivery mechanism. In this paper, we (1) explain how we develop the intelligent agents, which participate in a simple market model, and apply algorithms used for land transactions and land ownership partitioning in the MABEL client, and (2) describe the MABEL server infrastructure and the decision-reference processes that occurs between clients and the server containing Bayesian network agent profiles. We conclude by summarizing the main features of MABEL.

MABEL CLIENT

In the MABEL client, "base" agents own land, designated as parcels, on a landscape, the fundamental simulation environment. Land-use–based attributes are the main drivers of the simulation, and land-use–driven acquisition of land in a market model represents the basic framework for determining the actions of base agents. Currently, base agents in MABEL can be from any land-use category, such as farmer agents, resident agents, and forestry agents.

Each MABEL client represents a spatially defined area with various types of agents that simulate land-use changes over time and space. Like a person within a society, each agent makes decisions on the basis of information provided and interacts with other agents in the environment. MABEL client provides such an environment with specific policy rule regulations and communication interfaces with the MABEL server for agents to remote data acquisition and decision-making inference. Each MABEL client ideally represents an area that is under similar policy controls (i.e., a township in Michigan). Several modeling phases occur within the MABEL client.

MABEL Client: Initialization Phase

In the initialization phase, MABEL client first creates the simulation environment and the corresponding two-dimensional (2-D) geographic "world." The MABEL client then loads land-use, parcel, and socio-economic data for the specific area. Related data items in the land-use/ parcel and socio-economic data are linked together by an agent's parcel number, which represents and indexes a specific agent. Next, the MABEL client creates all agent objects with corresponding land-use/parcel attributes and socio-economic data organized by an assigned parcel number. The client then matches parcel locations to individuals in the socio-economic database through two variables in the Public Use Microdata Samples (PUMS) database that relate to parcel size and source of income (U.S. Bureau of the Census, 1995). Finally, each agent draws and updates itself in the geographic 2-D world and is ready to respond to user inquiries on GIS and socio-economic attributes associated with each parcel. The MABEL client can also be used to load the specific market model with policy rule models to control transactions among agents in the area.

MABEL Client: Multi-agent Interaction Phase

After initialization, each agent begins to act with other agents in the simulation area. During each time step, which can be predefined as a certain period, an agent runs its specific strategy based on the land-use type (e.g., farm) and other land-use/parcel and socio-economic data; agents then communicate with the server to inquire as to what the optimal transaction decision might be based on the Bayesian network agent profile. The decision requests that agents send an up-to-date state space of related GIS, human, and socio-economic variables, which are needed for the decision-making process in the corresponding Bayesian network profile database. The maximum reward decision received from the Bayesian network model includes action type (buy, sell, or keep/no action), action quantity, and appropriate agent types that match the transaction requirements. Finally, the set of agents that intends to make a transaction does so on the basis of most profitable deal within the market model within policy/rule regulations; agents then update their spatial/GIS/socio-economic attributes in the simulated world. The principle of a market model is to help a potential buyer agent make the most profitable deal with any corresponding seller agent. The degree of the profit in a transaction depends on how close the transaction quantity of the buyer agent meets that of the seller agent.

MABEL SERVER

The MABEL server acts as a bridge between the MABEL client, which represents agents in a specific area, and the external decision-making component stored in the belief network models in MSBNx (Microsoft Bayesian network). The MABEL server receives various decision requests from multiple MABEL clients and delivers them to different Bayesian network models for the decision inference using an agent's GIS/socio-economic attributes, such as land-use types. Finally, the MABEL server collects the resulting decision replies and sends them back to corresponding agents across different areas. To satisfy the reliability requirement of the communication between the MABEL client and MABEL server, a network connection is established with TCP/IP network protocols over the Internet.

For the design of the MABEL server, we use a multi-threaded technique to achieve parallel processing capability with high execution efficiency. Unlike the one thread in the single-threaded program, which executes tasks sequentially, each thread in multi-thread program executes only part of the task and synchronizes with other threads regarding the execution order of different parts in the task. In this way, a multi-threaded program can execute different tasks simultaneously with optimal execution efficiency. We explain the infrastructure of the MABEL server in two phases — the initialization phase and the execution phase.

MABEL Server: Initialization Phase

In the initialization phase, multiple proxy threads load their own configuration files and initialize the communication with corresponding external MSBNx proxy programs. Each external MSBNx proxy program is responsible for dealing with specific types of decision requests and communicates with a corresponding MSBN belief network model by calling routines in the MSBN3 API library. Each proxy thread in the MABEL server also has a specific working queue to buffer unprocessed decision requests. Every proxy thread watches its queue for incoming decision requests from specific types of agents that the proxy thread represents. When all of these initializations have been completed, the MABEL server is ready to receive the connect requests from MABEL clients.

MABEL Server: Execution Phase

For each client, at the beginning of each time step, all agents in a simulation area communicate with the server to inquire about optimal transaction decisions with the corresponding Bayesian network model. The decision requests that agents send include their up-to-date state space of related GIS, human, and socio-economic variables, which the Bayesian network model needs for the decision-making process.

Before the MABEL client sends agents' decision requests for inference, it first establishes the communication link with the MABEL server by a sending a connection request using the TCP/IP protocol. Once a connection request has been received from a MABEL client, the MABEL server assigns a communication link/socket for that client and launches two input/output (I/O) threads for the I/O operations with the client. The input thread is the request dispatcher thread, which is responsible for receiving different types of agents' decision requests from the MABEL client. The MABEL server then dispatches these requests, including the agent GIS/socio-economic data for inference, into corresponding working queues of the proxy threads by their GIS/socio-economic attributes. In addition, the request dispatcher thread attaches a client port number for every decision request, which represents the client/area information from which the request originated. Therefore, when the proxy thread finishes the decision-making process, it can send the results back to the corresponding client by this port number. On the other hand, the output thread — the decision collector thread — continues to wait for the decision results from proxy threads at the client port for a specific MABEL client; it then forwards the optimal decision to the MABEL client.

One of the main advantages of using a multi-thread technique is that the MABEL server executes different parts of tasks with different threads and synchronizes with each client about the execution order. Therefore, we can distribute and partition a large simulation task over different machines and coordinate the distributed working processes. Furthermore, the MABEL server can accommodate multiple clients simultaneously with optimal execution efficiency, which eases the work of result fusing and data analysis at the server.

SIMULATION

Using the MABEL environment, we simulated the land-use and transformation changes over time in different areas and scenarios within the State of Michigan. We ran the simulations on Traverse County in Long Lake Township and also in parts of Mecosta County. All simulated areas, which had various numbers of agents of different land-use types, are represented in different MABEL clients, which may run in different machines. All decision-making processes are routed through the MABEL server, which standardizes the inference interface with the Bayesian network agent profiles and provides server management and communication utilities.

CONCLUSION

This paper introduces a DMA framework that was used as part of the MABEL model. Some important issues are addressed as to how agent structure, a market model, and land partition strategy can be integrated within a client-server environment by using multi-threaded TCP/IP tools. We explain how the MABEL server acts as a bridge between specific simulation environments and general agent models interacting on clients.

The client-server architecture in the MABEL system allows simultaneous simulation of land-use change over large regions and does so efficiently. Moreover, the separation of simulation scenarios and agents' theoretic models simplifies the work of researchers and greatly eliminates the translation from the intent or conceptualization of a model to its implementation. Modelers need only to create, assess, and evaluate agents' theoretic decision-making models in their familiar model tool; the model builder can focus on organizing and scheduling the agents' activity in the specific simulation scenarios.

REFERENCES

Alexandridis, K.T., and B.C. Pijanowski, 2002, "Multi Agent-based Environmental Landscape (MABEL) — An Artificial Intelligence Simulation Model: Some Early Assessments," presented at the *AERE/EAERE: 2002 World Congress of Environmental and Resource Economists*, Monterey, CA, June 24–27, p. 26.

Alexandridis, K.T., B.C. Pijanowski, and Z. Lei, 2003, "Sequential Decision-making Process Simulation of Base-agents in a Multi-agent-based Economic Landscape (MABEL) Model," *Journal of Environmental & Resource Economics* (accepted for publication).

DeAngelis, D., L. Wolf, M. Mooij, M.P. Nott, and R.E. Bennetts, 2001, Individual-based Models: Tracking Variability among Individuals," in T.M. Shenk and A.B. Franklin, eds., *Modeling in Natural Resource Management: Development, Interpretation, and Application*, Washington, DC: Island Press, pp. 171–196.

Holland, J.H., 1975, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, Ann Arbor, MI: University of Michigan Press, pp. viii, 183.

Kohler, T.A., and G.J. Gumerman, eds., 2000, *Dynamics in Human and Primate Societies: Agent-based Modeling of Social and Spatial Processes, Santa Fe Institute Studies in the Sciences of Complexity*, New York City, NY: Oxford University Press, pp. xiii, 398.

Murch, R., and T. Johnson, 1999, *Intelligent Software Agents*, Upper Saddle River, NJ: Prentice Hall, pp. xiii, 210.

Parker, D.C., S.M. Manson, M.A. Janssen, M. Hoffman, and P. Deadman, 2001, “Multi-agent Systems for the Simulation of Land-use and Land-cover Change: A Review,” prepared for the *Special Workshop on Agent-based Models for Land Use/Land Cover Change*, Irvine, CA, October 4–7, 2001.

U.S. Bureau of the Census, 1995, *1990 Census of Population and Housing (A′-B′ Samples)* [Computer Laser Optical Discs], U.S. Department of Commerce, U.S. Bureau of the Census, Data User Services Division.

A TOPOLOGICAL APPROACH TO AGENT RELATIONS

T. HOWE,* The University of Chicago
M. DIGGORY, Harvard University

ABSTRACT

Agent-based modeling depends on rectangular grid spaces to represent space and relationships between agents. This limited, rigid requirement is not conducive for complex simulations. Researchers are exploring the use of irregular spaces in agent-based modeling environments, but few focus on compatible, swappable representations of relationships that have significant semantic meaning. This paper describes a new library in the Repast toolkit designed to achieve these goals. The library is based on relation topologies and context and allows agents to function differently depending on the situation. Topologies need not be fully connected and can take on appropriate structure. Because topologies are Java interfaces, the flexibility to add and swap them is great. By combining agents and topologies into contexts, the new library provides a flexible way of handling agent relationships and lets users focus on agent behavior.

Keywords: Relational topology, modeling agent context, model topology library

INTRODUCTION

Agent-based modeling and simulation has long been dependent on rectangular grids to represent both spatial and social relationships. While this has proven productive for many kinds of simple simulations, such a limited and rigid requirement has failed to meet the needs of many modelers. As such, researchers have started to explore the potential of using irregular spaces in agent-based modeling environments, including cellular automata (Flache and Hegselmann 2001). Most of these solutions have been "one-off" type solutions lacking generalizability. As computing power allows for increasingly complex spaces to be used in agent simulations, an efficient, generic model for agent relationships becomes necessary.

Recently, several publications have discovered that many models are sensitive to the structure of the relationships in which they are engaged. Flache and Hegselmann (2001) found "substantively interesting implications of the irregular grid that could not be identified with a regular grid structure" when studying irregular grid effects on influence dynamics and migration dynamics. Similarly, Rojas and Howe (2004) found significant effects from network structure on popular opinion change with both intra-network and extra-network influence. Particularly compelling about these examples is that they cross domains. The work of Flache and Hegselmann examined a migration model that contained a distinctly spatial characteristic, while Rojas and Howe focused on social networks. These examples demonstrate that many models, across domains, are sensitive to changes in relational structure.

Agent based modeling toolkits, such as Repast, Swarm, MASON, and Netlogo, have long treated spatial relationships as an entity distinct from social relationships. As such, reusing a

* *Corresponding author address*: Tom Howe, The University of Chicago, IL; e-mail: turtlebender@gmail.com.

model's logic with a different type of relationship can be quite involved. Gulyas (2002) described the difficulty of working with such libraries. As such, sensitivity testing using the approaches provided by most agent-based modeling frameworks is impractical.

The framework that we have designed is similar to previous attempts to design "swappable" spaces (Gulyas 2002). The basic units of the framework are "relation topologies" and "contexts." Abstractly, relation topologies simply represent collections of relationships between agents. As such, by examining a relation topology, one should be able to determine all of the relationships of a particular type between a given set of agents.

An example of a fairly common relation topology is the relationship between stores and their customers. The relation topology defines and maintains this "store-customer" relationship. In this type of scenario some agents represent stores and some agents represent customers. The "store-customer" relation topology maintains which customers a store has and which stores a customer patronizes. Note that agents can (and usually will) be involved in multiple relation topologies. Realistic agents will have very complex collections of relation topologies, which will comprise all of the relevant pieces of an agent's domain. As such, an agent might participate in the "store-customer" relation topology, the "parent-child" relation topology, the "employee-employer" relation topology, and so on. By keeping each of the relationships in a distinct relation topology, maintaining an agent's state with respect to all other agents should be much simpler.

Relation topologies are defined by two functions: distance and range. Distance yields the domain-specific distance between two points. For a grid-style relationship (where location is represented by a set of integer coordinates), these functions are quite straightforward. However, other types of relations may not be as simple to compute. For relationships in a continuous space, while a simple distance function may be sufficient, a modeler may wish to use another, more efficient representation for the relationship. When dealing with social relationships, no obvious representation of distance and range may exist. Any library for representing a topology must allow the modeler to describe his or her own representation.

For example, in a social relation, the distance between two individuals might be defined as the sum of the strengths of the edges that form the shortest path between the two individuals. It is important to note that the distance between individual "a" and individual "b" may not be equal to the distance between individual "b" and individual "a." In graph theory, this is a non-symmetric relationship. A range query returns all of the other members of the topology that are within a given distance, where, again, the distance is defined by the domain.

Individuals usually engage in multiple types of relationships. Krackhardt's (1987) work in the high-technology industry provides a good example of this. He looked at three sets of relationships: friends-with, reports-to, and asks-advice. Each one of the sets should be represented by a relation topology. This allows for simple and uniform swapping between relationship types.

The underlying motivation surrounding the creation of a library for supporting multiple topologies is the desire to treat all relationships in a consistent manner. In doing so, creating truly modular agents that can be tested in multiple modeling environments becomes much simpler. Also, by separating multiple types of relationships from one another, the maintenance of agent relationships becomes less complex.

CONTEXTUALIZED TOPOLOGIES

In order to make agent behavior more realistic, all relationships must be interpreted in terms of context. Context can have several meanings and as such we must define the term carefully here. Context is, essentially, the state of the model. From the context, agents can determine what the world looks like at any given point in time. Of course, the context will contain both global and internal information. The global information is accessible by agents and might be used by them to make decisions. The internal information should not be accessed by normal agents, but might play a role in how the relationships are determined.

Many situations have very different interpretations based on the state of the model. As such, agents might see the current state of a model and view their relationships with others differently depending on the set of circumstances. For example, the act of sitting in a restaurant may have distinctly different meanings depending on the surrounding circumstances. Normally, an agent may view this act as a leisurely moment, consisting of ordering food, eating, and paying for the meal. However, if a robbery attempt occurs in the restaurant, the agent's behavior will be quite different. In addition, the relationships between the actors shift with the context shift. While in one context the various actors may simply fill their roles as a waiter, patron, manager, etc., when the robbery occurs, the waiter may choose to confront the robber, changing his relationship with all of the other actors. In order to support these kinds of interpretations, the agents must be able to access information about the situation when examining their relationships and the Relation Topologies should be contextualized.

The topics of context and situated agents go far beyond the scope of the current discussion. However, when examining issues of topology and relationships for realistic agents, it is important to keep in mind that relationships can be context dependent.

THE LIBRARY

Building on ideas presented by others and the ideas contained herein, we have built a java library for the Repast agent toolkit (although the classes are certainly not limited to use with Repast). The entire library is based on three interfaces: Context, RelationalTopology and ModifyableTopology. These three interfaces support the behaviors outlined above.

The RelationalTopology interface is fairly simple. It contains a "RelationshipType" property, a range query, and a distance query. The RelationshipType property is a string that provides a meaningful description of the relationship described. The real meat of the RelationalTopology is provided with the other two methods. They determine how relationships are determined and maintained for agents.

The distance query is described by the following method signature:

double distance(Object element1, Object element2);

This method returns the distance between two objects as determined by the RelationalTopology. As such, implementers of this interface are required to determine how distance is calculated between agents. For relationships representing a grid, this may be a very simple arithmetic calculation, while for a geographic information system (GIS) style relationship, this may involve calculating a complex graph.

The range query is described by the following signature:

Collection getRelations(Object o, double d);

Given an agent, this method retrieves all of the other agents whose distances are less than or equal to the supplied distance. Most simply, this could loop through all of the relationships to find those with a distance less than the provided distance. But usually, implementers will want to provide a more efficient approach.

Some kinds of topologies cannot be altered directly. Most relationships involving a grid fall into this category. For example, adding a Von Neumann relationship is fairly nonsensical, since that type of relationship is determined by location. However, most topologies do require the ability to add and remove relationships. Social networks would not be very interesting if there were no way to add and remove relationships. To support this, we provide ModifyableTopology. This interface extends the RelationalTopology interface and provides two additional methods. The first method is described by this signature:

addRelation(Object element1, Object element2, double distance)

Not surprisingly, this method allows the user to create a new relationship of the type described by the RelationalTopology between element1 and element2, with the given distance. Similarly, the method

removeRelation(Object element1, Object element2)

removes the relationship between element1 and element2. The removeRelation method does not necessarily remove the complementary relationship, however. That decision is left up to the implementing class.

CONCRETE IMPLEMENTATIONS

As a first pass for this library, we provide concrete implementations for two sets of relationships, grid and network. We chose these two because they seem to be the most widely used and because of the amount of existing code that has been built to handle this type of data.

Two dimensional discrete spaces or grid spaces support agents whose location can be described using standard Cartesian coordinates [of the type (x, y) where x and y are integer values less than equal to the width and height of the space, respectively]. This type of space is very commonly used as it is easy to understand, fairly simple to implement, and efficient, performance wise. However, most implementations tend to be very brittle and depend on the semantics of the grid. As such, most grid implementations are not particularly compatible with other types of relations.

In order to support two-dimensional discrete spaces (which we are calling grid spaces) we need to handle several types of relationships. We decided to support three: a Von Neumann relationship, a Moore relationship, and a hexagonal relationship. The Von Neumann relationship is described as the set of objects to the north, south, east, and west. The Moore relationship is the same as the Von Neumann relationship except that it adds northeast, northwest, southeast, and southwest. The hexagonal relationship is determined by a hexagonal grid. The neighbors are the six contiguous spaces surrounding the agent.

We have noticed that previous attempts to develop a generic method to represent relationships for grids have suffered from performance problems. This was one of the motivating factors for creating this library. The challenge was to treat grid relationships in a way that would be consistent with graph relationships without accruing the cost of creating relationship objects.

To solve this problem, we needed to create a mapping between the semantically meaningful relationship structure and the efficient array structure. For the grid, we created a new object, the location object. A location object is just that: an object that represents a location in space. It knows its coordinates and can contain one or more objects. The location objects can then be stored in an array that is queried for its coordinates. So, agents (or other objects) will have an “at” relationship with a location. When the agents want to find, for example, their Von Neumann neighborhood, they ask their location object for the neighborhood. The location object then uses the array to find the neighbor locations (keeping in mind that the location object has a set of relationships with other locations) and returns the collection of neighbors. This allows us to treat all of the connections using a uniform concept of relationship, while maintaining a reasonable performance level.

This approach has a cost, however. Location objects need to be created and stored in memory. The creation time is fairly acceptable, though, given the number of other objects that need to be created at the beginning of a simulation. The memory usage is also small since this object only requires two integers and a reference to one or more resident objects. In addition, these objects can be lazily created so that they are only constructed once it is determined that they need to hold an object.

The implementation for network style relationships was a simpler task. Most network libraries already support the methods required by the topology library. Since most networks are represented by graphs, they already contain methods for retrieving and creating relationships. So, it was a fairly trivial task to hide the implementation of Repast’s existing network library behind the RelationalTopology interface. Because agents needed to support multiple types of relationships, slight alterations needed to be made support multiplex networks, but those changes were simple and required little additional code.

Future work for the topologies should include support for agents that are GIS aware. A similar approach to the implementation of the grid space can be used to maintain relationships in real world spaces. Some GIS software has already used a topology approach to index object relationships in geo-spatial data. There are a couple of differences in the specific implementation of a GIS topology system compared to a grid-based system, but the approach is the same. An array-based data structure is not applicable to GIS data because vector-type data is continuous. Advanced algorithms, such as Delaunay triangulation, have been developed to index spatial data. It is beyond the scope of this paper to examine these algorithms in depth, but they allow one to essentially create a graph of relationships between objects in space. So, agents, or other objects, have an “at” relationship with a location object. In this case, the location object contains a set of double precision coordinates. The location object can then query the RelationalTopology to determine its neighbor locations. Of course, because the coordinates are represented by a pair of doubles, these location objects do require more memory, but, again, because this approach provides uniformity in terms of handling relationships.

CONTEXT CONCRETE IMPLEMENTATION

A context is comprised of a set of agents and a collection of RelationalTopologies that affect those agents. The context is primarily responsible for maintaining the state of the environment in which the agents exist. It should expose situational information publicly. At a minimum, this information should include a representation of time and the types of relationships in which an agent can engage. In addition, though, any information or tools that should be accessible to all agents should be made public by the context. This might include some utility like a random number generator or part of the world environment, like the weather. The context is responsible for keeping all of the various parts of the environment synchronized between the various agents. The basic context interface is quite simple. It provides methods for retrieving the types of relationships that exist in that context as well as methods for working with those relationships. Another method is included to return the time as a double.

The ModelContext is the root context for all models. It provides the methods described above for the whole model. The ModelContext maintains the master time clock, as well as all of the agents. For many models, the agents will handle all of their relationships. The ModelContext can also contain other contexts. Each of these contexts represents a certain situation. They each can maintain their own set of relationships and public information. For example, in the example described earlier of the meal at a restaurant, before a robber enters the restaurant, the agents could be functioning in the ModelContext, but after the robber enters, they may shift into a new context called the CrisisContext. The CrisisContext would have its own set of relationships and might have its own sense of time. In addition, it would hold more specific information about the situation such as whether the police had been called. Once the crisis is resolved, the agents might switch back to the ModelContext.

By combining different sets of topologies and agents with contexts, we are able to create a very modular and rich set of tools for interacting with other agents. Because each of the agents can be contextually aware, they can interpret relationships differently in different sets of circumstances. Also, because contexts can be combined, context libraries can be built to allow for maximum modularity amongst model components.

CONCLUSIONS

Creating a library which can support uniform treatment and swappability of topologies is an important challenge to agent-based simulation toolkit developers in order to provide the maximum modularity and structure to users' models. Such a library is a requirement when the modeler wants to run experiments comparing the results of a simulation across multiple types of topologies. When a researcher wants create a highly realistic model with agents acting in a situated and contextualized manner, the modularity provided by this approach simplifies the task.

The library structure provides a space for future work, as well. New types of topologies can be created and plugged in. Another direction of future work would be to create a method of context shifting, possibly combining aspect-oriented programming with the context structure here.

REFERENCES

Flache, A., and R. Hegselmann, 2001, "Do Irregular Grids Make a Difference? Relaxing the Spatial Regularity Assumption in Cellular Models of Social Dynamics," *Journal of Artificial Societies and Social Simulation*, **4**:(4) [http://www.soc.surrey.ac.uk/JASSS/4/4/6.html].

Krackhardt, D., 1987, "Cognitive Social Structures," *Social Networks* **9**:109-134.

Rojas, F., and T. Howe, 2004, "Structure Matters: An Analysis of Consensus Formation in Social Networks," *American Sociological Association Annual Conference*, San Francisco, CA.

DISCUSSION:

TOOLKITS AND TECHNIQUES

(Thursday, October 2, 2003, 3:30 to 5:15 p.m.)

Chair: *Micheal North, Argonne National Laboratory*
Discussant*: Pamela Sydelko, Argonne National Laboratory*

MASON: A Java Multi-agent Simulation Library

Claudio Cioffi-Revilla:. I'm going to start off with a very few slides and then Sean Luke, my colleague, will continue and present the core of this work. And we should say that this has been a very collaborative project with graduate students, Gabriel [Balan] and Livew [Panait], and to a lesser extent, also Sean Paus.

So not coincidentally, MASON is a multi-agent simulator of networks and neighborhoods, and we'll begin with a few general features, and I'm going to give you mostly the big picture about this project that we're very excited about. This is the very first time we present this, so pardon the rough edges here and there.

[Presentation]

Sean Paus: What I'm going to talk about mostly is some of the details of how MASON works. To give you kind of a feel of it, I'm going to take you through this and you're going to see a lot of things, and you're going to say, "Well, that's got a Repast feel to it." And to some degree that was intentional. Repast did a lot of things right, so it was one of the models we based this work on. But I'm a computer scientist, especially in artificial intelligence and evolutionary computation, and I only know Repast mostly through the A-life world. My specialty has usually been robotic simulators in other kinds of areas. So some of the similarities you're going to see are more due to the fact that we borrowed ideas from other areas, which happen to coincide with Repast. And some of them are just wild chance.

[Presentation]

Unidentified Speaker: Since your visualization is a separate layer, would it be possible for the visualization to be on a different computer, through RMI or another method.

Paus: The answer is yes, but you'd have to write the code for it to communicate over the separate layer. I mean, of course, you can always do it over X-host. But the Java code that we've currently got written is intended to load the model in and display on the same platform, but there's no reason why you couldn't set up an RMI remote display on another machine.

Wesley Stevens: How many agents is the most you've run?

Paus: The most I've done is a million and a half.

Stevens: I like that answer.

Paus: But let me say, that's kind of slow because Java has memory limits on certain machines, like a two gigabyte limit on a lot of machines. So it all depends on the size of your agent.

But I will tell you, reasonable numbers usually are in the 10,000 to 100,000 range at most, until it starts getting a little slow for *my* tastes, for the kinds of things I need to do.

Unidentified Speaker: You mentioned you had the neophyte program in the high school. How long did it take them to get up to speed?

Paus: Rather little, actually. Dan Kuberich is a student who doesn't know a whole lot about Java. In fact, he wrote three of the demos, and you go through and you find out he didn't know the instance of operators, so he actually went to the reflection library to figure out the identification, the class, and then the instance of operators. And he did three of our demos: the L systems, the little soccer players, and the light cycles game. He did all three of those, and the media work for us over the course of maybe a month. We also have the port that was recently done from the swarm; the anthrax model was done by somebody who knows Java reasonably well and learned the system in about two days, and then spent about three days doing the port.

Unidentified Speaker: Are you doing any open source licensing on this?

Paus: This is all going to be open source, and it's going to be under a modified BSD Netscapish, MITish kind of license.

Unidentified Speaker: How do you see MASON evolving, particularly with regard to the noncore layers, and how do you see the community possibly driven to that?

Paus: My original vision of this, and I think Claudio has a very similar vision of this, was that MASON was a core that [had] enough tools that you could do lightweight simulations like these in it, but it was specifically designed for people to wrap larger simulators over it, for example, to build Teenbots II on top of it, and then distribute those as open source. And we also have a lot of commercial entities that are interested in putting their simulators on there for going after DARPA contracts, etc.

So the answer is that's kind of the model we've been going for. More, we're trying to make this as easy as possible for people to put their stuff on top of it rather than for us to really put those on top ourselves; although I will tell you, the Anthrax one, for example, we hooked up JChart to it without any problem.

Unidentified Speaker: Are you doing any things in particular to stimulate people to create those other layers?

Paus: Yes, but most of my efforts so far are because of my interests in the robotics category. But you have to understand, MASON is nearly brand new. This was developed, starting in January, on and off for about six months, and then, well, continuing on actually through the summer. So it's not been around for very long. This is almost a brand new system, and nobody has seen it. You are almost the first people outside GMU to see it.

Simulation and Distributed Architecture of Multi-agent-based Behavioral Economic Landscape (MABEL) Model within Swarm

Brian Pijanowski: We are working on some National Science Foundation and U.S. Environmental Protection Agency-funded projects in which we are taking sophisticated regional atmospheric models, like rams, and coupling them to crop forced net primary production models, which eventually feed into agent-based simulations. We are also working with some infectious disease models, some in East Africa. Some of our colleagues are working on biogeochemical dynamic models, so as landscapes change, nutrients loadings also change. We have also worked with some process-based hydrologic models.

These agents and models are actually responding to many changes. The models are generally process-based ones, but at this time, we are wrestling with the idea of how to actually couple these and also how to look at feedbacks. The scale issues are tremendous. When you have a climate model that is running on 15-second time steps and 120-kilometer grids, how does it talk to an agent-based model where you have an individual interested in making decisions about its parcel?

Looking at tipping points is another area of interest. We are very interested in learning how changes in the environment change decisions. We are looking at issues of uncertainty. Some of what Zen [Lei] is going to talk about really focuses on cyber infrastructure problems. When these are coupled together, things slow down very quickly. We would like to bridge qualitative and quantitative approaches by conducting well-planned simulations to obtain some broader impacts.

The model that is going to be presented is called MABEL. MABEL stands for Multi-Agent-Based Environmental Landscape. We have incorporated a behavioral component based on some beliefs and an economic market model. It is a landscape-level model. We are here for the conference, of course, but we are also here because MABEL is here and we have FRED. FRED is actually MABEL’s father.

I’d like to briefly introduce four people. I’m here, of course, from Purdue. Zen Lei is at Michigan State. Costas is at Purdue, and Snayhill is living out of a U-Haul truck somewhere between East Lansing and West Lafayette, heading in our direction [laughter].

[Presentation]

Sydelko: We have time for three or four questions.

Unidentified Speaker: What kind of thing do you try to replicate in terms of realities?

Zen Lei: Is there a data set that you try to replicate in terms of your transactions?

Unidentified Speaker: I work with the data. We have digitized historical data. These data go back to the beginning of the century, especially some or our data from Michigan. We have until 1992. We are attempting to simulate — with our database as input — how the landscape would look, given the decision-making properties of the agents and the people, because each person is not just a person. Why do we want to see how the landscape will look 20 or 30 years from now? It is a very good tool because decision makers will know how the decisions that are made today will affect the landscape 20 or 30 years later. That is a very powerful decision-making tool.

Lei: Have you made any real comparisons yet?

Unidentified Speaker: Yes, and they are real data experiences.

Unidentified Speaker: I'd like to add to what Lei has said. We are using things such as role and gaming and expert judgment to blend the qualitative with quantitative approaches, your Turing-type tests, to be able to determine the value against a real data set. We are also trying to incorporate that idea. And we also have got, as part of the interface to the gaming simulation, the Microsoft Belief Network software, which allows players of the game to construct the Bayesian network and assign initial probabilities. We use that as part of the proxy server construction.

Unidentified Speaker: My comment has to do with a talk by a person at Warwick in England. He was doing modeling for the British government on foot-and-mouth disease. One of the requirements was for the models to assume that Welsh farmers exaggerated their numbers of sheep by 40% for the purposes of a government subsidy or something, because otherwise the disease would have spread much faster through Wales. Making this assumption was the only way to get the models to [work]. The point of my comment is that sometimes having a lot of government data on farms is a dangerous thing.

Unidentified Speaker: Since I am familiar with Europe because I worked for the European Union, I am aware of the problem, and it has to do with the high European Union subsidies. The Union gets more revenue from overestimating your yields or the herds or the flock rather than selling it and doing business with [the proceeds]. But, yes, there is a data quality and availability issue. We would like to see more data and digital policy-based data, because we're talking about, in some instances, 30-by-30-meter data that is digitized from satellite data. So this process is slow, and sometimes it requires a lot of resources.

Unidentified Speaker: I have a question. When using multiple clients, quite a few different [ones], these are separate PCs that each one is running on?

Lei: Yes, a separate PC.

Unidentified Speaker: Does it have to be a manual intervention, or do you have a batch facility to start up the clients of various PCs?

Lei: Actually, we have each client on a different PC in the lab, and the server servicing our other machine.

Unidentified Speaker: Oh, no, I understand that, but to start up the clients and get those running. . . .

Lei: Actually, the server is running first.

Unidentified Speaker: It's a manual process. We're moving towards a remote method.

Unidentified Speaker: And why are you using the clients right now? Is that for a higher performance or is that just for experimentation now?

Unidentified Speaker: Mostly for visualization, but also for simulation.

A Topological Approach to Agent Relations

Sydelko: Next, Tom Howe will talk about “A Topological Representation of Agent Relations in Repast.”

Tom Howe: Actually, the title has been changed slightly to “A Topological Approach to Agent Relations.” The reason for this change is that, while there is an implementation in Repast of agent relations, the ideas and the way in which I will present it is more generic than that and more of an approach rather than a specific implementation. Having said that, it should become obvious that this talk is going to be somewhat different than the two previous presentations, which presented systems, where we are presenting a methodology.

[Presentation]

Sydelko: We have time for some questions.

Unidentified Speaker: Tom, it seems that, if you talk about social structure, that you would have two structures: a friendship structure and a family structure. Personally, I would think you would put that on the edges, that there would be a connection between agents — say, they’re friends — and then this other one is a type of a family edge.

I would have thought the context would have been more temporal. In other words, it might be easier to think of it as either endogenous or exogenous, but, for example, in a time of famine, everyone resorts back to the family. Can you comment on that?

Howe: These are handled by the topology. In a time of famine, however, the strength of the family relationship might be weighted higher.

Unidentified Speaker: ESRI is one of the largest companies in the GIS field at this time. They are moving toward developing agent-based simulation environments *in* GIS. It seems that those of us who are working in spatial environments love our GIS because of its efficiency. We have all the tools necessary. The most difficult thing that we do in our lab is move the data out of a GIS and into a simulation environment that was not built for it.

Therefore, getting closer to the GIS environment, to the point of possibly including the tool within that environment, seems the way to go because you have all of spatial relationships. I am Agent A. What township am I in? A GIS can easily handle a really complex spatial relationship.

Howe: Yes, that is correct.

Unidentified Speaker: When it gets to visualization, the tool is right there.

Howe: Yes, I think that is true. Now, of course, ESRI is basically a monopoly, and building things into ESRI's products is difficult because of their dependence on Visual Basic.

To that end, one of the best days of my life was the first time I imported a shape file into Repast in a successful and easy way and discovered that all of a sudden I had access to any kind of spatial structure I wanted. I could do complex networks, or I could do just a grid.

There is still this problem, however, that we need to be able to get to those spatial operators. There has been a movement among several people, and I have spearheaded one of those movements, to integrate some good GIS facilities into Repast. I know it seems like reinventing the wheel to put it into Repast instead of putting it into ESRI, but we have already talked about some of the problems with that.

With help from several people (including James MacGill, who is sitting here right in front), I have been working on a sort of integrated method with the geo-tools, open-source GIS limitation. This is back-ended at the moment by the Java Topology Suite, which has all of the basic spatial operators, such as Lizon, adjacent touches, all of those things, built into it. Once that is stronger, it becomes very easy to do the kind of relationships that you mentioned. Actually, there will be a paper on Saturday about this very topic

One of the things that I have noticed about the ESRI's concept of topology is that it is very limited. It does not seem to have a lot of different topology-building tools. While I am not a GIS expert, I haven't seen things like Delaunay Triangulation built into ESRI.

ESRI's basic topology builder will build either an adjacency matrix or a congruency matrix; I don't remember which. There is a benefit, though, to having this feature. Having the basic graph library setup makes it easy to implement various topologies, which can then be interfaced easily with an ESRI-type tool. So the goal with GeoTools is to make it very easy to go back and forth between the import and the export of data because the flood view utility in ESRI is still one of the coolest things I have ever seen, and that will just not be implemented in Repast.

Unidentified Speaker: In view of the fact that you delete an old location and add a new location and in view of the distributed architectures that we've seen here today, are you going to support transaction processing and have it commit in a rollback on the location change?

Howe: Yes, but the goal would be to make that transparent. That is sort of an artifact of the way grid operations have been handled in the past, which was removing yourself from a space and adding yourself to a space. It seems quite legitimate to think that the approach of removing your location and adding your location will change to hide things like transaction support because I do not want people to start a transaction, commit, catch, roll back, and so on.

Jesse Voss: My question has to do with your notion of topology and how you use it. I am interested in topology from the standpoint that Kurt Lewin takes in topological psychology, where you can have multi-dimensional psychological space. A simple example would be political parties compared to regular relationship networks. Let's say that someone has a tie to a person in politics, but belonging to a political party is not spatially oriented. Can your plans support multi-dimensional and pan-dimensional topological structures, which (based on different topologies) could build nested topological relationships for individuals or groups?

Howe: Yes, you bring up an interesting point. I don't see why that couldn't be done, although I haven't thought about a sort of nested topological structure. It should be do-able in the sense that you could have a topology. Your individual objects in one context could be actual contexts in themselves.

In your situation, for example, you have political parties, and each of the political parties exists in a large-scale national context. Then inside each of those political parties, you have a localized context, which perhaps consists of individual members. The challenge is to have nested contexts, where an inner context interacts with an outer context, which it seems necessary in your particular situation. There is definitely room to explore that, although it becomes rather complex, so I think that is in the future. In and of itself, however, I don't see any reasons why that wouldn't be possible.

Sydelko: We have listened to three very interesting presentations. They covered a wide range of topics. In particular, we heard about two different toolkits and also a discussion on topology.

In terms of coming up with a summary or a synthesis, I first heard multiple implementations of common underlying concepts. There are different ways of implementing the same basic concepts, and some of the concepts presented related to management of time. Simulations run things forward, creating focus points for agency. We had said these would be agents, but different ways of creating focus points. Regarding the management of relationships between agents, we might ask: How do you manage the space, or the relationships, that connect agents? I am sure that the audience recognized other concepts as well, and those would be great things to suggest as we go into the discussion.

Each implementation has unique strengths and weaknesses. All of the speakers who discussed some variation or implementation talked about the things that the implementation did better, but I think they would have to admit that there are a few things that they did not do as well when compared to the other toolkits. My conclusion would be that there is no one perfect way to do these things, but there are multiple ways to do them. Those various approaches are appropriate for different types of problems.

One thing that not discussed quite as often as I had hoped is that well-thought-out modularity seems to be one of the real keys. I think this idea was mentioned briefly, but someone would need to decide which modules to use for this system and how those modules should be factored. People hit on that individually, but as an overall focus, I would think that would be very important to discuss. It would also be important to talk about the substantial tension that exists between what should be the modules and what should be the core of the system and how those modules should be factored. In particular, one person's core future turns out to be another person's optional module in that no one actually wants to be the optional module in this world.

At this time, modularity seems to be defined within a given toolkit rather than between toolkits. I think this is a practical thing. There is very little, if any, discussion of taking a part from one toolkit and putting it into another. I think that is a very sophisticated thing to do; it is a very hard thing to do; and first we need to get the toolkits to work in all the directions we want them to work. That goal, however, is for the future.

Finally, in a perfect world, we would probably have a high-level toolkit-independent language to describe these things. This language would be transparent so that researchers could understand exactly what their simulation is doing and not always have to rely on programmers. You're laughing, because it is hard to do. Actually, it would be a high-level language, perhaps declarative, as David Sallach has suggested during earlier discussions. But the language could then be used to instantiate or create a model in any given toolkit, which means you would inherit the strengths and weaknesses of that toolkit. This is a long-range vision, not something I would see as practical immediately. But that would be a great thing to do because each of these toolkits has some very unique strengths. I saw distributed computation; I saw some efficient operation and 3-D visualization, topology's relationships, these types of things — all unique strengths. It would be great to develop an overall model and have a sense of transparency about that model, so you understand what it does, and then instantiate in the toolkits' different strengths and weaknesses. Ultimately, you could even perform a type of docking. That is, you would run the model in the various toolkits and see how those strengths and weaknesses do or do not influence the robustness of the results. So that, to me, would be an advantage.

I'm sure people in the audience have questions, but first I have the prerogative to ask all of you a question. Multiple agent-based toolkits, boon or bane? What would your view be on multiple-agent toolkits?

Paus: On multiple agent toolkits, boon or bane, I think for the time being the answer is boon, although there may be people in certain narrow fields that say it's bane, because you're dividing resources, but I think for the time being, very much boon.

We have a lot of discussion about unification and standardization, but the truth is that this is probably 10 years premature. Swarm and Ascape and Repast are all Generation I and Generation II toolkits. All are about 30-K lines, which is small. They are the kinds of things that are hungry to get replicated very rapidly.

Before we can really start talking about the problem of standardization, I think we are still at the point where we will be seeing several more major toolkits coming out, becoming very popular in the system before this even becomes an issue.

Unidentified Speaker: Boon, definitely. I agree with a great deal of what Sean just said regarding this being a growing field. Toolkits have a long way to go. However, I am not as against standardization as you are. When I say "standardization," I don't mean that all the toolkits do the same thing because I think that each toolkit brings a unique perspective to the field. As new toolkits are developed, new ideas and new approaches are going to be developed. Still, I believe that there is, at the base, an abstraction of what an agent-based model is. Taking the time to figure out what an agent-based model is at its core and then building up the various toolkits as they are created around some of those ideas — not to say that those ideas won't evolve or change — gives many benefits in terms of docking for validation, as we're talking about. I also think that they give benefits in terms of cross-pollination and affording the opportunity to explore new ideas within the various toolkits.

Sydelko: Anyone for a bane?

Lei: Yes. In my point of view, the multi-agent-based modeling is a booming field. I am interested in the large-scale simulation and also in simplifying. We need to find a way to simplify the simulation, and we need to combine some modeling — existing modeling — to build the external brain. The body is in the simulation, so in that way, we can combine some existing artificial intelligence tools into the multi-agent-based modeling and achieve a way to quickly translate the theoretical modeling into the implementation.

Sydelko: Are there questions from the audience?

Pijanowski: I have to ask a couple of questions. What is an agent-based model? What is it allowing us to do that we were not able to do in the past?

I see similar models that have been around for 20 years, and the question is … What does it allow us to do? What does it allow us to explore that we could not look at in the past? I don't know if there is an answer yet.

Joanna Bryson: I would like to follow on from that question and also from a question that was asked during the talk. How do you tell how good these things are? This is not a new question. I think we are looking at it the wrong way. Building a tool base is like building a chip set. I think the only thing you can do is benchmark. I do think you have to say, how quickly can we replicate well-known results? That brings up two sets of questions. How quickly can a novice do it, and how quickly can an expert do it because is it worth becoming an expert? So I think that there are ways to evaluate these things, but that there are new things because of the power.

Unidentified Speaker: Everything that we do now could have been done in the 1970s and was done in the 1970s.

Unidentified Speaker: Almost all of the MAS stuff has been done by the robotics community for 15 years.

Unidentified Speaker: They do some Fortran and C++ and it cooks.

Unidentified Speaker: So what it really boils down to here is software architectural design. I think this is a multi-objective problem, though. For example, I have an evolutionary computation system, a machine-learning system, that is probably the best in its field, but it has a very steep learning curve, and it's huge. It is very, very sophisticated, it can do almost anything, but it takes forever to learn it. On the other hand, Ken DeYoung, two doors down from me, has a very, very small, simple machine that is used by many people. This machine can't do everything, but it can do 80% of what people want, and you can learn it in five hours or less.

We are finding systems that are addressing the same kind of simulation functions that have been done for a long time, but for a different community with very different needs, and the needs are quite unusual. A large chunk of the community is our relative programming novices — I don't mean that with any disrespect — for which we need to have tools that are really easy to use. Then another segment of the community, such as some of the AI people and other people that are coming to this, are people who need very sophisticated capabilities and are willing to put in the effort. Both moved in both directions from what the kinds of tools in the 1970s and the 1980s were able to do.

Unidentified Speaker: Yes, I want to further that just slightly because I agree that fundamentally we are not doing anything we couldn't do. However, I also agree with Joanna

[Bryson] that working with robots takes a long time. But if you think about how people are doing their agent-based simulations, I would venture, without having any hard data, that the vast majority, or at least the slight majority, of people are rolling out their own solutions.

My point is that we are not doing anything that you can't do, that other people haven't done. The challenge in the near future is going to be to try to make it so that the "novice" programmers can build substantial realistic models that are not little toy models that you've done in the original Logo or in the StarLogo but make a way, construct a way, that people can (a) build those models and (b) have some form of validation, whether it be docking or whatever. I think that is where the big value added will come from. It will be from all the work we've done in the past five or six years.

Unidentified Speaker: I would like to add one thing as background. I would say that the focus may be a little bit too software-oriented because this meeting is not necessarily a software-oriented one; rather, it's a modeling meeting. I think that is the difference. In principle, we could go further back, not 20 years, but back to John Von Neumann architecture or to Alan Turing perhaps. We could work our way back to what seems like thousands of software years.

The idea, though, is to keep our work in perspective .In principle, ever since either Turing or Von Neumann or maybe Stephen Wolfram, who, I'm sure, would credit Nolan, the basic idea has been that you can back it off to some great distance and reach the point where nothing new has been invented in computers. The fundamental idea is that we're trying to model something — human behavior, animal behavior, or social systems. This conference talks more about social systems, so you'll see a variety of different things that we're trying to capture and model. That, I think, is different and new to some extent.

The multi-agent community was in some sense trying to replicate certain parts of human behavior, but it had a very different focus (than the AI community). This focus was to create an artificial human intelligence. Here, the idea is to capture social systems and understand them better, or capture other types of systems and understand them better. I think that to a large extent quite a bit of that *is* new. So, in that sense, the software structures may be very similar, even identical, but how they are being used is very unique and innovative. In fact, that is the reality of all software. It's similar to having the Von Neumann architecture, and then everything else is an application. The point is that it depends on the level that the application is being developed for.

Cioffi-Revilla: I understood the question somewhat differently. From the point of view of applications in social science in substantive models, there are several important problems in economics, in sociology, in political science that had been intractable by statistical and mathematical modeling approaches that have now become tractable. We understand a lot more about the way revolutions break out, for example. We understand a great deal more about the way in which the political geography of the world evolves. Bob Reynolds has explained in ways that other approaches have failed to explain how is it that communities of chieftains and hunter-gatherers form states. All of that is new, but it is substantive science. It is all new and important because of the recent generations of agent-based models. There is no turning back on that. New science is being done now, just like radio-astronomy allowed new science to be done and other similar tool innovations. My view, therefore, is more optimistic.

David Sallach: I agree with Claudio to some extent. I think, however, that the response is perhaps that something critically new needs to be done, and that is part of the problem. We need to do new things. We are in a situation where probably 80% of substantive social scientists do not see the relevance of agent modeling because it's so terminally simplistic in its assumptions. There's a higher-dimensional, rich, dynamic, flux-oriented, interpretive process.

We need to have a genuine engagement with the AI community. When you want to talk about the difference between building AI in units of one versus building it in units of 100,000, and then actually building social processes, dynamic social processes into that, being necessary — I mean, you can dismiss a certain part of the substantive social science thing in terms of disinterest or lack of technical background or ideology, close-mindedness, what have you, but there's a big hunk of it that you can't. That is, that people are immersed in area studies, people are immersed in qualitative, rich, historical or other kinds of data, and therefore, see the assumptions as being simplistic. Part of this has to be a dialogue between computational and social sciences in terms of how to bring that richness into computational modeling. When we do that, we will be doing something different than what was done 20 years ago.

Unidentified Speaker: So you have to understand that I'm looking at this from a computer science perspective, that the tools have been available for a long time, but there is an important feature that has not been available for a long time, and that is the big iron.

I think why you're seeing this blossoming that's come up, you know, it started with the A-life community and Santa Fe, and now you're seeing it move into social complexity and bio-complexity communities, etc. What's making this feasible for these models is that previously the only way you could do these models is (1) on the very, very small scale, and (2) you had to be a high priest of computer science to be able to understand how to write the low-level assembly necessary to get it run without taking a year of computational time.

Now we're seeing large, inexpensive machines, and we're seeing that software that runs on these large inexpensive machines is able to do these kinds of things in a more reasonable fashion for everybody. Essentially, in the last 10 years, I think that's where we're seeing these toolkits that are coming out that have been making these fields possible in the first place. You know, you're not dealing with simple, finite element analysis things; you're dealing with very large numbers of interactions, and that wasn't possible unless you had a lot of money and a really good programming team 15 years ago.

James MacGill: Assuming what you're saying about the big iron and the new toolkits, you said at one point in your talk, in the answer to our first question, you see several more toolkits down the line in the future. In each one of those, we're going to see an ever-prettier version of heat bugs running ever faster with ever more bugs. What is critical at this point is that we get a way to describe some of our simple models. I'd love to throw a challenge to you two guys to write an XML document that describes heat bugs that both of you can read.

I think it's no coincidence that Swarm, that Repast, that Ascape and MASON, are all open-source projects. We're scientists, and the way we communicate our findings and our research is by passing source code to each other. It's our language. We're not mathematicians; we don't have mathematical expressions that we can pass to each other, so we throw source code at each other. But that's why a lot of what we're talking about is to say, "Well, this is a modeling meeting, so why do we keep talking about software engineering?" It's because our language in which we express ourselves is software at the moment, and it needs to step up from that to something more abstract so that we can capture what we're talking about, and that when the next mega-thing comes out, the first task somebody doesn't do is write heat bugs to load the heat bug definition file.

Unidentified Speaker: I agree with you completely, with one exception. You know, I think that getting a way that we can come up with a language that's not code is exactly the way to go. I really don't think it should be XML, though. I think we need to have something that's somewhat serious — not just easier to work with, but also more semantically meaningful for what we're trying to do.

Unidentified Speaker: I'd add, just as a note: David Sallach said that if he did things over again and developed a new, say, declarative system or something like that, that he would not do heat bugs, as a political statement, actually.

Unidentified Speaker: I've never seen that one, so I'm not sure. I've seen Ascape. So he said he would not do it as a political statement. And also, there's an interesting comment that was made as well concerning David's protesting his heat bugs. That is to say, you know, the Internet Engineering Task Force has a policy where they don't accept proposals for new standards or other new policies unless they see working code, unless they see an actual implementation. That certainly is a higher standard, and it's probably a good one. So the next time people develop either a toolkit or a major new technique — yes, this means you — I'm kidding, Tom, I have to pick on you — that first people show a working model, substantive model in some domain that uses that feature effectively, so it might be a way to do things.

We have time for maybe one or two more questions. Do people have questions? I know you have an outstanding one, Steven. Does anyone else?

Unidentified Speaker: You may or may not have the answer today, but I think we, as researchers, are working with the multi-agent or agent-based simulation as a tool just for ourselves to understand the complexity and understand each individual agent's behavior, or just like you guys are working on building the toolkits to provide or to facilitate the researcher who works in this field to understand and extend the work for in different scale. But I think in another point of view, is it possible? I think it's possible, but I just try to leave this point. Why don't we do it in, you know, a parallel way? Try to build some kind of tools for those individual agents to learn and understand, because we never predict precisely what is going to happen in the reality in the future, because that emerges from the individual. So if it can do that, let them learn.

Sydelko: Does anyone want to comment? Good. Thanks. That's a very good point, actually. A cry, set them free, yes. It is an outstanding question.

Approaches to Validation

MODELING PLAYGROUPS IN CHILDREN: DETERMINING VALIDITY AND VERIDICALITY

W.A. GRIFFIN,* L.D. HANISH, C.L. MARTIN, and R.A. FABES
Family and Human Development, Arizona State University

ABSTRACT

Over the past two years, we have been developing an agent-based modeling program, called PlayMate, that simulates playgroup formation in children. In the fall of each year, new and returning children come together in our child development lab; many eventually settle into groups of semi-stable play partners. Factors contributing to the formation of these playgroups are currently unknown, but the dynamics of this evolution appear to have similar characteristics as other structures in social organizations (e.g., agent actions appear to be rule based). The children's social environment differs slightly each year because of variations in playgroup formations, and these formations derive from the stability of who plays with whom. Both the groupings and the resulting structures evolve as the year progresses. Agent-based modeling provides a mechanism for simulating this type of evolution. In modeling the emergent behavior, we assume that individual child attributes influence the quality and subsequent likelihood of peer interactions. Analyses comparing the simulated and the realized data indicate that the current implementation of PlayMate effectively captures the general formation of specific groups within the classroom. We illustrate and discuss how the strength of this interpretation is qualified when model veridicality is probed in depth and across time.

Keywords: Agent-based model, playgroup, model veridicality, ABM

INTRODUCTION

How do young children, each with unique preferences for an ideal play partner, form semi-stable playgroups that evolve as each child matures? This question addresses a fundamental problem in contemporary social science: how do disparate entities, through some unknown process, emerge as self-organized clusters that embody well-known, but poorly understood, social processes (Watts and Strogatz, 1998; Macy and Willer, 2002)? The ontology of each child influences the quality, duration, and frequency of time spent playing with other children, and in turn, this engagement alters the developmental trajectory of each child. Hence, a model of dynamic reciprocal influence characterizes the immediate social and physical worlds of children as they change and adapt. At the center is the playgroup, and although the formation of playgroups is well studied, it is unknown how this critical socio-developmental context develops and changes over time (Rubin, et al., 1998; Hartup, 1999).

Children's social networks are characterized by multiple morphologies. Consequently, investigators have classified the networks according to their gross structure; they typically distinguish between dyadic relationships (e.g., friendships; Hartup, 1996), relationships among

* *Corresponding author address:* William A. Griffin, Department of Family and Human Development, Box 872502, Arizona State University, Tempe, AZ 85287-2502; e-mail: william.griffin@asu.edu.

small groups of peers (e.g., membership in social networks or cliques; Cairns, et al., 1998), and relationships with large groups of peers (e.g., acceptance by classmates, social status; Coie, et al., 1982). Categorization in this manner provides no information about the mechanisms or processes involved in the initiation, maintenance, and evolution of playgroups. It is critical that social scientists move beyond categorization by static structure and begin utilizing the burgeoning and innovative work being done in social network analysis (Newman, 2003).

Although the focus is on young children, at the core of our endeavor is an attempt to understand and model the reciprocal evolutionary dynamics ubiquitous to all social processes (Conti, et al., 1998). As such, the research is informed by multiple scientific disciplines ranging from economics (Arthur, 1994), political science (Cederman, 1997; 2002), and sociology (Gilbert and Troitzsch, 1999) to computer science (Feber, 1999), physics (Rocha, 1999), and applied mathematics (Newman, 2003). During the last decade, the traditional boundaries among these disciplines have been broached by a general scientific methodology — agent-based modeling (ABM). ABM is a common language, and with it, comes common assumptions (Axelrod, 1997; Casti, 1997). Among the relevant assumptions, one is most pertinent to this research: social processes are complex and continuously evolving entities that adaptively configure themselves according to basic rules that, in turn, modify the environment housing the agents that comprise the entities. This reciprocal relationship among the individual agent, time, and the emergence of social phenomena is illustrated in Figure 1. Moreover, Figure 1 illustrates the individualism at the agent level, the ordered grouping of agents, and the resulting macro-level

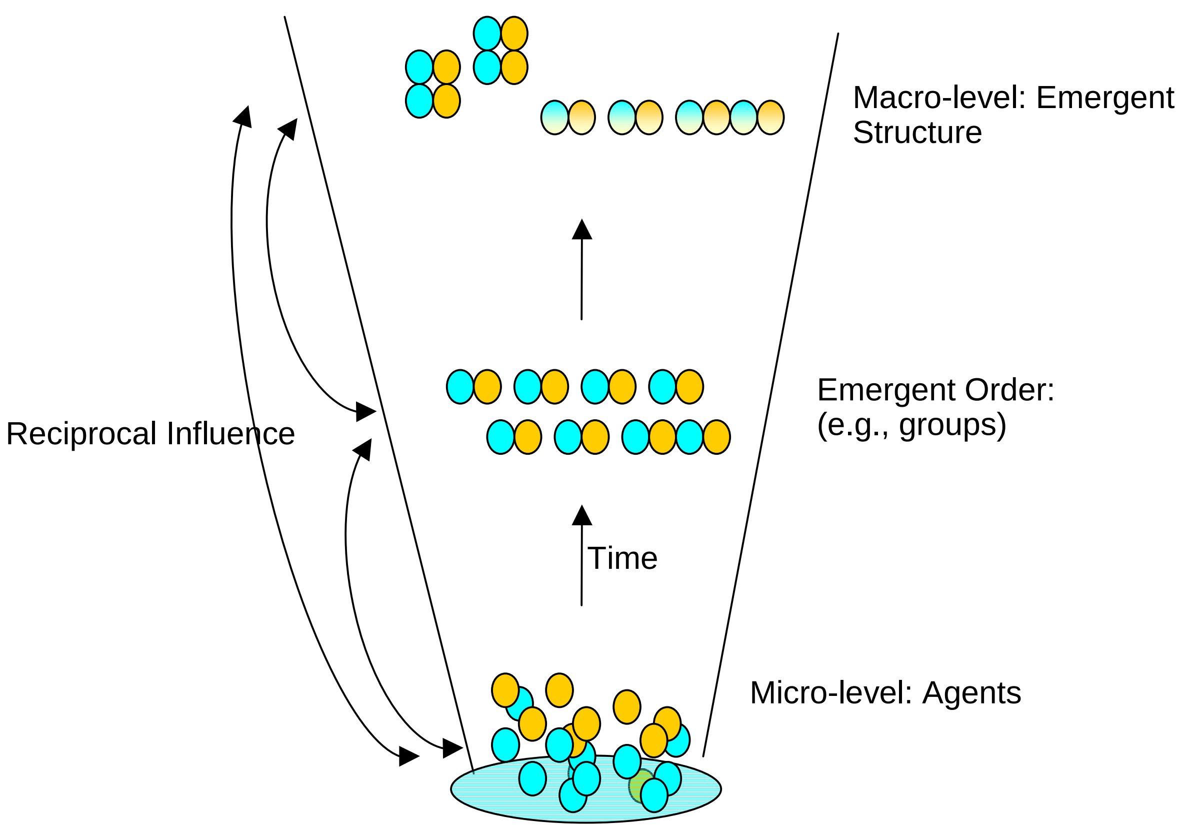

FIGURE 1 Emergence within a dynamic system derived from the interaction of individual agents over time

structures that emerge from the patterned collection of agent behavior. A contemporary debate in ABM is conveyed in this illustration: at what level of social structure (agent, grouping, emergent phenomena) does the modeler construct a model and, within the hierarchy where should validity be assessed (Conte, et al., 2001; Cederman, 2002)?

Our work resides at the center of this discussion. We have agents that cluster together because of individual- and group-level attributes, and these groupings, in turn, modify the environmental field that permits the expression of these attributes. We assume that individual children playing together co-create an environment that allows the emergence of a larger, more general social process — peer preferences. Peer preferences, and the process of peer group formation, simultaneously modify child attributes and the environment within which this metamorphosis occurs.

This assumption is the foundation for our work; at the general level, we investigate social processes, and at the specific level, we want to know how young children's playgroups form and change. From this foundation, we have several broad questions that focus our work:

- How are groups emergent (Axelrod, 1997)? That is, why is the aggregate not obtainable from simply examining its constituent parts (Humphreys, 1997)?

- What is the developmental value of groups that exceeds the socialization value obtained in simple parings? What is the limiting or carrying capacity of groups of children in this age range (Jin, et al., 2001)? Stated differently, is there an optimal group size that maximizes the assumed benefit of group play?

- How does the group entry criterion, as determined by the best attribute or set of attributes, change with group size or heterogeneity? Does group heterogeneity change with group size or density relative to the number of other groups with a given classroom size? And how does this rule change over the course of a year (Girvan and Newman, 2002)?

As evident from these questions, our research objectives extend beyond merely studying playgroups: we address questions about social processes that are germane to all human interactions involving micro-exchanges of social rewards and the diversity of shifting reinforcers — and how these crucial processes change over time. Finally, all computational modeling in service of these objectives must demonstrate validity across the period of emerging friendships. This paper illustrates some of the methods we use to determine if our ABM is valid, is veridical, and is sensitive to the intra- and interagent evolutionary changes that occur during friendship development among five-year-old children.

PLAYGROUPS

The Importance of Peers and Playgroups

Early childhood is an important starting point for the development of peer relationships. In particular, it is a critical time for the development of skills and expectations related to interacting within larger groups of children (Fabes, et al., 2003b). For many children, the earliest opportunity to interact with a consistent set of peers occurs in preschool as many same-age peers are brought together to socialize on a regular basis. Children embedded in a social network have more opportunities to develop their cognitive, behavioral, social, and emotional competencies, which facilitates later adaptation. These networks provide the socialization opportunities that children need to develop the nuances required for social negotiation. In contrast, children who have only limited interactions or who are rejected by peers do not get the opportunities to experience positive peer socialization. Consequently, these children are at increased risk for psychosocial and academic maladjustment (Hanish and Guerra, 2002). Understanding how young children develop positive peer relationships is critical to understanding the conditions that contribute to successful socialization and adaptation to life experiences. To date, however, most of the research examining the impact of peers has focused on older children and adolescents. Relatively little research has focused on young children even though it is well established that early peer relationships foreshadow the quality of later peer relationships (Hartup and Laursen, 1999).

Playgroup Formation and the Selection of Playmates

One of the most significant features about children's playgroups is homophily; that is, they are characterized by a high degree of within-group similarity (Berndt, 1982). Peer groups form around similarities in propinquity, sex, race, and behavioral dimensions, such as aggression (Cairns, et al., 1998). The notion of homophily is well established, but the processes accounting for homophily are not (Espelage, et al., 2003).

In young children, one of the most obvious dimensions of similarity in playgroups is sex. Preferences for same-sex peers emerge around 30 and 36 months and increase across childhood (Serbin, et al., 1994). By most estimates, more than one-half of all young children's peer interactions involve play with same-sex peers, approximately one-third involve mixed-sex peers (playing with both a boy and a girl), and less than 10% involve play only with other-sex peers (Fabes, et al., 2003). Same-sex peer preferences are stable over time (Martin and Fabes, 2001); they are stronger when activities are unstructured and when adults are not immediately present or involved in children's play (Thorne, 2001); and they are resistant to change (Serbin, et al., 1977).

Gender can serve as a primary basis for selecting social partners, but it does not explain the multiplex of peer relationships. Same sex peers are not selected indiscriminately. Choices about which boys or girls to play with are also influenced by behavioral compatibility, such that children seek out peers who exhibit similar behaviors or who have similar interests (Rubin, et al., 1994). For instance, aggressive peers tend to congregate together, and the social networks that surround aggressive youngsters often consist of other aggressive youths or children who actively encourage bullies' aggressive behavior (Espelage, et al., 2003). Children also are attracted to peers who share other characteristics, including prosocial behavior, and interest in academic activities (Fabes, et al., 2003a).

Peer Influence

Peers have the potential to be powerful socialization forces. For young children, this idea has been examined in research on sex segregation. Because of the high levels of sex segregation in children's play, children have more exposure to, and thus obtain more practice with, the styles of interaction characterizing their own sex. And, because the sexes play in very different ways, peer experiences can be described as separate cultures for boys and girls. Boys' groups are larger, and they tend to play in more public places with less proximity to adults than do girls (DiPietro, 1981; Fabes, et al., 2003a). Boys' play also tends to be rougher and more active than girls' play. Boys quickly establish a hierarchical pecking order, which remains stable over time (Maccoby and Jacklin, 1987). In contrast, dominance hierarchies in girls' groups are more fluid and less stable. Girls emphasize cooperation and use enabling forms of communication that promote group harmony. Compared to boys' groups, girls' groups are more likely to select activities that are governed by strict social rules (Leaper, 1994). Because boys' and girls' groups promote different styles of interacting, it is not surprising that they show different patterns of peer experiences. Experiences gained within boys' and girls' peer groups foster different behavioral norms and interaction styles. Over time, repeated exposure to these different behavioral and motivational norms and interaction styles has been hypothesized to promote the development of different skills, attitudes, motives, interests, and behaviors (Leaper, 1994; Maccoby, 1998).

Recently, evidence on the effects of peer socialization was demonstrated in a study of preschoolers' sex-segregated play. Martin and Fabes (2001) examined how individual differences in the "social dosage" of same-sex peers over several months influenced children's behavior. The results showed that both sexes became more gender typed in their behavior over time (e.g., boys became more aggressive; girls increasingly played near adults), and these differences were evident above the initial differences that may have led them to play with same-sex peers. The effects of peer socialization have also been identified for a range of behaviors, including specific interaction styles. These effects can be seen, for instance, in the ways in which exposure to particular kinds of peers affects children's own behavioral and emotional tendencies. Analyses of an extensive observational data set suggest that spending time with aggressive peers increases the likelihood that children will escalate in their own aggressive and disruptive behaviors, particularly girls. In contrast, spending time with prosocial peers resulted in increases in positive emotionality and decreases in negative emotionality (Fabes, et al., 2002; Hanish, et al., 2003). Furthermore, peer socialization effects are bi-directional and complex; exposure to particular interaction styles modifies children's own behaviors, and children become more alike over time as they interact (Kindermann, 1998).

The Dynamics of Young Children's Playgroups

Most studies on peer relationships have approached playgroups as static entities that classify children into groups or categories as they exist at a single point in time. Even if a static view is not presumed, difficulties of measurement often provide only a single point in time assessment of peer relationships. This approach has been crucial in building extant models of peer group formation, but with recent developments in methodologies that are capable of capturing dynamic shifts in social phenomena, it is now possible to assess dynamic changes in peer groups.

A dynamic approach also can be used to compare competing ideas in the literature, namely, whether the homophily seen in groups is due to selection of peers who are similar or whether it is due to the processes of influence that occur in peer groups. This issue has been central in the study of groups for many years from a variety of disciplines. For instance, criminologists have long noted the strong connection between delinquent adolescents and association with delinquent friends. Is this similarity due to the influence of peers (Sutherland and Cressey. 1974) or to the inability of the adolescent to make friends with nondelinquent adolescents (Hirschi, 1969)? A dynamic approach can incorporate both theories by allowing that selection features come into play by influencing who a child is initially interested in (and who may be interested in playing with the target child), and by proposing that these selection features likely change over time and depend on the range of available options. Furthermore, central to a dynamic model is the assumption that peers influence each other, and that this influence varies depending on the social dosage, or exposure, that a child has to specific children. This exposure, in turn, may change a child's selection criteria and/or desirability as a play partner.

Simulating Playgroups: PlayMate

PlayMate is an agent-based model constructed to simulate the formation of playgroups in children ages four to six years (Griffin, 2003). To keep the model simple and results tractable, PlayMate uses static (e.g., sex) and dynamic (e.g., sociability) child attributes to modify the likelihood of interacting with another child (Griffin, et al., 2002; 2003). The effects modeled for these traits or attributes can be modified to represent postulated developmental shifts.

PlayMate is constructed as a multi-threaded, object-oriented, agent-based platform where each child, as an agent, is assigned a separate thread and is derived from a parent-child class. Written in Python, a high-level, interpreted scripting language, PlayMate is framed around a state transition model, where a child is always in one of four states:

1. Playing with another child,
2. Playing with an adult (a teacher),
3. Playing alone after playing with another child, or
4. Playing alone after playing with an adult.

Early in our work, it became obvious that solitary play, either item 3 or 4, occupies about 20%–25% of a child's time, and the propensity to enter and exit this state varies according to whether the child plays with another child or with an adult.

Two key components are used to model the shifts in play likelihoods between and among children across the four states. The first is Play Propensity, the likelihood that any specific pairing of children will occur. The second is Arousal, a behavior proxy (of a child's internal configuration of cognitions, affect, motivational, and behavioral tendencies) that externally characterizes the propensity to shift states. This latter component does not imply a change in physiological systems (e.g., central nervous system); rather, it is a descriptive term to indicate the current level of a child within each state as he or she moves toward shifting states.

The underlying mechanism PlayMate uses is briefly described as follows. At each observed epoch (analogous to a single, real, 10-second playground observation), a child is in one of four discrete states (noted above). Although in a particular state, the child has a cumulating value in each of the four states that is used to allow spontaneous state transitions (excluding those logically not permitted; e.g., solitary [3] following solitary [4]). In round-robin fashion, a child is selected to play with another child from the available pool (one is randomly removed to simulate a "sick" day), and upon pairing, child *i* assesses child *j* on several dimensions determined by the investigator; minimally, these include sex and the relevant attribute (e.g., aggression, anxiety; see details below) being examined. Arousal, and thus the propensity to exit the child-playing state, increases proportional to play partner dissimilarity. The greater the homophily, as assessed by closeness on the variables in the model, the less likely the child is to exit the child-playing state and to continue playing with other children. This reduces the amount of solitary play and increases the likelihood of playing with other children as long as they are similar. After each play episode, two things happen using the assessed attribute level difference: (1) the arousal level of each state is updated according to a set of transition rules and values associated with those rules, and (2) the degree of similarity in attribute level plus the assigned value for sex similarity plus a memory value (higher value assigned to having played recently) is converted to an integer value associated with an investigator-determined curve (e.g., exponential), and this value is entered into an adjacency "tally" matrix. This matrix is a proxy to the observation matrix containing real data. After each run, the simulation tally matrix is converted to a child-to-child play probability matrix and compared to a similar matrix derived from actual data.

Following the admonishment of Carley (1996) and others (e.g., Leik and Meeker, 1995; Rykiel, 1996) regarding the necessity of model validation, and her work on veridicality or truthfulness in the model (Carley, 2002), throughout the evolution of PlayMate, we have tied its output to real data. Real data were collected via intensive 10-second observations of children's naturally occurring behaviors and interactions at preschool or kindergarten. Each year, a large group of observers were trained to record the activities, actions, and play partners seen in each observation. Data were recorded in real-time into handheld computers. This procedure was repeated for a randomized list of children in each classroom. We typically get 2,000–3,000 observations in a month of data collection. Assessment of the reliability of each coder was conducted and was consistently found to be high (see Martin and Fabes, 2001, for an example). For model validation, PlayMate generates numerous quantitative indicators of the structure and composition differences between the simulated and real data; these include difference measures of Euclidian distance, Mean cell values, Entropy, Uncertainty reduction (a measure of mutual information), Solitary play, and row (i.e., child) signal-to-noise ratios. Each measure is assumed to provide slightly different information about the characteristics of the matrix structure.

Data Simulation

To approximate a typical month of child observation data, a simulation run was defined as consisting of allowing each child in the class to play 50 rounds in the round-robin fashion. This routine is performed 50 times, and we generally obtain the appropriate 75–125 play episodes, characteristic of the number generated for each child during a month of observations. The 8-month school year was reduced to five data periods because coder training and reliability acquisition occur during the first month of school, and children are not available during the holiday period from mid-December to mid-January.

Response Space: The three primary factors influencing state shift propensities and play partner likelihoods are the influence of sex, attribute difference level, and recency of play

(Memory). Within PlayMate, each factor is weighted according to empirical or theoretical justification. In aggregate, these three factors determine the magnitude of the increase in play propensity of one child toward another; however, only attribute difference is used to modify the likelihood of the existing play with child state. Although, in principle, the weighting of these factors should have an empirical or theoretical justification, in practice we ran the simulations using parameter sweeps across each factor. Sex influences the model by allowing the preference for matching on sex to be higher for boys than girls, the girl-to-girl play being a percentage of the maximum of boy-to-boy play. Values ranged from 0.5 to 0.9 in 0.1 increments. This same-sex differential is consistent with a substantial body of literature (e.g., Martin and Fabes, 2001). As a proxy for a child's memory, a list is maintained of all recent play pairings. PlayMate currently maintains a list of five pairings; this list can be modified to correspond to the developmental level of the children (e.g., older children have better memories). Assuming that children tend to play with other children who are similar, and that the preference should be evident in the ordinal ranking within the list, integer values are assigned according this ranking. Similarly, the distance between classroom attribute ranking is also assigned an integer value. For the initial analyses, to assign values for Memory and Attribute distance, we did a parameter sweep using a gamma distribution, modifying the shape and scale parameters. Interestingly, the best fitting curve(s) reduced to an exponential distribution, implying that reinforcement for these two factors falls off at a constant rate. Although not in the current version of PlayMate, it is possible to allow an optimization method (e.g., genetic algorithm, see Mitchell, 1996) to maximize the correspondence between the value associated with the index location on the specified curve and the realized data. As we discuss below, however, this introduces a brittleness that optimizes the model to a particular group of children, and yet, it may not produce an optimal general model (see Bankes, 2002).

GENERAL MODEL VALIDATION

To date, PlayMate has been through two revisions. The first version was an agent-based model and equation-based hybrid (Griffin, et al., 2002). More recently, PlayMate was rewritten to be completely agent based (Griffin, et al., 2003). Our latest analyses indicate that the current implementation of PlayMate, although simple, effectively captures the formation of specific groups within the classroom (Griffin, 2003). Among the various indices used to compare simulated data with real data, the two most sensitive to children's attribute differences were Euclidian distance and mean cell difference. Not surprisingly, these were highly correlated ($r = 0.92$). The child attributes used in the analyses (obtained via teachers' and observers' reports) were (1) prosocial behavior, (2) activity level, (3) aggression, (4) social inhibition, (5) temperament, (6) anxiety, (7) physical attractiveness, and (8) social competence.

Analyses of the data consisted of running each attribute individually using parameter sweeps for Sex, Memory, and Attribution distance. We estimated an overall measure of association between the simulated and real data matrices by using the quadratic assignment procedure (Hubert and Schultz, 1976; Krackhardt, 1988) as implemented in UCINET 6 (Borgatti, et al., 2002). Taking the best fitting models that minimized Euclidian distance and mean cell difference, the attributes of prosocial and social inhibition produced matrices that were nonsignificantly different for periods 1, 2, 3, and 4 ($p < 0.01$). The mean cell differences were 0.025, 0.026, 0.025, and 0.029, respectively. This indicates that the simulation produced an average per cell (i.e., *ij*) play expectancy within about 2.5–2.6% of the actual data. Note however that at Period 4 the value moves up to 0.029, and although still significant, it does suggest the model fits less well over time. This suspicion was confirmed with Period 5; the mean cell difference was 0.034 ($p > 0.05$), indicating a significantly different matrix configuration than the realized data. In short, we were able to adequately simulate playgroup formation for Periods 1–4 using the attributes of prosocial and social inhibition (physical attractiveness also provided a significant model but only for Period 3); however, as the year progressed, our model fit less well.

PUNCTUATED VERIDICALITY AND DEPTH OF CORRECTNESS

Having shown that we have produced a pretty good model — in the general sense — we more closely examine how the model performs under greater scrutiny. Veridicality can be considered along at least two dimensions: patterns over time and depth of correctness. If we punctuate time into discrete intervals, we can assume that model veridicality is invariant within the specified window of time, and that this stability of truthfulness may or may not continue into the next interval. Such discretization permits exacting tests of the model as it attempts to capture processes that invariably evolve in a dynamic system (Casti, 1997), and it allows us to ask very specific questions of the simulated data. For example, can we predict clustering of children over time, and can we determine the depth of peer preference at each point of assessment? By dividing the time year into five periods, the analyses presented thus far have addressed the initial question. The latter question (addressing preference strength) is just as important as the former.

Envision time running horizontally, where the model has been divided into approximately equal segments, either for analytic or theoretical reasons, and depth of correctness running perpendicular to time. In the perpendicular plane, degrees of correctness are demarcated (one could use a 0–1 range to indicate percent of correctness) for one of several categories of correctness. We illustrate this concept by examining three depth categories per punctuation point: (1) correct classification of clustering, (2) correct classification of within- and between-cluster preferences, and (3) correct classification of strength of preference. In the first category, theory suggests a strong gender affiliation, and this characteristic was built into the model. As noted above, parameter sweeps for the influence of sex on subsequent play propensity did not drastically modify the fit to the data; the data fit well as long as the model specified that boys moderately prefer the company of other boys more than girls prefer the company of other girls (approximately 60–80% of the preference of boys). Consequently, the general fit in this category can be considered good. This finding is evident by comparing the realized data in Figure 2 with the simulated data in Figure 3. Specifically, these figures show the web of connections where the number of interactions is greater (i.e., 5) than the class mean (real data: $\underline{M} = 4.335$, $\underline{SD} = 4.45$ and simulated data: $\underline{M} = 4.344$, $\underline{SD} = 2.819$) for Period 2. Period 2 is used for illustration because its fit to the data is approximately the same as Periods 1, 3, and 4. The teacher report data were also collected during this period.

As can be seen by the connections between vertices (each being a child; blue = boy) the simulated data captures most of the same-sex interactions and several critical between-sex interactions. However, additional questions, at greater, more microscopic depth need to be addressed to determine model truthfulness at this punctuation point. First, does the model capture pertinent same-sex versus between-sex clustering; that is, are we identifying boys and girls that play with each other? (Recall that same-sex play is configured tightly in the model.) Second, can we identify and predict peer preference strength, an assumed critical indicator of playmate longevity?

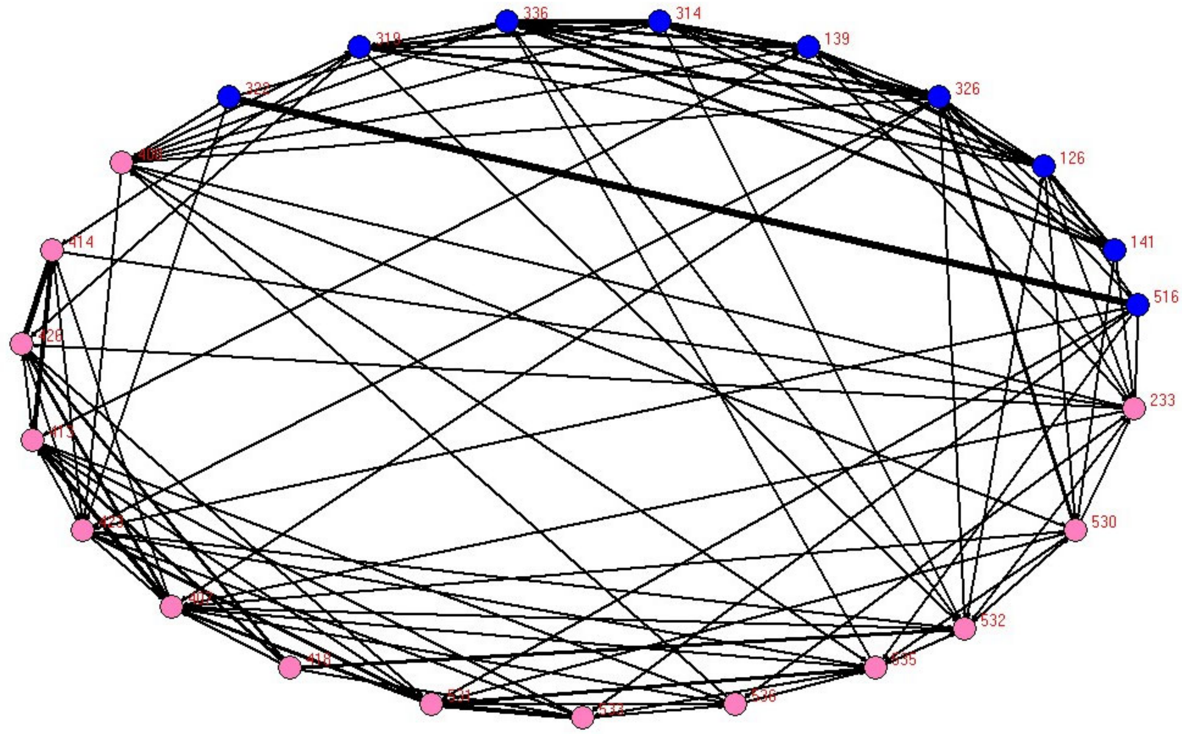

FIGURE 2 Realized data with edges greater than 4 at Period 2

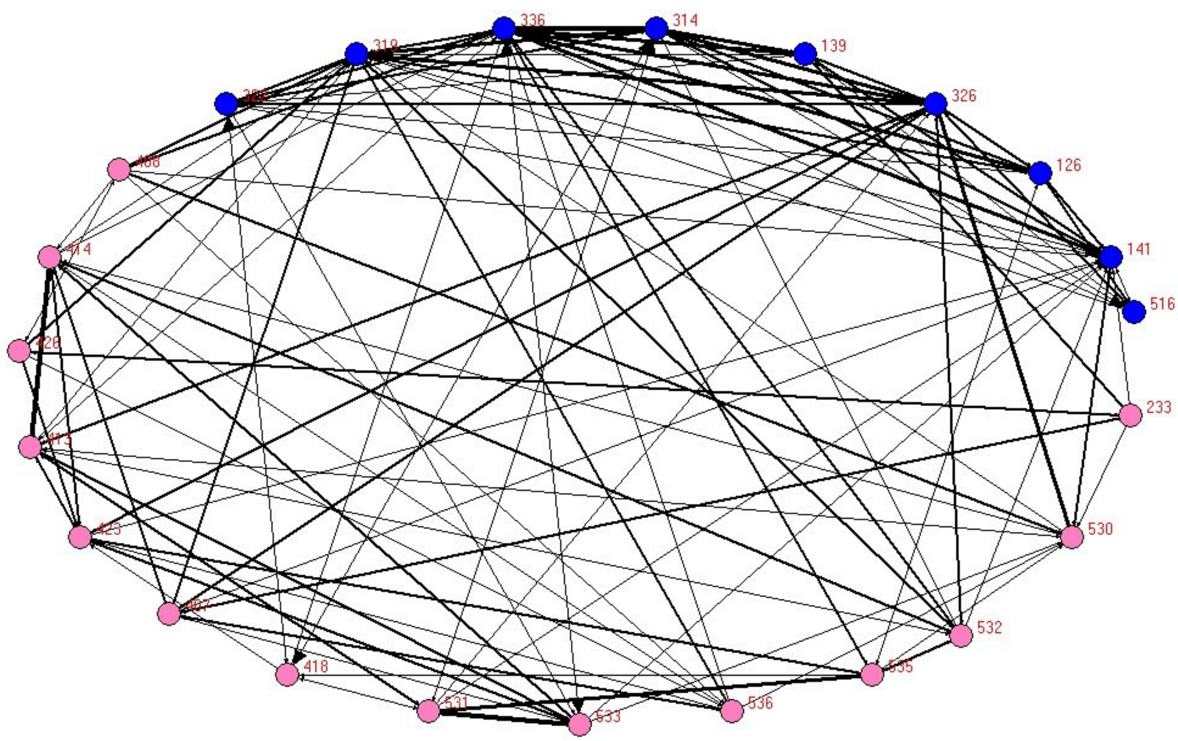

FIGURE 3 Simulated data with edges greater than 4 at Period 2

To address these questions, we again raise the threshold of criteria for an edge to seven or greater interactions. This change provides an increase of about 50% over the mean, thus reducing the number of edges in the web and providing a better visual presentation of if and how the simulated data differ from the realized data. Figure 4 shows the realized data for Period 2, and Figure 5 shows the simulated for the same period. It is immediately apparent that the simulated data capture many of the same-sex interactions but fail to identify several between-sex playmates. It is clear from Figure 4 that several boys and girls played together on a regular basis, and although the simulation suggests this cross-sex play occurs at a rate comparable to the realized data, it generally failed to identify the correct couplings. In addition, information about the strength of particular couplings that were missed is evident in Figure 6. This figure shows the edge differences between Figures 4 and 5. Edge width reflects the strength of the interaction; a wider band indicates more frequent interactions. Not surprisingly, Figure 6 shows that several cross-sex interactions were missed; it also shows that several significant same-sex playmates were not found in the simulated data — for both boys and girls.

This finding of approximate fit at the meso or pattern level (Casti, 1997), and a weaker fit at the agent level was consistent across the first four periods; not surprisingly, Period 5 had the poorest fit at all levels. It is clearly evident that as we punctuate time into discrete windows looking for veridicality within each, and as depth progresses — either at the agent level or the process sequence level — fit between the realized data and the simulated data lessens. *Does this mean model veridicality depends on the level of examination, or does determining model truthfulness of dynamic processes within a complex evolving system require some latitude for the system's inherent variability and possible nonreplicability*? We are trying to address these questions as another school year of data collection begins, and the model is again being revised.

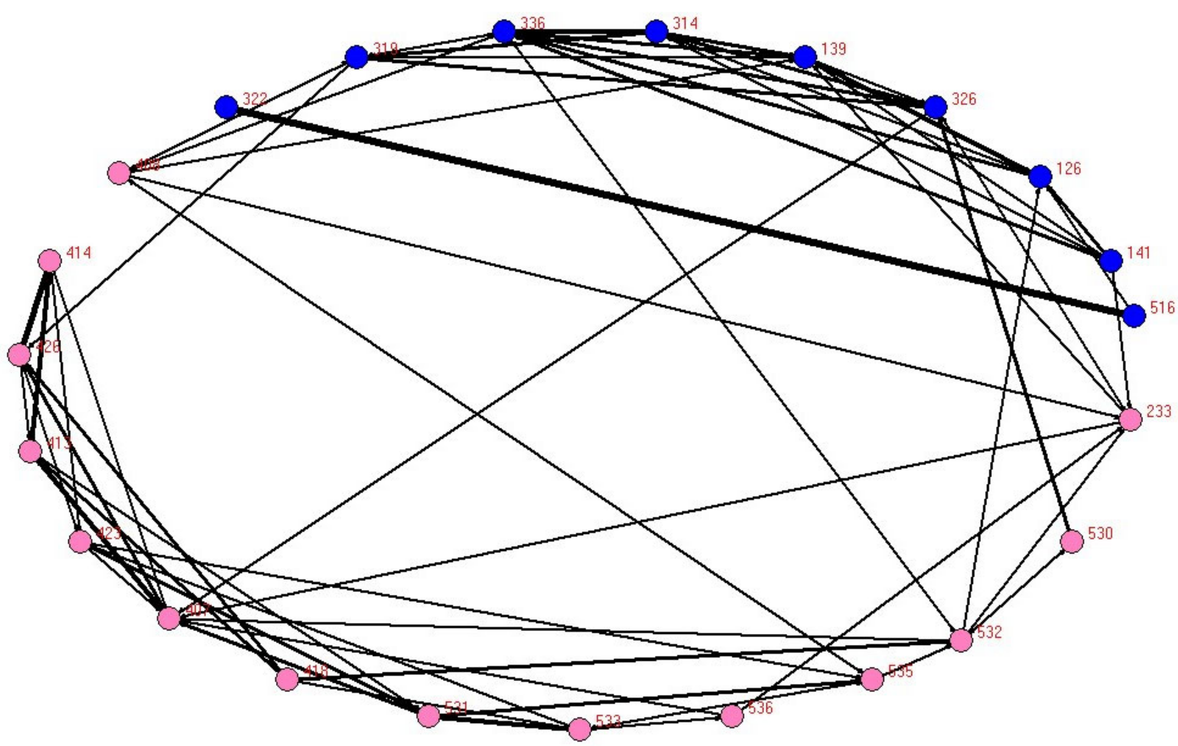

FIGURE 4 Realized data with edges greater than 6 at Period 2

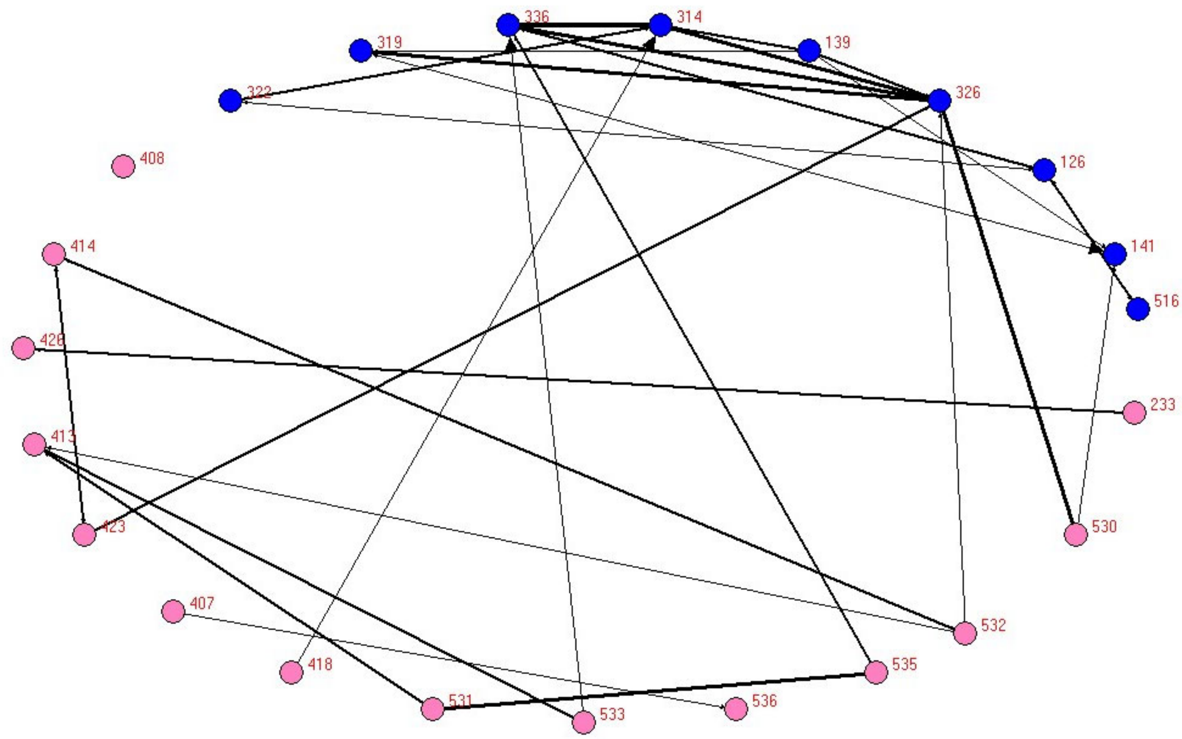

FIGURE 5 Simulated data with edges greater than 6 at Period 2

The web diagrams in this section indicate that the simulated data are more democratic in their distribution of playing time, and that real data cluster at higher rates and are maintained longer (i.e., more occurrences) than our current model generates. The simulated data encompass most of the connections at levels around the mean, but they fail to lock-in on unique relationships. This suggests that something beyond mere homophily is operating to maintain a friendship. Sex clearly maps onto the realized process, and in conjunction with prosocial behavior, we get a generally good fit to the model. Although at a level less than prosocial behavior, other attributes, especially social inhibition, and to a lesser degree, physical attractiveness, improved fit to the data. We are currently examining methods of creating vector variables, with and without element weighting, consisting of these attributes in the hope that unique combinations might generate simulated data that capture the depth and complexity seen in the realized data. This idea is discussed in greater detail below.

DISCUSSION

Proposed Refinement of PlayMate

Although analyses indicated that the model generally was adequate, several prominent shortcomings of PlayMate were revealed. First, in its current implementation, the model does not provide a mechanism to modify the child's attribute level as a function of interactions with other children (see Figure 1). This reciprocal modification among interacting children is key to

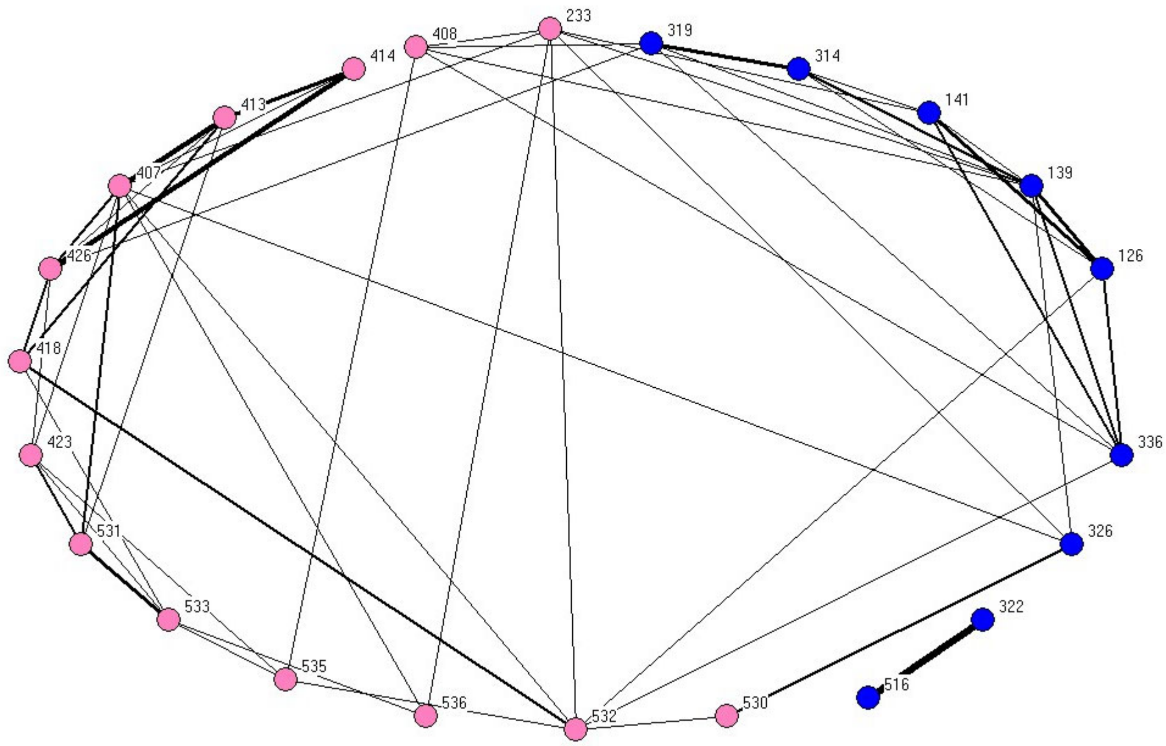

FIGURE 6 Edge differences between realized and simulated data for edges greater than 6 at Period 2 (line width indicates strength [frequency] of interaction)

modeling the evolution of change (e.g., co-evolution; Fogel, 1993). If this was in the model, we may not have seen the fit drift toward the end of the year. An initial, though nonsystematic, attempt to modify PlayMate by having each interaction slightly modify the child's rated attribute level was not successful (all children drifted toward a single attribute level); further refinements of mutual influences across children will be explored in the future. Such agent characteristic modification mechanisms, however, are common in the ABM literature and should be incorporated (Axelrod, 1997; Conte, et al., 1997; Goldspink, 2002). This shortcoming is confounded by our lack of temporally relevant data to use in the model. In our current data, children's individual attributes were measured in mid-December, and the simulation accurately modeled play during this time (Periods 1–3); as we attempted to model behavior farther from the assessment period, model accuracy diminished. In fact, it is possible that with periodic assessments of attributes throughout the school year, the existing model may not show the drift found the current analyses.

Second, there are several *prima facie* assumptions within the code structure of PlayMate that need to be modified to better capture the ontological complexity of playgroups emerging from simple homophilous partners. These assumptions can be separated into two areas: individual children and groups. For children, the assumptions are (1) each has perfect attribute knowledge about other children; (2) attributes are equally important developmentally, and this importance does not vary over time; (3) the within-child attribute level is stable and is not modified by play; (4) no costs are associated with play (e.g., social standing or energy expenditure); and (5) arousal levels are uniform across all children. At the group level, we assume that (1) groups form around homophilous attributes and attribute levels, and this formative mechanism is not affected by group or class size; and (2) group composition and entry criteria may or may not evolve. Clearly, these assumptions are untenable, yet at PlayMate's initial stages of development, they were necessary to ensure tractability. Ideally, the refined simulation would address each of the aforementioned assumptions, either singularly or in a configuration that would allow us to track the dynamics of the groups.

Third, PlayMate is currently configured to systematically simulate a child modifying his or her interaction with another child contingent on sex, memory, and a single attribute on the assumption that similarity of attribute level, combined with sex, establishes the requisite homophily. In effect, each child is represented as a two-dimensional agent. In reality, children probably evaluate each other in *n*-dimensional space. Although the length of dimensionality, its configuration, and possible differential weighting of each dimension are unknown, it is possible to construct a vector score or an amalgamated index score using the combinatorial methods developed by Griffin (2000) and then use these scores in the simulation.

Finally, in PlayMate, children play with each other via an algorithm (i.e., random assignment within a round-robin format). Although the children — as agents — are heterogeneous across sex and attribute variation, they are not imbued with the ability to evolve beyond simply reacting and responding to other children on the primary putative factors (e.g., sex) assumed to foster group formation and adaptation to changing environments. Although agent diversity is present (Page, 2002), PlayMate fails to maximize it in the service of the research question. This limitation is not unusual in many ABMs, but in PlayMate, with its basis in the simulation of socio-developmental processes, lacking a mechanism for intra-agent recognition, learning, and evolution restricts the validity, robustness, and generalizability of the model. Two prominent methods are being used in ABMs that would address this lack of intra-agent communication. The first is Holland's tagging method (Holland, 1995; Riolo, 1997). Tags are a form of primitive communication that involve signals. They indicate a property that an agent has, and other agents can view the tags and take action, making assumptions based on the information. The second approach is via reputation systems (Alt and King, 2002; Mui, et al., 2002; Sabater and Sierra, 2002). Within these systems, each agent possesses a reputation based on group affiliation, direct exchanges with other agents, and information obtained indirectly from other agents. This setup mimics the plausible process that children may use to determine with whom and why they play. Moreover, coupling diverse agents with rules allowing variability in response to each exchange (based on attributes and rule variations) generates better modeling of the richness and complexity underlying human engagement, reaction, and change. Integrating this method of intra- and interagent behavior into PlayMate would add realism to our model. In turn, this capability would allow us to attest to the veridicality of the model.

REFERENCES

Alt, G., and K. King, 2002, "A Dynamic Reputation System Based on Event Knowledge," in S. Rabin, ed., *AI Game Programming Wisdom*, Hingham, MA: Charles River Media, Inc., pp. 426–435.

Arthur, W.B., 1994, "Inductive Reasoning and Bounded Rationality," *American Economic Review* **84**:406–411.

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

Bankes, S., 2002, "Tools and Techniques for Developing Policies for Complex and Uncertain Systems," *Proceedings of the National Academy of Sciences* **99**(3):7263–7266.

Berndt, T.J., 1982, "The Features and Effects of Friendship in Early Adolescence," *Child Development* **53**:1447–1460.

Borgatti, S.P., M.G. Everett, and L.C. Freeman, 2002, *UCINET 6.0*, Columbia: Analytic Technologies.

Cairns, R., H. Xie, and M. Leung, 1998, "The Popularity of Friendship and the Neglect of Social Networks: Toward a New Balance," in W. Bukowski and A. Cillessen, eds., *Sociometry Then and Now: Building on Six Decades of Measuring Children's Experiences with the Peer Group*, *New Directions for Child and Adolescent Development* **80**:25–54.

Carley, K.M., 1996, *Validating Computational Models*, Carnegie-Mellon University working paper. See http://www.casos.ece.cmu.edu/casos_working_paper/validating_abstract.htm.

Carley, K.M., 2002, "Simulating Society: The Tension between Transparency and Veridicality," in C. Macal and D. Sallach, eds., *Proceedings of the Workshop on Social Agents: Ecology, Exchange, and Evolution*, Chicago, IL: Argonne National Laboratory and University of Chicago, pp. 103–113.

Casti, J.L., 1997, *Would-be Worlds: How Simulation Is Changing the Frontiers of Science*, New York City, NY: John Wiley and Sons.

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

Cederman, L.-E., 2002, "Levels of Complexity: Endogenzing Agent-based Modeling," in C. Macal and D. Sallach, eds., *Proceedings of the Workshop on Social Agents: Ecology, Exchange, and Evolution*, Chicago, IL: Argonne National Laboratory and University of Chicago, pp. 175–184.

Coie, J.D., K.A. Dodge, and H. Coppotelli, 1982, "Dimensions and Types of Social Status: A Cross-age Perspective," *Developmental Psychology* **18**:557–570.

Conte, R., R. Hegselmann, and P. Terna, 1997, *Simulating Social Phenomena*, Berlin: Springer.

Conte, R., B. Edmonds, S. Moss, and K.R. Sawyer, 2001, "Sociology and Social Therapy in Agent-based Social Simulation: A Symposium," *Computational & Mathematical Organization Theory* **7**:183–205.

DiPietro, J.A., 1981, "Rough and Tumble Play: A Function of Gender," *Developmental Psychology* **17**:50–58.

Espelage, D.L., M.K. Holt, and R.R. Henkel, 2003, "Examination of Peer-group Contextual Effects on Aggression during Early Adolescence," *Child Development* **74**:205–220.

Fabes, R.A., L.D. Hanish, C.L. Martin, and N. Eisenberg, 2002, "Young Children's Negative Emotionality and Social Isolation: A Latent Growth Curve Analysis," *Merrill Palmer Quarterly* **48**:284–307.

Fabes, R.A., C. Martin, L.D. Hanish, M. Anders, and D.A. Madden-Derdich, 2003a, "Early School Competence: The Roles of Gender-segregated Play and Temperamental Regulation," *Developmental Psychology* **39**:848–859.

Fabes, R.A., C.L. Martin, and L.D. Hanish, 2003b, "Young Children's Play Qualities in Same-, Other-, and Mixed-sex Peer Groups," *Child Development* **74**(3):921–932.

Feber, J., 1999, *Multi-agent Systems: An Introduction to Distributed Artificial Intelligence*, Reading, MA: Addison-Wesley.

Fogel, A., 1993, *Developing through Relationships: Origins of Communication, Self, and Culture*, Chicago, IL: University of Chicago Press.

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

Girvan, M., and M.E.J. Newman, 2002, "Community Structure in Social and Biological Networks," *Proceedings of the National Academy of Sciences* **99**:8271–8276.

Goldspink, C., 2002, "Methodological Implication of Complex Systems Approaches to Sociality: Simulation as a Foundation for Knowledge," *Journal of Artificial Societies and Social Simulation* **5**(2).

Griffin, W.A., 2000, "A Conceptual and Graphical Method for Converging Multi-subject Behavioral Observational Data into a Single Process Indicator," *Behavior Research Methods, Instruments, and Computers* **32**(1):120–133.

Griffin, W.A., 2003, "Agent-based Models and Computer Simulation as a Methodological Aid for the Study of Micro-social Behavior," presented at the *2003 Biennial Meeting of the Society for Research in Child Development*, Tampa, FL, April 24–27.

Griffin, W.A., W. Cree, C. Martin, R. Fabes, and L. Hanish, 2002, "Emergent Structure in Children's Play Group Formation," presented at the *First Lake Arrowhead Conference on Computational Social Science and Social Complexity*, UCLA Center for Human Complex Systems, Lake Arrowhead, CA, May 9–12.

Griffin, W.A., C. Martin, R. Fabes, L. Hanish, M. Anders, S. Leonard, and M. Herzog, 2003, "A Multi-agent Computational Model of the Evolution in Children's Playgroup Formation," presented at the *Second Lake Arrowhead Conference on Computational Social Science and Social Complexity*, UCLA Center for Human Complex Systems, Lake Arrowhead, CA, March 19–23.

Hanish, L.D., and N.G. Guerra, 2002, "A Longitudinal Analysis of Patterns of Adjustment Following Peer Victimization," *Development and Psychopathology* **14**:69–89.

Hanish, L.D., C.L. Martin, R.A. Fabes, S. Leonard, M. Anders, L. Dinella, M. Herzog, and D. Madden-Derdich, 2003, *Peer Contagion Effects on Young Children's Externalizing Symptomatology*, Arizona State University (unpublished manuscript).

Hartup, W.W., 1996, "The Company They Keep: Friendships and Their Developmental Significance," *Child Development* **67**:1–13.

Hartup, W.W., 1999, "Peer Experience and Its Developmental Significance," in M. Bennett, ed., *Developmental Psychology: Achievements and Prospects*, Philadelphia, PA: Psychology Press, pp. 106–125.

Hartup, W.W., and B. Laursen, 1999, "Relationships as Developmental Contexts: Retrospective Themes and Contemporary Issues," in W.A. Collins and B. Laursen, eds., *Relationships as Developmental Contexts*, Mahwah, NJ: Erlbaum, pp. 13–35.

Hirschi, T., 1969, *Causes of Delinquency*, Berkeley, CA: University of California Press.

Holland, J.H., 1995, *Hidden Order: How Adaptation Builds Complexity*, Reading, MA: Addison-Wesley.

Hubert, L.J., and J. Schultz, 1976, "Quadratic Assignment as a General Data Analysis Strategy," *British Journal of Mathematical and Statistical Psychology* **29**:190–241.

Humphreys, P., 1997, "Emergence, Not Supervenience," *Philosophy of Science* **64**:337–345.

Jin, E.M., M. Girvan, and M.E.J. Newman, 2001, "The Structure of Growing Social Networks," *Phys. Rev. E*. **64**:046132.

Kindermann, T., 1998, "Children's Development within Peer Groups: Using Composite Social Maps to Identify Peer Networks and to Study Their Influences," *New Directions in Child Development*, pp. 55–82.

Krackhardt, D., 1988, "Predicting with Networks: Nonparametric Multiple Regression Analysis of Dyadic Data," *Social Networks* **10**:359–381.

Ladd, G.W., and S.M. Profilet, 1996, "The Child Behavior Scale: A Teacher-report Measure of Young Children's Aggressive, Withdrawn, and Prosocial Behaviors," *Developmental Psychology* **32**:1008–1024.

Leaper, C., 1994, "Exploring the Consequences of Gender Segregation on Social Relationships," in C. Leaper, ed., *The Development of Gender and Relationships*, San Francisco, CA: Jossey-Bass, pp. 67–86.

Leik, R.K., and B.F. Meeker, 1995, "Computer Simulation for Exploring Theories: Models of Interpersonal Cooperation and Competition," *Sociological Perspectives* **38**(4):463–482.

Maccoby, E.E., 1998, *The Two Sexes: Growing Up Apart, Coming Together*,” Cambridge, MA: Harvard University Press.

Maccoby, E.E., and C. Jacklin, 1987, “Gender Segregation in Childhood,” *Advances in Child Development and Behavior, Vol. 20*, Orlando, FL: Academic Press, pp. 239–287.

Macy, M.W., and R. Willer, 2002, “From Factors to Actors: Computational Sociology and Agent-based Modeling,” *Annual Review of Sociology* **28**:143–166.

Martin, C.L., and R.A. Fabes, 2001, “The Stability and Consequences of Young Children’s Same-sex Peer Interactions,” *Developmental Psychology* **37**:431–446.

Mitchell, M., 1996, *An Introduction to Genetic Algorithms*, Boston, MA: MIT Press.

Mui, L., A. Halberstadt, and M. Mohtashemi, 2002, “Notions of Reputation in Multi-agents Systems: A Review,” in *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, Bologna, Italy, July, pp. 280–287.

Newman, M.E.J., 2003, “The Structure and Function of Complex Networks,” *SIAM Review* **45**:167–256.

Page, S.E., 2002, “The Interplay of Differences,” in C. Macal and D. Sallach, eds., *Proceedings of the Workshop on Social Agents: Ecology, Exchange, and Evolution*, Chicago, IL: Argonne National Laboratory and University of Chicago, pp. 311–327.

Riolo, R.L., 1997, “The Effects and Evolution of Tag-mediated Selection of Partners in Populations Playing the Iterated Prisoner’s Dilemma,” in T. Bäck, ed., *Proceedings of the 7th International Conference on Genetic Algorithms (ICGA97)*, San Francisco, CA: Morgan Kaufmann, pp. 378–385.

Rocha, L.M., 1999, *From Artificial Life to Semiotic Agent Models*, LA-UR-99-5475, Los Alamos National Laboratory, Los Alamos, NM.

Rubin, K.H., D. Lynch, R. Coplan, L. Rose-Krasnor, and C.L. Booth, 1994, “Birds of a Feather…”: Behavioral Concordances and Preferential Personal Attraction in Children,” *Child Development* **65**:1778–1785.

Rubin, K.H., W. Bukowski, and J.G. Parker, 1998, “Peer Interactions, Relationships, and Groups,” in W. Damon (Series ed.) and N. Eisenberg (Volume ed.), *Handbook of Child Psychology: Vol. 3. Social, Emotional, and Personality Development*, 5th Ed., New York City, NY: Wiley, pp. 619–700.

Rykiel, E.J., 1996, “Testing Ecological Models: The Meaning of Validation,” *Ecological Modelling* **90**:229–244.

Sabater, J., and C. Sierra, 2002, “Reputation and Social Network Analysis in Multi-agent Systems,” in *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, Bologna, Italy, July, pp. 475–482.

Serbin, L., I.J. Tonick, and S.H. Sternglanz, 1977, "Shaping Cooperative Cross-sex Play," *Child Development* **48**:924–929.

Serbin, L., L.C. Moller, J. Gulko, K. Powlishta, and K.A. Colburne, 1994, "The Emergence of Gender Segregation in Toddler Playgroups," in C. Leaper, ed., *Childhood Gender Segregation: Causes and Consequences*, San Francisco, CA: Jossey-Bass, pp. 7–17.

Sutherland, E.H., and D.R. Cressey, 1974, *Criminology (9th ed.)*, Philadelphia, PA: Lippincott.

Thorne, B., 2001, "Girls and Boys Together…But Mostly Apart: Gender Arrangement in Elementary School," in R. Satow, ed., *Gender and Social Life*, New York City, NY: Wiley, pp. 153–166.

Watts, D.J., and S.H. Strogatz, 1998, "Collective Dynamics of 'Small-world' Networks," *Nature* **393**: 440–442.

ALIGNMENT AND VALIDATION IN AN AGENT-BASED MODEL OF ON-LINE B2C AUCTIONS

KUMAR MEHTA,* University of Connecticut, Storrs, CT
SIDDHARTHA BHATTACHARYYA, University of Illinois, Chicago, IL

ABSTRACT

This paper discusses the development and validation of an agent-based computational model of the on-line B2C auction marketplace. The model incorporates all of the relevant elements of the environment (i.e., consumers, retailers, and auctioneers), allowing investigation of various aspects of B2C auctions. A specific focus is on the development methodology that ensures alignment of the agent models with previously proposed analytical models, as well as its subsequent validation against field-observed price formations. Such alignment is critical to ensure that the agent-based model embodies known economic and behavioral principles and produces known or field-observable results, so that it can be subsequently used for studying different aspects of the B2C auction marketplace and also to aid in the design of such auctions.

Keywords: On-line auctions, agent-based model, alignment, validation

INTRODUCTION

The environment in which on-line auctions operate raises numerous research questions, ranging from issues dealing with the design of these auctions to issues of social welfare. Taking into account the complexity of the environment, factors that can influence the participation and outcome of an auction include the nature of the Internet, prices in the retail market, demographics and behavior of the participating consumers, and the design of the auction itself. Development of an "all-encompassing" single analytical model of the market is not feasible given the level of complexity involved and the degree to which one component may directly or indirectly affect outcomes in the B2C auction market. For example, the revenue outcome in an auction could be determined by (1) the consumers' ability/inability to search the retail market; (2) the nature of the retail market itself, in terms of the number of retailers offering the product, the posted prices, etc.; (3) the auction mechanism, in terms of duration of the auction, number of consumers demanding the product, quantity being auctioned, auction's format and rules, etc.; and (4) the demography of the consumers participating in the auction, their search of price-related information, and their bidding behavior. Although a separate theoretical model could be developed for each of these specific cases, it would prohibit understanding the potentially complex interactions between one or more factors that simultaneously could be at work in the marketplace.

This paper discusses the development and validation of a computational agent-based model (ABM) of the electronic auction marketplace. Such a model allows investigation of the various aspects of B2C auctions by incorporating all of the relevant elements of the environment

* *Corresponding author address:* Kumar Mehta, Department of Operations and Information Management, 2100 Hillside Road, University of Connecticut, Storrs, CT 06269; e-mail: Kumar.Mehta@business.uconn.edu.

(i.e., consumers, retailers, and auctioneers). By undertaking a bottom-up approach to develop models of individual agents based on existing knowledge and findings, the paper specifically focuses on the development methodology that ensures alignment of the agent models with previously proposed analytical models, as well as subsequent validation against field-observed price formations. Such alignment is critical for two reasons. First, it ensures that the ABM embodies known economic and behavioral principles and produces known or field-observable results so that it can subsequently be used for studying various aspects of the B2C auction marketplace. Second, it aids in the design of such auctions.

A simple validation of output from the ABM vis-à-vis field-observed data is insufficient to prove the adequacy and appropriateness of models used for modeling agents' behaviors. We propose a methodology that ensures an alignment in model selections, correspondence in conditions for output generation, and final validation by means of output comparisons. A caveat for the reader: the proposed methodology is for a particular class of problems that aim to build an ABM of a real-world phenomenon with the objective of utilizing the model for normative as well as predictive research. Our objective is to use the observable parameters of the real-world marketplace to model the properties of the agents (i.e., retailers, consumers, and auctions) so that the computational model essentially provides a "synthetic test bed" for simulating the market, allowing for future normative and/or predictive studies.

The general nature of the class of problems is best illustrated by using the typologies outlined in Axtell (2000) and Tesfatsion (2002). Axtell (2000) presents three distinct uses for adopting agent computation in the social sciences:

1. Traditional simulation of operations research problems,

2. Research areas where mathematical models can be formulated but not completely solved, and

3. Inability to mathematically model the problem, except at the rudimentary level in a piecewise manner.

This research fits into the second category of problems, which are in an analytical sense only partially soluble. In this class of problems, the theory or theories (as the case may be) informing the problem yield mathematical relationships, but these relationships are not directly soluble. A problem can resist detailed analysis in various ways, most commonly when no appropriate solution concept is available; stability of the equilibrium is uncertain; and in an analytical sense, it is not possible to readily solve for the dependence of the equilibrium on parameters of interest. Understanding B2C auctions resists full solubility because of each of these reasons. Tesfatsion (2002) categorizes the agent-based computational economics (ACE) research into the eight application areas shown in Table 1. In addition to belonging to the category of problems that resist full solubility as defined by Axtell (2000), in the context of the ACE application areas described above, the primary objectives of the ABM presented here are to replicate the real B2C auction market (parallel experiments with real and computational agents) to provide predictive capability, and to use bottom-up modeling of market processes to enable future testing of auction design (design of market protocols).

TABLE 1 Eight application areas of ACE research

1. Learning and the embodied mind
2 Evolution of norms
3. Bottom-up modeling of market processes
4. Network formation
5. Intra-firm organization
6. Use of ACE laboratories to test the design of market protocols
7. Use of ACE laboratories to test the design of computational agents for an automated market
8. Parallel experiments with real and computational agents

Because research areas that use ABM to tackle complex problems resist solubility, typically when the agents in the model are heterogeneous, Axtell (2000) suggests first building an agent-based computational model where the agents can be made heterogeneous. The ABM then can be docked with the analytical model by imposing constraints in the simulation that are identical to those in the analytical model (usually homogeneity of agents). The docked ABM should then reproduce the known analytical results providing the first-order validation of its agents. Following the validation of the ABM, the assumptions can be relaxed for a systemic study. In designing agent-based systems intended to mimic the real world, the issue of docking becomes especially significant. In such scenarios, given the lack of a theoretical model that is soluble, direct avenues to validate the ABM do not exist. Our ABM falls in this real-world category of problems, and no earlier research exists in docking the ABMs in such cases.

The ACE research area most relevant for this study aims to mimic the real-world market through bottom-up modeling of market processes. To the best of our knowledge, no ACE research has been conducted with the objective of replicating a real-world system. The research objective for mimicking a real-world system is to understand the underlying dynamics of the observed emergent phenomenon by modeling an "equivalent" phenomenon in the laboratory using agents (human or computational) in market conditions equivalent to "real-world" settings. In contrast, the objective of this research (seeking to mimic the price formation in a B2C auction market) is not only to aid in the understanding of underlying dynamics, but also to design and validate an ABM of the B2C auction market that possesses predictive capabilities. This objective necessitates that we replicate the market (rather than just mimic) to prove the robustness of models used for agent's behaviors, allow the system's use for predictive purposes, and use as a "synthetic test bed" for evaluating and designing auctions. This objective of replicating the market raises additional research issues with respect to verification and validation of the computational model as discussed in the next section.

We adopt a multi-stage approach for constructing the agent-based system and its validation. A first part of this research proposed (and empirically using field-observed data) a revenue model for the auctions, based on interaction between the retail and auction market by way of consumers' search for price-related information (Mehta and Lee, 2003). This paper identifies the relevant agents and develops detailed specifications for each of the agent's behaviors in alignment with the environment and broad constructs of the earlier model. These specifications for agent models are obtained through deconstruction and specification of lower-level processes, while maintaining the theoretical alignment at the aggregate level, for example, consumers' search the posted-price market until expected savings from additional observations are unlikely to offset the marginal cost of search. The introduction of newer models at detailed specifications of these behaviors introduces new variables and assumptions that require further validation. In contrast to the validation of the theoretical model conducted using the field-observed final revenues, the ABM is validated by using the entire price-formation data from the auctions and posted-price data from retailers for the same and related products.

RESEARCH METHODOLOGY

Hales (1998) illustrates the methodological frameworks adopted in research dealing with artificial societies (ASoc). The methodologies illustrated by Hales (1998) include existence proof, behavior modeling, theory testing, theory building, and explanation finding. Because our research cannot be strictly categorized as a typical ASoc work but does aim to build an ASoc equivalent to a real marketplace, we borrow the elements of the framework to develop and illustrate the proposed methodology (Figure 1).

Research using advanced computer modeling (ACM) can be considered as a set of theories *T* informing the formulation of a set of agent-based models *M*; a set of runs *R*, comprising the execution of simulations that embody *M*; and a set of observations *O* obtained from the runs *R*. Axtell, et al. (1996) align the two computational models based only on establishing equivalence of their outputs. For modeling some real-world phenomenon, because of the flexibility accorded by ACM, one risks building an overly complex model, and mere equivalence of output does not provide a sufficient guarantee of appropriateness or adequacy of *M*. An overly complex model under certain settings, however, can produce an equivalent output; significance established using various statistics from comparison of outputs does not imply validation of the model. To overcome such pitfalls of building overly complex models and to aid in selection of *M*, it is necessary to first engage in the testing of theory *T* that inform the

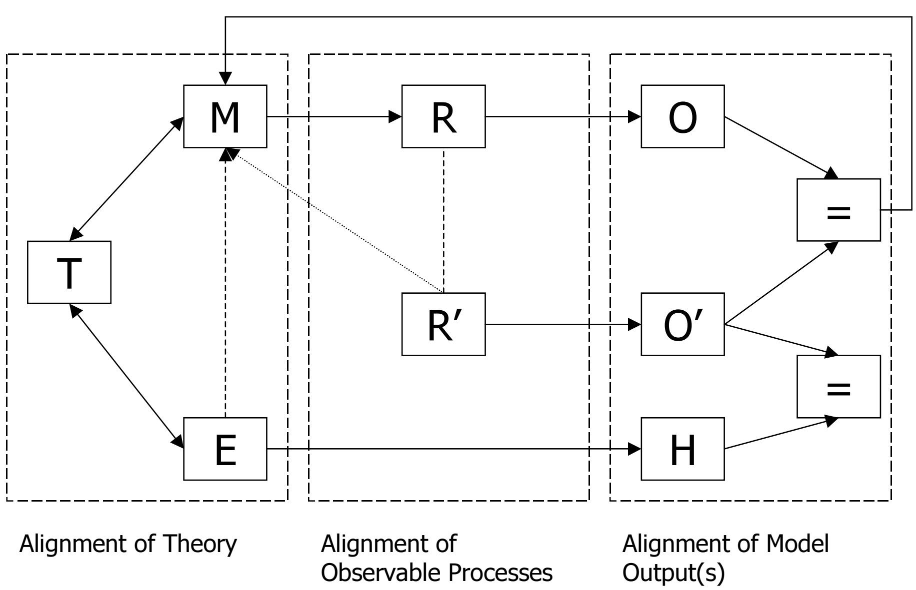

FIGURE 1 Methodology for alignment and validation of ACM

models M. Theory testing involves translation and/or abstraction of some existing T pertaining to a set of real-world process R' into a set of explanations E. To obtain support for E, a set of hypotheses H is formulated and tested against a set of measurements O' yielded by R'. Validation of H thus implies validation of abstraction E and by induction gains support for applicability of T for understanding R'. Alignment of the ACM to the real world and its validation thus requires: alignment of theory, alignment of observable processes, and alignment of output. These three alignments together constitute a successful validation of the ACM.

Alignment of Theory

In alignment of theory, the T used for abstracting E are now used in the construction of M. To leverage the richness in modeling available in ACM, the T previously abstracted in E can now be decomposed into multiple models during construction of M while maintaining corresponding equivalence with E. In this context, traditional numerical simulations involve direct translation of theoretical formulations, followed by relaxations of assumptions and constraints that were necessary for maintaining analytical tractability. Such direct translation implicitly maintains an alignment. In the case of ACM, however, the explicit formulation of lower-level processes may bring to the surface models that were nonexistent at higher levels of abstraction, thus making a direct comparison through mapping infeasible.

We adopt the following four approaches for alignment of M:

1. *Qualitative approach* – A qualitative approach is utilized when no suitable formulation exists for a given behavior, but evidence of the behavior has been reported. In such cases, we draw on existing domain knowledge to build the simplest possible formulation that can adequately represent the known behavior.

2. *Direct mathematical formulation* – Direct mathematical formulation is adopted for proven behavior models.

3. *Higher-level theoretical principles* – Only higher-level theoretical principles are maintained because underlying assumptions of T and E have been relaxed, resulting in some of the lower-level models coming from Approach 1, thereby making Approach 2 infeasible. For example, consumers' search of the retail market yields posted prices not only for the item being auctioned, but also for items that can be considered as substitutes. Since no known models exist for incorporating price-related information of substitutes, we introduce a model based on "degree of similarity" to allow for assimilation of all related information into formulation of "willingness to pay," while maintaining the theoretical principles of search models (i.e., search is costly and consumers stop searching when expected savings from additional searches cannot be offset by the cost incurred).

4. *Abstraction of models to a level where they are replaced by exact values (states) observable from the real world* – While this approach might seem counterintuitive considering that the principles of ACM emphasize decomposition rather than abstraction, it is necessary in cases where

> specification of the actual ABM is not critical to the objective of the research and introduction of the model can destabilize the alignment process. For example, in our model, the retailers' product offerings and pricing models are abstracted away and replaced by field-observed, product-price offering information. Though introduction of retailers' models for product offering and pricing strategies would lend richness to the overall model, they are not critical to the objective of this research, and the increase in complexity of the overall model would prevent proper validation of the model. In fact, replacement with the actual product price information provides a point of alignment with the real world and a better comparison environment for alignment of observable processes and in turn outputs by ensuring that any convergence/divergence between outputs results from adequacy/inadequacy of *M* of primary importance.

Alignment of Observable Processes

Alignment of observable processes requires equivalence in conditions of *M* to conduct *R* such that a meaningful correspondence with *R′* allows for the most direct comparison of *O* and *O′*. However, the conditions producing *R′* may be only partially observable, and every effort should be made to replicate the observable conditions in *M*. For example, as mentioned earlier, when the retailers' product offering and pricing models are replaced with actual posted prices, the settings of the auction in terms of product offered on auction, duration, bidding rules, etc., are replicated in *M* to maintain a correspondence between *R* and *R′*.

Alignment of Outputs

The comparison of outputs *O* and *O′* constitutes the final step. If an alignment of theory and process is complete, no significant differences should be observed between *O* and *O′*. However, since the real world is not entirely transparent to the researcher, complete fore-knowledge of model specifications (values of certain parameters) must be estimated. If the alignment of the theoretical models and formulations of the processes are deemed adequate, any observed divergence between *O* and *O′* can be assumed to be a result of an incorrect estimation of these parameters. The results from comparison of *O* and *O′* can thus be used to revise parameter values iteratively until an equivalence is established. Since adequacy of formulations cannot be guaranteed by merely conducting alignment of theory and processes, it is possible for parameter estimates to compensate for any shortcomings and provide a false sense of "validation." As a result, follow-up testing is necessary to establish that parameter estimates were the only source of observed divergence and the estimates obtained did not compensate for inadequacies in *M*. The final validation of the model is thus conducted using independent sets of *R*2′ and *O*2′ and parameter estimates obtained from *O′*. A successful comparison of *O2′* with *O2* yielding equivalence constitutes completed validation of the ACM. Additional support for the validation can also be obtained through sensitivity analysis of the parameter estimates.

AGENT-BASED MODEL

In the defined context, the relevant agents identified are the retailer, consumer/bidder, auction, and product. The level of detail specified for each agent is limited to actions that are directly relevant to the events in a single set of auctions. Thus, auctions are examined in a somewhat static setting where agents do not learn from one auction to the next. Figure 2 provides an overview of the environment and interactions between agents.

A consumer desires to purchase a product and is willing to accept some perfect and imperfect substitutes for the desired product. Before making the final purchase, the consumer must select a channel (auction or retail), seller (auctioneer or retailer), and the product-price combination offering the best "deal" (utility maximization). To make this decision, the consumer searches through the retail market to gather price-related information for desired and related products (perfect and imperfect substitutes), and evaluates product-price observations based on the similarity with the desired product to form the highest willingness to pay. For example, assume the consumer desires a product for which the lowest retail price observed is $100. If a similar product is available, but it provides only 80% of the utility because it lacks some of the features of the desired product, the consumer will be willing to pay $80 for the similar product.

Thus, upon arrival in the market, the consumer assumes a search state and engages in search for price-related information from the retail market for the desired product and the acceptable substitutes. Following the consumer's search of price-related information, the consumer "visits" the auction, participating only if the product being auctioned is an acceptable

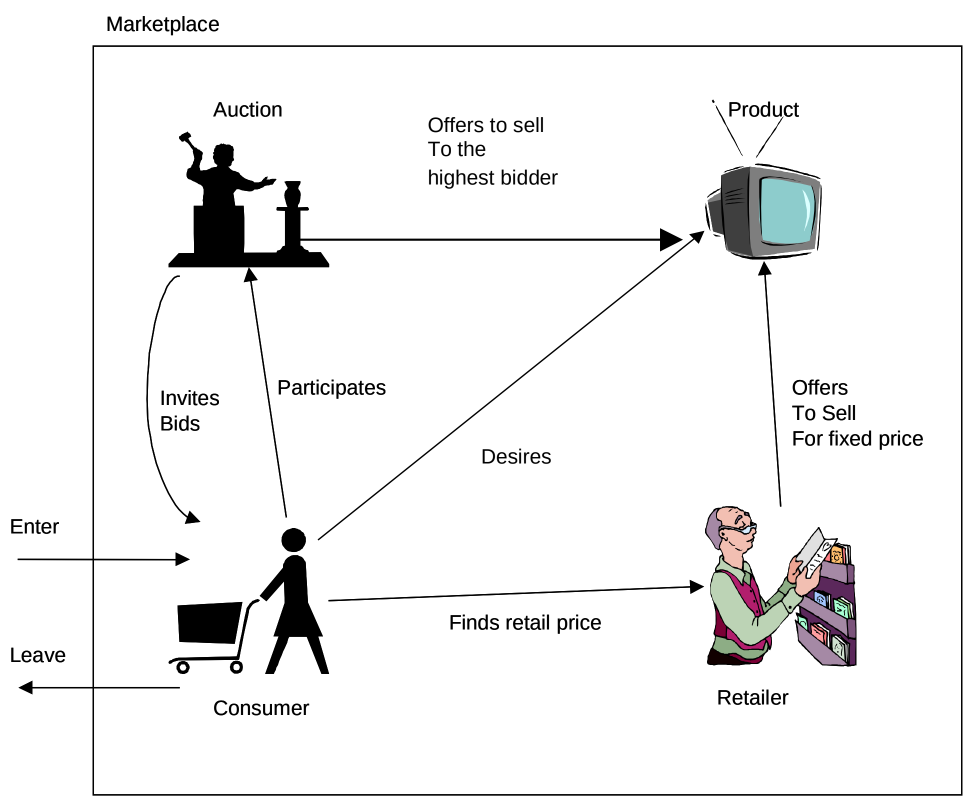

FIGURE 2 Schematic of the implemented agent-based simulation model

substitute for the desired product and continues to participate as long as the required bid for winning the auction is less than consumer's willingness to pay for the product. The consumer leaves the auction if the product being auctioned is not acceptable or the required winning bid exceeds the consumer's willingness to pay.

In any time period, when participating in the auction, the consumer assumes one of five states:

1. *Watch* – Monitor the progress of the auction,

2. *Sleep* – Remain dormant,

3. *Bid* – Place a bid in the auction,

4. *Evaluate bid* – Evaluate the success of a bid placed in the previous time period, or

5. *Leave* – Leave the auction if the willingness to pay has been exceeded.

At each time period, during participation in the auction, the consumer undertakes actions according to the current state and decides the state for the next period. This process continues at each time period until either the consumer decides to leave or the auction closes. Depending on the state assumed, the consumer agent obtains the necessary information to execute actions for that state and for deciding on the state to assume in the next time period. This decision is determined by the state transition rules and the state↔behavior relationships model as shown in Figure 3 (state determines behavior in time t, the behavior in turn determines the state in $t + 1$). The dependence of states and actions allows for each consumer to act independently, obtaining and processing potentially different information in each time period. In a given time period, two consumer agents possessing identical information and in the same state can also decide differently owing to differences in their attribute values, such as search efficiency, risk profile, and desired product.

The retail market, as modeled, consists of various retailer agents, each of whom offers to sell a product (not necessarily the same product) at a fixed posted price. The only consumers that observe this price are those whose search for price-related information leads them to this retailer.

The auction (auctioneer) is modeled as an agent who offers to sell q quantity of a product to the highest bidders. The auction keeps track of the time elapsed and knows when to close the auction. The auction also advertises the current winning bids and the minimum required bid to displace the current winners. For each time period until close of the auction, all participating consumers are invited to submit new bids. After collecting the response from all the consumer agents, the auction agent combines the list of submitted bids with the list of current winning bids. The highest q bidders from the combined list are chosen as the new winners of the auction. In the case of a tie, the tie for the q'th position is broken by using arbitration rules, giving preference to the bidder whose first bid was placed earlier. If the two arrivals are simultaneous, the tie is broken by using a random draw with each of the tied bidders having an equal probability of winning. After completion of bid evaluation, the successful and unsuccessful bidders are informed of the results.

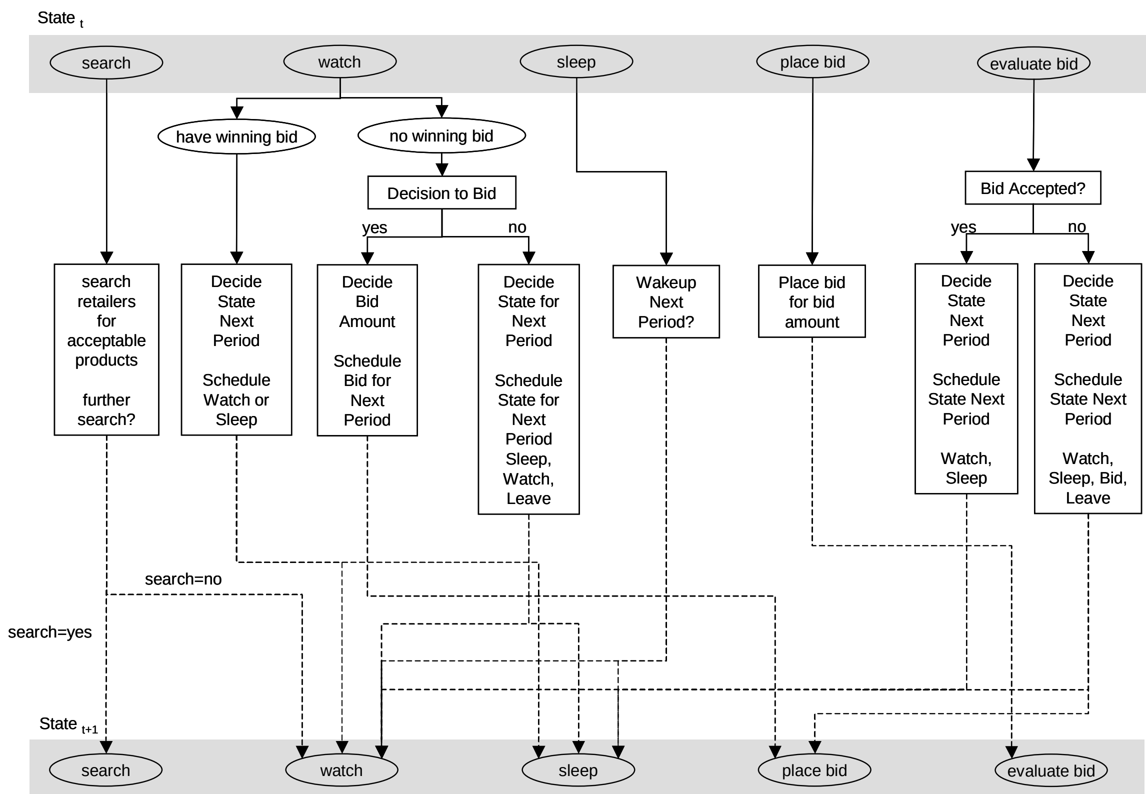

FIGURE 3 State transition rules and state actions dependence for the consumer agent

The auction also keeps track of closing time for ending the auction. The closing time is either static (auction closes at declared closing time) or dynamic (auction closing time is extended such that a fixed length of time must elapse after the last bidding activity before closing the auction). In the latter case, the auction does not inform the consumer agents about either the extension or the duration of the extension.

In their roles, only the auction, consumer, and retailer agents are modeled as “animate” in the sense that they are able to act autonomously and interact with other agents. Although the retailer agents interact with the consumer agents (reveal posted price), they do not change their state; retailer’s price offerings and specific product’s properties do not change for the duration. The product agents are modeled as “inanimate” because they lack the capacity to initiate an interaction with other agents. (For the detailed formulations of the individual models, see Mehta and Bhattacharyya [2003].)

Alignment of Observable Processes: Utilizing Field Observations

As discussed earlier, the alignment of observable processes involves setting conditions in the ACM such that a suitable correspondence between conditions for simulation runs R and the real-world conditions that define the context of R'. In this context, Figure 4 illustrates the class diagram of the ACM and identifies agents whose parameters are directly observable (at the level

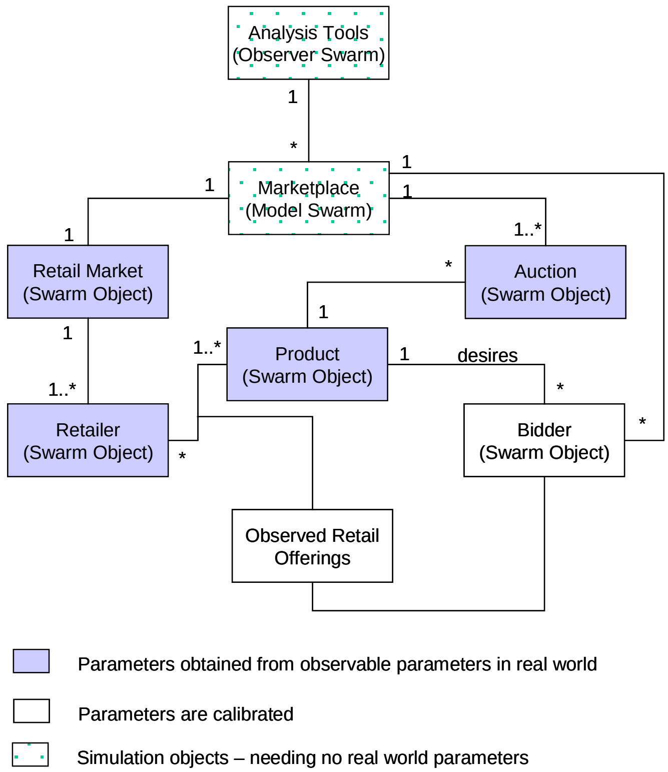

FIGURE 4 Class diagram of ACM for B2C auction

of modeled abstraction) and can be used to bring about the needed correspondence. The data collection for each of these is described in the next section.

Retail Market: Retailer and Product Data

Retail price information was collected for the hard drive market by using two shopping agent websites: pricescan.com and pricewatch.com. While many other websites offer comparative prices from multiple retailers, the goal was to collect enough data to model a sufficiently representative retail market. The search yielded a total of 1,436 posted prices from various on-line retailers for 183 different makes and models of hard drives. The hard drives were organized in order of similarity in their technical characteristics: storage capacity, rotation speed (rpm), and data transfer rate. The technical specifications were gathered from the manufacturer for each of the 183 hard drives. Hard drives that were technically identical in all respects (excluding make and model) were assigned the same product ID (integer value), with the IDs being in increasing order of storage capacity. This process resulted in a total of 77 products.

The final data set used for modeling the retail market in the simulation includes 1,436 retailers offering to sell one of the 77 products at a posted price. The product offered for sale through the auction is assigned an ID; it is the same as the one used for identifying the corresponding retail offerings for the product. In absence of any specific information regarding demand, the consumer agents are randomly assigned a desired product ID with equal probability.

Auction Market Data

Data were collected from UBid over a 3-day period for more than 40 different auctions belonging to the “hard drive” category. The auction settings data recorded were (1) product offered mapped to product IDs used in simulation, (2) duration of auction, (3) start bid, and (4) bid increment. In addition, the price-formation data were also collected for these auctions for use in alignment of outputs and final validation of the model (discussed in next section). All of these were single-unit auctions belonging to the hard drive category and, given the volatility in retail prices for computer-related components, the data collection period was limited to 3 days to prevent changes in the retail market from impacting the auction market.

Alignment of Outputs

Price-formation data were collected at 1-minute intervals from the start to the end of the auction for each of the auctions mentioned above. Collection of data was restricted to single-unit auctions to prevent inclusion of resellers who tend to participate in multi-unit auctions and bid for multiple quantities of the same item to reduce the per-unit shipping costs.

The objective was to obtain price-formation data from multiple auctions with identical settings in terms of auction parameters (i.e., product auctioned, duration of auction, and bidding increments). To validate the models of the underlying processes in the ABM, it is essential to replicate the auction market in general rather than replicate the events of a single auction. Comparison with multiple price-formation series from auctions with identical settings O' provides us with the price formation in general (mean of these series) along with variations between auctions because of other environmental uncertainties. For validating the ABMs, the results O yielded from multiple runs R of the simulation should reflect the price-formation series observed in the B2C auctions.

From the data collected, the price formation series identified for use in parameter estimation and validation of the ABM includes three auctions each for two different product items (three auctions of Western Digital 30 GB hard drive with manufacturer part number WD300AB and three auctions of Western Digital 40 GB hard drive with manufacturer part number WDC400BB).

To ensure proper validation of outputs and to avoid an over-fit solution to the observed price formations, the final validation of output of the ABM is conducted in two stages — calibration (parameter estimation) and validation. Calibration of the parameters is conducted by using the first group of data (three auctions of WD300AB) to obtain the simulated price-formation series, and parameters are revised to fit the simulated price-formation to the field-observed data. By using the calibrated parameters along with the auction settings for the second product (WDC400AB), results from the simulation are obtained and validated against a second group of data. The two items were selected to ensure that the product IDs were sufficiently unique to prevent any confounding unobserved effects during parameter estimation from also affecting the validation runs.

RESULTS AND DISCUSSION

Parameter Estimation and Validation

During the calibration process, three different random seeds were utilized, and the outcome of the simulations was compared against the observed price formation series of product WD300AB. The parameter estimates were refined until the simulated price formation O converged to field-observed data O', and statistical tests indicated predictive capability of the model at a better than 95% significance level. Following the calibration, 10 new random seeds were chosen, and simulations were carried out for validation against the 3 price formation series from auction of WDC400BB, which constituted the holdout sample for validation of the ABM. New random seeds were chosen for the simulation to prevent biases in parameter estimates that could have been caused by conditions generated by specific random number seeds. The parameters for auction settings were set to those used by the auctioneer for WBC400BB, and the product offered ID was set to the ID used in the retail market model to denote WDC400BB. The retailer agent's parameters are modeled using the retail market posted prices for various hard drives and do not change from calibration to validation. The consumer agent's parameters are based on the ones obtained from the calibration.

For both calibration and validation, to test for similarity between the simulated and field-observed price formation data, consider a field-observed, price-formation data from j'th auction, $Rj = \{B_{jt}\}$, and the simulation results with identical auction settings from i'th run, $R_i = \{B_{s,it}\}$, where $B_{s,it}$ and B_{jt} are bid levels at time t in the simulation results and field data, respectively. Given the nonlinear nature of price formation in auctions, mean bid levels were compared at 0, 2, 4, 6, 8, 10, 20, 40, 60, 80, 100, and 120 minutes. By using the mean bid levels from multiple runs of the simulation and the mean bid levels from field-observed B2C auctions, a paired t-test is conducted to statistically test the similarity of the two price-formation series. The calibration runs of the simulations (for product WD300AB) were only for parameter estimation. These estimates are then used to model the auction of another product (WDC400BB). The validation of simulation results thus obtained, against the field-observed, price-formation data constitutes the final proof of the ABM's ability to replicate price formation in the B2C auction market. The average bid levels obtained from the 10 simulations, along with the price formation data from the 3 B2C auctions, are shown in Figure 5.

A paired t-test comparison of the bid levels at the above-mentioned times indicates the difference between the two price-formation series to be statistically insignificant from 0 at 10%, proving that the ABM adequately captures the underlying processes at play in the field-observed B2C auctions. Examination of the residuals from comparison of the price formations (difference of the two series) also indicated no significant correlation of the residual with the field-observed price formation. Ideally, one should also test for equivalence of the variance at the same time periods between the simulated and field-observed data. Because of the paucity of field-observed data in terms of price-formation series from identical auction settings, we were unable to do so; however, we examined other aspects of the output to establish equivalence in the dynamics of the process. The comparison of simulation output with field-observed data and implications from an analytical model are summarized in Table 2.

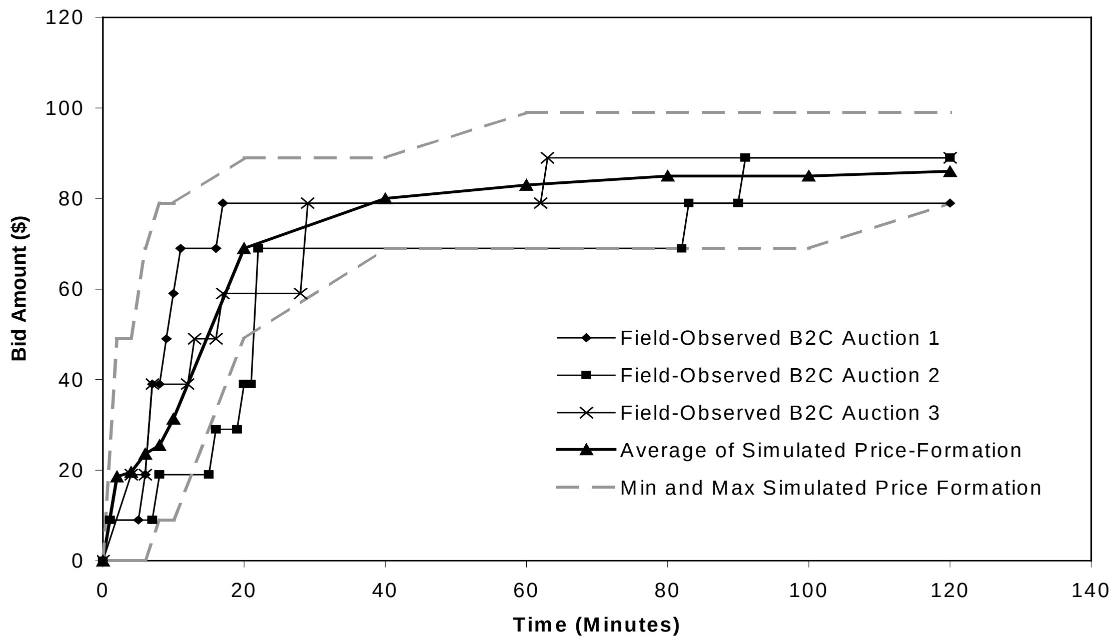

FIGURE 5 Comparison of simulated and field-observed, price-formation processes

TABLE 2 Comparison of simulated and field-observed, price-formation processes

Aspect of Simulation Output	Compared Against	Remarks
Mean bid levels at geometric time intervals	Field-observed data	Difference between mean bid levels is statistically insignificant from 0 at 90% confidence level.
Range of number of bids placed in auction	Field-observed data	Both data indicate 6 to 7 bids placed.
Number of bidders placing bids	Field-observed data	Field observations indicate 3 to 4 bidders bidding, whereas simulation indicates 3 to 5 bidders in 9 out of the 10 runs, and 7 bidders in 1 case.
Number of total participants needed to obtain the approximately 10% premium over the lowest posted price for the item being auctioned	Analytical model proposed by Mehta and Lee (2003)	Analytical model indicates approximately 20 participants, each observing more than 5 posted prices for the exact item. Simulation model indicates an average of about 17 participants desire the same item or are willing to accept it as the perfect substitute and observe 5–11 posted prices for the desired item and perfect substitutes.

CONCLUSION

Agent-based models offer a suitable mechanism for developing a realistic, all-encompassing model of the B2C auction marketplace. Given the flexibility accorded the modeler, however, one runs the risk of building overly complex models. Such models, even when generating output identical to that of the system being studied, would not appropriately represent the characteristics of individual agents and their actions. We propose that obtaining an agent-based computational model that adequately captures the system under study requires alignment at each of the following stages: (1) model selection, (2) observable processes, and (3) final output produced. The application of the proposed multi-staged methodology is illustrated in the context of the design, development, and validation of the ABM of B2C auctions.

Results from the agent-based simulation demonstrate the usefulness of this approach for replicating the dynamics of the auction market. The model is useful for investigating various aspects of B2C auctions, including the following:

- Examination of market characteristics, such as alternative distributions of posted prices, demand for items, and degree of product differentiation in the retail market;

- Consumer characteristics in terms of their search efficiencies and bidding behaviors; and

- Auction parameters related to the design of the auction, such as start bids, bid increments, and number of units on auction.

The methodology highlighted here is applicable across a range of areas adopting agent-based modeling of real-world systems/markets, including network pricing, bandwidth allocation, and dynamic routing in packet-switched networks.

REFERENCES

Axtell, R., R. Axelrod, J.M. Epstein, and M.D. Cohen, 1996, "Aligning Simulation Models: A Case of Study and Results," *Computational Mathematical Organization Theory* **1**(2):123–141.

Axtell, R., 2000, *Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences*, Working paper No. 17, Center on Social and Economic Dynamics, Brookings Institution.

Hales, D., 1998, "Artificial Societies, Theory Building and Memetics," presented at Memetics Symposium in *Proceedings of the 15th International Conference on Cybernetics*, Namur, Belgium: International Association for Cybernetics.

Mehta, K., and B. Lee, 2003, "A Model of Market Interactions: Auctions vs. Posted Price," Working Paper, University of Connecticut.

Mehta, K., and S. Bhattacharyya, 2003, "Design, Development and Validation of an Agent-based Model of Electronic Auctions," *Information Technology and Management* (forthcoming).

Tesfatsion, L., 2002, "Agent-based Computational Economics: Growing Economies from the Bottom Up," *Artificial Life* **8**(1):55–82.

A MULTI-MODEL DOCKING EXPERIMENT OF DYNAMIC SOCIAL NETWORK SIMULATIONS

J. XU,* Y. GAO, J. GOETT, and G. MADEY
University of Notre Dame, Department of Computer Science

ABSTRACT

Axtell, Axelrod, Epstein, and Cohen have described a "docking" or alignment process and an experiment for verifying simulations. By comparing simulations built independently with different tools, the docking process can be used to discover bugs, misinterpretations of model specifications, and inherent differences in toolkit implementations. When the behavior of multiple simulations is similar, verification confidence increases. North and Macal reported on such an experiment in which they used Mathematica, Swarm, and Repast to simulate the Beer Distribution Game (originally simulated using system dynamics simulation methods). This paper presents the results of docking a Repast simulation and a Java/Swarm simulation of four social network models of the open source software (OSS) community. Data about the SourceForge OSS developer site have been collected for more than two years. Membership in projects is used to model the social network of developers. Social networks based on random graphs, preferential attachment, and preferential attachment with both constant and dynamic fitness are modeled and compared with collected data. Furthermore, this paper describes how properties of social networks, such as degree distribution, diameter, and clustering coefficient, are used to dock Repast and Swarm simulations of four social network models. The simulations grow artificial societies that represent the SourceForge developer/project community. A by-product of the docking experiment is a set of observations concerning the advantages and disadvantages of the two toolkits for modeling such systems.

Keywords: Dynamic social network, docking, agent-based modeling, open source software

1 INTRODUCTION

Agent-based modeling (ABM) has become a popular computational methodology in recent years because researchers can simulate complex systems in a relatively straightforward way. Unlike traditional mathematical simulation tools, ABM simulates artificial worlds on the basis of components called *agents* and defines rules to determine their interactions. Although commonly used in simulations, ABM does not guarantee an accurate representation of a particular empirical application (Axelrod, 1997). In this context, Axtell, et al. (1996) claimed, "It seems fundamental to us to be able to determine whether two models claiming to deal with the same phenomenon can, or cannot, produce the same result."

* *Corresponding author address:* Jin Xu, 214 Cushing, University of Notre Dame, Notre Dame, IN 46556; email: jxu1@cse.nd.edu.

An agent-based simulation is validated in several ways:

- Simulation output is compared with the real phenomenon. This method is relatively simple and straightforward; often, however, complete real data cannot be obtained on all aspects of the phenomenon.

- Results of agent-based simulation are compared with results of mathematical models. The disadvantage of this method of validation is that mathematical models must be constructed, and these models can be difficult to formulate for a complex system.

- Other simulations of the same phenomenon can be docked. Docking is the process of aligning two dissimilar models to address the same question or problem. The objective is to investigate their similarities and differences, but, most important, to gain new understanding of the question or issue (Burton, 1998).

Axtell, et al. (1996) have described a docking or alignment process and experiment for verifying simulations. By comparing simulations built independently with different simulation tools, researchers can use the docking process to discover bugs, misinterpretations of model specifications, and inherent differences in toolkit implementations. When the behavior of multiple simulations is similar, confidence in verification increases. North and Macal (2003) reported on such an experiment in which they used Mathematica, Swarm, and Repast to simulate the Beer Distribution Game (originally simulated using system dynamics simulation methods). Ashworth and Louie (2002) performed docking by comparing results of the canonical Garbage Can Model with those of the NK model. Xu and Gao (2003) used Repast and Swarm to dock a random network model of the open source software (OSS) phenomenon. Although the above experiments show the importance and advantages of docking, only a few studies have been performed, and none has used topological properties of social networks as parameters.

This paper presents the results of docking a Repast simulation and a Java/Swarm simulation of four dynamic social network models of the OSS community. These results are part of a study of the OSS by a number of researchers.[1] Data regarding the SourceForge OSS developer site have been collected for more than two years. Developer membership in projects is used to model the social network of developers. Social networks based on random graphs, preferential attachment, and preferential attachment with both constant and dynamic fitness are modeled and compared to collected data. Properties of social networks, such as degree distribution, diameter, and clustering coefficient, are used to dock Repast and Swarm simulations of four social networks. The simulations grow artificial societies that represent the SourceForge developer/project community. As a by-product of the docking experiment, we provide observations on the advantages and disadvantages of the two toolkits for modeling such systems.

The remainder of this paper is organized as follows. Section 2 provides background on our OSS study and simulation. Section 3 discusses docking simulations using Repast and Swarm. Section 4 gives experiment results and comparisons, and Section 5 presents conclusions.

---

1 Researchers include Madey, et al. (2002a,b), Madey, et al. (2003a,b), Gao, (2003), Gao, et al. (2003a,b), and Xu and Gao (2003).

2 SOCIAL NETWORK MODEL

Social network theory is a conceptual framework through which we can view the OSS developer movement. The theory, built on mathematical graph theory, depicts interrelated social agents as nodes or vertices of a graph and their relationships as links or edges drawn between the nodes (Wasserman and Faust, 1994). The number of edges (or links) connected to a node (or vertex) is called the index or degree of the node.

Of special interest are the evolutionary processes and associated topological formations in dynamic growing networks. Early work in this field by Erdos and Renyi (ER) (in Barabasi, 2002) focuses on random graphs, those in which edges between vertices were attached in a random process, called ER graphs in this paper). The distributions of index values for random graphs, however, do not agree with the observed power law distribution for many social networks, including the OSS developer network at SourceForge. Other evolutionary mechanisms include the following:

- Watts-Strogatz (WS) model (Strogatz and Watts, 1998),
- Barabasi-Albert (BA) model with preferential attachment (Albert, et al., 1999; Barabasi and Albert, 1999; Barabasi, et al., 2000),
- Modified BA model with fitness (Barabasi, et al., 2001; Barabasi 2002), and
- Extension of the BA model (with fitness) to include dynamic fitness based on project life cycle (Gao (2003); Gao, et al. (2003a,b); Madey, et al., 2003a).

The WS model captures the local clustering property of social networks and was extended to include some random reattachment to capture the small world property but failed to display the power-law distribution of index values. The BA model added preferential attachment, while preserving the realistic properties of the WS model *and* displaying the power-law distribution. The BA model was extended with the addition of random fitness to capture the fact that sometimes newly added nodes grow edges faster than previously added nodes (the “young upstart” phenomenon).

The OSS development movement is a classic example of a dynamic social network; it is also a prototype of a complex, evolving network. Previous research suggests that the OSS network can be considered a complex, self-organizing system (Faloutsos, et al., 1999; Adamic and Huberman, 1999; Barabasi, 2002). These systems are typically composed of many locally interacting elements.

The OSS community can be described as a dynamic social network. Our model of the OSS collaboration network has two entities — developer and project. The network can be illustrated as a graph. In this network, nodes are developers. An edge is added if two developers participate in the same project, and the edge is removed if they no longer participate in that project. The study of the OSS collaboration network can help us to understand the evolution of the social network’s topology, the development patterns of each individual object, and the impact of the interaction among objects to the evolution of the overall network system.

We use ABM to simulate the OSS development community. Unlike developers, projects are passive elements of the social network. Thus, we define developers only as the agents that encapsulate a real developer's possible daily interactions with the development network. Our simulation is time stepped rather than event driven (one day of real time = one time step). Each day, a certain number of new developers are created. Newly created developers use decision rules to create new projects or join other projects. Further, each day existing developers can decide to abandon a randomly selected project, to continue their current projects, or to create a new project. A developer's selection is determined by a Java method based on the relative parameter and the degree of the developer.

3 DOCKING OSS COLLABORATION NETWORK SIMULATION

This section describes the docking of our OSS collaboration network simulation by two ABM tools — Java Swarm and Repast. Simulation details are compared between these two models.

3.1 The Docking Process

The docking process is an important stage of the OSS project (Freeh, et al., 2003). The initial simulation was written using Swarm. Docking is necessary in this project for two reasons:

- Testing the correctness of the Swarm implementation and
- Providing the Repast version of the OSS simulation that we would like to use in our future research.

Repast has several advantages for this project: it is written in pure Java, which makes debugging easier; it provides a graphic representation of the network layout; and, most important, Repast 2.0 provides a distributed running environment (Collier and Howe, 2003).

As shown in Figure 1, both Swarm and Repast simulations are docked for four models of the OSS network. Our docking process began when the author of the Swarm simulation wrote the docking specification. The Repast version was then written on the basis of the docking specification. Simulations are validated by comparing network attributes generated by running these two simulation models.

3.2 Swarm

Originally developed at the Santa Fe Institute (Minar, et al., 2002), Swarm is a software package for multi-agent simulation of complex systems. In the Swarm model, the basic unit is called an agent. Modelers can define a set of rules to describe the interaction of agents. Furthermore, Swarm also provides display, control, and analysis tools.

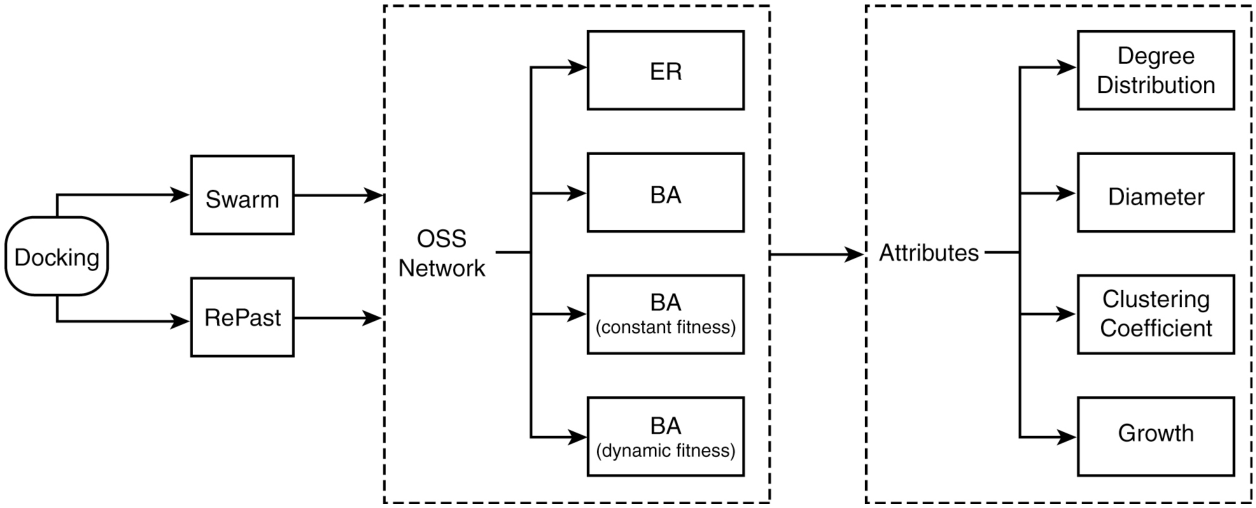

FIGURE 1 Docking process

Our Swarm simulation has a hierarchical structure that consists of a *developer* class, a *modelswarm* class, an *observerswarm* class, and a *main* program. The modelswarm handles the creation of developers and controls their activities. In modelswarm, a schedule is generated to define a set of activities of the agents. The observerswarm is used to implement data collection and draw graphs. The main program is a driver to start the entire simulation.

The core of a Swarm simulation consists of a group of agents. Agents in our simulation are developers. Each developer is an instance of a Java class. A developer has an identification, a degree that is the number of links, and a list of projects participated in by this developer. Furthermore, a developer class has methods to describe possible daily actions: create, join, abandon a project, or continue the developer's current collaborations. A separate Java method models each of the first three possibilities. A fourth method encapsulates a developer's selection of one of the three alternatives. Here, three model parameters appear. Each represents the probability of one of the three developer activities. Comparison of a randomly generated number to these probabilities determines which behavioral method the agent enacts.

3.3 Repast

Created by Social Science Research Computing at the University of Chicago, Repast is a software framework for agent-based simulation (Repast Home Page, 2003). Like Swarm, Repast provides an integrated library of classes for creating, running, displaying, and collecting data from an agent-based simulation (Collier, 2003). In addition, Repast is written in pure Java, which has better portability and extensibility than Swarm. Furthermore, Repast provides some different library packages that provide such features as network display, QuickTime movies, and snapshot.

Our Repast simulation of OSS developer network consists of a model class, a developer class, an edge class, and a project class. The class structure of the simulation differs from that of the Swarm simulation, in part because Repast has a graphic network display feature. The model class is responsible for creating and controlling the developers' activities. Furthermore, information collection and display are also encapsulated in the model class. The developer class is similar to that in the Swarm simulation. An edge class is used to define an edge in the OSS network. We also create a project class with properties and methods to simulate a project.

4 EXPERIMENT RESULTS AND COMPARISONS

This section describes docking of Repast and Swarm simulations on four OSS network models: ER, BA, BA with constant fitness, and BA with dynamic fitness. The results and a comparison are also presented.

4.1 Docking Procedure

The objective of our docking process was to verify our Repast migration against the original Swarm simulation. The process began with a comparison of degree distribution between corresponding models. Upon finding differences, we compared each developer's actions.

The first attempt at docking compared the degree distributions between these two simulations. The Swarm simulation used its built-in random number generator. The Repast simulation used the COLT random number generator from the European Laboratory for Particle Physics (CERN). From a graphic comparison of degree distribution for projects and developers over multiple runs of Swarm and Repast, we observed systemic differences between the two simulations' outputted data. Over one subdomain of the developer degree distribution, Swarm had a higher count than Repast. Over another subdomain, Swarm had a lower count. The next step in the docking process determined that the random number generators did not cause this difference. We ran the two simulations using exactly the same set of random numbers: each simulation used the same random number generator with the same seed. The developer and project degree distributions from these runs, however, revealed similar systemic differences between the two simulations.

To determine the exact reasons for the differences, we had the simulations log the action that each developer took during each step. Through comparison of these logs, two reasons emerged to explain the differences.

First, one simulation occasionally threw an SQL exception (our data are stored in a relational database for post-simulation analysis). To recover from such an error, the simulation does not log the developer's action: it moves on to the next developer. Because the developer's previous actions affect its future actions, one error can cause more discrepancies between the two simulations at future time steps. The cause of this error was a problem with the primary keys in the links table of our SQL database (this problem is a programming bug). Further inspection of the data logs showed that a simulation's data snapshots, which are used in analyzing macroscopic graph properties, were out of phase by one unit of time. Even if the corresponding simulation ran identically, this extra time step would prevent the output data from matching. The Swarm scheduler begins at time step 0, whereas the Repast scheduler begins with time step 1. Thus, when snapshots were logged at time step 30, Swarm had actually performed one extra time step.

With these two problems corrected, the corresponding logs of the developers' actions matched. Using the same sequence of random numbers, the Swarm and Repast simulations produced identical output.

4.2 Comparisons of OSS Parameters

Degree distribution, diameter, and clustering coefficient are frequent attributes used to describe a network (Newman, 2001a,b) and have been used since the foundation of random network theory. To study the validity of our simulation, we compared these attributes in Swarm and Repast simulations. We observed matches of these attributes between corresponding Swarm and Repast models, which indicate a clean docking.

Degree distribution *p(k)* is the distribution of the degree k throughout the network. The degree k of a node equals the total number of nodes to which it is connected. Degree distribution was believed to be a normal distribution, but Albert, et al. (1999) recently found it fit a power law distribution in many real networks. Figure 2 illustrates developer distributions in four models implemented by Swarm and Repast. The *X* coordinate is the number of projects in which each developer participated, and the *Y* coordinate is the number of developers in the related categories. The figure shows that the ER method does not have a power law distribution. Rather, the distribution looks more like the mathematically proven normal distribution. Developer distributions in the other three models match the power law distribution. Slight differences occur between the Swarm results and the Repast results; however, we believe these differences are caused by various random generators associated with Swarm and Repast.

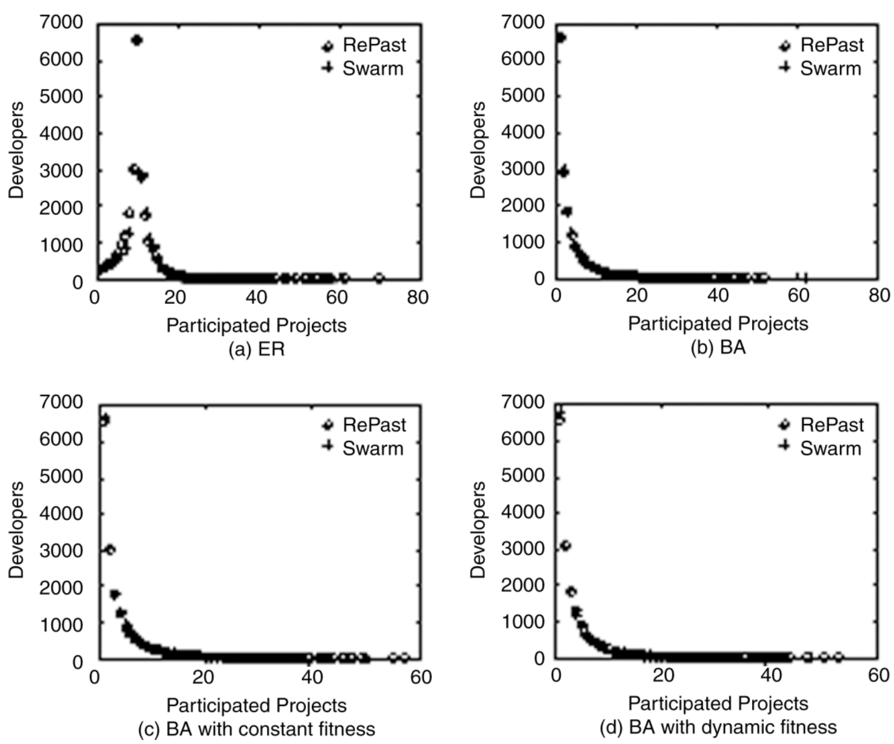

FIGURE 2 Developer distribution

Degree distribution was the diameter of a network or the maximum distance between any pair of connected nodes. The diameter can also be defined as the average length of the shortest paths between any pair of nodes in the graph. We use the latter definition because the average value is more suitable for studying the topology of the OSS network. Figure 3 shows the evolution of the diameter of the network. We can see that Swarm and Repast simulations are docked. In the real SourceForge developer collaboration network, the diameter of the network decreases as the network grows. In our models, we observe that the ER model does not fit the SourceForge network, whereas the other three models match the real network.

The neighborhood of a node consists of the set of nodes to which it is connected. The clustering coefficient of a node is a fraction that represents the number of links actually present relative to the total possible number of links among the nodes in its neighborhood. The clustering coefficient of a graph is the average of all the clustering coefficients of the nodes represented. Because clustering is an important characteristic of the topology of real networks, we also investigated the clustering coefficient, which is a quantitative measure of clustering. Figure 4 shows the clustering coefficients for the developer network as a function of time. All models are docked very well. We observe the decaying trend of the clustering coefficient in all four models. The reason is that, with the evolution of the developer network, two co-developers are less likely to form a new project because their ongoing projects are approaching their limits.

Figure 5 shows the total number of developers and projects relative to the time period in four models, which describe the developing trends of size of developers and projects in the network. The size of developers and projects is very similar for Swarm and Repast simulations.

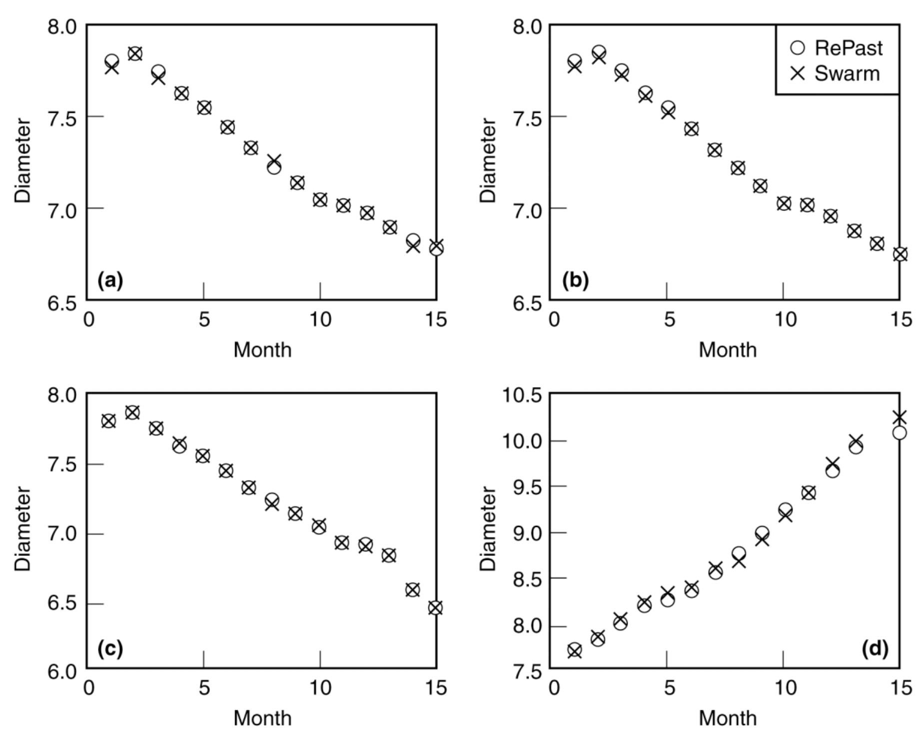

FIGURE 3 Diameter of the network

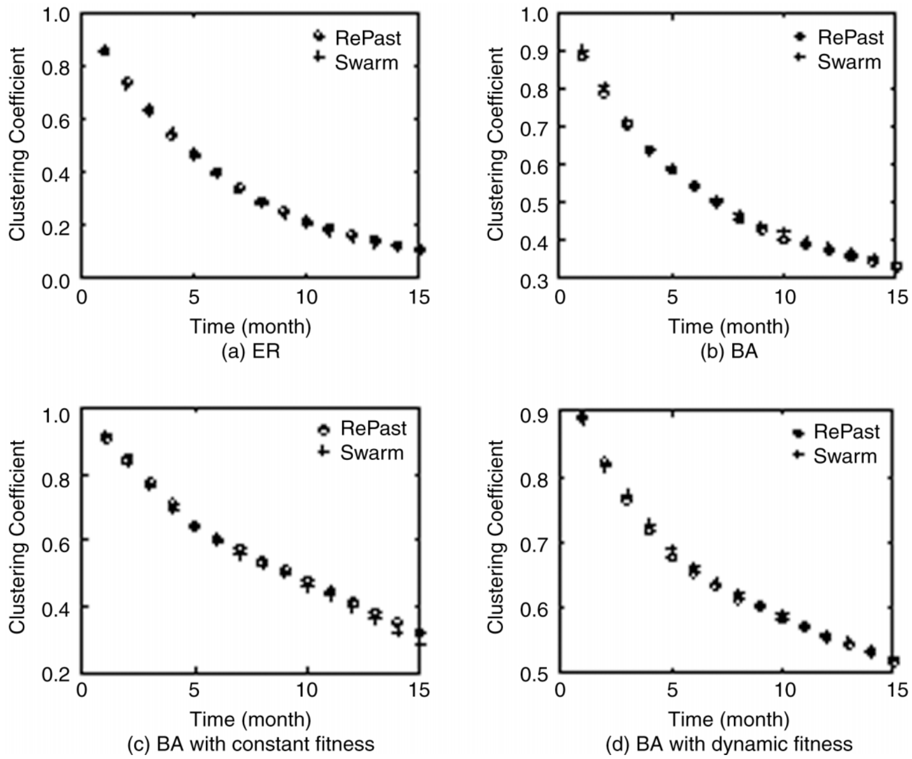

FIGURE 4 Clustering coefficient

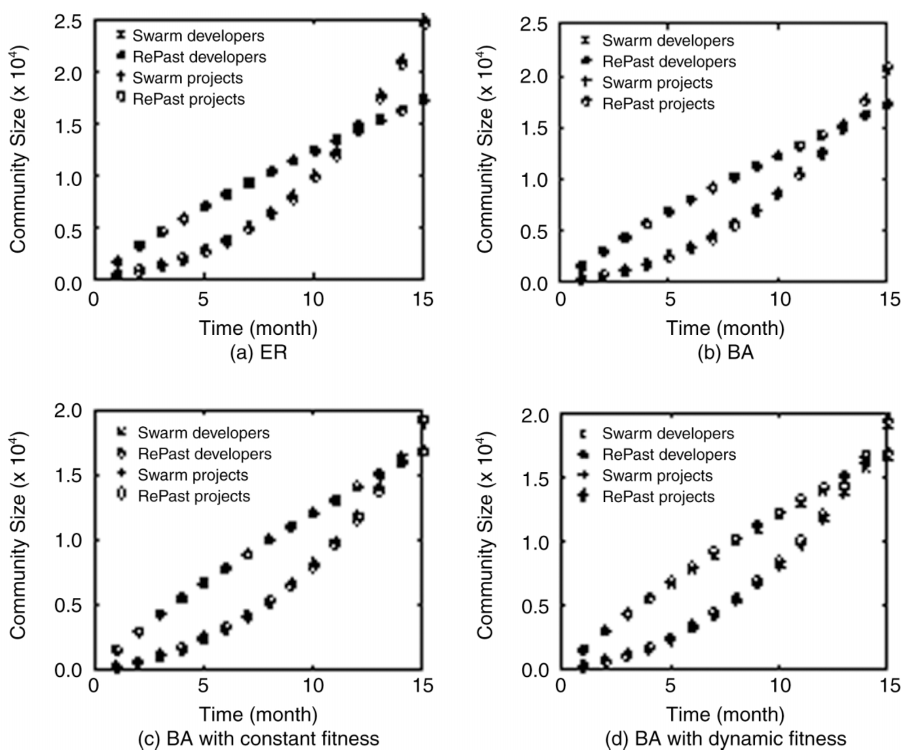

FIGURE 5 Community size development

5 CONCLUSION

This paper discusses the validation of agent-based simulation by using the docking process. It describes four simulation models of an OSS developer network using Swarm and Repast. Properties of social networks, such as degree distribution, diameter, and clustering coefficient, are used to dock Swarm and Repast simulations of four social networks. Results show that docking two agent-based simulations helps to validate a simulation. A docking process can also be used to validate a migration of a simulation from one software package to another. In our case, the docking process helped with the transfer to Repast to take advantages of its features. The Repast simulation runs faster than the Swarm simulation because Repast is written in pure Java, whereas Swarm is originally written in Object C, which causes some overhead for Java Swarm. Furthermore, Repast provides more display library packages, such as a network package, which help users perform analyses.

6 ACKNOWLEDGMENT

This research was funded in part by the National Science Foundation Award-0222829, from the Digital Society and Technologies Program, Directorate for Computer and Information Science CISE/IIS.

7 REFERENCES

Adamic, L., and B.A. Huberman, 1999, "Internet: Growth Dynamics of the World Wide Web," *Nature* **401**:131, Sept. 9.

Albert, R., H. Jeong, and A.L. Barabasi, 1999, "Internet: Diameter of the World Wide Web," *Nature* **401**:130–131, Sept. 9.

Ashworth, M.J., and M.A. Louie, 2002, *Alignment of the Garbage Can and NK Fitness Models: A Virtual Experiment in the Simulation of Organizations*. Available at http://www.casos.ece.cmu.edu/casos working_paper/GCandNK_Alignment_abstract.html.

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

Axtell R., R. Axelrod, J. Epstein, and M. Cohen, 1996, "Aligning Simulation Models: A Case Study and Results," *Computational and Mathematical Organization Theory* **1**(2):123–141.

Barabasi, A.L., and R. Albert, 1999, "Emergence of Scaling in Random Networks," *Science* **286**:509–512.

Barabasi, A.L., R. Albert, and H. Jeong, 2000, "Scale-free Characteristics of Random Networks: The Topology of the World Wide Web," *Physica A* 69–77.

Barabasi, A.L., H. Jeong, Z. Neda, E. Ravasz, A. Schubert, and T. Viscek, 2001, *Evolution of the Social Network of Scientific Collaborations*. Available at www.lanl.gov/arXiv:cond-mat/0104162v1 (accessed April 10).

Barabasi, A.L., 2002, *Linked: The New Science of Networks*, Boston, MA: Perseus.

Burton, R., 1998, "Aligning Simulation Models: A Case Study and Results," in *Simulating Organizations: Computational Models of Institutions and Groups,* Cambridge, MA: AAAI/MIT Press.

Collier, N., 2003, "Repast: An Extensible Framework for Agent Simulation." Available at http://Repast.sourceforge.net/projects.html (accessed in April 2003).

Collier, N., and T.R. Howe, 2003, "Repast 2.0: Major Changes and New Features," in *Seventh Annual Swarm Researchers Conference (Swarm2003),* Notre Dame, IN.

Faloutsos, P., C. Faloutsos, and M. Faloutsos, 1999, "On Power-law Relationships of the Internet Topology," in *SIGCOMM'99*, Cambridge, MA: Association for Computing Machinery, pp. 251–262.

Freeh, V., G. Madey, and R. Tynan, 2003, "Research Project on the Free/Open Source Software Development Phenomenon." Available at http://www.nd.edu/~oss (accessed in January 2003).

Gao, Y., 2003, "Topology and Evolution of the Open Source Software Community," in *Seventh Annual Swarm Researchers Conference (Swarm2003),* Notre Dame, IN.

Gao, Y., V. Freeh, and G. Madey, 2003a, "Conceptual Framework for Agent-based Modeling," in *North American Association for Computational Social and Organizational Science (NAACSOS 2003),* Pittsburgh, PA.

Gao, Y., V. Freeh, and G. Madey, 2003b, "Analysis and Modeling of the Open Source Software Community," in *North American Association for Computational Social and Organizational Science (NAACSOS 2003),* Pittsburgh, PA.

Madey, G., V. Freeh, R. Tynan, and C. Hoffman, 2003a, "An Analysis of Open Source Software Development Using Social Network Theory and Agent-based Modeling," in *The 2nd Lake Arrowhead Conference on Human Complex Systems,* Lake Arrowhead, CA.

Madey G., V. Freeh, R. Tynan, Y. Gao, and C. Hoffman, 2003b, "Agent-based Modeling and Simulation of Collaborative Social Networks," in *Americas Conference on Information Systems (AMCIS2003)*, Tampa, FL.

Madey, G., V. Freeh, and R. Tynan, 2002a, "The Open Source Software Development Phenomenon: An Analysis Based on Social Network Theory," in *Americas Conference on Information Systems (AMCIS2002),* Dallas, TX, pp. 1806–1813.

Madey, G., V. Freeh, and R. Tynan, 2002b, “Understanding OSS as a Self-organizing Process,” in *The 2nd Workshop on Open Source Software Engineering at the 24th International Conference on Software Engineering (ICSE2002),* Orlando, FL.

Minar, N., R. Burkhart, C. Langton, and M. Askenazi, 2002, “The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations,” Santa Fe Institute, NM. Available at http://www.santafe.edu/projects/ swarm/swarmdoc/swarmdoc.html.

Newman, M.E.J., 2001a, “Clustering and Preferential Attachment in Growing Networks,” *Physics Reviews* **64**(025102).

Newman, M.E.J., 2001b, “Scientific Collaboration Networks: I. Network Construction and Fundamental Results,” *Physics Reviews* **64**(016131).

North, M., and C. Macal, 2003, “The Beer Dock: Three and a Half Implementations of the Beer Distribution Game.” Available at http://www.dis.anl.gov/msv/cas/Pubs/BeerGame.PDF.

Repast Home Page, 2003. Available at http://Repast.sourceforge.net/.

Strogatz, S.H., and D. Watts, 1998, “Collective Dynamics of Small-world Networks, *Nature* **393**:440–442.

Wasserman, S., and K. Faust, 1994, “Social Network Analysis: Methods and Applications,” in *Structural Analysis in the Social Sciences*, Cambridge, U.K.: Cambridge University Press.

Xu, J., and Y. Gao, 2003, “A Docking Experiment: Swarm and Repast for Social Network Modeling,” in *Seventh Annual Swarm Researchers Conference (Swarm2003)*, Notre Dame, IN.

DISCUSSION:

APPROACHES TO VALIDATION

(Thursday, October 2, 2003, 5:45 to 7:15 p.m.)

Chair and Discussant*: Steve Bankes, RAND Graduate School*

Modeling Playgroups in Children: Determining Validity and Veridicality

Bill Griffin: Let me tell you a little bit about what we're going to talk about for the next 25 to 30 minutes. I'm going to break this into two parts. The first part is just a quick review of the model that was used. That is not the primary focus of the talk, but now that the model's a little older, a little more mature, we're looking at what's wrong with it, where does it work, where does it not work, and how did we come to decide that it doesn't work?

[Presentation]

Steve Bankes: We have about one minute. Let's take one question.

Joanna Bryson: I'm very interested in your methodology question, but unfortunately, since this is the only question [permitted], I actually had a question about your model. You're talking about these strong things, but you also said you were missing some of the gaps. So you're missing some of the aversions as well as the fondnesses. You said that one of the great advantages of agent-based modeling, of course, is memory. If you cranked up the memory, you can't get replication of that?

Griffin: Well, yes. The length of the memory is easy to modify, obviously. But we're trying to decide what is the age-appropriate memory link, because we're trying to think of these kids' cognitive abilities.

Bryson: Well, it's not only about length, it's also about importance.

Griffin: Well, I can talk to you about that later. Actually, once there's a parameter sweep attached to that, which I ran on both the effects of gender and memory, how much value is associated with that on the link of the memory? Right now, we're letting them hold for five memory places; it basically follows an exponential curve.

Is that what you're asking?

Bryson: I meant the prioritization. So if every day you're going in and playing with the person you saw the previous day, because you had a good time and you saw them yesterday or whatever, or every day you're avoiding the other person. That may be more important than their gender to you, the fact that you played with them last, as long as you can remember you've been playing with them, or avoiding them.

Griffin: Well, yes. In fact, I talked to the individual on our team who's the gender expert about that exact same question. At what point does something like memory override, or is there some sort of ratio, is there some sort of multiplicative thing going on where you could do something exactly like that? And she really holds onto gender, and so she wants me to work memory through gender.

Bryson: Yes, because it sounds like internal state. So you only observed the children in the playground?

Griffin: Yes.

Unidentified Speaker: Okay, so there could be things going on in the classroom, they may be neighbors, they may have siblings that are friends, in Sunday school, they may have assigned seating in the classroom that could influence future actions.

Griffin: One of the things that resulted from the initial model two years ago, then this one last year is that we changed, or *they* changed their data collection procedure. We gathered much of the data in the single data point here in mid-December. Now we're taking monthly updates on the teacher attributes, we're getting the coders to get an attribute rating, plus the observational data, and we're looking for any outside connections like you're referring to, because they're the same sort of questions.

I just wanted to get that out, that you're exactly right. But, see, it's those kinds of things that are sort of generating these results. We went into meetings, and I said, "We need more data than these 2,000 or 3,000 data points per month. We need a different kind of data."

Alignment and Validation in an Agent-based Model of On-line B2C Auctions

Bankes: The next talk is "The Alignment and Validation in an Agent-based Model of Online B2C Auctions," and the speaker is Kumar Mehta.

Kumar Mehta: I've been attending talks since the morning, and I didn't know this world of agent-based modeling existed. It's very, very radically different from anything I've seen. But what I've seen has been only for three to four years. So it's been kind of refreshing. I come from a very quantitative technical background, so this is an incredibly refreshing point of view. What I'm going to present is a small part of a rather large stream of work. This is part of my dissertation.

[Presentation]

Bankes: We'll take a quick question.

Unidentified Speaker: One thing I noticed in the final output results: You showed the bounds or range of the simulation output and also the average. What about the mean? I mean, was it a typical sort of straight curve through there, or did it vary a lot?

Mehta: No, there was a solid line in between, which is the mean of the three.

Unidentified Speaker: Oh, no, I saw the mean, but I mean like a typical trace. What would a trace look like, one of the runs?

Mehta: It would look like in the steps, exact step manner, yes.

A Multi-model Docking Experiment of Dynamic Social Network Simulations

Bankes: So the next talk is going to be given by Jen Xu.

Jen Xu: The topic I will be talking about today is “A Multi-model Docking Experiment of Dynamic Social Network Simulations.” This work was done with Yung-Chi Gao, Jeff Goett and Gregory Madey. This research was partially supported by the National Science Foundation. I will first give a brief introduction of our docking experiment.

[Presentation]

Bankes: Questions for the speaker?

Unidentified Speaker: I was curious. I have a vague recollection at Lake Arrowhead that somebody had an agent-based model of open-source software development. I forgot who it was, but maybe instead of using the Erdos algorithm, or the Barabasi stuff, if you looked at that and compared that to the real data from Source … It seems like it would be a neat thing to do. Does anyone remember?

Unidentified Speaker: That probably was me. I’m a co-author.

Unidentified Speaker: I’d like to hear more about differences between Swarm and Repast in terms of which was a bigger pain to program in and other differences.

Xu: Actually, the conception is a little bit similar, but Repast has some extra features.

Unidentified Speaker: So Repast is better.

Unidentified Speaker: Yeah.

Unidentified Speaker: I guess what I’m hearing is that a 10% difference in performance isn’t very convincing. So I’m looking for even more reasons to go with Repast. And as you’re saying, it’s the features.

Xu: I heard somebody say that in some applications Repast outperformed Swarm, but in some applications, Swarm is better. But in our simulation, we found that Repast is better. And we also want to transfer to Repast because Repast implemented a distributed feature that may improve our performance in the future. We want to increase the speed, the running speed of our simulations.

Bryson: Just to follow up from what he just said . . . . So there was no difference in development? It’s sort of an unfair question, because you’d already prototyped and had all the hard conceptual work while you were doing Swarm. But do you think it would have been as fast to rewrite it in Swarm as it was to write it in Repast?

Xu: Maybe Yung-Chi can say how long it took him to write in Swarm, but for me, I used one week to study Swarm simulation and transfer it into Repast, so I think it's pretty quick.

Unidentified Speaker: I'm not sure what you asked about comparing Swarm with Repast based on the programming cost or the running performance.

Xu: Running cost, I think it should be similar, because actually the ideas behind Swarm and Repast are kind of similar for me. So if you write a program in Swarm already, it will just take a little jump to migrate to Repast.

Jesse Voss: What metric specifically are you using to make the choice between one or the other? Is it just run-time speed, or is there something else that you're using that's better? Because I've heard it said that some particularly complex kinds of relationships that you're trying to model can't really be modeled in either Swarm, Repast or Ascape. So I just wanted to know if you looked at that at all?

Xu: Actually, I like the network display. I think Repast provides a better network display than Swarm. And also it has some distributed architecture that we will use in the future.

Panel Discussion

Bankes: Okay, if we could have all the speakers move forward. As discussant, I'll go ahead and ask the first question, or make my comments in the first volley and then open it up. This is the last session, so as long as you guys have energy and interest, this can go on all night.

You know, our conversations are always shaped by the terminology we've adopted, and I observe that the word "validation" has caused a lot of mischief to a lot of simulation communities going back many decades. I once did an exercise of going off and trying to find a definition of "validation" and found four or five in various documents that had an Aristotelian turn, where they really felt the need to come up with a formal definition. And one thing that's remarkably true, anytime anybody or any committee's tried hard to carefully define their terms, they end up in the middle of this long legalistic bunch of stuff, saying, "Valid for a particular purpose." And one of the ways that validation and the implication of the phrase that "a model is either valid or it's not, and if it's not valid, what good is it?" is this tendency to drop out for what purpose?

And so what I intend to do is give a challenge and say, "You guys, I did a validation. For what purpose is your work valid?" And "docking" is a much more modest phrase and not near as pernicious, but nonetheless, I invite the third speaker as well to talk about the limits of the exercise and the extent to which it looks like we've got two models that are really almost, you know, a re-implementation more than a document in the hard sense. But is there an edge past which it wouldn't work? And to avoid this being a really hard snap quiz, I observe by reading the papers and listening to the talks that there's a variety of things that we accomplish by comparing models to models or models to data, ramping from a kind of verification, where when you see a difference you bore into your model, you discover places where you goofed up the implementation, or some choice you made produced an artifact that is unwelcome. So you're able to get rid of it to the next kind of phrase where you see differences, and it caused you to think about phenomenology more. So you can climb a hill in model space, you get a better model.

Then there's this next tier, and I claim this is a validated model in the sense it actually replicates the real world, which is an aggressive claim epistemologically for almost any model, I think. But to the extent one wants to make it, it invites, then, the question, "For what range of phenomena in the real world have you established validity and where's the edge past which you have to say, 'It's not validated for data classes or cases that don't have this characteristic, and so forth?'" And so, not to make it real hard, but just a brief statement from all the speakers about what is the edge of your work? What delimits what you've accomplished?

Unidentified Speaker: That's a good question. With ours, I'm much more conservative about how I would define validity. One definition I have looks just like the previous definition on verification, all the way from does the code do what it's supposed to do internally? And the validity, does the model in some way grossly represent the physical data that you actually possess? Does a validation, using the very strict sense of the word, capture the data, including the process, not just the outcome?

At least in our work, I'm much more conservative in thinking we've got one run of data under this particular model that we've got running. I'd want several years of data, so that when I did drill down I was able to capture most of the phenomena most of the time. Now, what is most of the data most of the time? I'm not sure yet. But where I can say that, in general, as I drill down, do the data still map onto the simulation? Any variation past where these spontaneous eruptions, like that one slide I put up? At some point, I'll have to say that, "How well," and this is the phenomenological idea of how well can we capture a dynamical process with a single model that replicates itself over and over again, but never in the same form.

I don't know if that answers your question.

Unidentified Speaker: There are quite a few limitations to what we've done. What we are following is more of a spiral development methodology where we are starting with a core in a sort of a constrained set, as we build more and more behaviors into [the model]. The one limitation we have right now is we are looking at products which can be collapsed into a series of integer values. So essentially on one dimension you can map the similarities, and that is the choice of hard disks comes into the play. It also currently does not take into account any reputation effects of retailers, which is why we picked hard disks again as a category, because no matter what I buy from X or Y, I'm going to get the same exact item. It's the branding which is the issue.

Memory is the third one, and this is the first thing we have just built in. So essentially those were the ranges within which [we worked]. And we are looking at only one unit option to prevent any resellers currently. So as we sort of step one further, we are adding memory to it right now, then we'll add the reseller, the guy who's spotting to buy bargain items and sell it to you again. Those behaviors are all we are working.

Xu: The limitations in our work is just to compare one rung of the result. But we can still see, there's still some difference between those simulations, so we want to do more statistical analysis to check if they are really matched, because in theory they should match exactly.

And the second limitation I think is that in both our simulations is the speed of performance is not very much faster than when we simulate a large number of developers. So … this will make our documentation very difficult, because we need to wait a long time to finish the simulation and to compare results. So we want to improve our simulations maybe by some remote procedure call for migration some part of simulations distributed them on several machines. So that's our next plan.

Brian Pijanowski: Actually, I just want to follow up with your question. It seems to me that whenever you talk about model validation, you also have to consider what the assumptions are of the model, because you have to make them. I mean, that's what a model is. And so you have to have a correspondence between the assumption, the nature of the assumption, and the validity of the model. So unfortunately, as a modeler, we oftentimes have to communicate our results to people that think they're kind of suspicious, they don't quite understand them, and it seems to me that oftentimes we kind of stumble in our communication because we don't state what our assumptions are.

And so they start pointing to the model and they say, "It's invalid because of" whatever. And you think, "Well, I made this assumption over here because I didn't want to consider it in my model. I'm trying to simplify a complex system." So I think that whenever you talk about validating model, you also have to consider what the assumptions are.

Unidentified Speaker: In fact, I had another whole slide of assumptions, a group level, but at the individual level I had to make assumptions, for example, that the attributes didn't change within a certain window of time, and that the modification of an attribute is the same across all attributes. And we know that's probably not true, but we have to get more data to find out, more multiple data points to find out if they measured at say a 1.5 on an attribute at the beginning in September, and there at 2.7 in May, is that progression standard across all kids or is there some unique combination when you see that progression? So, yes, there are assumptions all over this thing.

Unidentified Speaker: I think there's expert judgment.

Unidentified Speaker: With all the constraints in your models, you have to validate them one at a time. But there's a pitfall to that, in the interaction effects, when two models are switched on at the same time, you have no way of knowing whether it is still valid or not. But there is too much synergistic action.

We've been running some recent experiments [in which we] keep increasing the number of units sold. What we found was slightly counterintuitive, which is true in general, but there are small cases where actually our revenues jump up. We had seen this in the market, and then it popped up on — I'm going to toss in the name, just for the heck of it — Swarm. It popped up in Swarm. And we started digging this. We wanted to trace back what was going on. Because there are more number of seats that people could have taken up before. Instead of one seat, now you have five, and they would take them in different order. And so there are counterintuitive things that pop up which will not come up on your analytical model, because it has abstracted those concepts away.

So there is a pitfall to talking just to the analytical model. But in absence of real data, that's possibly the only option one has.

Unidentified Speaker: Do you think that 23 is enough for your sample size?

Unidentified Speaker: The absolute sample size is 23, but when you get 2,000 to 3,000 observations per month, over a given year, now, of course, as I mentioned a while ago, then you take the subsequent year's students and then you take the subsequent year's students and you do this over and over again.

Unidentified Speaker: And it seems like you simulated the data and at the same time, you used real data to compare the visual simulated.

Unidentified Speaker: No. I wrote the model independent of having the data. I just wrote that by myself in a little booth. And when I just went in and said, "What variables do you have, what variables does the literature suggest? And what is the assumed relationship?" And I wrote the code never having seen the data.

Unidentified Speaker: I'd just like to follow up. It seems to me that the thing that is from a social perspective, the most problematic of the way that you've collected the data is the teacher attributions. And I think it might be interesting, if you do it too much, it could be intrusive in its own right, but if every day, like the first thing in the morning every day, or maybe even just once a week, if you kind of reframed those questions so that the children answered those questions, and you said, "Who is most helpful?" you know, "Who plays rough?" and a couple things like that and get *their* perspective, since the teacher's perspective is quite a bit different.

Unidentified Speaker: That's why this year — that's actually one of the things that they do with older children. They do what's called sociometric ratings where you say, "Who's my best friend? Who do I like to spend time with?" etc.

And with regard to the teacher data, this year they're implementing it over each time period. We're running into pragmatic problems, because we're having to pay teachers extra to fill out all the forms. We're now using the coders who before were just doing the palm pilots and doing the behavioral counts. We are now asking them to do similar ratings as the teachers do so we get two, multi-raters' data. So we're trying to expand the width of the data band coming in to see how it compares. But then, of course, you think, do we have a composite score? Once you do that you have all these other problems.

Bryson: You'd expect also to get a huge impact — well, maybe not huge with kids that age, but whenever you make people bring things into declarative, then that totally affects their behavior.

Unidentified Speaker: That actually came up in the meetings, yes. And these are the same people who do the data counts. Yes. And one of the things we're actually asking for in a present submission to NSF is enough money to have those as independent parties. You know, money takes care of a lot of issues because we could pay our teachers more, they're more willing to be helpful, and we have a separate group of coders, one for the attributions summaries and the other one for the data counts.

Konstantinos Alexandridis: It's important to make a distinction which, I don't know about the other researchers that developing and implementing agent-based modeling, but to make the distinction that it's different to ask, for example, if that specific agent at a specific time makes accurate decisions and another thing to say what are the persistent properties that we observe? And that's where people that they're not familiar with agent-based modeling don't easily understand.

I've been running a lot of situations where people are asking, for example, "Well, is that specific farmer there in that parcel?" that simulation implies that he will make that decision at that specific time. And that kind of accuracy is not a part of agent-based modeling approach. And that kind of validation is not applicable, I think, and that has to be clear.

Unidentified Speaker: You know, that's how I started out talking. Ed and I talked about this at lunch. When I give this talk, people say, "Does that mean Johnny will play with George?" I go, "No. It doesn't mean that. It means they have to share the same characteristics of this cluster that tended to play together, but that doesn't mean a specific child." The prediction to a person or an agent, I just don't know.

Alexandridis: And that also mean that we have to acknowledge, in terms of validation, that this kind of validation is not complete validation.

Unidentified Speaker: Well, it's a validation to the process and not to the person, or to the agent. It's a validation — I'm shooting from the hip here. The idea is that we're validating, or possibly validating, how an end result came vis-a-vis this process that we've coded the rules for. That's not all we can say, and it maps on fairly well. That's not all we can say.

Bryson: Wouldn't you expect to get — so you wouldn't be able to say for sure two kids are going to play together, or that two programs are going to program together, right? But wouldn't you expect to get a probabilistic result?

Unidentified Speaker: Yes.

Bryson: So then you could say, "I predict that these two people are likely to, so 60% of them actually will.

Unidentified Speaker: No, I think you're right, but the flip side of that is if you're wrong, some individuals assume that the model is not valid.

Bryson: Hypothesis sets in, right, so you have a 95% chance of being wrong or something. Yes.

Unidentified Speaker: Yes, you know. And I was thinking about that. I mean, we want to put a confidence band around some of the outcomes, yes.

Friday, October 3, 2003

Welcome:

Stephen Gabel

WELCOME

STEPHEN GABEL, Associate Provost, The University of Chicago

On behalf of The University of Chicago, I want to welcome you to the *Agent 2003 Conference on the Challenges of Social Simulation*. Since my academic training is in literature and is entirely non-technical, and since I teach subjects like Homer and Aristotle, my presence here may require a bit of explanation. For the past year or so, in the Provost's Office at The University of Chicago, I have been working with old colleagues in the university and new colleagues at Argonne National Laboratory — Tom Wolsko, Chick Macal, Mike North, and others — to help build new collaborations and foster exchanges between the social scientists on campus and the scientists at Argonne who are active in computational social science. In the process, I have had to try to understand what a complex adaptive system is and what in the world folks mean by agent-based simulations. I have to admit that I am still trying.

Yesterday, I made what I think is a small step in understanding agent-based simulations. I realized that I had read an account of the special value of simulation as a mode of discovery in the work of an author familiar to all, that is, the Greek philosopher, Aristotle, who lived more than 2,300 years ago.

Aristotle devoted a treatise to simulations: what they are, how they differ from other products of the mind, and what the standards are for evaluating them. The treatise is the *Poetics*, Aristotle's analysis of how the human propensity to imitate what we observe can eventuate in complex symbolic simulations. The simulations Aristotle had in mind were ancient Greek dramas.

Plato, Aristotle's teacher, was a philosophical idealist, and tended to see simulations as merely imperfect images of reality and of no intrinsic interest. This attitude is one that I would guess some of you have encountered in one guise or another.

But Aristotle understood that simulations — properly performed — offer a unique way of gaining knowledge about the world. Or, as Aristotle put it: "The poet's function is to describe, not the thing that has happened [that is, empirical or historical data], but a kind of thing that might happen, i.e., to describe what is possible as being probable or necessary (1451a36–40)."[1] That is why, he goes on to argue, a simulation such as a drama "is something more philosophical and of graver import than history, since its statements are of the nature of universals…. By a universal statement I mean one as to what such and such a kind of man [or agent!] will probably or necessarily say or do (1451b6–10)."

I could continue with this exercise of discussing Aristotle's reflections on drama (which are very much in a scientific spirit) and argue further that his reflections reveal that he understood drama as essentially a simulation, and that he believed simulations can yield a kind of knowledge that is available to us in no other way. But if you did not already accept the proposition, you probably would not be here.

---

[1] Aristotle, *Poetics*, trans. Ingram Bywater, in *Introduction to Aristotle*, R. McKeon, ed., Chicago: University of Chicago Press, 1973. Citations are to standard line numbers which are the same in all editions.

My real reason for drawing on Aristotle this morning is twofold. First, it seems that it is a good thing for all of us to remain aware of the intellectual genealogies of our disciplines. Today's science and scholarship — even cutting-edge science — are like a branch of a tree that is very old, with deep roots. And some of the problems we try to understand today are problems humans have been thinking about for a long time. Second, it seems quite likely that the tools discussed will have a great deal of resonance and utility for scholars in fields other than those represented here today. You should get ready to encounter other visitors like me, aliens from the library who are intrigued by the work you are doing. Please be patient with us. I wish you all a stimulating and productive day.

Invited Speaker:

Steven Bankes

IMPROVING THE UTILITY AND THE RIGOR OF AGENT-BASED MODELING THROUGH ENSEMBLES OF MODELS

STEVEN BANKES,* Evolving Logic, Los Angeles, CA

ABSTRACT

Agent-based modeling (ABM) has demonstrated great promise, but it also faces significant challenges. Central among the latter are the need for greater levels of rigor and of demonstrating important applications. This paper argues that both these challenges can be met, at least in part, by adopting techniques of reasoning over ensembles of alternative versions of models.

Keywords: Ensembles, rigor, robust inference, agent-based modeling

INTRODUCTION: THE PROMISE AND CHALLENGE OF AGENT BASED MODELING

Agent-based modeling (ABM), and computational science based on simulation more generally, has demonstrated great promise, but it also faces significant challenges. ABM provides new representational options to allow inference from theory and data that did not fit into previous formalisms. It can thus provide important theoretical findings that would not previously have been possible to achieve. It can augment the literary methods of much of social science with a more formal framework and simultaneously augment descriptive models with related dynamic ones.

But, if agent based modeling is to make a significant contribution to science, much greater rigor in its use will be required. A large fraction of ABM research to date has been exploratory and suggestive, featuring hypothesis generation with little hypothesis resolution. Definitive studies that have been validated against data are rare. The need for greater rigor has been expressed by some leaders in the field as a need for more "prediction." While predictive accuracy is a powerful attribute to establish, if it can be achieved, the emphasis on prediction is somewhat misleading, as I will argue below.

Related to the need for rigor is a shortfall in developing important applications of this tool. In the 1950s, the newly minted tools of operations research that had proven their value in military settings were deployed to industry, with substantial documented benefits in cost reduction and improved profits. To my knowledge, no similar examples of direct financial benefits of ABM have yet been documented. Similarly, there is not yet any example of a major public policy problem that has been met through ABM studies. It can be argued that these two problems, rigor and applications, are very closely related.

* Author's address: Steven Bankes, Evolving Logic, 3542 Greenfield Ave., Los Angeles CA 90034; e-mail: bankes@evolvinglogic.com

WHAT IS RIGOR?

Standards of rigor vary in their expression across various fields of science depending on their traditions, and upon both the nature of the problems being addressed and tools that are used to address these problems. In a new and hybrid field such as computational science, discussions about rigor sometimes become confused as a consequence. So, it is useful to return briefly to first principles in analyzing what the actual challenges of rigor are for ABM research.

It is sensible to speak of rigor in the use of models that make no use of data. In particular, it is important that the claims made by computational research are supported by the modeling experiments that were conducted. This requirement of internal consistency is similar to standards of proof in mathematical reasoning, though different in that deductive inference does not have a central role.

However, most appeals for greater rigor in ABM research are fundamentally appeals for more studies that compare models to measurements. These appeals often take the form of insisting that models must be more "predictive." Unfortunately, the use of the word "prediction" as a synonym for rigor introduces yet more confusion, as this word again has multiple definitions in different fields arising from different applications to different types of problems. There are at least three different definitions of prediction:

1. Correct forecasts of future events

2. Correct model-based inference of new knowledge from available knowledge and data

3. Sufficient similarity of model outputs to data not used in its construction (cross validation)

All three forms of prediction are good properties to achieve, but they are not at all the same thing. In particular, definition #3 is neither necessary nor sufficient for definition #1. And none of these definitions are necessary for a model to provide utility in solving problems. In order to sort this all out, we must return to first principles in thinking about how models relate to data.

The formalism of statistical inference provides us the machinery for thinking about this. The approach founded by R.A. Fisher is based on distinguishing model specification from model estimation. Model specification is a step that happens outside the frame of statistical inference. In model specification, the researcher asserts (assumes) that the available data were generated by one of a parameterized family of models plus a source of noise. Once this is done, model estimation is the mathematical problem of computing the parameter vector (picking a single model from the family) that maximizes the likelihood that the resulting model generated the data. This is the so-called maximum likelihood estimator (MLE). Thus, in a simple form of statistical inference, where the model is a linear equation relating several predictors, model specification is the selection of the predictors to include in the equation, and estimation involved solving for parameter values that minimized the residual squared error that results from comparing the "predictions" of this equation to the actual data.

Within this framework, formal machinery for uncertainty analysis can be erected. For simple models, it is possible to calculate the probability distribution of estimated parameters given an assumed noise process. A variety of goodness-of-fit measures (i.e., an F test) can be been devised to assess how well the model explains the pattern seen in the data. In the culture of statistical modeling, a good score on an appropriate measure of fit suffices to demonstrate that the model is "good." (Sometimes this will be called "prediction" under definition #2.) Typically, this outcome is used to validate the model specification, though there are other ways to screen for misspecification, such as correlated residual noise.

These fundamental details are widely known, but it is useful to emphasize their logical basis. R.A. Fisher's formal structure does not depend on asserting that the specified model family contains the correct model. And if it did, this does not mean that the process of estimation would accurately identify it. Rather, model specification can be understood as an analytic device that reveals patterns in data. And model estimation, given that the specification is correct, is a process of minimizing the expected difference between the estimated and true models, given the limits imposed on our reasoning by the presence of noise. Thus, while often interpreted idealistically, the framework of statistical inference can be understood as highly pragmatic. That is, this approach serves the question "How can we best solve specific problems given available data?" where "best" is defined within the pragmatic constraints of limited information about a noisy universe. This pragmatic stance can provide important benefits when we turn to thinking about comparing agent-based models to data.

The framework of statistical inference was developed in a period of computational poverty, where the computation involved in a single model estimation using linear models could be significantly expensive. With increasing computational resources, there has been growing interest in doing lots of estimation experiments, automating specification search, as well as using more complex non-linear models. This trend is fundamentally virtuous, as it brings the previously ad-hoc process of exploring across model specification into an analytic framework. It also has entailed various problems, as the assumptions behind model estimation can be easily violated with naïve specification search (Miller, 1990).

Initially, attempts at specification search received the pejorative label of "data mining" for the bulk of the statistical community, and any procedure that tried out lots of model variants was viewed as highly suspect. Simply searching through many alternative specifications and keeping the one with the best goodness of fit is a practice that can lead to very bad results. If it is done without penalizing complex models, the result can readily be a procedure guaranteed to select a highly complex model that over-fits the data. This will usually result in a highly biased model. Even where model complexity is properly penalized, a specification search can still manage to "model the noise." To make matters worse, many of the elementary goodness-of-fit statistics cannot be properly used to compare models from different families.

While these problems were used to condemn automated specification search in the early days of computational statistics, the same problems can occur in connection with the ad hoc specification search that occurs when researchers revise their modeling approach iteratively by hand, seeking "good" results. Automation can make foolish mistakes more likely, but can also more readily reveal the misspecification of first guesses that had acceptable goodness-of-fit statistics. And with growing computing power, specification search by some means or another was inevitable.

During the past few decades, a broad stream of research in the statistical community has provided a variety of tools for addressing these problems (Draper, 1995; Hastie et al., 2001; Mendes and Billings, 2001). While considered advanced and less widely taught than classical methods, they are centrally important to the use of highly complex computational models such as ABM.

Most fundamentally, cross entropy or relative entropy, also known as the Kullback-Leibler (KL) metric (Hastie et al., 2001) can be used to compute an effective distance between models drawn from different families. Further, it can be proved that even in the case where there is misspecification, where the family of models does not include the true model, the member of that family that is maximally likely given the data, is also the member of that family with the minimum KL-metric to the (unknown) true model. This provides a theoretical justification for using maximal likelihood as a criterion for model estimation given that model misspecification is nearly inevitable for complex models.

The Akaike information criterion (AIC, Akaike, 1973) combines the KL metric with a penalty for the number of parameters employed, and provides a measure by which specification search can be pursued with greater care. Subsequent work has extended these initial steps, for example by combining the AIC with hypothesis testing to establish whether the difference between two models is statistically significant.

The comparison of ABM to data requires the sophistication of the portfolio of statistical tools. But agent-based models are much more complex than are the data models of essentially all statistical practice. They thus present special challenges that merit yet further consideration.

FITTING ABM TO DATA

Statistical practice first developed using linear models with a small number of predictors. As our sophistication and computational resources have grown, ever more complex and nonlinear models are being used. Currently, Bayes Nets are an example of some of the most complex models being routinely fit to data, and the most complex of them have parameter complexity equal to many simulation models. That said, any simulation that has a non-linearity at a given time step will present a highly non-linear response surface due to the iteration of that non-linearity through time. And of the simulation models, ABM is perhaps the most deeply non-linear due to the combination of rule-based descriptions of agent behavior and complex trajectory bifurcations driven by agent interaction. Our instincts regarding data analysis are informed by experience with linear, generalized linear, or linearizable models. Highly non-linear models present problems in data analysis that these instincts do not serve well.

For linear models modeling a data table with a limited number of columns, model specification is a relatively contained exercise of deciding which predictors to include in the model. When the phenomenon being modeled is indeed relatively linear, model specification is a simple determination of the most important predictors to include. Here model misspecification amounts to a modest amount of unmodeled pattern in the data. Further, model estimation is framed as an optimization problem. For linear models, the likelihood surface is smooth and unimodal.

For highly non-linear models such as ABM, the situation is very different. The universe of alternative model formulation, varying agent attributes, and behavioral rules is vast. Thus, some amount of model misspecification is highly likely, and experimentation with alternative model structures is inevitable. More crucially, the high degree of non-linearity means that the likelihood surface will be multi-modal, and it may be quite rugged. In general, there can be many models whose ability to explain the data cannot be distinguished, and these models can differ greatly in both structure and parameter values.

While the literature is not large, there are examples of studies that have estimated simulation models from data, fit these models to data, or perhaps “calibrated” their models to existing data (Chang and Delleur, 1992, for example). The techniques for doing this vary in detail, but they amount to sensitivity analysis combined with hill climbing to find a vector of model parameters that are locally maximal. The technique is useful, but as the landscape becomes increasingly rugged, the strength of claims made for the outcome of a hill-climbing exercise must be correspondingly weakened. And for many agent-based models, the likelihood landscape may be quite rugged indeed.

For those experienced with ABM, the assertion of rugged landscapes may appear quite reasonable, but for others an example may be useful. Figure 1 displays a response surface from a quite simple agent-based model. The model in this case is a reimplementation of a classic work in the combat modeling literature, the demonstration by Dewar et al. (1996) of the possibility of chaos in combat models. Here there are two combatants, Red and Blue, each with an initial number of troops that commit a fraction of their forces to a battle where losses occur according to a Lanchester formula (Engel, 1954). The commanders of the two sides reinforce their forces in this battle out of their reserves according to a rule with two parameters, one for the force level at which to reinforce, the other the size of reinforcement to send. The resulting model has several other parameters, including the initial force levels of the two sides. Figure 1 displays the ultimate winner of the war, as a function of two of the parameters that determine the behavioral rules used by the two combatants. There is a region in which the outcome is quite nonmonotonic, and indeed is nonmonotonic along nearly every parameter. In particular, leaving all else constant, but adding incrementally more Blue forces, the outcome flips back and forth many times. There are thus counterintuitive situations where giving Blue more capabilities causes Blue to do worse. This phenomenon is a product of the delicately balanced (indeed formally chaotic) dynamics emerging from the interaction of the two reinforcement rules. This particular model was crafting as a theoretical demonstrator. But consider the problem of fitting or tuning it to match data from an actual war, should we choose to do so. As Figure 1 suggests, there may be numerous different parameter combinations that could explain the observed behavior equally well. And in general, alternative explanations of data regarding an emergent phenomenon can easily interact to create complex borders such as this one.

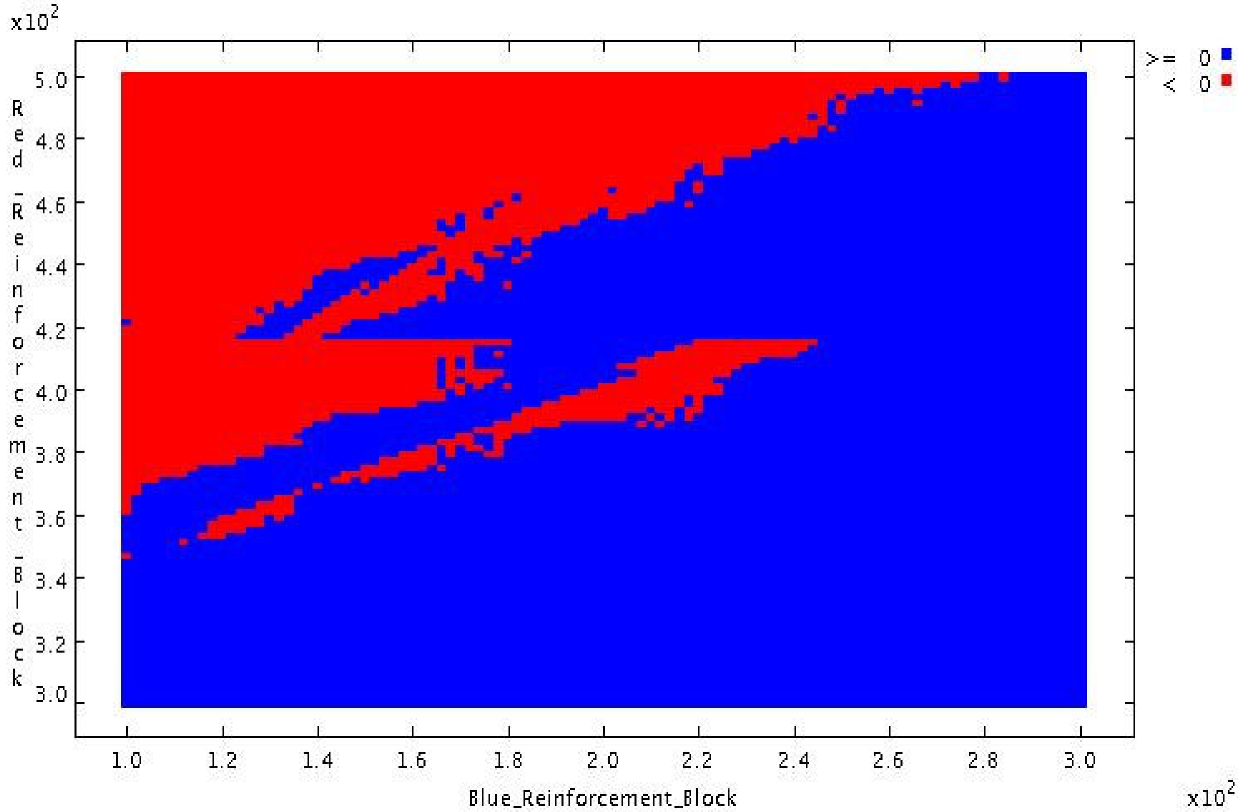

FIGURE 1 Complex borders can emerge from agent interactions

Indeed, the landscape can easily be so rugged that the discovery of the global MLE is computationally intractable. As a consequence, the large corpus of theorems regarding desirable properties of MLEs are not relevant, as the MLE cannot be determined, and even if one had the MLE in hand, one might not be able to prove it maximally likely. Estimation of non-linear models must involve some sort of non-linear optimization algorithm, which may return an answer that is only locally maximally likely. Even if the global maximum was discovered, there may be “second place finishers” whose likelihood is essentially equivalent to the MLE but whose structure or parameter values are very different from the MLE.

When multiple alternative models or parameter vectors are effectively equivalent in explaining the data, and the global maximum likelihood may not be effectively computable, it is questionable what significance should be attached to the most likely model that can be discovered. Instead, we may define a threshold in likelihood that is sufficient that any model of greater likelihood has explanatory power, and investigate the set of models with this property. For a linear model, such a “level set” on the likelihood surface is an ellipsoid that is completely characterized by its center (the MLE) and the variances that characterize its axes. But for highly non-linear models such as agent-based models, the level set may tend to be non-convex and perhaps non-contiguous. For such a set, identifying a single point, even if it is a local maximum in likelihood, does little to characterize the properties of the set as a whole. In this context, neither non-linear optimization of parameters interpreted as model estimation nor specification search across alternative models can be thought of as discovering the “correct” interpretation of the data. Instead, it is reasonable to think of experiments that sample from the likelihood surface as data-driven inference, where multiple alternative models can capture more information from the available data than can any single model.

Thus, in reasoning jointly about ABM and data, we may learn more by viewing the data as constraining the range of plausible model variants (including constraining that range to be the null set when data are disconfirming). Non-linear optimization of likelihood is challenging, and may in the end provide nothing more than another type of hypothesis generation.

This line of reasoning leads us to approaches to understanding data using ensembles of models rather than single models. This has a clear connection with recent developments in the field of data mining, where various approaches to developing and using ensembles of models (albeit much simpler models than ABM) have been under investigation for several years. Examples include the practice of bagging (bootstrap resampling of the training data set to generate an ensemble of alternatives, followed by model averaging), boosting (iteratively modeling the residual of previous modeling steps), and techniques specific to a given modeling approach, such as random forests (Breiman, 2001).

The techniques used by the data mining community do not directly solve the problem of fitting ABM to data. But all of the foregoing suggests the feasibility of developing ensembles of agent-based models that reflect knowledge and assumptions about the structure of the model and data from the system being modeled. The most important property that techniques for doing so should have is that the collection of models generated to represent the actual (typically infinite) ensemble be as diverse as possible while being constrained by the data.

While most applications of ensemble approaches have used model averaging to combine model predictions, ensembles of models can have many other uses. This topic will be explored next.

ROBUST INFERENCE FROM ENSEMBLES OF MODELS

Once an ensemble of models is created that represents the combination of our knowledge, theories, hypotheses, and data, there are a diversity of ways this ensemble can be used. Fundamental to all these uses is the assertion that all these models are plausible, that is, they are all consistent with what we know. Thus, the diversity of models is a resource for uncertainty analysis. Further, while it is difficult or impossible to establish that nothing has been left out, those properties shared by all members of the ensemble do represent a derived fact (albeit one conditioned upon assumptions inherent in the method for generating the explicit members of the ensemble).

Thus, an ensemble of models generated from data can be used as a challenge set to support robust inference. A hypothesis can be assessed against the ensemble to see whether it is true for all members, or whether there is a minority that contradicts it, meaning that it must remain a hypothesis. Even in that case, the hypothesis has been informed by the discovery of the circumstances under which it would fail. Averaging the responses from all members of the ensemble is sometimes a useful way to summarize (and can be viewed as an expectation if the ensembles are thought to be drawn from a probability distribution over our knowledge), but this approach does not exploit all the information available.

I and my colleagues have been particularly interested in applying this framework to policy analysis, where it can be very useful to identify plausible models that, if they were true, would make a policy fail (Bankes, 1993, 2002; Lempert et al., 2002 2003). An example will help to make this clear.

Figures 2 and 3 display results from previously unpublished research in which we explored the use of neural networks to model patterns of terrorist activity. Data on precursors of terrorist activity and corresponding terrorist acts had been collected from public sources, coded, and modeled using a classical two-layer feed-forward neural net, with promising results.[1] We replicated the neural net modeling, but with the twist that we performed bootstrap resampling on the training data in order to create an ensemble of neural net models. Each of the models trained on resampled data predicted the cross-validation data nearly as well as the original neural net. Further, model averaging demonstrated that the ensemble contained more information in the sense of making forecasts that are equal or better than those of the original model. More importantly, the ensemble provided an indication of the certainty in this prediction across the ensemble. Figure 2 shows the percentage of members of the ensemble forecasting each of four categories of terrorist action for three different test cases. As can be seen, in one situation there is 100% uniformity in predicting an assassination attempt, while in another each category of terrorist activity receives at least a small amount of weight from some model using some protocol for making predictions.

The agreement or divergence of predictions across the ensemble gives some sense of the certainty of the forecast, which is clearly more useful than a single forecast would be. Moreover, we can take a next step and use the ensemble of models we have developed from the data as a challenge set to use to develop robust policies. For demonstration purposes, we asserted a payoff matrix that gives a utility associated with the combination of a terrorist act and an associated counter-terrorism strategy, shown in Figure 3. This allows us to explore issues of type 1 vs. type 2 errors, the desirability of portfolio strategies, and so forth. (A full description of this mock analysis is beyond the scope of this paper.) Figure 4 displays a landscape in which the color-coded expected outcome is displayed against two dimensions of uncertainty or choice.

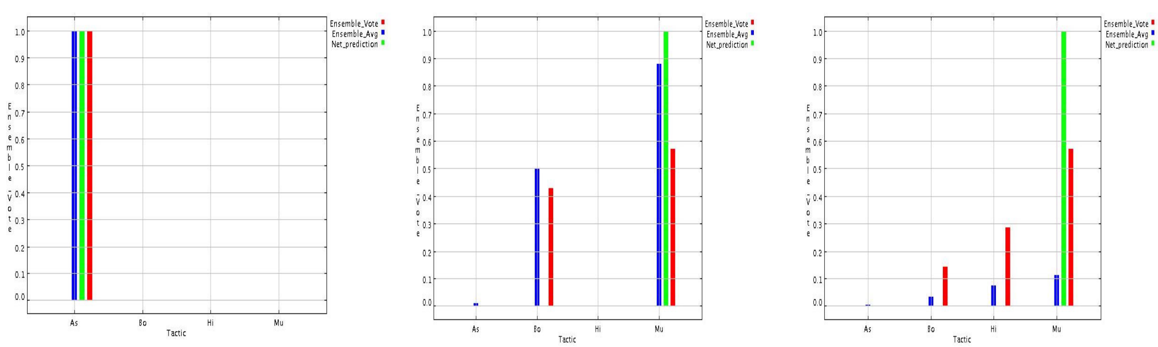

FIGURE 2 Responses of an ensemble of neural nets to three new terrorism test cases (The three colored bars represent different rules for using the ensemble of networks to make a forecast, e.g., winner take all versus ensemble average.)

[1] The collection of data, its coding, and the original neural net modeling was performed by colleagues at the American Institute for Research as part of DARPA-sponsored research.

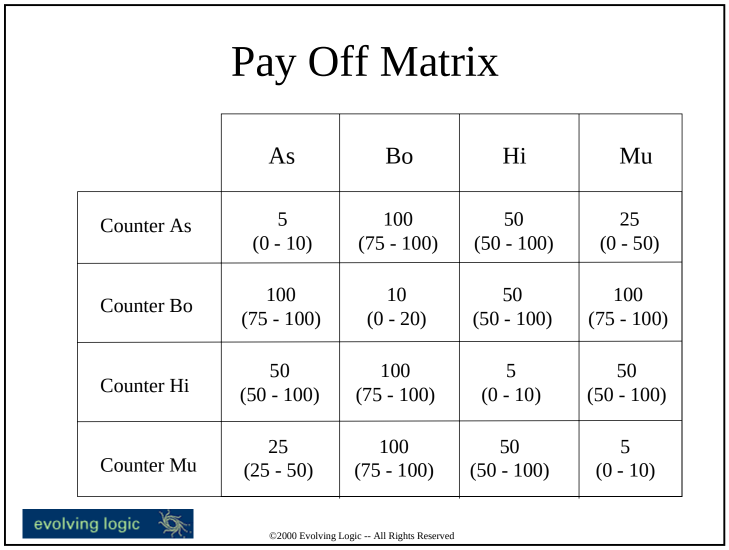

Pay Off Matrix

	As	Bo	Hi	Mu
Counter As	5 (0 - 10)	100 (75 - 100)	50 (50 - 100)	25 (0 - 50)
Counter Bo	100 (75 - 100)	10 (0 - 20)	50 (50 - 100)	100 (75 - 100)
Counter Hi	50 (50 - 100)	100 (75 - 100)	5 (0 - 10)	50 (50 - 100)
Counter Mu	25 (25 - 50)	100 (75 - 100)	50 (50 - 100)	5 (0 - 10)

evolving logic

FIGURE 3 Assumed payoff matrix for terrorist acts versus counter-terrorism measures

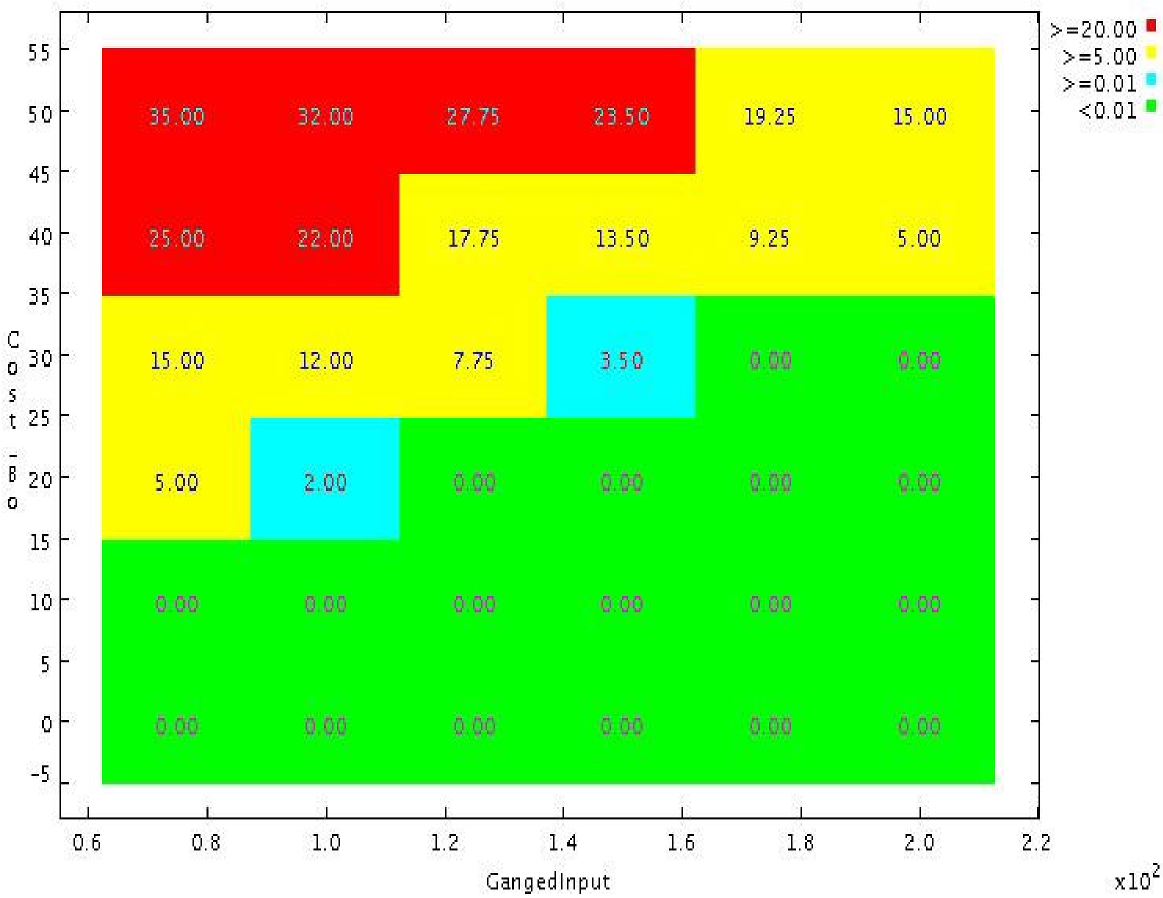

FIGURE 4 Expected outcomes of possible strategy as function of two uncertainties

In more recent work, we have explored decision theoretic approaches to compensating for the risk that the method for constructing a given ensemble may have biased the following analysis. In Lempert et.al. (2003), we made the process of ensemble construction an iterative one, in which a tentative conclusion is used to seek additional plausible models that might invalidate it. The result is a co-evolutionary dynamic in which computer and human resources are used in parallel to seek (1) strategies that are robust across the ensemble being used as a challenge set and (2) members of the ensemble that will be more challenging for the leading candidates. This approach appears to be very promising.

CONCLUSIONS

The state of ABM reflects both significant promise and significant challenge. Research strategies based on developing and exploiting ensembles of alternative model instances can help meet the challenge of both incorporating data in the construction of agent-based models and in making them more useful in problem solving. The two problems confronting the fitting of ABM to data, likely specification error and the computational complexity of estimation, can both be met in part by pruning ensembles of models to be consistent with the data instead of seeking the single best model. And ensembles can readily serve as a challenge set against which to ask questions. For science, one can seek statements robust (invariant) across the ensemble. And for policy analysis, one can search for policies that perform well for any member of the ensemble, that is, any plausible model.

REFERENCES

Akaike, H., 1973, “Information Theory and an Extension of the Maximum Likelihood Principle,” in *Second International Symposium on Information Theory*, 267–281.

Bankes, S., 1993, “Exploratory Modeling for Policy Analysis,” *Operations Research*, **41**:435–449.

Bankes, S., 2002, “Tools and Techniques for Developing Policies for Complex and Uncertain Systems,” *Proceedings of the National Academy of Sciences*, **99**:7263–7266.

Breiman, L., 2001, “Random Forests,” *Machine Learning*, **45**(1):5–32.

Chang, F.J., and J.W. Delleur, 1992, “Systematic Parameter Estimation of a Watershed Acidification Model,” *Biological Processes*, **6**:29–44.

Dewar, J., J. Gillogly, and M. Juncosa, 1996, “Non-monotonicity, Chaos, and Combat Models,” *Military Operations Research*, **2**(2):37–49.

Draper, D., 1995, “Assessment and Propagation of Model Uncertainty,” *Journal of the Royal Statistical Society, Series B (Methodological)*, **57**(1):45–97.

Engle, J.H., 1954, “A Verification of Lanchester’s Law,” *Operations Research*, **2**:163–171, Hastie, T., R. Tibshirani, and J. Freidman, 2001, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, New York, NY: Springer.

Lempert, R., S. Popper, and S. Bankes, 2002, “Confronting Surprise,” *Social Science Computing Review*, **20**:420–440.

Lempert, R., S. Popper, and S. Bankes, 2003, *Shaping the Next One Hundred Years: New Methods for Quantitative, Longer-Term Policy Analysis*, RAND Report MR-1626-RPC.

Mendes, E., and S. Billings, 2001, “An Alternative Solution to the Model Structure Selection Problem,” *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, **31**(6):597–608.

Miller, A.J., 1990, *Subset Selection in Regression*, London, U.K.: Chapman and Hall.

Computational Organization Theory

DYNAMICS OF EXPERTISE IN ORGANIZATIONS: AN AGENT-BASED MODELING EXERCISE

K.C. DESOUZA,* S. BHATTACHARYYA, and J.R. EVARISTO,
University of Illinois at Chicago, Chicago, IL

ABSTRACT

Organizations are rational entities and only enlist individuals (as employees) as long as they provide some resources of interest. Such resources are in the form of the tacit knowledge that employees bring to the organization. The integration and synthesis of such expertise for the collective good of "the organization" is not yet fully understood. This paper takes a first look at understanding the dynamics of the allocation of expertise and movement among agents in the organization. Of specific interest is the way that expertise moves in the organization through the process of socialization.

Keywords: Socialization, tacit knowledge, expertise, crossover

INTRODUCTION

Organizations are rational entities and enlist individuals (as employees) only as long as they provide some resources of interest. These resources are in the form of the tacit knowledge that employees bring to the organization (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998). An organization's most valuable asset is the knowledge that resides in the minds of its employees (Nonaka, 1994; Grant, 1996). As often noted, organizations have a great deal of individual expertise; however, the integration and synthesis of such expertise for the collective good of "the organization" are not fully understood (Tiwana, 2003; Tiwana and McLean, 2002).

In this paper, we take a first look at understanding the dynamics of the allocation of expertise and movement among agents in an organization. Of specific interest is how expertise moves in the organization through socialization. Socialization is a key process in bringing tacit knowledge and expertise to bear on projects (Nonaka, 1994). Because of the lack of literature in the area of the dynamics of expertise, our model is simplistic and was developed as a result of our observations of behavior in organizations.

Our initial experiments garnered an interesting set of results. For instance, we found that an increase in the percentage of experts to nonexperts does not always lead to an increase in the overall knowledge of the organization. After a given point, increasing the number of experts results in a decline in overall knowledge in the organization. Similarly, we found that the initial disposition of experts on domains of knowledge affects the number of new agents that will become experts in these domains. These experiments have applications for work in organizational theory and strategic management. Specifically, we feel that uncovering the dynamics of expertise in organizations will help to set policy and better manage knowledge and expertise in these organizations.

---

* *Corresponding author address:* Kevin C. Desouza, Department of Information & Decision Sciences, University of Illinois at Chicago, 601 South Morgan Street, M/C 294, Chicago, IL 60607; e-mail: kdesou1@uic.edu.

MODEL

We conceptualize our organization as having an interest in a set of domains of knowledge. These domains can be areas in which an organization conducts activities, such as accounting, finance, legal, human resources, marketing, and operations. Each agent has a level of expertise (a score) in these domains ranging from 0 to 9. A score of 0 means that the agent has no expertise in that domain, whereas a score of 9 indicates that the agent possesses maximum expertise in that domain.

Agents are restricted as to the total amount of expertise they can possess across all domains of knowledge. The sum of all scores on the domains of expertise must be less than the limited cognitive capacity. As many studies have shown, agents have a cognitive capacity that governs how much they can store and recall at any point in time. In our model, the cognitive limit capacity is set at 50; thus, the sum of an agent's score on the 25 domains cannot exceed 50.

Two classes of agents are included in the organization — experts and nonexperts. Initially, experts have a higher level of expertise in selected domains of knowledge. Experts are defined as those who have high scores (>7) in five or more domains, while nonexperts have no score greater than 5 in any domain. Initially, we segmented our pool of agents into 20% experts and 80% nonexperts. Except for the initial endowments of scores in domains, no difference exists between the two classes of agents.

Agents increase their expertise through learning. Learning is defined as the acquisition of new knowledge or expertise within a given domain of specialization (Tiwana and McLean, 2002). This new knowledge can either decrease or increase the agent's level of expertise in the domain. People learn while interacting with their peers and working on tasks. Members in organizations need to exchange knowledge to accomplish tasks (Kaplan and Miller, 1987). We model two types of learning. *Interaction-based learning* is akin to the traditional crossover operator in genetic algorithms (Holland, 1975). *Communication-based learning* occurs when two agents interact. Once two agents are selected, they follow the rules of engagement. First, each agent determines its top three domains of expertise. Second, we select an agent and go through the top three domains as follows:

- *If the current expertise is an expertise of the other agent:* Have each agent perform observation noise checks on each other's value at the selected domain. Once these checks have been completed, each agent observes the true score or an artificially inflated or deflated score in the domain of expertise if observation noise was applied. The agent with the higher observed value retains its value at the selected domain, while the other agent copies the higher value from his peer subject to the copy noise. If two agents have the same observed value at the selected domain of expertise, each agent exchanges values at a random domain that is an expertise of neither. This case is also subject to copy noise.

- *If the agents do not have matching domains of expertise:* Conduct the exchange on a random domain that is an expertise of neither. Copy noise is applied to the exchange.

Some of the rationale for choosing the rules engagement are based on the literature. Stasser and Titus (1985, 1987) found that groups were much more likely to discuss information that had been previously shared than to talk about unshared information. Hence, when two agents are chosen for communication, they first look for commonalities in their domains of expertise. As asserted by Stasser, et al. (1995), members with expertise (experts) try to focus their search on other domains so that they can obtain relevant information rather than improve their areas of expertise. When members of the organization do not share the domain of knowledge, there is a potential for members to seek out and acquire new knowledge (Stasser, et al., 1995). Moreover, even when domains of knowledge are shared, members can acquire information that they have temporarily forgotten or cannot recall (Kaplan and Miller, 1987; Stasser, et al., 1995). This fact is captured via the rule that if two agents have similar scores in areas of expertise, they attempt to conduct an exchange in other domains.

Each agent also learns independently, which is modeled via a mutation operator. At every time step, an agent with a probability of 0.005 mutates five domains of expertise. Mutation can cause an agent to either increase or decrease its expertise in a given domain. The rationale is that an agent might either learn something new, thus increasing its expertise, or realize that its knowledge in a domain is outdated or obsolete, meaning its score declines.

It is difficult to observe what knowledge and expertise an agent possesses and to transfer such knowledge perfectly (Nonaka, 1994; Van den Bosch, et al., 1995). To account for those factors, we incorporated an observation noise and a copy noise. An *observation noise* is defined as the imperfection in an agent's perception of his/her peer's expertise in a given domain. A *copy noise* is defined as the imperfection in imitating or transferring expertise between two agents. Observation noise is the probability that each agent's score will be artificially inflated or deflated by a probability of 0.25. Observation noise stays consistent throughout any number of exchanges and the life of the simulation. Copy noise, however, decreases on the basis of the frequency of interactions between agents. The first time two agents meet, they are essentially strangers, with varied backgrounds and contexts, and hence copy noise will be high. If they meet for the second time, however, they have developed some aspect of a share context that will help to decrease the difficulty in the transfer of expertise.

In their study of the process of the socialization, Nonaka and Takeuchi (1995) have ascribed to some of the above phenomena. Many studies have also attested to the fact that members bring unique knowledge and expertise to a group or organization (Stasser, et al., 1995), but it is difficult to identify these unique knowledge areas (Stasser, et al., 1995). The copy noise is modeled as follows. If two agents interact for the first time, there is a 0.8 probability that copy noise will occur, which reduces the expertise transferred by 0.75. If two agents are communicating for the second time, the chance of copy noise is 0.5, which reduces the expertise transferred by 0.5. Agents that interact more than twice have no copy noise and can transfer expertise perfectly. Wegner (1986) asserts that groups who have a long history of working together can pass knowledge more easily and also value each other's areas of expertise. Moreover, communication and interpretation among members of such groups are very fluid. Wegner (1986) and Wegner, et al. (1985) articulate the role of transactive processes and memories. They argue that individuals working in groups construct and reconstruct separate memories to determine smoothness in information transfer over time and develop shared knowledge spaces. Wegner states, "The transactive memory system begins when the individuals learn something about each others' domains of expertise" (1986, page 191).

People work in groups or around projects in organizations, which implies frequent interaction with a few peers and rare but necessary, interaction with members from the rest of the organization. To model these interactions and associated intricacies, we selected the following approach. Each agent interacts for a given percentage of time with agents who are in its neighborhood; the remaining percentage of time, the agent interacts with agents in the community at large. Each agent interacts 60% of the time with agents two Euclidean distances from its placement on the grid. During the remaining 40% of the time, the agent interacts with anyone from the rest of the organization. Agents are placed on a grid on the basis of the affinity of their expertise. Agents with similar domains of expertise are placed close together.

RESULTS

Table 1 details the parameter setting for the simulation. Figure 1 displays the number of domains with nonzero scores and the number of domains with high scores (>7) for all agents in the simulation, for a total number of domains of 6,400 (256×25). Initially, we see that the number of domains with nonzero scores falls sharply, and the number of domains with high scores rises. In other words, agents initially increase their expertise in domains at the cost of having no knowledge in other domains. Because an agent's total expertise is constrained by the cognitive capacity (50), agents must move expertise between domains.

TABLE 1 Model parameters

ORG_KNOW_SPACE	25
EXPERTISE_RANGE	[0–9]
NUM_OF_DOMAINS	25
MAX_COGNITIVE_CAPACITY	50
MAX_SKILL_VALUE	9
MIN_EXPERT_SKILL_VALUE	7
MAX_AVERAGE_SKILL_VALUE	5
NUM_AGENTS	256

Figure 2, which is a continuation of Figure 1, depicts the state of simulation up to 2,250 cycles. As can be inferred, the rise in the number of domains with nonzero scores occurs at a faster rate than the rise in the number of domains with high scores. This result can be due to the fact, that by this time, agents have fixed domains with high expertise, and when exchanges occur between peers, they would rather focus on domains that are unknown to them. This artifact also occurs because agents interact more closely with their peers in the neighborhood. Many share their areas of expertise; hence, a level of expertise saturation is reached. As a result, they explore areas where they know nothing and learn new domains because all share common areas of expertise with very similar scores on the domains.

We also generated six maps to uncover patterns of spatial expertise. Each agent's position on the grid was highlighted with a color that represented its score on various attributes of interest. Table 2 depicts the coloring scheme along with the associated scores for each of the maps; Figure 3 depicts the maps at various time cycles.

In Map 1, we looked to see *how an agent's top three scores faired*, with a minimum of 0 and a maximum of 27. The top three scores of most agents ranged from 16 to 20 at the start of the simulation (*t* up to 200); a few agents' scores ranged from 25 to 27. Of interest is that by cycle 2,250, all agents' top three scores ranged from 16 to 20. Thus, even agents that had high scores or were very knowledgeable in their three domains of expertise, lost some of their

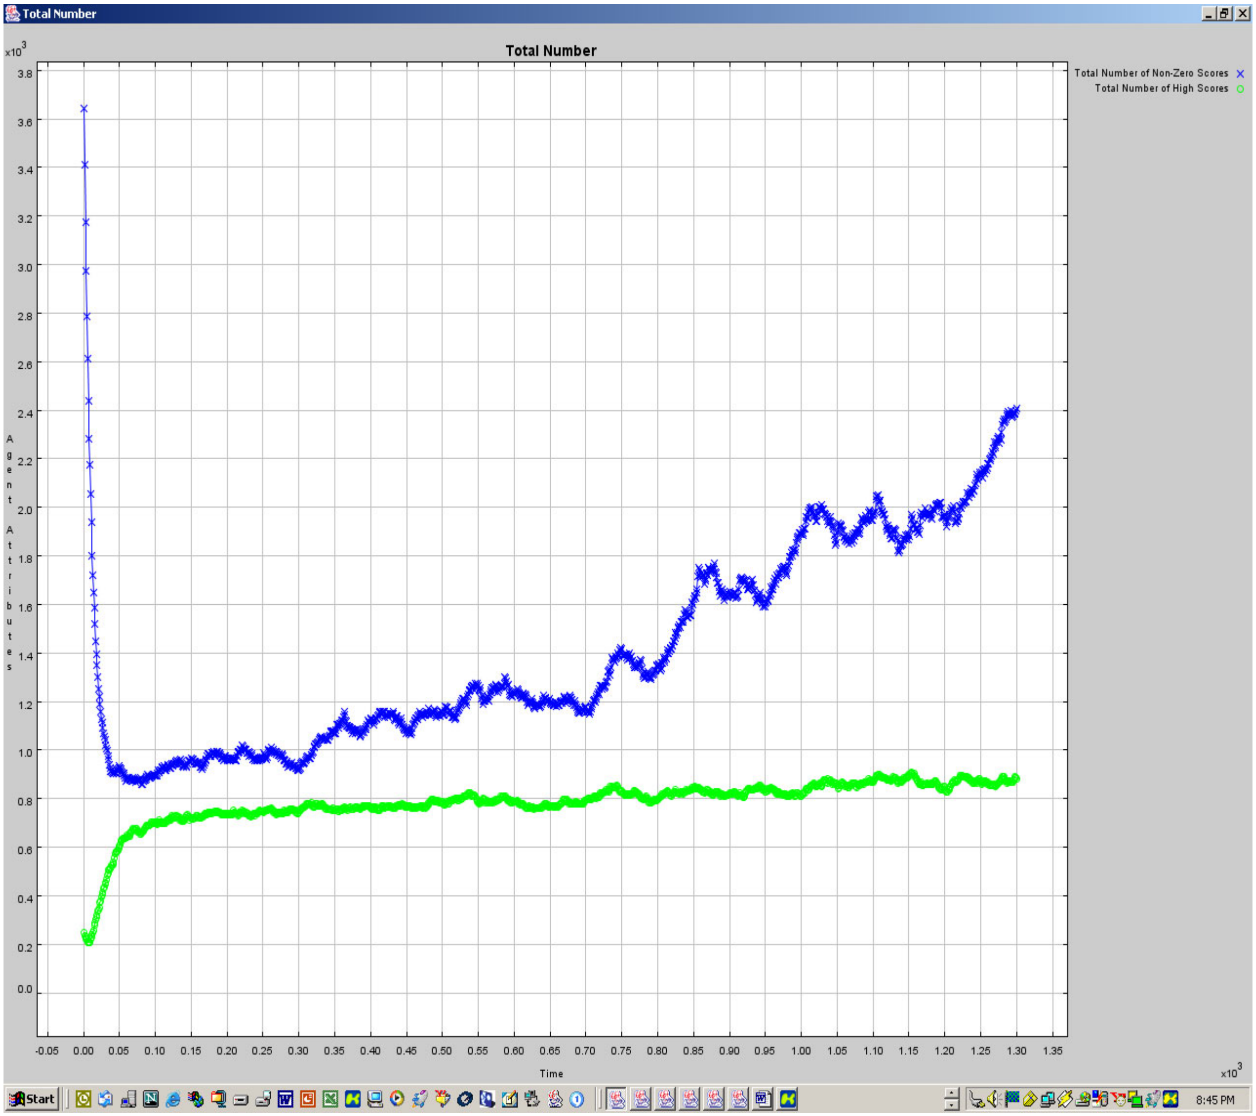

FIGURE 1 Number of nonzero scores (blue) and high scores (green), up to t = 100

knowledge. The reason for this behavior is the presence of a mutation operator, and the fact that over time agents learned in areas other than their domains of expertise and sacrificed some of the high scores to account for the cognitive limit.

In Map 2, we looked at the *total domain knowledge* an agent possessed, with a maximum of 225 (25×9) and a minimum of 0. Since we imposed a cognitive capacity, however, no agent could have a score greater than 50. In the beginning of the simulations, agents were widely distributed based on the total domain knowledge possessed. As time passed, however, all converged at their maximum capacity of 50.

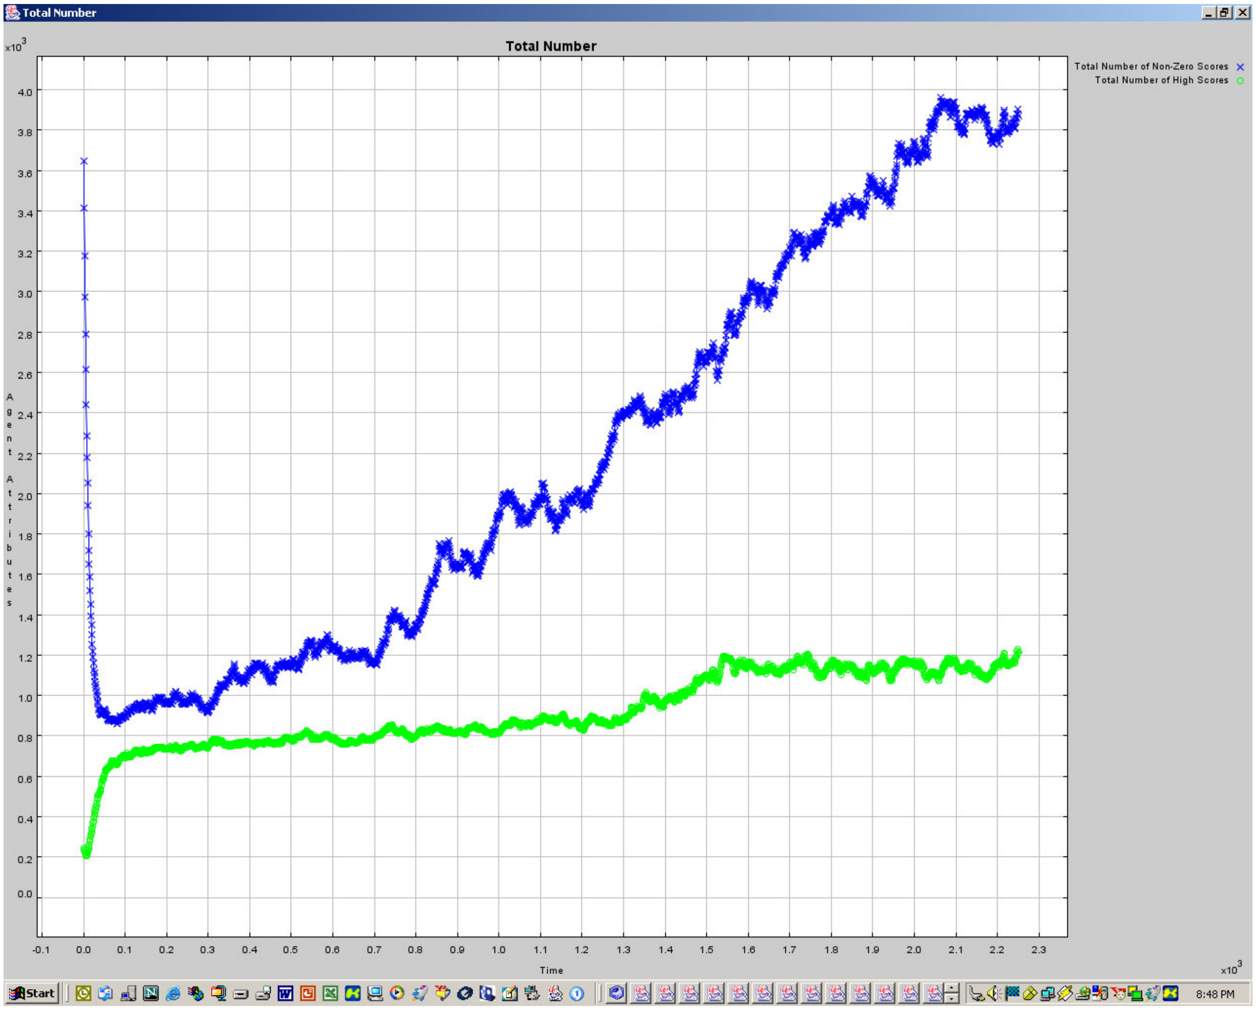

FIGURE 2 Number of nonzero scores (blue) and high scores (green), up to *t* = 2,250

In Map 3, we looked at the *number of domains in which zero scores appear* for an agent, with a maximum of 25 and a minimum of 0. At the start of the simulation, a large proportion of agents had more than 20 domains with a score of 0; few agents ranged from 16 to 20. As agents interacted with their peers, expertise was generated, and a form of *exploration* emerged in which agents started learning knowledge not in their domains of expertise. As expected, over time (up to 1,000 and 2,250 runs), most agents had from 4 to 7 domains with 0 scores. In addition, a significant number of agents had one-half or more of their domains with 0 scores (12–14).

In Map 4, we looked at the *number of domains with scores less than or equal to 3* for an agent. A vast majority of the agents had low scores on 21 or more domains in the initial runs of the simulation. Over time, this pattern persisted with a marginal improvement, where around 40% of the population had lowered the number of domains with a score of less than 3. At time step 2,250, most agents had scores of less than 3 in 1,220 domains. This fact indicates that agents have started to develop their core areas of expertise. It is interesting to note that a sizable group of agents have low scores in 15 or 17 domains. We can argue that this score is indicative of agents choosing domains of areas of specialization at the cost of these domains.

TABLE 2 Spatial maps

		Maps					
No.	Color	1[a]	2[b]	3[c]	4[d]	5[e]	6[f]
1	Blue	–	0–10	0–2	0–2	0–2	0–2
2	Cyan	–	11–15	3–5	3–5	3–5	3–5
3	Gray	0–3	16–20	6–8	6–8	6–8	6–8
4	Green	4–7	21–25	9–11	9–11	9–11	9–11
5	Magenta	8–11	26–30	12–14	12–14	12–14	12–14
6	Orange	12–15	31–35	15–17	15–17	15–17	15–17
7	Pink	16–20	36–40	18–20	18–20	18–20	18–20
8	Red	21–25	41–45	21–23	21–23	21–23	21–23
9	Yellow	25–27	46–50	24–25	24–25	24–25	24–25

[a] Sum of three highest scores.

[b] Sum of domain knowledge.

[c] Number of domains with scores of 0.

[d] Number of domains with scores less than or equal to 3.

[e] Number of domains with scores less than or equal to 6.

[f] Number of domains with scores greater than 6.

In Map 5, we looked at the *number of domains with scores less than or equal to 6 but greater than 3* for an agent. During the initial runs, most agents had few domains with an *average* level of expertise (most have only 0 to 2). This pattern persisted for most of the simulation. If the simulation runs to infinity, domains with average patterns rise only slightly — to between 6 and 8 (see Step 2,250); a select few have between 9 and 11 domains with average knowledge. This pattern shows that agents have a high degree of variance in their expertise. They are very strong in certain areas (as shown in Map 6) or have a large number of domains with below average knowledge. We can also argue that this pattern is due to the emergence of core competencies.

Finally, in Map 6, we looked at the *number of domains with scores greater than 6* for an agent. Most of the agent population had from 0 to 5 domains of expertise during the initial period. As time progressed, most of the population converged and had between 3 and 5 areas of expertise. What is interesting to note is that if the simulation runs to 2,250 steps, one-half of the agents increase the number of domains with expertise to between 6 and 8, while at the same time, many preserve their original number of domains.

Figure 4 displays the evolution of expertise among average and expert agents. All else being equal, expert agents have a slower learning rate than the average agents. We ran various simulations changing the proportion of experts in the population to see if the proportion of experts played role in the evolution of expertise in the organization. As seen in Figure 5, varying the proportion of expertise did little to change the overall level of expertise in the organization.

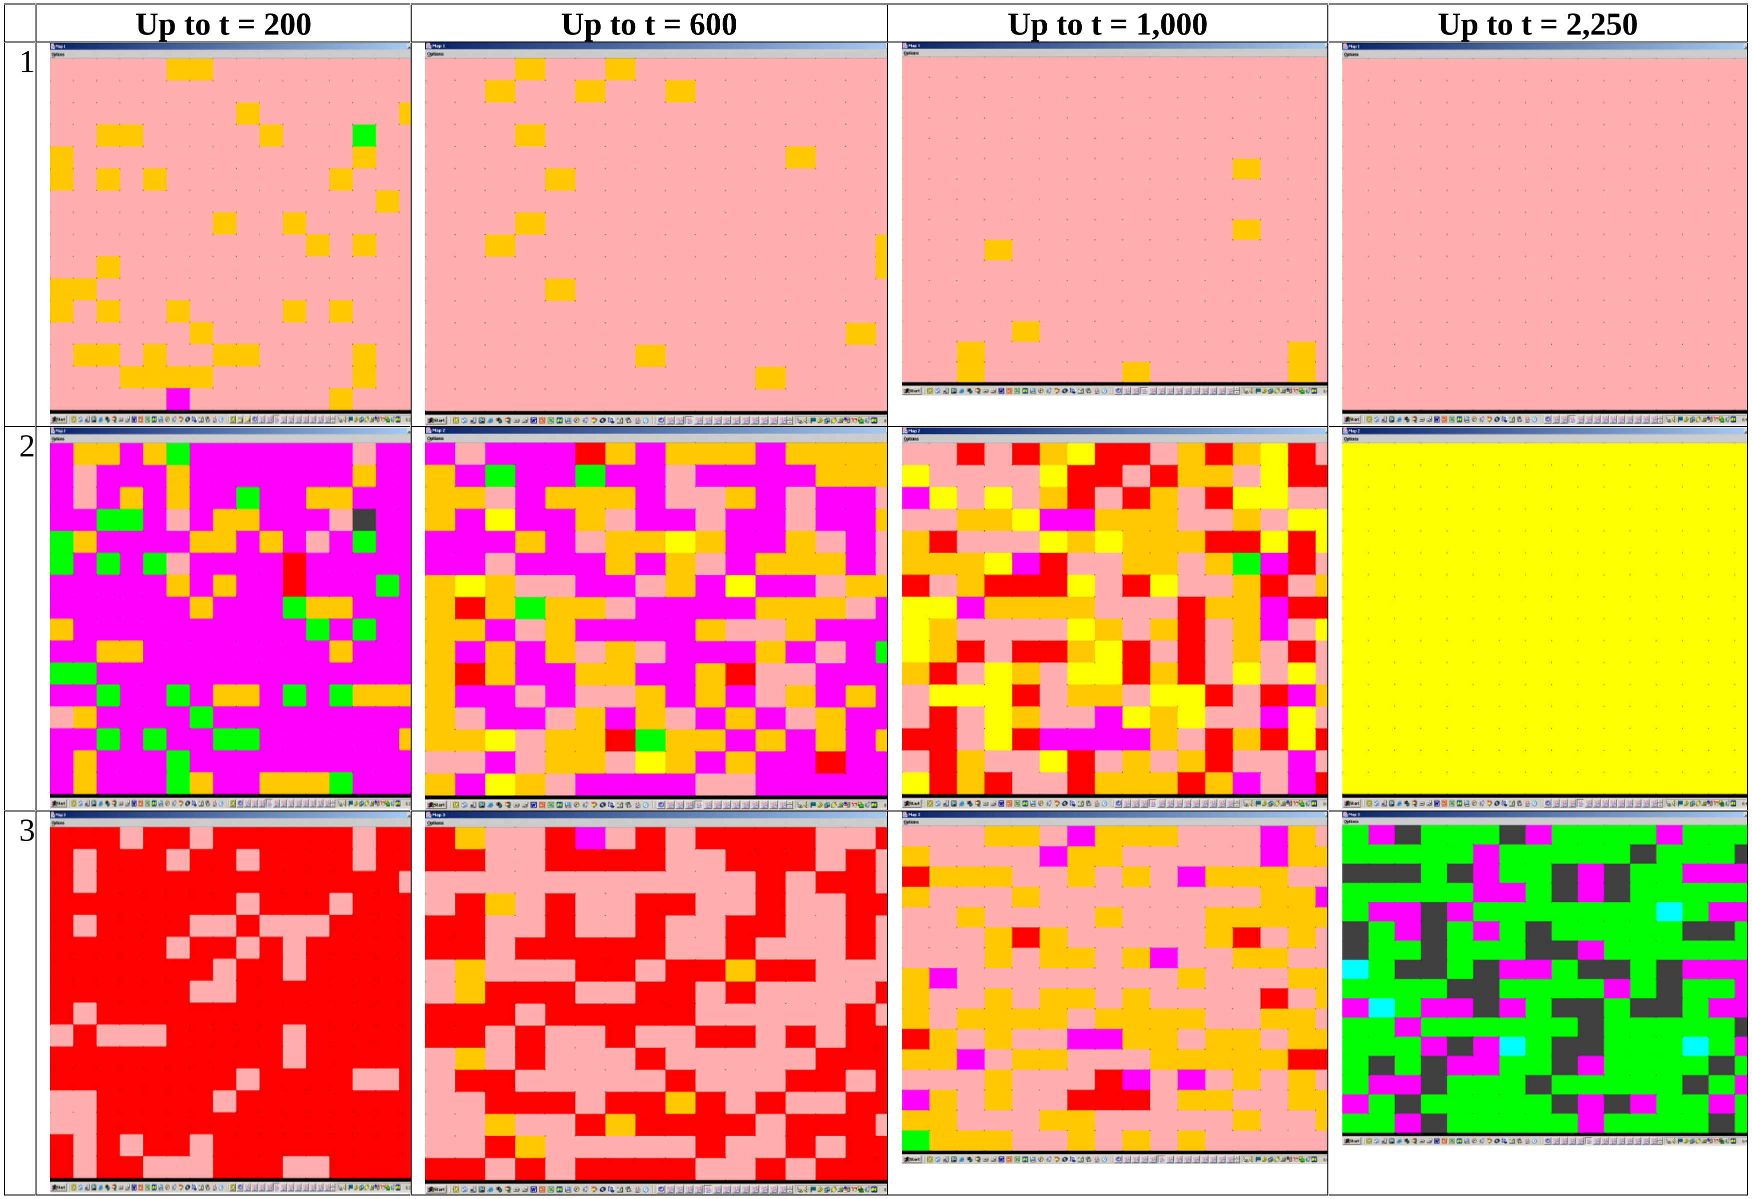

FIGURE 3 Maps used to uncover spatial expertise patterns at various time cycles (see Table 2 for explanation of colors)

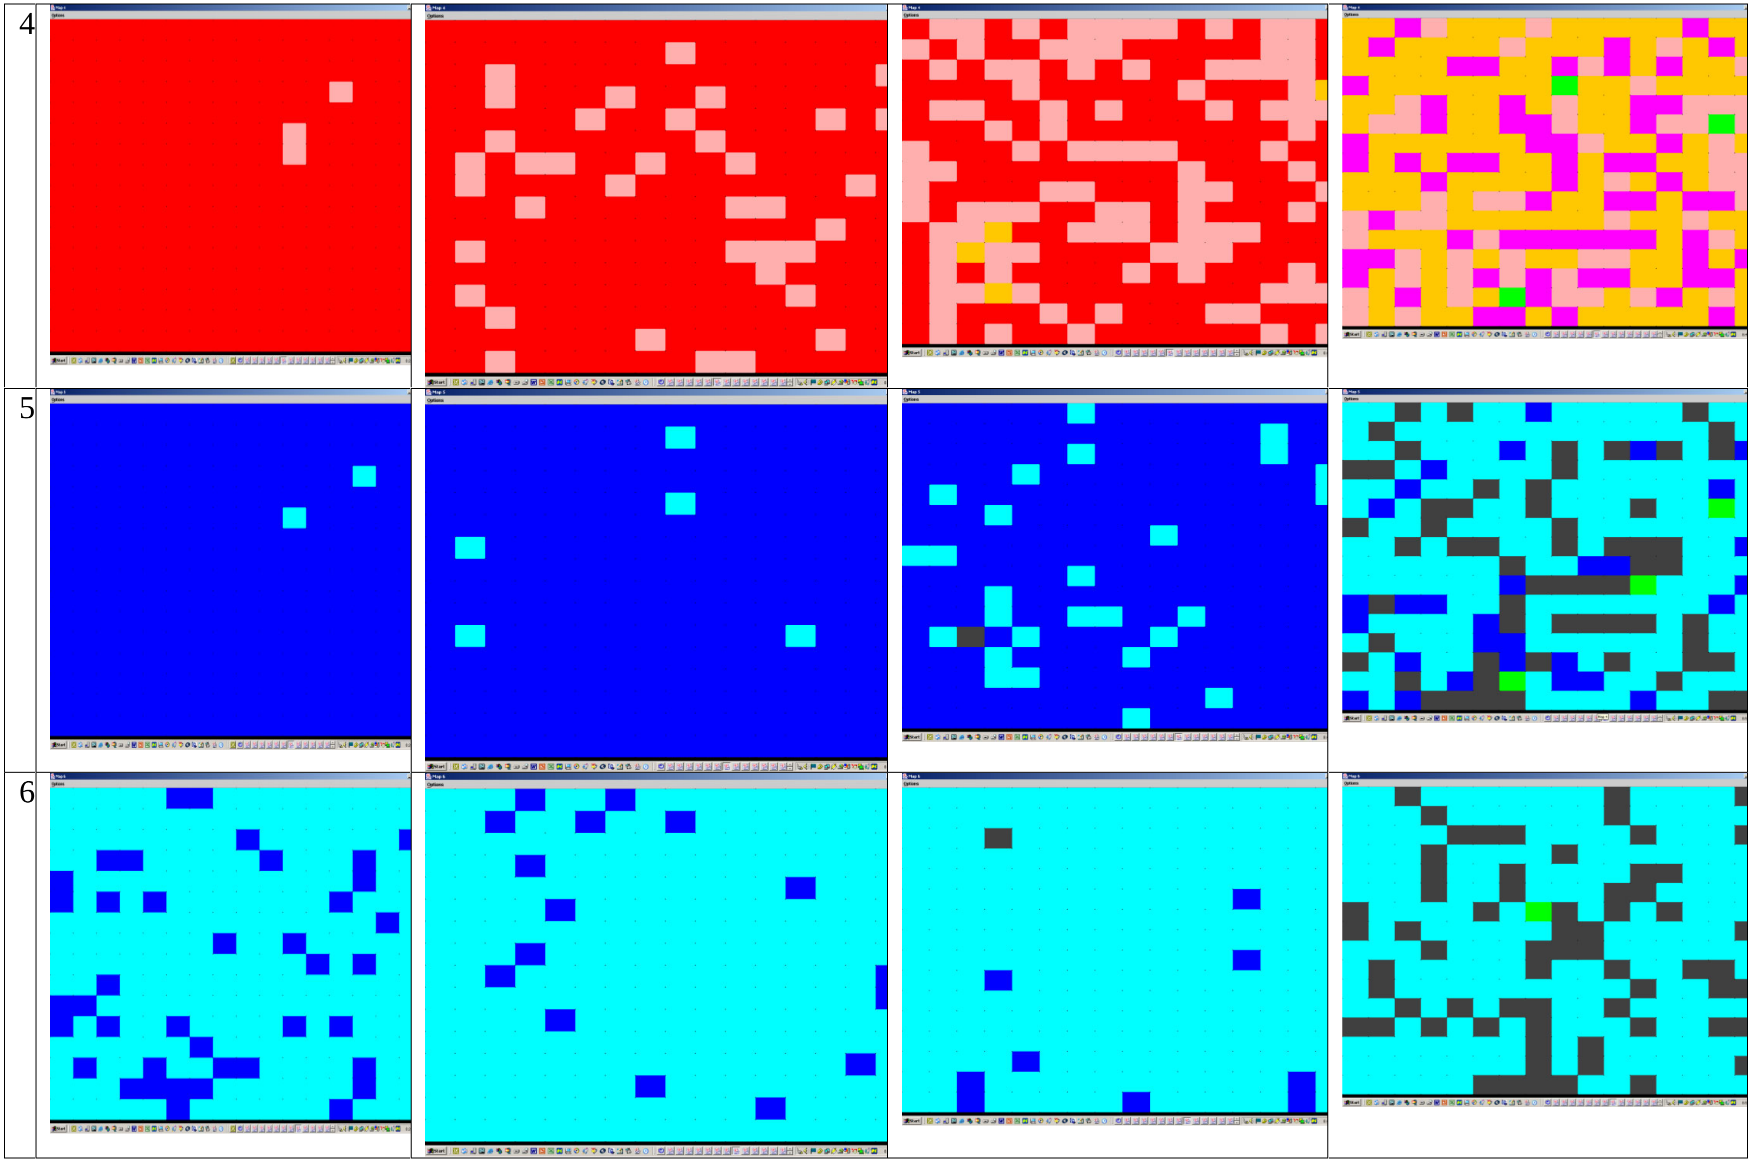

FIGURE 3 (Cont.)

CONCLUSION

This research investigated the dynamics of expertise in organizations. Our results are preliminary and must be viewed in that light. Much work can be carried out to study how different network topologies might affect expertise in organizations. Researchers are also well advised to carry out studies in which agents enter and drop out of the organization. This factor would enable us to capture a more realistic setting in which experts leave and join organizations.

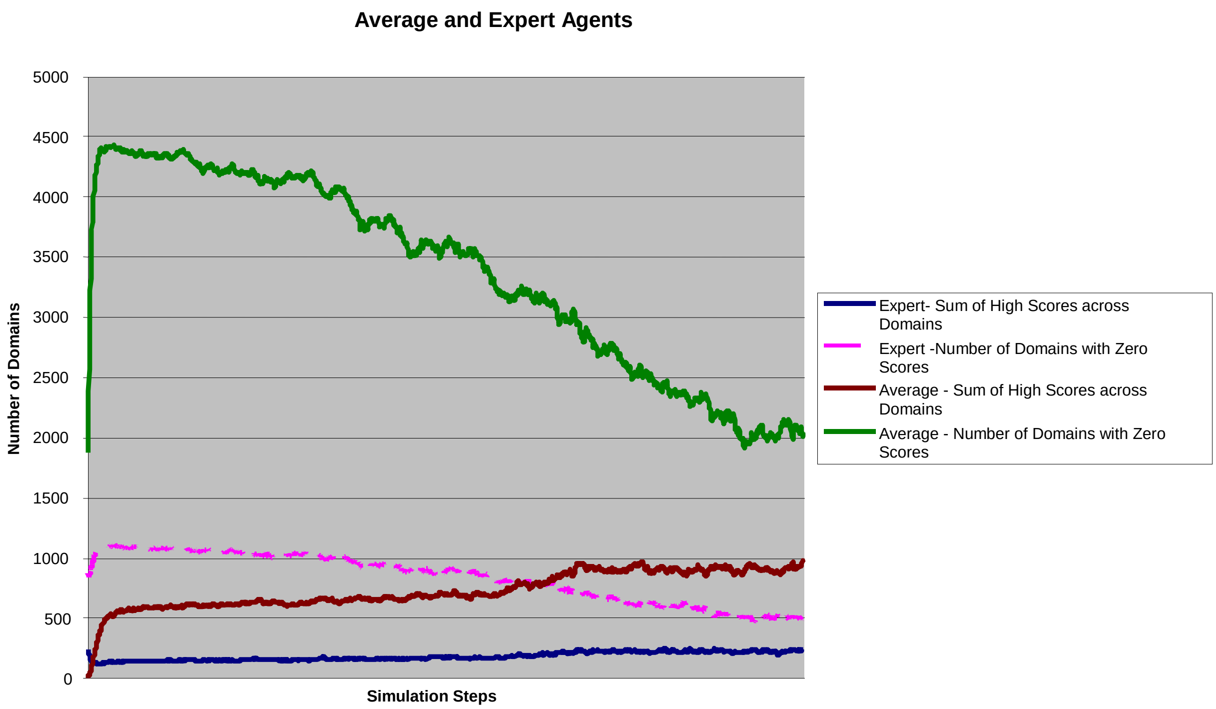

FIGURE 4 Expert versus average agents

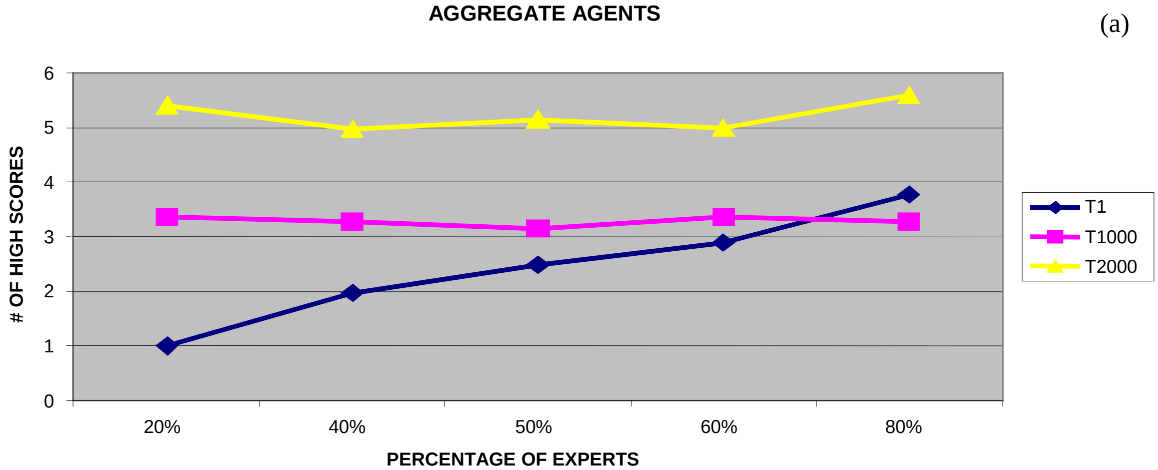

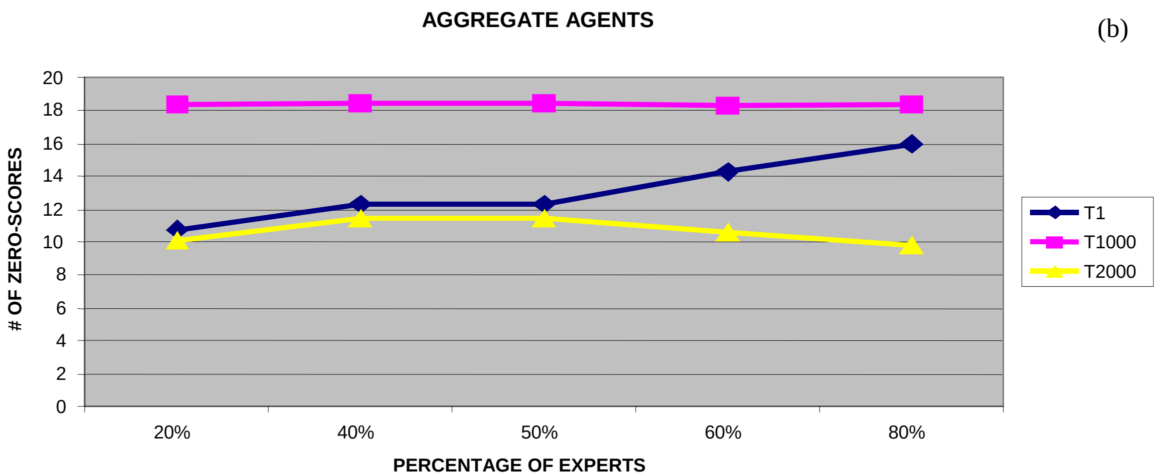

FIGURE 5 Varying proportion of experts with (a) number of high scores and (b) number of zero scores

REFERENCES

Davenport, T., and L. Prusak, 1998, *Working Knowledge*, Boston, MA: Harvard Business School Press.

Grant, R.M., 1996, “Toward a Knowledge-based Theory of the Firm,” *Strategic Management Journal* **17**:109–122 (Winter Special Issue).

Holland, J., 1975, *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press.

Kaplan, M.F., and C.E. Miller, 1987, “Group Decision Making and Normative versus Informational Influence: Effect of Type of Issue and Assigned Decision Rule,” *Journal of Personality and Social Psychology* **53**:306–313.

Nonaka, I., 1994, “A Dynamic Theory of Organizational Knowledge Creation,” *Organization Science* **5**:14–37.

Nonaka, I., and H. Takeuchi, 1995, *The Knowledge-creating Company: How Japanese Companies Create the Dynamics of Innovation*, New York City, NY: Oxford University Press.

Stasser, G., and W. Titus, 1985, “Pooling of Unshared Information in Group Decision Making: Biased Information Sampling during Group Discussion,” *Journal of Personality and Social Psychology* **48**:1467–1478.

Stasser, G., and W. Titus, 1987, “Effect of Information Load and Percentage of Shared Information on the Dissemination of Unshared Information during Group Discussion,” *Journal of Personality and Social Psychology* **53**:81–93.

Stasser, G., D.D. Stewart, and G.M. Wittenbaum, 1995, “Expert Roles and Information Exchange during Discussion: The Importance of Knowing Who Knows What,” *Journal of Experimental Social Psychology* **31**:244–265.

Tiwana, A., 2003, “Affinity to Infinity in Peer-to-peer Knowledge Platforms,” *Communications of the ACM* **46**(5):77–80.

Tiwana, A., and E. McLean, 2002, “Knowledge Integration and Individual Expertise Development in E-Business Project Teams: From the Pod to the Peas,” Association of Computing Machinery Special Interest Group on Computer Personnel Research Conference on the Information Technology Workforce in the Digital World: Managing IT Workers across Borders, Time, Distance, and Culture, Kristiansand, Norway, May 14–16, pp. 120–123.

Van den Bosch, B., and C.A. Higgins, 1995, “Executive Support Systems and Learning: A Model and Empirical Test,” *Journal of Management Information Systems* **12**(2):99–131.

Wegner, D.M., 1986, “Transactive Memory: A Contemporary Analysis of the Group Mind,” in B. Mullen and G. Goethals, eds., *Theories of Group Behavior*, New York City, NY: Springer-Verlag, pp. 185–208.

Wegner, D.M., T. Giuliano, and P. Herterl, 1985, “Cognitive Interdependence in Close Relationships, in W.J. Ickes, ed., *Compatible and Incompatible Relationships*, New York City, NY: Springer-Verlag, pp. 252–276.

A QUANTUM MODEL OF CONTROL FOR MULTI-AGENT SYSTEMS

W.F. LAWLESS,* Paine College
J.M. GRAYSON, Augusta State University

ABSTRACT

The major unsolved problem of social interaction, studied with social psychology from the 1920s and game theory from the 1940s, is to distinguish a group of individuals from its disaggregate. Apparently, social interactions cannot be simulated efficiently with traditional methods. The failure to solve this problem efficiently likely will preclude agent autonomy, especially with multi-agent systems using reinforcement or adaptive learning for control. In contrast, the quantum perturbation model has made progress in understanding social interaction with field evidence and a mathematical model of the two factors of action and observational uncertainty based on the entangled members of a group. We have extended our findings to organizational and argument theory. We begin to extend our work, a work-in-progress, to control theory.

Keywords: Quantum agents, perturbations, organizations

INTRODUCTION

Computational social models predicated on traditional social learning theory (e.g., game theory) assume that action information *I* and observation *I* are equivalent — similar to the assumption of perfect *I* in game theory, where interdependence is crafted through the configuration of arbitrarily valued, forced choices. The general result of these models underscores the value of cooperation (Axelrod, 1984; Nowak, et al., 2000) to forcibly seek consensus in decision making; the greater value of an individual compared with a group rational perspective (Stroebe and Diehl, 1994); and the lack of trust from the competition or conflict inherent in the majority rule of democratic decision making (Worchel, 1999). Yet, traditional models have been contradicted by the persistence of, or even the necessity for, tension from competition to uncover hidden *I* to solve ill-defined problems (*idp*s), recognized by Kuhn (1977) as the essential ingredient for scientific inquiry (see also Von Neumann, 1961). Luce and Raiffa (1967) concluded that the rational individual perspective mathematically was likely unable to comment on social processes, and Wendt (1999) concluded that paradoxically, trust did not arise from cooperation. Further, over the years, consensus decision making has been criticized for political (European Union, 2002), experimental (Janis, 1982), and theoretical reasons (Lawless and Schwartz, 2002).

For example, the European Union justified its recent switch to majority-rule decision making by noting that consensus-seeking in a political context can hold hostage the solutions to difficult problems of governance (European Union, 2002, p. 29). In other words, the more ill-defined a legislative problem, as the number of participants (here as nation-states) who must

* *Corresponding author address*: William F. Lawless, Paine College, 1235 15th Street, Augusta, GA 30901-3182; e-mail: lawlessw@mail.paine.edu.

forcibly cooperate to achieve a solution agreeable to all increases, the weaker the solution becomes.

At a regional level, Lawless (2004) studied these two decision processes by contrasting citizen groups making decisions to accelerate the U.S. Department of Energy's (DOE's) nuclear waste cleanup at its Savannah River Site (SRS) in the State of South Carolina, where majority-rule decisions were used, with its Hanford Site in the State of Washington, where consensus decisions were used. He found that citizen decisions for SRS were quicker (based on interviews, about 1:4), more specific to site cleanup, and more helpful in accelerating cleanup than similar decisions made for Hanford (specifically, decisions regarding transuranic and high-level radioactive wastes). Unexpectedly, he also found that decisions for SRS were made with less conflict among participants (citizens, scientists, and managers) and were more broadly based than those made for Hanford, which apparently generated more conflict and were more aligned with special interests.

A model of the time to reach a decision, and tentative support for the quantum perturbation model (QPM), is estimated by time $t \approx exp\ (N\ \Delta V)$, where $\Delta V = V(B_0) - V(B_{1\ or\ 2})$ is a potential energy E "barrier" to be overcome, B the choice that represents the attractiveness of either alternative 1 or 2 to neutral agents, and N the minimum number in each group required to reach a decision. Lawless found that the group using consensus (the Hanford Advisory Board in Washington State, with $N = 31 - 4$ as their minimum consensus) struggled to reach environmental cleanup decisions in about 2 hours, giving $\Delta V \approx 0.0257$ units. The contrast came from a second group using majority rule (Savannah River Advisory Board in South Carolina, with $N = 25/2$, rounded to 13), with an average t of 1/2 hour to reach majority-rule decisions, giving $\Delta V \approx -0.053$ units. Thus, the effort expended by the citizens' group at SRS to overcome its potential E barrier for majority-rule decisions was significantly lower than the barrier imposed by the citizens at Hanford with its consensus-rule process.

QUANTUM PERTURBATION MODEL VS. GAME THEORY

On the basis of the above phenomena, we make a strange proposition. The traditional view of cooperative decisions derives a consensus solution to problems, but this cooperation makes use of competition to forcibly squelch dissent and drive sequential *I* transfer. Further, in QPM, cooperation among neutral agents is maximized when agent "operators" compete fiercely to win, but as a consequence, driving the group of neutral states into a special state of cooperation (superposition, characterized by the lowest entropy state possible among agents) to randomly explore the landscape of alternative solutions. These two counterintuitive insights allow us to propose that consensus seeking disguises an underlying competitiveness, while overt competition between two agent operators drives neutrals into an enhanced state of cooperation.

Game theory has never been validated for any social, psychological, or economic phenomenon, including its use as a model of the social interaction, its *raison d'être* (Lawless and Chandrasekara, 2002). Nonetheless, we find many agreeable points of contact between game theory and QPM (Lawless and Chandrasekara, 2002). Our most serious contention with game theory is its identification of forcible cooperation with social welfare. (Both Hardin [1968] and Axelrod [1984] believed that the value of cooperation to social welfare outweighed the need to coerce it.) In our view, social welfare should not be a goal, but rather, the end product of operators driving neutrals who in turn provide feedback to forge a limit cycle that controls both society and operators.

Game theory is based on social learning theory. However, social learning theory, which is based on the individual rational perspective and static in the interaction, has been unable to establish the fundamental shift from a disaggregated collection of individuals to a dyad, group, business organization, political faction, culture, or nation (Allport, 1962; Jones, 1990). This shift is fundamental to "emergence" processes, the existence of which is rejected by traditional modelers (Epstein, 1999, reviewed by Sallach, 2003). And by concentrating on the positive aspects of social learning theory (specifically, to reinforce "cooperation") to avoid the negative imputation of cognitive or social dissonance, but which Kuhn considered essential for the practice of science, computational agent models based on social learning theory have been mostly restricted to reinforcement learning among nearest and next-nearest neighbors in order to reduce communication costs between agents, consequently producing computational agent systems with power too low to solve *idp*s. As Tambe and his colleagues discovered with their computational agent system designed to simply manage the schedules of faculty and graduate students (Pynadath, et al., 2001), current agent models are unable to achieve sufficient autonomy even for the solution of well-defined problems (*wdp*s).

In contrast, QPM brings formal methods to the study of social interaction and perturbation in agent systems across a broad spectrum of social, psychological, and economic phenomena. While game theory usefully introduced an interdependence between the choices participants are forced to make in a given game configuration, there has been no theoretical justification offered by game theorists for its static *independence* between action and the observational uncertainty involved in these choices. By comparison, our QPM begins with the *interdependence* between action uncertainty Δa and observational uncertainty ΔI to link within our model the uncertainties that occur naturally in an interaction (from Bohr, producing $\Delta a \, \Delta I \approx c \approx \Delta E \, \Delta t$; Lawless, et al., 2000). In addition, as humans manage interaction uncertainty, feedback cycles arise with outcomes that at best can be roughly predicted, initiating a limit cycle. Instead of the narrow feedback from forced choices between cooperation and competition, we have focused on how initial and subsequent decisions generate a limit cycle. Following the lead of conflict theorists (e.g., Simmel, 1964), with this model we were able to establish mathematically the problems created by cooperation (e.g., the corrupting influence of hidden I; Lawless and Schwartz, 2002), still not a consideration with traditional models (e.g., Wright, 2000). With our mathematical model, we have also studied organizational growth; business mergers during economic instability that resemble ant and slime-mold mergers during environmental instability, suggesting a scale-free model (Lawless, 2003); terrorism (Lawless and Chandrasekara, 2002); social responses to environmental disasters; and recently, with coupled Kolmogorov nonlinear equations, the wax and wane of knowledge K (e.g., expectations, predictions, beliefs, and algorithms; Lawless and Grayson, 2004).

QPM FOR CONTROL OF MULTI-AGENT SYSTEMS

Our quantum model is not meant to copy reality. In the sense that the atom constructed by quantum physicists is an abstraction that permits exact predictions to be calculated and validated, it matters less that the QPM we propose matches social reality than that it leads to new discoveries that can be validated, such as the control of multi-agent systems (MASs) (for reviews of model validation and social phenomena, see Carley [2002] and Bankes [2002]).

Generalizing our earlier conclusion that social debate among discussion leaders defending orthogonal positions produces superior decisions (e.g., in science, Bohr versus Einstein on quantum theory; in the courtroom, a defense attorney versus a prosecutor; and in business, the recently settled web-browser wars between America Online and Microsoft), we propose that a large computational parallelization derived from entangling *N* agents with pro-con beliefs simultaneously superposed during debate is a condition sufficient to control an MAS to resolve an *idp* and achieve a solution. (See Zlot, et al. [2002] for an example of the superposition of multiple robot interpretations employed to construct a single map to navigate the environment.)

In the traditional view of human or agent computations, computations can occur in parallel. Traditionally, *I* is shared sequentially among agents or broadcasted from a central command point, both slowing the computational process. In this view, evolution occurs (e.g., genetic algorithms) from the random transfer of *I*, generally by agent reproduction and within the constraints of a well defined problem (*wdp*). In contrast, Feynman (1996) showed that quantum-mechanical states could evolve from the action of operators. Then Deutsch (1989) showed that a superposition of quantum states could be explored simultaneously, producing parallel computations more powerful than digital ones.

Digital logic states can be either $| 0 >$ or $| 1 >$, with each known as a "bit." Digital logic gates can sequentially transform single bit inputs to output values. For the most part, digital computers can solve the same problems proposed for quantum computers. However, an exponential increase in classical information processing requires an exponential increase in the number of digital computers and physical space (i.e., $n \times n = n^2$). In contrast, quantum information processing logic states can be in a linear combination of ground $| 0 >$ and excited states $| 1 >$, known as a "qubit," with each qubit producing 2^1 values. Thus, with *N* agents in a superposition of 2^N states, an exponential increase in quantum computation occurs with only a linear increase in agents and physical space. These simultaneous states are a superposition, allowable for quantum information processing but without analog in classical digital computing. This difference has important implications for the time it takes to complete a computation, the relative and absolute power of quantum parallelization, the physical space occupied by the computer processors, and the heat generated by the respective processors. Additionally, quantum information processing opens new aspects of nondeterministic random executions of computations, producing a direct link with the Fourier elements in the biological computations proposed by May (2001) — efficient (fast) digital algorithms exist for addition, for example, but not for factoring, where an efficient algorithm has been discovered for quantum processing (i.e., Shor's fast quantum Fourier factoring algorithm).

Three traditional computational agents exist in three of eight possible states. To explore randomly these eight states, a traditional parallel model of an MAS would probe each possible agent group state sequentially. In contrast, three human agents (or three quantum qubits) exist in 2^3 or eight states simultaneously, and can be probed at once. Further, traditional agents do not rely on emotion per se. But human and quantum agents exist in ground and excited states. To extend this analogy to a quantum agent model of an organization (Lawless and Grayson, 2004), we propose that aspects of a quantum agent organization or system are in a superposition of ground and excited states. In our view of an MAS, at least two organizations are operators that drive the remainder of the system as superposed neutrals across a fitness landscape.

Many ways are possible for quantum information processing to attain an answer described by the amplitude of a quantum state (complex numbers that mostly cancel each other). Quantum information processing is efficient when only the correct answer survives with high probability and wrong answers cancel (Berman, et al., 1998, p. 21).

Compared with traditional computational states where one bit represents either 0 or 1, a single qubit is in a superposition of 0 and 1, with a register of N qubits being in a superposition of all 2^N possible values. Briefly, with $|\uparrow>$ representing the "pro" ground-state proposition and $|\downarrow>$ the "con" excited-state refutation, a single basis state can be represented as $|\Psi> = a|\downarrow> + b|\uparrow>$, where a and b are complex numbers such that $|a|^2 + |b|^2 = 1$. Parallelization for two agent qubits, each with the same single basis state, can be symbolized as $\{|\downarrow\downarrow>, |\downarrow\uparrow>, |\uparrow\downarrow>, |\uparrow\uparrow>\}$. This simple example illustrates the exponential growth of the state space with a linear increase in the number of agent participants. Superposed or entangled states have no classical analog and cannot be decomposed (supporting the notion of "emergence"); however, in that social dissonance is characterized by small numbers of strongly held but polar opposite positions witnessed by a larger group of mostly neutral participants, we characterize the strongly held positions as traditional social forces — Feynman's operators — driving the neutral agents to a solution, but with neutrals reflecting Deutsch's notion as the register of superposed states that are being driven to randomly explore the space of alternative solutions (Lawless and Schwartz, 2002). (For indirect support of our position regarding the general existence of neutrals among humans, the review by Eagly and Chaiken [1993] concluded that surveys can be worded to obtain almost any desired result; similarly, Tversky and his colleagues observed that the correlation between decisions and their subsequent justifications is negligible [in Shafir, et al., 1993]. Most experimental subjects have weak connections between their beliefs and the actions they enact, but in general, this is not true for the experts or operators who drive a system or debate for the benefit of neutrals, as in courtroom attorneys or scientists like Bohr and Einstein.)

Control extends the decision-making process: Predictions from the results of decisions produce an outcome with an error component, a larger error occurring during times of economic expansion or environmental stability (i.e., stable environments promote competition and, as a consequence, volatility and social evolution). Consequently, feedback about discrepancies has an effect on earlier decisions by causing another decision to be made to reduce the error or discrepancy, establishing an iterative process, the end result being a limit cycle to regulate or control a system (May, 2001). Three predictions from this model contrast with more traditional computational approaches: The optimum limit cycle occurs with an increase in the number of participants in the decision-making process (the best fit from increasing the number of "neutral" participants or Fourier components occurs once a solution is reached); attacking a group at a rate faster than its natural feedback response rate will produce panic (see Figure 1); and if the number of fluctuations across a social system is constant, when the overall system environment is stable the community becomes easier to control (a larger community matrix eigenvalue, representing a quicker return to stability) even as an increase in competition between constituent groups produces more unstable groups (and social evolution), increasing the innovation rate (where technology reduces the size of environmental fluctuations and gives a competitive edge to a group; see Ambrose, 2001).

For example, assuming in 2003 that the recession and aftermath of the al-Qaeda attacks in 2001 made 2002 more unstable than 2003, a recent PricewaterhouseCoopers survey (online.wsj.com) compared the change from 2002 to 2003 in risk-taking among corporate CEOs.

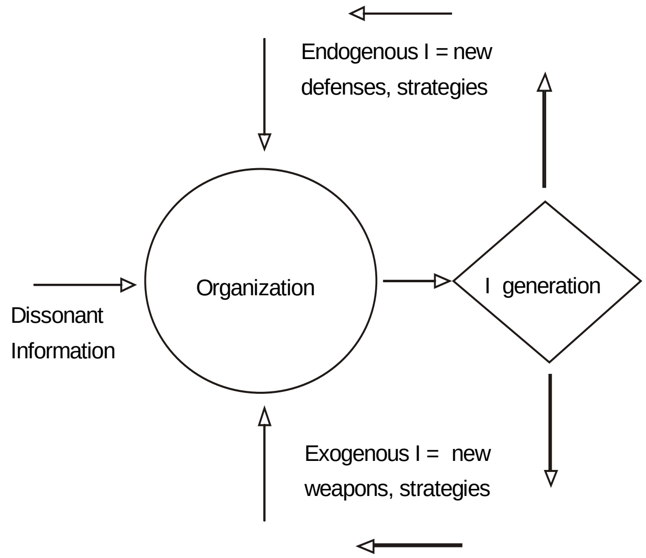

FIGURE 1 Quantum Perturbation Model for Organizations

The survey found that in 2003 48% of CEOs reported being more aggressive, 31% reported no change, and 20% reported being less aggressive. Thus, during stable times, organizations and individuals take on more *idp*s than in unstable times, when they consolidate and retreat to *wdp*s.

QPM for Organizations

In Figure 1, if we assume that an organization exists in a ground, excited, or combined state, then perturbations provide invaluable I about the structure and competitiveness of an organization (from Lewin, 1951), with each "measurement" limited by an uncertainty in action (Δa) and observation (ΔI). After a perturbation (from Lawless and Grayson, 2004), an organization's goal is to respond with endogenous feedback to dissonant I by creating new knowledge K to design new innovations, strategies, or technologies to defend the organization. (In general, K arises when $\Delta I \rightarrow 0$; here, $K_{new} = K_{alg} + K_{\chi}$, where K_{alg} is algorithmic knowledge and K_{χ}, is interaction knowledge, such as beliefs or expectations.) Conversely, using exogenous feedback, a competitor's goal is to devise innovations, strategies, or technologies to defeat the organization. In general, the quicker respondent determines which organization wins and evolves (in 2003 in the war with Iraq, for example). Coalition decision-making and implementation of those decisions occurred faster than that of Iraq's Defense Forces, causing the latter to panic and its organization to dissolve (i.e., in engineering control theory, late feedback is destabilizing; May, 2001, p. 5).

The Quantum Perturbation Model of Organizations offers a ready explanation for the premium on deceiving or bluffing opponents into thinking that intentions for an action or strategy may or may not be the one implemented, the tendency for terrorists to cooperate to hide their intentions (Lawless and Chandrasekara, 2002), and the greater ability of democracy to uncover hidden intentions, thereby reducing corruption in comparison to consensus or command governments (Lawless and Schwartz, 2002). This model also accounts for some of the underlying forces in mergers: Should an organization's execution of technology falter, producing weak operational results (e.g., AT&T Wireless's difficulties with enacting phone number portability in early 2004), the organization becomes the prey or acquired organization instead of the predator. In this model, power accrues to the winning organization and its chief strategist.

In general, from mathematical control theory (May, 2001), the stable control of an MAS should occur when its responses to error are faster than its natural response rate, which is much more likely under majority-rule decision-making than under consensus decision-making (Lam and Suen, 1997; Lawless and Schwartz, 2002). The quantum agent approach should assist in achieving optimum control by regulating the system to seek the best fit between a problem and its solution as neutral participants are added to the decision-making process (i.e., Fourier components) until a solution is found.

WORK IN PROGRESS

We are considering two possible approaches to operationalize, test, and explore the QPM of organizations shown in Figure 1: agent model and modified Markovian.

First, we consider the effect of uncertainty on decisions made in dynamic social structures. Axelrod and Cohen (2002) recognized that strategy space in the prisoner's dilemma game is stochastic, and not noiseless or error-free, leading them to incorporate noise into their experiments. Similarly, we recognize that the nature of our model depends on I that is not noise-free, but is subject to errors in receiving and processing.

The benefit of entangling multiple independent actors is that collectively, they can reduce the noise (or variance) of the central tendency of the information. This is analogous to the beneficial effect of obtaining multiple independent samples in estimating the mean of an unknown population. The variance estimate of the unknown mean is reduced by a factor of $1/(\sqrt{n})$ as the number of independent assessments is increased, giving a sampling distribution of the mean with variance/n. Since neutral actors are independent but entangled in their assessments, there is an analogous benefit of more actors (N) reducing the variance of the unknown "mean" of the reality of an *idp*. We expect to find that entangled agents will be in a lower state of variance or entropy than correlated agents, who in turn have lower variance than independent actors in games.

We will use an agent model to test our theory. Initially, we will validate our model against the analytic solutions of quantum game theory by Arfi (2003). Specifically, in the "Battle of the Sexes" game, analytic game theory solutions should be less correlated than quantum game theory solutions. Afterward, we will model and analyze MAS organizations and mergers to establish the costs and extent of interactions necessary between agent systems and humans to facilitate autonomy.

Second, we propose to test the concepts in Figure 1 with an MAS model of the competition between an agent model of GM, Toyota, and all other car manufacturers (as a group). We assume that the primary focus is from the perspective of GM. Then:

1. GM observes its current state and its context. Estimate the percentage of buyers who choose GM cars, Toyota cars, or other cars.

2. GM proposes a plan to capture more of the market. (GM's plan will contend with internal factors such as structure, personnel, talent, and costs, as well as external factors such as competitiveness, markets, price, and technology leadership.) Estimate the Markov transition matrix. GM implements its plan. Estimate the percentage of buyers who stay with GM or switch to Toyota or others.

3. Obtain feedback. Update the transition matrix.

4. Build a binomial tree: Determine the probability p of having correctly estimated the transition matrix and probability $(1–p)$ of the current state continuing unabated. (GM does not know for sure whether its plan will work, but perhaps with an estimate of p the company's plan will work, meaning a target progression of customers to GM according to the transition matrix.)

5. Some probability exists on the number of stages that the company progresses through on the way to a steady state until a counterattack by opponents. (Even if GM's plan is successful, some number of steps/stages/time periods occur before the market reaches steady state. Toyota will not wait, but will counterattack as soon as it can, possibly before the market reaches a new steady state. If so, then the new market status will either be the ground state — if there was sufficient time — or some intermediate excited state, depending on the number of stages completed before a counterattack occurs.) [Note: GM and Toyota are operators, driving the car-buying public of neutrals to randomly explore the alternative space of solutions by determining the optimum choice of, in this case, a car.]

6. At the point of counterattack the progress stage of percentages becomes the new current state.

7. Simulate Steps 1 through 6 to estimate the probability of success in reaching a minimally acceptable new state for GM.

ACKNOWLEDGMENTS

The lead author thanks J.A. Ballas Information Technology Division, Naval Research Laboratory, Washington, D.C., where most of this research was conducted with funds from the Office of Naval Research through an American Society for Engineering Education grant.

REFERENCES

Allport, F.H., 1962, “A Structuronomic Conception of Behavior: Individual and Collective,” *Journal of Abnormal and Social Psychology* **64**:3–30.

Ambrose, S.H., 2001, “Paleolithic Technology and Human Evolution,” *Science* **291**:1748–1753.

Arfi, B., 2003, “Ambivalence of Choice and Holistic Strategic Interaction: A Quantum Game-theoretic Proposal,” University of Illinois-Urbana Champaign, Unpublished manuscript.

Axelrod, R., 1984, *The Evolution of Cooperation*. New York: Basic.

Axelrod, R., and M.D. Cohen, 2002, *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, New York: Simon & Schuster.

Bankes, S.C., 2002, “Perspective. Agent-based Modeling,” *Proceedings National Academy of Sciences* **99**(3):7199–7200.

Berman, G.P., G.D. Doolen, R. Mainieri, and V.I. Tsifrinovich, 1998, *Introduction to Quantum Computers*, River Edge, NJ: World Scientific Publishing Co.

Carley, K.M., 2002, “Simulating Society: The Tension Between Transparency and Veridicality,” in *Social Agents: Ecology, Exchange, and Evolution*, Argonne, IL: University of Chicago and Argonne National Laboratory.

Cohen, M.D., R.L. Riolo, and R. Axelrod, 1999, “The Emergence of Social Organization in the Prisoner’s Dilemma,” Santa Fe Institute Working Papers (www.santafe.edu/sfi/publications/99wplist.html), NM: Santa Fe Institute.

Deutsch, D., 1989, *Proceedings of the Royal Society* **A425**:73–90.

Eagly, A.H., and S. Chaiken, 1993, *Psychology of Attitudes*, Fort Worth: Harcourt.

Epstein, J.M., 1999, “Agent-based Computational Models and Generative Social Science,” *Complexity* **4**(5):41–60.

European Union, 2002, *European Governance* (COM, 2001, 428 final), white paper, Brussels: European Union.

Feynman, R., 1996, *Feynman Lectures on Computation*, New York: Addison-Wesley.

Hardin, G., 1968, “The Tragedy of the Commons,” *Science* **162**:1243–1248.

Janis, I.L., 1982, *Groupthink*, Boston: Houghton Mifflin.

Jones, E.E., 1990, *Interpersonal Perception*, New York: Freeman.

Kuhn, T., 1977, *The Essential Tension*. Chicago: University of Chicago Press.

Lam, L., and C.Y. Suen, 1997, "Application of Majority Voting to Pattern Recognition," *IEEE Transactions Systems, Man, and Cybernetics* **27**(5):553–568.

Lawless, W.F., 2003, "Information Density Functional Theory (IDFT): A Quantum Approach to Knowledge Fusion and Organizational Mergers," in *AAAI Spring Symposium on Agent-Mediated Knowledge Management*, Stanford, CA: AAAI, pp. 71–73.

Lawless, W.F., 2004, "DOE's Policy of Consensus Seeking Versus Truth Seeking in Traditional American Democracy," in *Proceedings of the Waste Management 2004 Symposium*, Tucson, AZ, February 29–March 4.

Lawless, W.F., T. Castelao, and C.P. Abubucker, 2000, "Conflict as a Heuristic in the Development of an Interaction Mechanics," in C. Tessier, L. Chaudron, and H.J. Muller, eds., *Conflicting Agents: Conflict Management in Multi-agent Systems*, pp. 279–302, Boston: Kluwer.

Lawless, W.F., and R. Chandrasekara, 2002, "Information Density Functional Theory: A Quantum Approach to Intent," in *Proceedings AAAI Fall Symposium on Intent Inferencing*, pp. 26–31, Cape Cod, MA: AAAI Press.

Lawless, W.F., and M. Schwartz, 2002, "The Social Quantum Model of Dissonance. From Social Organization to Cultural Evolution," *Social Science Computer Review (Sage)* **20**(4):441–450.

Lawless, W.F., and J.M. Grayson, 2004, "A Quantum Perturbation Model (QPM) of Knowledge Fusion and Organizational Mergers," in L. van Elst, V. Dignum, and A. Abecker, eds., *Agent Mediated Knowledge Management*, pp. 143–161, New York: Springer-Verlag.

Lewin, K., 1951, *Field Theory in Social Science: Selected Theoretical Papers*, D. Cartwright, ed., New York: Harper Torchbooks.

Luce, R.D., and H. Raiffa, 1967, *Games and Decision*, New York: Wiley.

May, R.M., 1973/2001, *Stability and Complexity in Model Ecosystems*, Princeton, NJ: Princeton University Press.

Nowak, M.A., K.M. Page, and K. Sigmund, 2000, "Fairness versus Reason in the Ultimate Game," *Science* **289**:1773–1775.

Pynadath, D., P. Scerri, and M. Tambe, 2001, "MDPs for Adjustable Autonomy in Real-world Multi-agent Environments," in *2001 AAAI Spring Symposium*, Technical Report SS-10-03, Stanford, CA: AAAI.

Sallach, D.L., 2003, Panel presentation: "The simulation of Emergence: Mysticism, Reduction, or Social Complexity?," NAACSOS-03, Pittsburgh, PA: Carnegie Mellon University.

Shafir, E., I. Simonson, and A. Tversky, 1993, *Cognition* **49**:11–36.

Simmel, G., 1964, *Conflict and the Web of Group Affiliations*, New York: The Free Press.

Stroebe, W., and M. Diehl, 1994, “Why Groups are Less Effective than Members,” *European Reviews of Social Psychology* **5**:271–303.

Von Neumann, J., 1961, “The Mathematician,” in *Collected Works*, Pergamon Press.

Wendt, A., 1999, *Social Theory of International Politics*, Cambridge University Press.

Worchel, S., 1999, *Written in Blood. Ethnic Identity and the Struggle for Human Harmony*. New York: Worth Publishing.

Wright, R., 2000, *Nonzero. The Logic of Human Destiny*, New York: Pantheon Books.

Zlot, R., A. Stentz, M.B. Dias, and S. Thayer, 2002, *Market-driven Multi-robot Exploration* (CMU-RI-TR-02-02), Pittsburgh, PA: Carnegie Mellon University.

APPLICATION OF AGENT-BASED SIMULATION TO POLICY APPRAISAL IN THE CRIMINAL JUSTICE SYSTEM IN ENGLAND AND WALES

SEÁN BOYLE,* London School of Economics
STEPHEN GUERIN, Redfish Group, Santa Fe, NM
JULIAN PRATT, London School of Economics
DANIEL KUNKLE, Redfish Group, Santa Fe, NM

ABSTRACT

This paper describes an agent-based approach for constructing a model of criminal justice system operations in England and Wales. The primary purpose of the model is to assess the impact of policy variants across the entire criminal justice system. Because of the structure of this system, three separate government departments interact and deliver services. Decisions in one area of the criminal justice system can be crucial in determining what happens in another area. Our purpose was twofold. First, we needed to contribute to the Treasury's spending review by working with different groups in criminal justice agencies to reach a consensus on how things actually occur (i.e., linking behavior and actions of one group with another and with resources). Second, we needed to produce a model of the entire criminal justice system that would provide insights into questions related to capacity, case flow, and costs. We also needed to model the ways in which individuals go through the system. The result is a hybrid model that combines a simple system dynamics approach with an agent-based model. The distinctive approach used in this work integrated modeling with practical ways of enabling people to engage in strategic policymaking, while taking into account the complexities of the criminal justice system. The agent-based framework developed to meet these needs models the criminal justice system, provides the ability to assess policy across the system, and allows sharing of model output to improve cooperative efforts among departments.

Keywords: Agent-based modeling, criminal justice system, visualization, policy appraisal simulation

1 INTRODUCTION

This paper reports on an agent-based approach for constructing a model that shows the operations of the criminal justice system in England and Wales. The primary purpose of the model is to be able to assess the impact of policy variants across the entire justice system.

Because the model is designed to help people to think about what happens when things are changed in a deliberative manner, we provide some examples of policy changes for which the model is designed to provide help. We also discuss a visualization that represents what the model can do for different policy views. With a view of "model as icon for change," the reality of the visualization is not as important as how it looks.

* *Corresponding author address:* Seán Boyle, London School of Economics and Political Science, LSE Health and Social Care, Cowdray House, Houghton Street, London, U.K. WC2A 2AE; e-mail: s.boyle@lse.ac.uk.

Section 2 describes the context for the project — the structure of criminal justice in England and Wales. Section 3 discusses the purpose of the project, which goes beyond the mere construction of a model. Construction of the model included at least two aspects of interest: the way in which the problem was approached and the physical representation of some kind of solution, which we define as the model. These aspects are discussed in Sections 4 and 5. Section 6 provides concluding remarks.

2 CONTEXT

The criminal justice system in England and Wales is delivered by diverse government bodies; the same is true in many other countries. In England and Wales, these are not part of a single government department. Three departments are involved: the Home Office, which is by far the biggest financially and in terms of human resources; the Department of Constitutional Affairs; and the Crown Prosecution Service. Each of these departments has its own government minister, and in the case of the first two, has a range of responsibilities outside those considered in constructing a model of the criminal justice system. Thus, the Home Office is also responsible for immigration and for homeland security, whereas the Department of Constitutional Affairs also has the responsibility for civil and family law.

The Home Office criminal justice responsibilities include the Police Service, the Prison Service, and the Probation Service, but it is not a direct operational responsibility. Other agencies are responsible for delivering each service. Little direct financial accountability occurs (although all rely on central government funds), and there is only limited operational interference. Top-level targets are set for each service, but the utility of these is uncertain. Because operational control is divided across 42 areas of the country, determining what happens is a local matter.

The Department of Constitutional Affairs is responsible for both the courts and, via an executive agency, the provision of free criminal defense services (known as Legal Aid). The courts are divided between lower and higher courts: the former are called magistrates' courts and deal with lesser offenses; the latter are called the Crown Court and generally deal with more serious cases.

The Crown Prosecution Service is responsible for prosecuting criminal cases. It is the least complex of these bodies.

The functionality of the criminal justice system depends crucially on the way in which each of these bodies delivers services and on the interactions among them. Each part of the system has thousands of individual agents who act according to sets of rules. Some rules are fairly prescriptive, and others are rules of thumb, often undescribed.

Most of the funding for these service providers comes through the U.K. Treasury. Some other money flows through either local government sources or are private funds. For every government department, the U.K. Treasury has a system of spending reviews; these take place every two years and look three years ahead, therefore overlapping by one year.

Decisions in one area of the criminal justice system can be crucial in determining what happens in another. For example, how well the police functions may make the life of courts easier or more difficult, the workload of prisons more or less. This has been recognized by the

Treasury. Thus, in the 1998 spending review, the government undertook the first-ever review of the performance and management of the criminal justice system as a whole, cutting across all three government departments.

The 2002 spending review saw a cross-departmental review of the criminal justice system, which built on the work begun in 1998. However, the Treasury did not feel that the collective criminal justice system elements presented were sufficiently "joined up." Thus, for the 2004 spending review, the Treasury has required further development of the way in which all agencies bid, so that bids take into account what the other agencies are doing. The Treasury also requires that the bidding process be mediated through a model of the entire system. Our work is designed to address this need.

3 PURPOSE

Our primary task was to do something that would contribute successfully to the Treasury's spending review for 2004, and, beyond this, that could be used for assessment of future policy development across the whole of the criminal justice system. To achieve this goal, we worked at two levels. First, to establish a consensus about how things actually happen in the system, we worked with groups of people from different agencies in the criminal justice system. We gathered evidence of links between the behavior and actions of one person or group of people and another and through this process made arguments for the best use of resources. This new process was essentially about encouraging a change in the style of working of these core government agencies (see Pratt, et al. [1999] for a discussion of some ways of achieving this).

Second, we needed to produce a model of the entire criminal justice system that all actors in the system would recognize. To achieve this task, we worked with modelers and statisticians in the various government agencies and departments who were technical people interested in building better models. We acknowledge the extent of the contribution of the Criminal Justice Performance Directorate in this respect as well as various individuals in each of the departments and agencies of the criminal justice system. Our aim was to build on existing models of the system to produce an end-to-end computer model of the criminal justice system to provide insights, in particular into questions of capacity, case flow, and costs. This has the feel of a standard modeling problem, although ours was not a standard solution.

We needed to model how individuals — criminals or cases — move through the criminal justice system from the initial crime event to final disposition, culminating in receiving either a prison or a community sentence (including various forms of post-prison supervision), or in being released as a free member of the population. Moreover, these flows needed to be mapped against costed resources to meet Treasury requirements.

There was a third level of approach that we were only able to engage in tangentially. This involved the people doing the job, who are in fact those represented as agents in our model.

4 CONSTRUCTION OF THE MODEL

The project had two distinctive parts: (1) working with the people involved in making and delivering policy in the criminal justice system and (2) developing an adequate model of what the system does.

Working with people involved a range of activities:

- Determining user requirements through individual interviews and workshops, which culminated in the production of a user requirements report;

- Developing ways to assure the client that the model was really “them,” again through interviews and workshops, and culminating in a test suites report; and

- Recording what the system does and why in terms of processes, activities, and resources, again through interviews and workshops, and resulting in the production of what was called a modeled processes report.

Each of these activities was also of fundamental importance in delivering a successful model, which comprised the second part of our task. The model developed was based on agent behaviors.

4.1 Inputs

To provide inputs to the model, we asked each agency to consider the following types of questions:

- What resources are used in providing services (i.e., what police and types, courts, custody suites, etc.)?

- What does each resource do, how does it makes choices, and are there different rules that can be selected in making those choices?

- What happens when capacity limits are threatened; how does prioritization take place?

- What are the costs of each resource, and how does this vary as decisions are made?

4.2 Outputs

The model represents the flow of activity through the criminal justice system, which can be analyzed in terms of, for example:

- Number of crimes reported,

- Number of cases tried in magistrates’ courts,

- Cost of various types of resources, and
- Numbers waiting at different points in the system.

Each of these activities can be viewed at the minimal level of disaggregation (i.e., one agent doing one activity in one time slot), but each also can be aggregated over time, people, and activities to any required level.

5 AGENT-BASED HYBRID MODEL

We set out to produce a model that would engage people in the system. To meet this need, we adopted an agent-based approach. In the time allotted, however, it would not be possible to build an agent-based model for every part of the criminal justice system. The key question was, could we produce a model that would satisfy the needs of the client, while at the same time, take the client-system down the agent-based road by providing a model that the client could readily build on, and most important, would want to build on.

The result is a model that is a hybrid. It combines a simple system-dynamics model of flows through the criminal justice system — albeit with relatively complex interactions at each stage or node — with an agent-based model of individual agents that behave in ways that produce results that cannot be predicted from looking at the behavior of groups of the same agents.

Figure 1 represents the hybrid model concept. In some parts of the system, our model is more like process descriptions with high levels of agent homogeneity (superagents). In other parts, we have good descriptions of activities of individual agents with significant interaction among agents. Ultimately, the process and activity descriptions are mutually consistent (Bonabeau, 2002).

The model is structured in a way that allows the user to examine simple questions or more complex policy issues. "Simple" questions, however, are often only simple because the more complex issues they imply are ignored in that instance. Some typical policy issues are described in Box 1.

Increase the number of police by 10,000 (currently 130,000). Determine the impact on the system.

The impact depends on what activities the new police choose to do — more patrolling, more investigation, better case preparation, etc. All of these activities will have effects down the line for other service providers, and all will also affect how the agents themselves work.

Increase sentencing powers, for example, from 6 to 12 months, for certain offenses.

It may seem obvious that this policy will increase the prison population, but sentencers may choose to use the power differently. Moreover, defendants may react to longer sentences by appealing more or choosing a different court for the hearing. Any of these may result in different consequences from those that might have been supposed when the policy was first devised.

Box 1 Typical policy issues

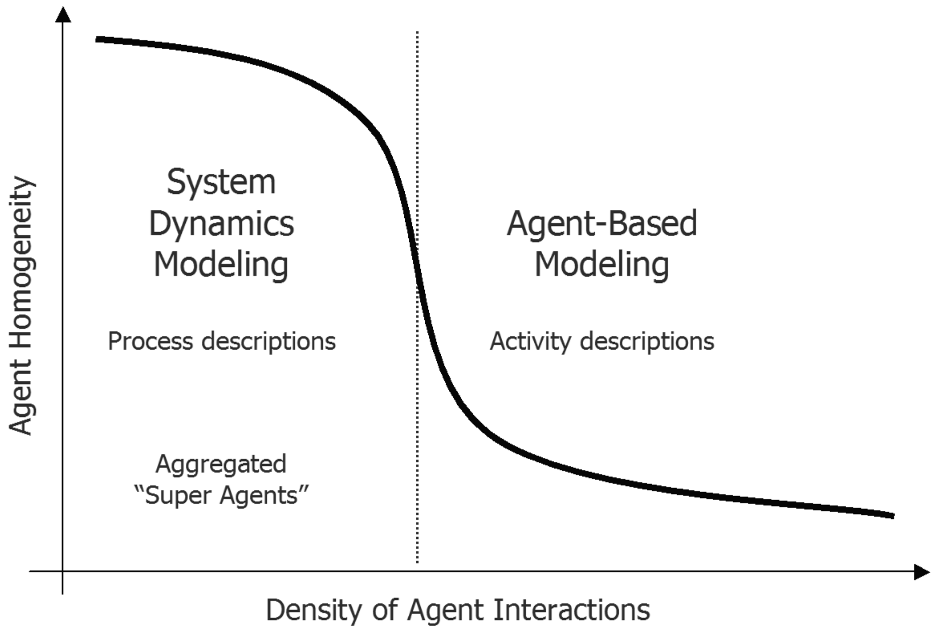

FIGURE 1 Hybrid modeling: process and activity-based descriptions

6 CONCLUDING REMARKS

At the end of the project, we have delivered an agent-based framework with the potential to model the impact of government policy on the criminal justice system. In addition to developing the model, we delivered the free-standing policy "tools" listed below. Each tool enables a practical application of systemwide thinking.

Thus, as part of the creation of the model we:

- Set up and facilitated a group known as the Spending Review 2004 Group, which helped to give ownership across government departments.
- Supported and developed the role of the Project Steering Group, which spanned agencies across the whole system.
- Developed a template for systemwide policy formulation called the Systemic Impact Statement.
- Offered a high-impact demonstration of flows across the system through computer visualization developed with the model, which is especially useful for nontechnical policy people. Box 2 provides further discussion of our use of visualization.

Another key aspect for the client was that all government departments and constituent agencies were signed up to the outcomes of the project. We interpreted this as meaning that each player had to be involved in the development of the model, and the method of working. This was certainly achieved.

7 REFERENCES

Bonabeau, E., 2002, “Agent-based Modeling: Methods and Techniques for Simulating Human Systems,” *Proceedings of the National Academy of Sciences of the United States of America* **99** Suppl. 3, 7280–7287.

Pratt J., P. Gordon, and D. Plamping, 1999, *Working Whole Systems*, London: Kings Fund.

We felt it was important to provide a visualization of the system that a wide range of users could relate to — to reach beyond those with a technical interest in the model to those who determine policy, such as high-level public servants and politicians.

Our use of visualization also allows the different service providers to see themselves as integral parts of a large whole. In a way, the visualization comes to represent the model as icon: it is almost as if people have something that they can touch while making their decisions.

This diagram shows a screen capture from the visualization. Our aim is for users to become more aware of the system and its parts, at the same time they see the size of flows along edges between nodes (e.g., the proportion of capacity used, timeliness between two nodes, or costs of providing services at each node.

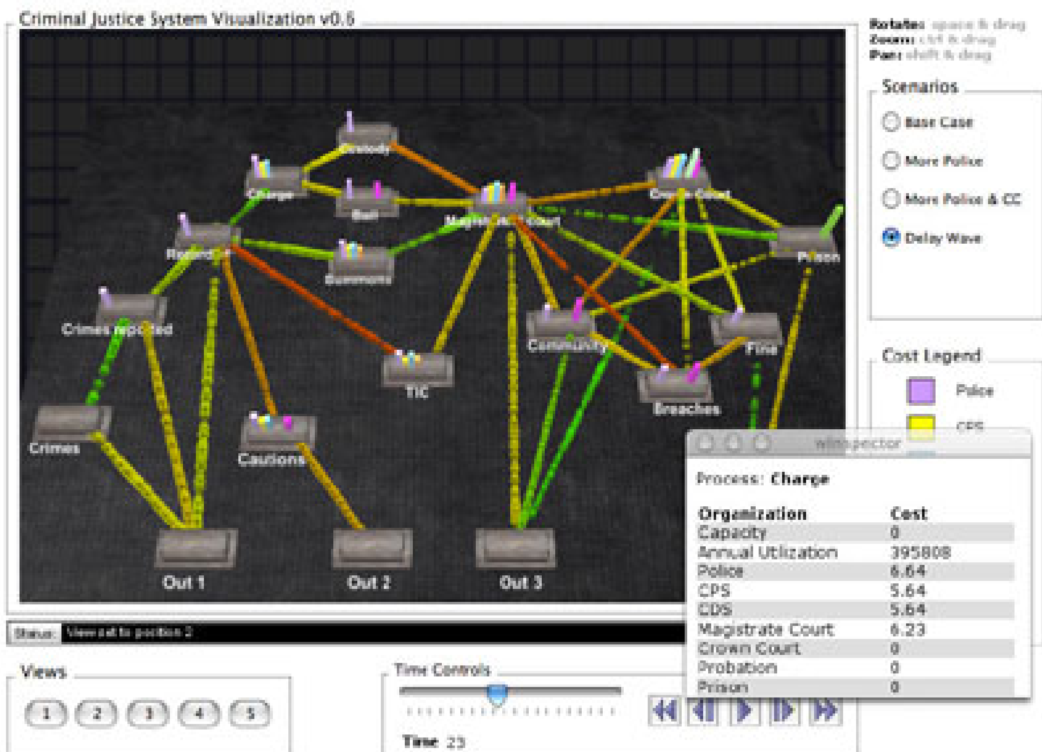

The visualization is decoupled from the model where the visualization reads the log files produced by the model. This approach allows us to easily switch between different scenarios produced by multiple-scenario runs. A second benefit is that it allows us to do early rapid prototyping to establish the scope of the project while the model is being constructed. We are able to use the same visualization for outputs of “scratch-pad” throw-away prototypes in various programming languages, then plug-in the actual model data when available. A third benefit to this approach, which cannot be overstressed, is the ability to more rapidly diffuse the model and its insights throughout the organization (because visualization with log files involves a much smaller memory footprint than deploying the model and all of its dependencies). The above benefits notwithstanding, one disadvantage of having a decoupled view is the inability to modify model parameters on the fly for interactive exploration by the user.

Box 2 Visualizing the system

AGENT-BASED SUPERVISION AND CONTROL OF COMPETITORS IN A HETEROGENEOUS ENVIRONMENT

E. TATARA, I. BIROL, F. TEYMOUR and A. CINAR,*
Illinois Institute of Technology, Chicago, IL

ABSTRACT

Much of the research in the area of agent-based modeling focuses on replicating observed behavior of the system of interest. The purpose of this paper is to illustrate how a multi-layered agent-based supervisory control system can interact and influence a physical model that is based on first-principles, an agent framework, or a combination of the two. A knowledge-based control strategy is implemented based on what is known about the behavior of the system. The command and control structure resembles that of a social organization. Local command agents determine the most appropriate course of action and the subordinate control agents execute the commands. Autonomy allows each level to act independently of others to some degree. Agents at the lowest level of the command structure can identify which techniques work best for solving different types of problems. The control agents must therefore adapt through trial and error. The degree of autonomy granted to individuals permits the emergence of highly complex behavior that cannot be anticipated. The primary focus of analyzing complex emergent behavior is demonstrating methods or combinations of methods that influence the behavior and capabilities of the agent-based control system. Simulation studies will illustrate scenarios of interest, especially those conditions that lead to emergent behavior. The behavior of competing autocatalytic chemical species in a continuous stirred tank reactor (CSTR) has been used as a model for more complex phenomena such as competing biological populations. This framework will be used to illustrate the structures and tools described.

Keywords: control of distributed systems, intelligent supervision, industrial process control

INTRODUCTION

Much of the research in the area of agent-based modeling focuses on replicating observed behavior of the system of interest. Often, the agents are designed to operate autonomously of each other and free from any type of external supervisory control structure. While software agents can be used to simulate physical systems, they can also be used to simulate hierarchical regulatory structures, such as the management of natural resources, where the fundamental equations governing the ecosystem may be known, but the overlying control mechanisms (human management) in place are far too complex to simulate with traditional methods. Additionally, agent-based control systems can be used to monitor and control real-world physical systems such as computer networks or industrial production facilities (Monostori and Kadar 1999; Jennings 2003).

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

Multi-agent control systems have several properties that make them particularly attractive for use with large, complex systems (Lesser 1999). The first, and usually most important in critical systems, is a high level of reliability. Modularity and scalability also play an important role in multi-agent systems. Agent-based monitoring and control systems often produce several different solutions to the same problem. *Solution multiplicity* arises when several agents, using completely independent methods, arrive at different conclusions based on the presented data. Negotiation between agents is therefore required to resolve the situation. Additional agent layers may be used to evaluate information and conclusions produced by diagnosis agents. Criteria such as performance history are used to determine the validity of agent diagnosis when solution conflicts arise.

The behavior of continuous stirred tank reactors (CSTRs) has been studied extensively over the past several decades. Fascinating static and dynamic phenomena have been observed for many classes of reactions including autothermal reactions, autocatalytic reactions, and polymerization reactions (Uppal and Ray 1974; Bilous and Amundson 1955; Uppal et al. 1976; Farr and Aris 1986; Lin 1981; Gray and Scott, 1983, 1984). Birol and Teymour (2000) have studied isothermal autocatalysis when two competing chemical species are introduced to the system. It was shown that with N species, although there may be as many as $2(2^{N-1})$ steady states, only one species can exist stably in the reactor as $t \to \infty$.

If several identical CSTRs are connected, such that material is exchanged between them, the system becomes spatially heterogeneous, compared with a single larger CSTR. The degree of heterogeneity can be varied by manipulating the interconnection flow rates between reactors. Taylor and Kevrekidis (1993) have studied the effects of reactor coupling. Oscillatory states tend to synchronize as long as the frequency of the oscillations in each reactor are not too different. Weak coupling of the reactors is usually employed, since strong coupling of a reactor network causes the network to behave as one large reactor.

Recent work on two- and three-reactor configurations with cubic autocatalytic reactions has demonstrated that spatial heterogeneity enlarges the boundaries of chemical species survival (Birol et al. 2002). Furthermore, detailed analysis has shown that networks of reactors with autocatalytic replicators produce highly complex bifurcation structures and that the number of steady states increases exponentially with size of the system (Tatara et al. 2003). With the autocatalytic reaction scheme, larger networks permit more steady states and spatial combinations thereof than smaller networks. In addition to the single-reactor states, which are omnipresent in larger networks, networks exhibit states that are unique to configurations of more than one reactor.

The high degree of nonlinearity provides a challenging obstacle in the control of autocatalytic reactor networks. In the case of a single CSTR with competing autocatalytic species, a nonlinear control scheme is necessary to first remove the invading species and then return the host species to the original steady state (Chaivorapoj et al. 2002, 2003). This concept is extended to control autocatalytic species in a network of multiple reactors. Systems of more than one reactor require multiple controllers and may require transients through several operating regimes to achieve the desired operation. Furthermore, when the system contains multiple reactors, situations arise such that a global control objective can be satisfied by several different combinations of local control objectives. This leads to the requirement of a hierarchical control structure whereby local control objectives can change dynamically in order to achieve the global control objective of the system.

AGENT-BASED CONTROL SYSTEM ARCHITECTURE

Networks pose a tough challenge from a control perspective because multiple control strategies may be necessary to achieve the desired operational or organizational goal. Intelligent supervisory control systems (Kendra et al. 1994) provide adaptive capabilities that facilitate the control of such systems. Furthermore, the operation of highly nonlinear systems like autocatalytic replication networks benefit from *evolutionary control* because the optimal operating regime or control strategy may not be known *a priori*. Agent-based control systems (Jennings and Bussmann 2003) provide the capability for localized and global control strategies that are both reactive in controlling disturbances and proactive in searching for better operational solutions.

The agent-based control system architecture consists of several sub-systems, each of which is highly modularized (Figure 1). The network layer comprises interconnected nodes that represent a discrete physical entity. Each node in the network is monitored by an *observation* agent that is responsible for maintaining a communication channel to the node and acquiring data. These agents are also responsible for sampling data requested by other agents, as well as storing the data in a history for some specified time. Manipulation of the interconnections between nodes is handled by an actuation agent which receives commands from superior control layers.

The next layer in the control hierarchy is the *local decision* layer. Local decision agents are responsible for monitoring control functions and proactively improving the overall performance of the network. Due to the number of control responsibilities of decision agents, each agent may be further modularized into several sub-agents and so on. For example, the local control decision agent requires information regarding the state of the network. A sub-agent is therefore tasked with checking whether the network is at steady state or if the network is oscillating or behaving chaotically.

During network operation, local decision agents will attempt to improve the performance of the node they are controlling at the expense of another node's performance. Thus, disputes will arise as to the value of the interconnections between neighboring nodes. *Arbitration* agents are tasked with negotiating disputes between neighboring agents. Without a means to negotiate, the

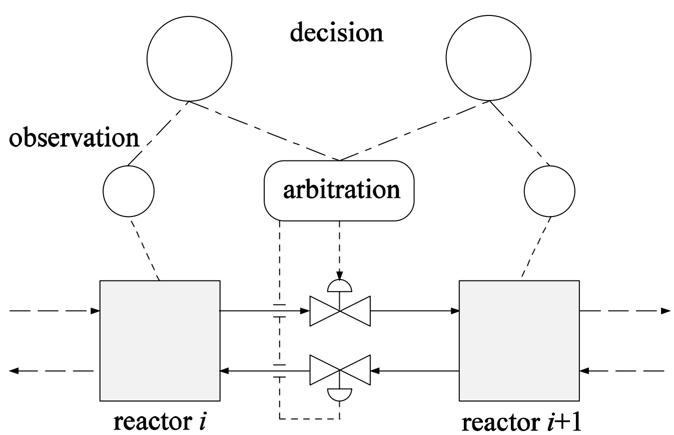

FIGURE 1 Control agent architecture

decision agents would not have an open communication channel between them. The arbitration agents receive requested operational procedures from the local decision agents and then presents a solution to them. *Supervision* agents function as the topmost layer in the control system hierarchy. This layer is responsible for setting the desired global operating conditions for the entire network.

SOFTWARE IMPLEMENTATION

The agent-based supervision and control system development and deployment are performed using G2 software (Gensym 2003). G2 is a graphical knowledge base (KB) development environment for creating intelligent real-time applications. It has a hybrid KB paradigm with classes and objects for knowledge presentation and rules for inferencing. Applications developed in the G2 environment are called knowledge bases (KBs) and contain graphical workspaces upon which object code is organized. Workspaces are arranged in a hierarchical structure such that the source code can be easily managed. Workspaces contain all the rules, variables, and objects that constitute a KB. Furthermore, workspaces can be organized further into modules for increased reusability and scalability. The G2 programming language is very similar in structure to the English language, allowing for rapid software development. Graphical representation of data in G2 is performed through several different types of real-time charts that are customizable by the user. The displays may be placed on any workspace, and the shape, position, and colors of the display may be modified as needed by the user. Customizable graphical user interface (GUI) controls such as buttons and text entry boxes are available in G2.

G2 provides an excellent platform for the development of supervisory KBs. However, complex numerical calculations, such as process simulations or the matrix manipulations needed for analysis, cannot be efficiently implemented in G2. One may take advantage of more sophisticated programs for numerical analysis and simulation by linking those programs to G2. Several options for networking and passing data to and from external programs are available. The G2 Standard Interface (GSI) bridge allows the KB to access remote procedures written in C. The functions contained in the C files are callable by GSI bridge code that communicates with G2 via TCP/IP.

The network node class definition in G2 contains basic specifications for network objects including instantiation, deletion, cloning and connections. Node attributes and methods are specified by the user for the particular system of interest. A software bridge links the G2 node objects to a numerical ordinary differential equation (ODE) solver (Figure 2). CVODE (Cohen and Hindmarsh 1994) is a numerical ODE solver package written in C and is capable of solving stiff large-order systems very efficiently. The solver code is linked to the system via a custom GSI bridge. With a high-level object based language like G2, it is possible to construct an interface between the solver and software agents in order to produce a highly flexible communication pathway between the nodes and the simulator. The executable bridge code does not need to be executed on the same CPU on which G2 is also currently running. This allows the KBS to run on one machine while the bridge code can be executed on a remote machine.

Node attributes are mapped to the specific array location in the ODE solver. When the simulator is initialized, the user may specify the size of the network, parameters such as the connection values between nodes, and initial conditions. The appropriate number of nodes are

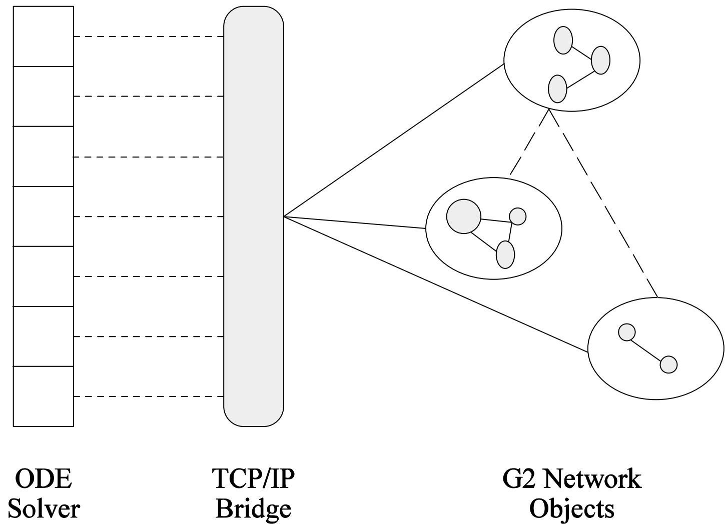

FIGURE 2 Software architecture

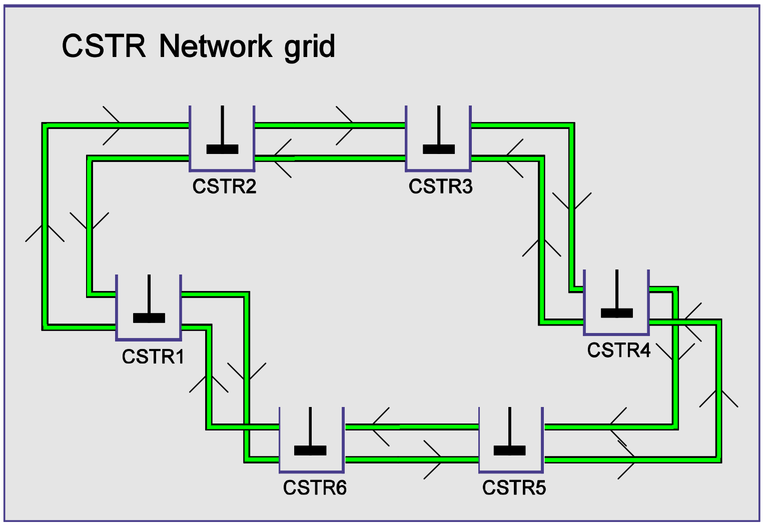

FIGURE 3 Graphical representation of reactor network nodes

automatically created by the agent-based system and the node objects are modified to include the initial conditions. Figure 3 shows the graphical representation of reactor nodes in G2. The CVODE solver simply requires the initial states and parameters to describe the system. The solver then dynamically creates the appropriate equations internally and returns the output to G2, where the states are then mapped back to the node objects. The benefit of using this type of architecture is clear when considering the natural language programming syntax of G2. Software agents can interact with the node objects by invoking methods and procedures using such syntax as "change the connection between node-1 and node-2 to 3."

NETWORK MODEL EQUATIONS

A network of I interconnected reactors is modeled by specifying the mass balance for an individual reactor at position i in the network, where $i = 1\ldots I$. Figure 4 shows the schematic for a system of four CSTRs. The cubic autocatalytic reaction for a single autocatalytic species is

$$R + 2P \xrightarrow{k} 3P$$

$$P \xrightarrow{k_d} D$$

R is the resource concentration, P is the species concentration, D is a dead (inert) species, k is the species growth rate constant, and k_d is the species death rate constant.

The production rates of the resource and species concentrations for a network of size I are

$$V\frac{dR_i}{dt'} = -VkR_iP_i^2 + F(R_0 - R_i) + G(R_{i-1} + R_{i+1} - 2R_i)$$

$$V\frac{dP_i}{dt'} = VkR_iP_i^2 - P_i(F + Vk_d) + G(P_{i-1} + P_{i+1} - 2P_i)$$

where R_0 is the resource concentration in the feed, R_i is the resource concentration in reactor i , P_i is the species concentration in reactor i, F is the feed flow rate, G is the interconnection flow rate and V is the reactor volume. The feed stream contains only resource. The state equations can be written in dimensionless form as

$$\frac{dr_i}{dt} = -kr_ip_i^2 + f(1 - r_i) + g(r_{i-1} + r_{i+1} - 2r_i)$$

$$\frac{dp_i}{dt} = kr_ip_i^2 - p_i(f + d) + g(p_{i-1} + p_{i+1} - 2p_i)$$

by redefining the variables as $r_i = R_i/R_0$, $p_i = P_i/R_0$, $f = F/(VR_0^2)$, $g = G/(VR_0^2)$, $d = k_d/R_0^2$, and $t = R_0^2 t'$. For $i > 3$, analytical solutions become intractable, although it should be noted that a single trivial steady state ($r_i=1$, $p_i=0$) exists for all i for every combination of model parameter values. This state represents total extinction in the system. The feed flow rates and interconnection flow rates are treated as manipulated variables. Constraints on the reactor flow rates ensure that material is conserved.

EFFECTS OF CONTROL STRATEGY

Since nonlinear systems like reactor networks are capable of producing highly complex static and dynamic behavior, the system behavior cannot be easily predicted when coupled with typical control strategies. Within a multi-agent control system, decision agents are given the task of increasing the concentration of the autocatalytic species in their reactor by manipulating the

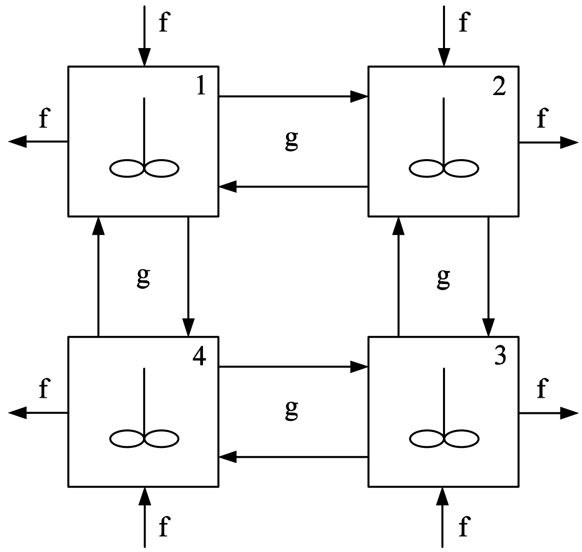

FIGURE 4 4-CSTR network schematic

interconnection flow rates between neighboring reactors. This can be achieved by local exploration of the parameter space. Without a cooperative control strategy, this task is nearly impossible to accomplish because the decision agents will aimlessly search the entire parameter space, frequently resulting in inferior reactor performance.

The result of using a purely competitive strategy on neighboring reactor concentrations is shown in the time series charts of Figure 5. The competitive control strategy used by the local decision agents allows each to agree on a new operating regime only if both agents see a improvement in their reactor. As shown in Figure 5, the agents are able to improve the operating conditions (increase in species concentration) only slightly before the performance begins to degenerate once again. The control strategy gets stuck in the parameter space and will oscillate indefinitely unless halted by the supervision agents.

A more robust approach to this problem is to permit the local decision agents to suffer some performance loss, but only for a small number of control moves. Figure 6 shows the concentrations in neighboring reactors when the decision agents are designed to cooperatively optimize the network performance. The rules governing the arbitration agent permit a decision agent to make a control move that is detrimental to its neighbor for only one iteration, otherwise it must return to its previous state. This strategy proves to be very effective at improving the network performance. Although the performance of CSTR 2 occasionally suffers while the performance of CSTR 1 improves continuously, this loss is only relative to the previous move and ultimately results in a net gain in performance for the whole network.

CONCLUSIONS

Nonlinear physically distributed systems like reactor networks may produce complex dynamic and static behavior. System behavior and performance cannot always be determined *a priori*, making selection of control variables and procedures very difficult. These uncertainties

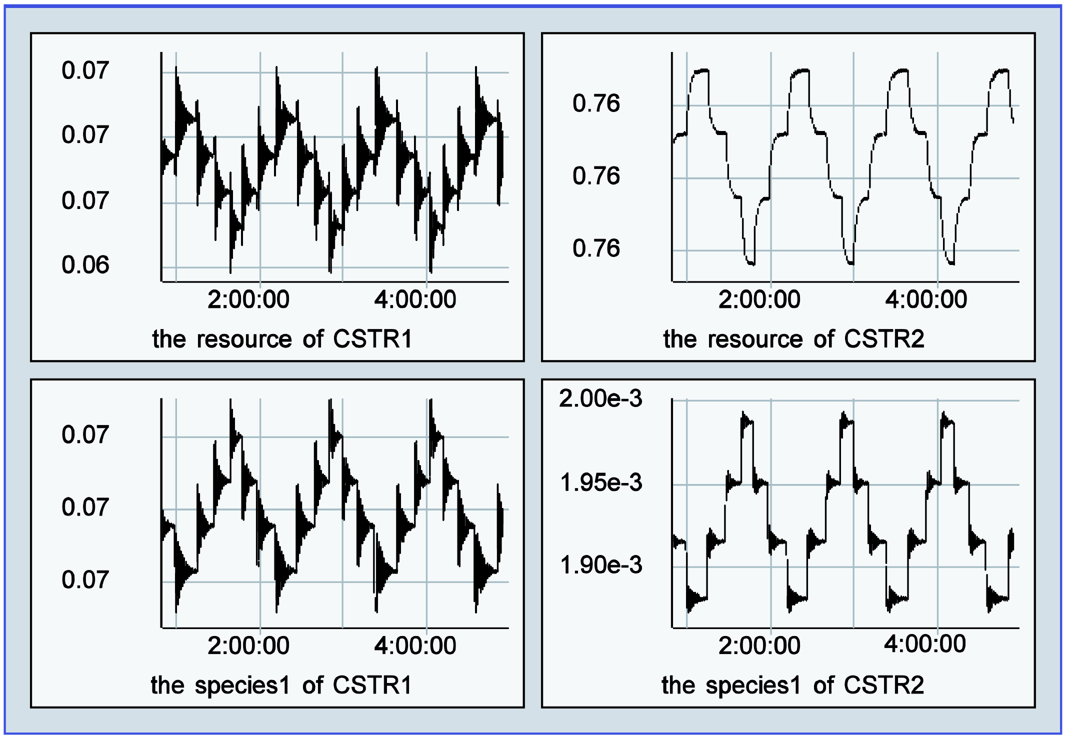

FIGURE 5 Non-cooperative agent control strategy performance

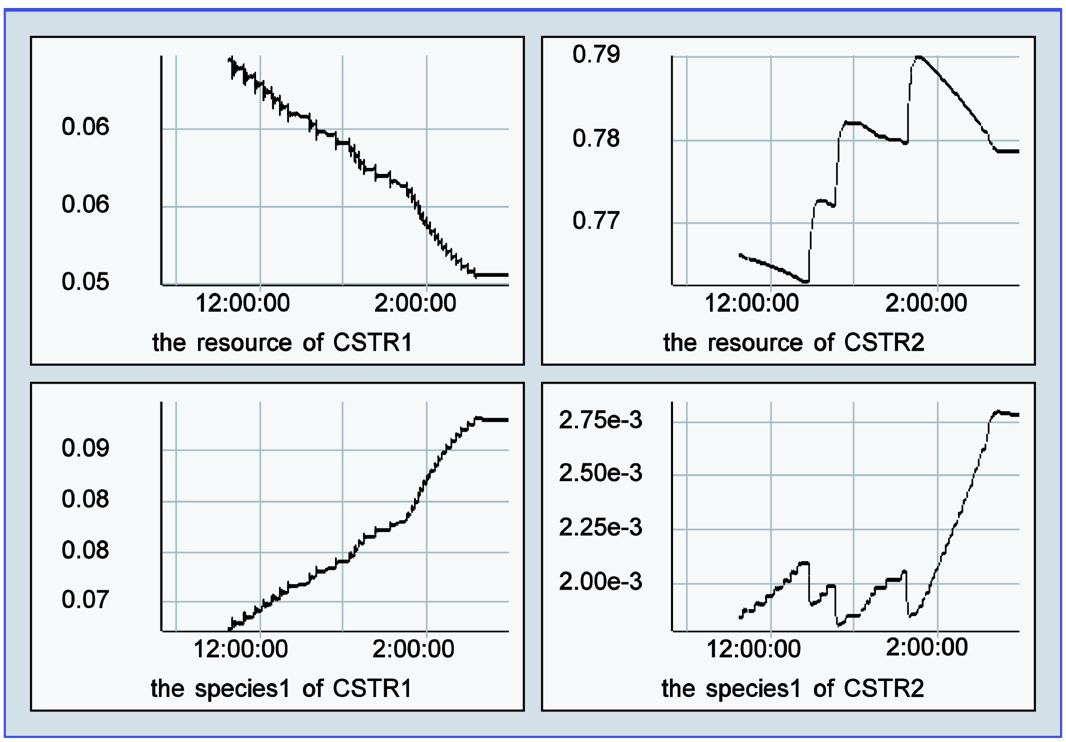

FIGURE 6 Cooperative control strategy performance

naturally result in complications in controlling the system, either in disturbance rejection or changing the operational regimes of the system.

A multi-layer agent-based control system has been developed and applied to a reactor network to improve the overall performance of the system. The multi-agent control system is able to explore the parameter space of the network and intelligently manipulate the network interconnection flow rates such that the specified goal is achieved.

Furthermore, it was shown that cooperative relationships between the decision agents provide a more effective paradigm for improving the network performance than selfish relationships for networks of autocatalytic replicators. While individual reactors may temporarily suffer a performance degradation, the overall performance of the reactor network is improved.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. CTS-0325378 of the ITR program.

REFERENCES

Bilous, O., and N. Amundson, 1955, "Chemical reactor stability and sensitivity," *AIChE Journal* **1**(4):513–521.

Birol, I., S. Parulekar, and F. Teymour, 2002, "Effect of environment partitioning on the survival and coexistence of autocatalytic replicators," *Physical Review E* **66:**051916.

Birol, I., and F. Teymour, 2000, "Statics and dynamics of multiple cubic autocatalytic reactions," *Physica D* **144**:279–297.

Chaivorapoj, W., I. Birol, A. Cinar, and F. Teymour, 2003, "Feedback control of a continuous-flow stirred tank reactor with competing autocatalators," *Industrial and Engineering Chemistry Research* **42**:3765–3785.

Chaivorapoj, W., I. Birol, and F. Teymour, 2002, "Competition between robust and nonrobust autocatalytic replicators," *Industrial and Engineering Chemistry Research* **41**:3630–3641.

Cohen, S., and A. Hindmarsh, 1994, *CVODE User Guide*, LLNL Report UCRL-MA-118618.

Farr, W., and R. Aris, 1986, "Yet who would have thought the old man to have had so much blood in him? - reflections on the multiplicity of steady states of the stirred tank reactor," *Chemical Engineering Science* **41**(6):1385–1402.

Gensym, 2003, *G2 Reference Manual*, Boston, MA: Gensym Corporation.

Gray, P., and S. Scott, 1983, "Absence of restrictions on permissible paths for changing the efficiencies of operation of flow reactors," *Journal of Physical Chemistry* **87**(11):1835–1838.

Gray, P., and S. Scott, 1984, "Autocatalytic reactions in the isothermal, continuous stirred tank reactor: Oscillations and instabilities in the system $A + 2B \rightarrow 3B ; B \rightarrow C$," *Chemical Engineering Science* **39**(6):1087–1097.

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

Kendra, S., M. Basila, and A. Cinar, 1994, "Intelligent process control with supervisory knowledge-based systems," *IEEE Control Systems,* **14**:37–47.

Lesser, V., 1999, "Cooperative multiagent systems: A personal view of the state of the art," *IEEE Transactions on Knowledge and Data Engineering* **11**(1):133–142.

Lin, K., 1981, "Multiplicity, stability and dynamics for isothermal autocatalytic reactions in CSTR," *Chemical Engineering Science* **36**(9):1447–1452.

Monostori, L., and B. Kadar, 1999, "Agent-based control of manufacturing systems," in *The Second International Conference on Intelligent Processing and Manufacturing of Materials*, Vol. 1, pp. 131–137, Honolulu, HI, July 10–15.

Tatara, E., I. Birol, F. Teymour, and A. Cinar, 2003, "Static and dynamic behavior of autocatalytic replicators in heterogeneous reactor networks," *AIChE Annual Meeting*, San Francisco, CA, November 16–21.

Taylor, M., and I. Kevrekidis, 1993, "Couple, double, toil and trouble: A computer assisted study of two coupled CSTRs," *Chemical Engineering Science* **48**(11):2129–2149.

Uppal, A., and W. Ray, 1974, "On the dynamic behavior of continuous stirred tank reactors," *Chemical Engineering Science* **30**:967–985.

Uppal, A., W. Ray, and A. Poore, 1976, "The classification of the dynamic behavior of continuous stirred tank reactors-influence of reactor residence time," *Chemical Engineering Science* **31**:205–214.

ASSURANCE OF SOLUTION CORRECTNESS IN TASK ALLOCATION WITHOUT STIPULATION IN AN AGENT SOCIETY

HAMID HAIDARIAN SHAHRI, Amirkabir University of Technology, Iran*

ABSTRACT

A common problem in multi-agent systems is finding a way to assign tasks to other agents. It is very desirable to be able to guarantee the error rate of a solution in a multi-agent system. This paper introduces a novel approach for solving this problem. We prove that the multi-agent system can achieve any desired and predetermined threshold of correctness for the final solution, regardless of the performance of selfish and unreliable agents in the society or any stipulation about their honesty, which is extremely critical and problematic in the design of open and real-world, multi-agent systems.

Keywords: Multi-agent system, fault tolerance, agent-based simulation, agent society

1 INTRODUCTION

In today's complex and distributed systems, the fitting solution for development of a practical application is to effectively use a flexible multi-agent system (MAS). An MAS is a system in which a large number of self-interested autonomous agents with different design objectives take advantage of other agents' skills to solve a problem. Many issues are involved in engineering an agent society and integrating separately designed agents so that they can work together in an unreliable and distributed environment.

An important point that must be considered in a large agent society is how to choose among the huge numbers of options available in the allocation of tasks to agents. Selection of options is of particular interest in an environment where agents are totally unknown, and there is no central agent or general source of information, such as a directory for querying and learning about the characteristics of agents in the system. To date, this case has not been addressed thoroughly because, in most real-world applications, the agents in an MAS have been designed and implemented by the same group of people. For example, in the production of an MAS for discovering an unknown planet, the designers are aware of the capabilities and the honesty of agents working in the system. With the growth of geographically distributed agents, however, which communicate through a wide area network like many e-commerce applications, there is a very urgent need to build mutual trust in agents' societies.

This paper presents three different fault tolerance mechanisms devised for assuring a correct result in an MAS. An allocating agent is totally unaware of the society in which he lives. The agent learns about the honesty of his partners gradually. Nevertheless, he can

* Corresponding author address: Hamid Haidarian Shahri, Faculty of Computer Engineering and Information Technology, Amirkabir University of Technology (Tehran Polytechnic), Tehran, Iran; e-mail: hhaidarian@aut.ac.ir.

guarantee any level of accuracy for the final solution, which consists of accumulating intermediate results from untrusted and self-interested agents in the environment.

Section 2 provides the mathematical model used for creating an agent society. Section 3 explains the three different fault tolerance mechanisms. The third method provides a general and flexible framework for modeling agent behavior. Section 4 presents the related work of ensuring correctness. Section 5 consists of the conclusion and a discussion of some directions for future work.

2 MATHEMATICAL MODEL

Consider that the agent is going to solve a large problem. The problem is divided into batches; each batch is made up of tasks. The problem-solving agent starts to distribute the tasks of one batch from the pool of tasks, between untrusted agents in the environment, until no tasks remain. Once the results of tasks from one batch have been gathered, the agent begins to distribute the next batch. Distribution of the second batch can use the results of the previous batch. The final solution of the large problem consists of the total of these results.

Error rate, e, is the ratio of the incorrect results accepted to the total number of results. For simplicity, it is assumed that batches are independent, and the results of one batch do not affect the next one. The aim is to design a mechanism that will ensure a predetermined nonzero acceptable error rate, e, for the solution of the large problem. This rate depends greatly on the type of application for which the system is used. For example, if the acceptable error rate for a problem is 1%, and the problem consists of 10 batches with 10 tasks each, the probability of one result failing is $e/(10*10) = 0.0001$. Although this number might seem small, Section 3 shows that it can be achieved with a minor increase in redundancy.

A fraction of malicious agents, m, of the total agent population returns wrong results without actually doing the assigned tasks. The master agent that allocates tasks might know m or can assume an upper bound for it. If m is higher than the assumed upper bound, the master does not guarantee any degree of correctness for the final solution. Because no assumption is made concerning the speed or computational power of other agents, each result could come from any of the agents with equal probability. Therefore, the original error rate would be m, without using any preventive mechanism.

The malicious agents are modeled as Bernoulli processes having a probability of s (sabotage rate) for producing a bad result, which is constant in time and the same for all malicious agents. It is assumed that the worker agents do not communicate with each other and cannot agree on when to give a bad result. When malicious agents return wrong results, however, they agree on the incorrect result to allow voting for malicious agents. If this assumption does not hold, we would expect a number of correct results that are better than the guaranteed rate.

The criteria for comparing the efficiency of different preventive mechanisms are redundancy and slowdown. Redundancy is the ratio of the total number of tasks assigned to the original number of tasks. Analogous to that, slowdown is defined as the ratio of the execution time using the mechanism to not using the mechanism. Although the two criteria are related, if agents can leave the computation or become blacklisted in the middle of a batch, the slowdown would increase, but the redundancy would be the same. Fault tolerance mechanisms should generally minimize or reduce the final error rate to an acceptable level, as well as minimize redundancy and slowdown.

3 FAULT TOLERANCE MECHANISMS

3.1 Voting

The agent allocates the tasks from a pool using eager scheduling in a "round robin" fashion. Tasks could be reallocated to different agents redundantly. As soon as v equal results (votes) for one task are returned, the task is marked "done." When all the tasks in a batch have been completed, the master accumulates the results and begins to distribute the next batch. The redundancy in this method is $v/(1-m)$, where m is the fraction of malicious agents. By approximation, the error rate is $c \cdot m^v$, which shrinks exponentially with v. This method performs well if the fraction of malicious agents m is low, and malicious agents cannot form a majority. In this case, the error rate can be reduced to very low levels with a small increase in redundancy.

The weakness of this method is that m can be high, which can occur in an agents' society. More important, redundancy cannot be reduced to less than two, even when m is zero (no malicious agents). Therefore, it can only be used in cases in which we have an honest society and the population is very large (i.e., we have many trustworthy resources).

3.2 Verification

In the method we call "verification," rather than asking for the result of a task from at least two agents, we might ask them to solve a task for which the result is known. The agents, however, are not aware that at some time they will be tested with a task that has a result known to the master. This method can be compared to the following analogy: the teacher scares students with the possibility of a quiz but does not always plan to give one. In this way, the teacher ensures that the students are working correctly, without incurring much slowdown in class progress. If a malicious agent replies with a wrong result to a known task, its results would be *backtracked* to the beginning of the batch. It could also be *blacklisted*, and no more tasks would be assigned to it. In this case, if the probability of verifying an arbitrary agent is p, redundancy will be $1/(1-p)$, which has a lower bound of one rather than two as in the previous method.

Through blacklisting, only agents cause errors; these agents survive until the end of the batch. If the agent can change its sabotage rate s in time, the error would decrease *inversely* with the length of the batch and in time. Nevertheless, that is a worst-case scenario, and the malicious agent does not have the required information to set s in such a way that maximizes the error rate. Generally, it is advantageous for the master to make the batches longer.

It is not always possible to enforce blacklisting because the malicious agent might hide or forge its identity and IP in a network. In this situation, the error rate would decrease *inversely* with the length of time l that the agent participates in the computation and would be significantly higher than with blacklisting. The reason is unreliable agents, which contribute to performing the tasks and leaving the computation early before getting caught. Therefore, they increase the error rate, and it would be desirable to somehow increase l. Malicious agents could return in the same batch, and the situation would then worsen. To prevent the return of malicious agents, sign-in delays could be enforced. If the sign-in process were delayed until the next batch, the malicious agents would not benefit from leaving the computation early.

3.3 Honesty

In addition to verification, voting could be used to exponentially reduce the already linearly reduced error rate and to achieve more accurate results for the same redundancy. This section explains an effective framework for fixing the problem of blacklisting and combining any other fault tolerance mechanism similar to the two mentioned above.

A parameter called "honesty" is assigned to every agent in the society, analogous to *human societies*. If the result of a task is only accepted when the conditional probability of that result being correct is higher than p, the probability of accepting a correct result, averaged over all tasks, is at least p. That is, accepting results from agents who are honest ensures a correct final solution for any desired nonzero error rate. The honesty of an agent is determined by monitoring behavior and administering tests. Obviously, newly arriving agents are unreliable and have a low degree of honesty. Honesty might also depend on the agents' society and the master's prediction about the fraction of malicious agents in the population.

The honesty of an agent determines the probability of whether the result passed by that agent is correct. The master might receive different results from agents in the society, and a result group consists of all *matching* results for the same task. Each task could have several result groups. The highest probability of accuracy for the result groups determines the probability that the result will be accurate for a particular task. The result group will be accepted only when we have reached a desired threshold for correctness of the result.

Four types of objects are in the system: agent, result coming from an agent, result group, and task. While the master is assigning tasks in a batch, the probability of accurate results in the system increases if the agent passes verification tests, if a matching result arrives from another agent (vote), or both. In time, the probability of correctness for a task will reach the desired threshold and is marked "done"; the master will not reassign it. This method is very efficient because if the answer comes from an honest agent, it will not require voting or verification (i.e., it reduces the redundancy in the system).

4 RELATED WORK

Research on the conceptual and practical tools needed for building dependable agent systems, resilient to errors and other unexpected situations, is at an early stage. Better methods are needed to develop multi-agent systems that can guarantee correctness, reliability, and robustness in an agent society. This fact has given rise to the arrangement of specialized agent conferences with the theme of agent societies; for example, see *Proceedings of the Agent 2003 Conference on Challenges in Social Simulation* and *Proceedings of the Fourth International Workshop on Engineering Societies in the Agents World.*

Turner and Jennings (forthcoming) recently initiated a project on agents for mobile communication environments and are currently exploring integrity and correctness issues. Using formal transformation systems for multi-agent system synthesis is one way to meet this growing need (Sparkman, et al., 2001). Defining boundaries in the behavior space by leveraging the goal hierarchy has also been used for validating complex agent behavior (Wallace, 2003). Ramamohanarao and Bailey (2001) discuss the development of an agent system with a computational model with correctness criteria. Rather than hardwiring robustness and fault-tolerant behavior into agent plans, notions of correctness are embedded at the semantic level. Verification can then be undertaken at the desired level of abstraction. Overeinder, et al. (2003) describe a design for the integration of AgentScape, a multi-agent system support environment, and DARX, a framework for providing fault tolerance in large-scale agent systems. Although these works are related to fault tolerance, they have all used different approaches for achieving it and do not capture the spirit of a real multi-agent community (i.e., they do not try to model the agent society).

5 CONCLUSION AND FUTURE WORK

In an open and geographically distributed multi-agent system, malicious agents might exist. There is no way to guarantee that they actually perform the tasks allocated to them because the agents have made no stipulations regarding this issue. They are usually autonomous, self-interested, and selfish; even worse, they could have spiteful intentions. It is very critical to use fault tolerance mechanisms to make the error rate tractable. This paper introduces setting up the assumption for producing an open multi-agent system, modeling the society of agents without any stipulation on agent reliability, and introducing the generic and effective honesty framework for modeling a society.

The first mechanism introduced is voting: the more agents that return the same result, the more likely it is to be true, assuming that the good agents form a majority. The assumption limits the use of this mechanism, and it is highly redundant and inefficient. The second mechanism is verification. The agents are sometimes tested to see that they are working properly. In this way, redundancy is mitigated, and, in effect, agents are scared into doing the right thing. The human society acts in an analogous fashion.

Finally, the generic honesty mechanism is proposed to create the required infrastructure in a society so that agents can act legitimately. The honesty attribute is neither stipulated in advance, nor is it a characteristic of the agent. Rather, it is computed gradually on the basis of the agent's behavior in society. It represents the faith of others toward the agent. The generic approach of solving this problem culminated in the acquisition of a better understanding of the agents' interactions in an inaccessible environment.

These mechanisms are essential, and some of their applicability is for software agents living in a competitive society like the web environment; agent simulation applications (Winoto, 2002) for research related to sociology, political science, economics, business, and ecology; and creation of a framework for negotiation and collaboration among agents. Considering the many possible real-world applications that can harvest these fault tolerance mechanisms, there is a need for a more detailed mathematical and simulation analysis of these mechanisms, and we are following this idea with promising results. The next logical step, and future direction for research, is to deploy the mechanisms in a real-world application similar to the ones mentioned above.

6 ACKNOWLEDGMENTS

The author deeply appreciates the cooperation and continuous support of Mr. Hossein Taleby and Ahmad Tavallayi at the Management Information Systems Division of the National Iranian Oil Company.

7 REFERENCES

Sallach, D., and C. Macal, forthcoming, *Proceedings of the Conference on Agent 2003: Challenges in Social Simulation*, Chicago, IL, October 3–4, 2003.

Proceedings of the Fourth International Workshop on Engineering Societies in the Agents World, 2003, Imperial College, London, U.K., October 29–31 (forthcoming in *Lecture Notes in Computer Science*).

Jennings, N., and P. Turner, forthcoming, *Software Agents for Future Mobile Communication Environments*. Available at http://www.ecs.soton.ac.uk/~nrj/mobilevce.html.

Overeinder, B., F. Brazier, and O. Marin, 2003, "Fault Tolerance in Scalable Agent Support Systems: Integrating DARX in the AgentScape Framework," *Proceedings of the Third International Symposium on Cluster Computing and the Grid*, Tokyo, Japan, May 12–15.

Ramamohanarao, K., and J. Bailey, "Transaction Oriented Computational Models for Multi-agent Systems," *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'01)*, Dallas, TX, November 7–9.

Sparkman, C.H., S.A. DeLoach, and A.L. Self, "Automated Derivation of Complex Agent Architectures from Analysis Specifications," *Proceedings of the Second International Workshop on Agent-oriented Software Engineering (AOSE-2001)*, Montreal, Canada, May 29.

Wallace, S., 2003, "Validating Agent Behavior," *Proceedings of the 23rd SOAR Workshop*, Ann Arbor, MI, June 23–27.

Winoto, P., 2002, "A Multi-agent Based Simulation of the Market for Offenses," *Proceedings of the AAAI-02 Workshop on Multi-agent Modeling and Simulation of Economic Systems*, Edmonton, Canada, July 29.

DISCUSSION:

COMPUTATIONAL ORGANIZATION THEORY

(Friday, October 3, 2003, 10:00 a.m. to 12:30 p.m., Session 1)

Chair and Discussant: *Noshir Contractor, University of Illinois at Urbana-Champaign*

Dynamics of Expertise in Organizations: An Agent-based Modeling Exercise

[Presentation]

Noshir Contractor: Thank you. Questions, comments, clarifications? Yes.

Joe Jeffrey: Have you had a chance to look at whether anything interesting results from the modelling of experts and domain acquisition. There might be things like a given expert, say in computer science, might easily acquire knowledge in mathematics, but have a much more difficult time acquiring knowledge in English literature or accounting or software management. So on a more complex map and profile of expertise …

Desouza: We have not tried that. It would be a good thing to try, but as of now we have not tried this. But, yes, that's an interesting thing to try, where we can give agents capabilities to learn.

Unidentified Speaker: Something else that would be interesting to see would be what happens with varying degrees of receptivity. We all know experts who are not interested in learning anything new, new domains; whereas, other experts are always eager to branch out.

Desouza: Yes. We actually thought about changing the expert capabilities, but one of the problems we ran into is that the experiment becomes so complex, exponentially difficult. That's why right now we are just at a first step of changing parameters.

Other questions?

Contractor: When we get time, we'll come back and revisit some of these papers later on.

A Quantum Model of Control for Multi-agent Systems

[Presentation]

Lawless: And the conclusions I won't go through. And if there are any questions about future research, I can address those. And that's it.

Thank you very much.

Contractor: Okay. So why don't we go ahead and take some questions here.

Lawless: Well, one of the things that Mike North, who's speaking in the other session right now, mentioned at a conference we were at last weekend was that they're considering revising Repast from its current Swarm model to a quantum model, and that's in my final slide here.

One of the problems that I mentioned to him is that if you come up with a quantum-based model, how in the world are you ever going to go out and test it? And so that's pretty much where we were. And I thought about it on the ride home, and I've come to a different conclusion.

If you look at that last slide, you'll notice that Badredine Arthey, who used to be at the University of Illinois, and I think he's recently moved to Southern Illinois, and others have come up with quantum game solutions already. Quantum game theory has been out since about 1996, and in Badredine's paper, and in others, he has a full solution for, say, a game, Battle of the Sexes. He's got a solution for game theory, he's got the best, the optimum solution for correlated games, and they've also got already analytical solutions for quantum games.

So it seems to me that this should be a logical next step, a really good step, if we're going to create quantum games or quantum agents [where] you could actually use these analytical solutions and establish that the agents can equal those exact proofs and then go from there. The work from there is a very difficult next step. I'm not sure how that will be done.

Application of Agent-based Simulation to Policy Appraisal in the Criminal Justice System in England and Wales

Contractor: Our next presenter is Stephen Guerin, armed with a bunch of colleagues that I'm sure he will introduce. The title of the paper is "An Application of Agent-based Simulation to Policy Appraisal in the Criminal Justice System in England and Wales."

Stephen Guerin: Great, thank you. This is looking at the criminal justice system in England and Wales. It was a collaboration with the London School of Economics and Redfish Group. We're a small consulting company down in Santa Fe. And Daniel Kunkle's with me today, as well as two joint authors, Sean Boyle and Julian Pratt.

The talk's going to be broken into two pieces. One is just kind of a case study of the criminal justice system. It's kind of agent-based modeling in the wild, some of the issues that come up. I can give you a little background there. And then, given the context of this workshop, a little bit more theoretical on modeling organizations and some of the research we're doing.

[Presentation]

Contractor: We have plenty of time for your comments and questions …

[No discussion was recorded.]

Agent-based Supervision and Control of Competitors in a Heterogeneous Environment

Contractor: Eric Tatara and his colleagues are talking about agent-based supervision and control of competitors in a heterogeneous environment.

Eric Tatara: Thank you. I just wanted to point out initially that the URL for our group is at the bottom, and we'll put the slides up shortly. As things go, you'll probably copy half of it before I go to the next slide. So it'll also show at the end.

[Presentation]

Unidentified Speaker: My question is whether in reality that is possible to keep the kind of guarantee you referred to.

Unidentified Speaker: Well, there's also a number of constraints on the system that take this into account in the arbitration. For example, if you look at these flows into the reactors, we have to balance it. You can't send out more than you're getting in. And if one guy closes his up a little bit, he can take in more from a neighbor, and he may restrict what's coming in from the other neighbor. So they will make a move, and then maybe going back to the previous move is less favorable than staying where they are.

Unidentified Speaker: In some models you also have interspecies relationships, correct?

Unidentified Speaker: Right. In this model, there is one species only, and they do have a relationship through the resource. But one does not eat the other or take another one's land or something of that nature. It's just simply through the resources.

Contractor: One of the issues to think about is that this session was focusing on organizational issues, but it has also, in *my* mind, and perhaps [in my] somewhat biased opinion, also focused quite a lot on various aspects of networks of one kind or another, whether it was looking at networks where you're trying to address getting expertise of other people or in a very dyadic sense the networks of cooperation and conflict that you talked about, as well as, in the case of the Redfish presentation, if I could call it that, and the London School of Economics presentation, it was actually multi-level networks, where it was networks at the interorganizational level, but also lower down, amongst the individuals involved in it, etc. And, of course, certainly the last presentation was also looking at networks.

One should not automatically equate computational models as being necessarily *network* models. But one theme that I've found today, and I think it's sort of gratifying, is to see that interest in looking at computational network models in particular.

That being said, I think one of the issues that my own background as a social scientist, even though my undergraduate is from a different IIT — not the Illinois Institute of Technology, but the one from India, the Indian Institute of Technology — but one of the things to think about when we start thinking about looking at this network computational models is there's a tremendous amount of literature in more the traditional social sciences that actually have a lot to say about the different motivations for why we create, maintain, dissolve network linkages.

And one of the things that's I've been involved in during the last couple of years, and this is of course a shameless plug for a book we've just completed — a colleague of mine, Peter Mongey, and I — on theories of communication networks. And one of the things we do in it is talk extensively about the use of computational models for this purpose, but specifically about all the different families of traditional social science theories that we could look at. Most people think of networks in terms of a contagion model, and that is the infection sort of notion, the epidemiological metaphor for networks, which says that if *I* have a certain opinion, through contagion it will flow through my network, or a disease would flow through the network or a behavior would flow through the network.

But in fact looking at also provides theoretical explanations based on theories of self-interest. So a lot of the economic models can be looked at as, say, I want to create a network link with Eric because it is in my self-interest to create it. There's something that Eric has that I want.

A separate mechanism would focus on exchange, which says, "Well, it's not likely that I'm going to have a sustained interaction with Eric if I simply go to Eric on the basis of self-interest." So a second mechanism is to say, "I'll go to someone where I have something, where I need something that person has, but in exchange that person wants something that I have," and so it sets up a social exchange mechanism, and there's a huge body of literature that many of you are familiar with in social exchange literature.

A third distinct mechanism is in terms of collective action, which says, "We're not going to go to each other because we need something from one another, but collectively we can get something from a third party." And so a lot of the models of collective action are in fact based on that particular mechanism. Each of these provides sort of different mechanisms that we could look at together or separately when we're trying to model these. So there's a family of about 10 mechanisms that we've described in our own work in this area.

There's a very interesting body of literature which they [Eric Tatara et al.] do cite in the paper on transactive memory systems, in a theory of transactive memory and social psychology, which says, looking at how experts are distributed in groups is largely being shaped by whether there is one expert and whether other people in the group can say, if they know who the expert is, no one else needs to be the expert. So you don't necessarily have the simple combination models that were described today, but instead to say, "If I am the expert on this topic, no one else in the group needs to be the expert, because it's a waste of their time. If they need help, as long as they know I'm the expert, they'll come to me. If they find something that is of interest, as long as they know I am the expert, they will give it to me." Likewise, if you are the expert, Bill, on a particular topic, the same would apply to you. And so you get this specialization of experts within groups, which works on the basis of each one being an expert on a particular topic and the rest of the people knowing who the experts are and then using that logic as a way of creating knowledge networks within these communities.

In any case, I just wanted to sort of again encourage all of you who have very sophisticated mathematical treatments to consider going back to some of these traditional social science theories and seeing ways in which a lot of what they have to say can be formulated in ways which *they* can't fathom, but we as people interested in computational modeling will be able to deploy in many of these contexts.

Unidentified Speaker: When we talk about democratic decision-making, it's different. I've seen this in many decision-making groups, where the scientist or the experts actually hide from democratic decision-making.

Contractor: Right.

Unidentified Speaker: But, on the other hand, I have seen experts who disagreed with each other in public, like a prosecutor and the defense attorney, serving valuable functions, where their expertise is used to educate neutrals to that decision-making. And they actually make the decision. So that's an alternative.

Contractor: So this is actually the point that I'm making, and that is, you cited Dick Moreland and Levine's work on small groups. Dick Moreland is sort of the big proponent of transactive memory. But the larger point here is that there are different theoretical mechanisms, and part of what is such a wonderful advantage of working with computational models is that you can juxtapose multiple theories and look at these differences in ways that allow us to decide whether they are context-dependent models that apply to a jury situation, as opposed to some specialized medical treatment, etc. And, again, who's to say that transactive memory systems are in fact the most effective. It's simply a theory. And all the work that's been done empirically in this area is typically done at the dyadic level. There are very few people who are even looking at it in terms of larger networks, which we have the ability to do out here.

So it is exactly the sort of multi-theoretical, multi-level approaches that I think would be so valuable, given the kind of computational modeling that we do.

Finance and Markets

AGENT-BASED MODELING OF LOTTERY MARKETS

SHU-HENG CHEN and BIN-TZONG CHIE[1]
AI-ECON Research Center, National Chengchi University

ABSTRACT

The lottery market modeled in this study cannot be explained by the conventional rational expectation approach. Clarifying lottery market behaviors is a daunting task. We apply an agent-based computational modeling technique in which each agent is modeled as autonomous with his or her own perceptions and actions. The objective is to use three empirical observations in lottery markets — the halo effect or lottomania, conscious selection of betting numbers, and aversion to regrets — to examine the effects of the lottery takeout rate on its revenue. Initial results show the Laffer curve, which indicates the existence of an optimal lottery takeout rate or range. This finding provides some insights to the empirical averaged rate for the 25 lottery markets examined.

Keywords: Lottery markets, agent-based computational modeling, Laffer curve, fuzzy system, genetic algorithms

1 INTRODUCTION

Economists find lottery market behavior to be an interesting subject. Many studies have used demographic and socioeconomic data to estimate lottery sales or demand. The standard econometric approach, however, primarily treats the demand decision as an individual rational choice problem. Within this framework, the number of tickets purchased by an individual is determined only by his or her personal profile; this choice has nothing to do with how other people would act. To model this aspect, an agent-based computational modeling approach is used to capture some aspects that cannot be described by using an analytical model.

Modern agent engineering techniques offer more advantages for capturing the idea of autonomous agents. Over the past years, these insights have extended to the economics analysis arena in some areas, such as the artificial financial market. As an extension of our earlier studies with an artificial stock market (Chen and Yeh, 2001; 2002), this paper addresses an agent-based model of lottery markets.

We survey the takeout rate of some lottery markets, which shows a wide distribution. This rate ranges from a low of 40% in Taiwan to a high of 68.4% in Brazil. Although these data are helpful in reaching a design, we observe that the takeout rate is only one dimension of the complex lottery design. Scoggins (1995), Hartley and Lanot (2000), and Paton, et al. (2002) have discussed this issue.

Empirical observations of psychological studies of the lottery market have motivated us to use an agent-based modeling approach. Actually, we find that gamblers are not so concerned

---

1 *Corresponding author address:* B.-T. Chie, AI-ECON Research Center, National Chengchi University, Wenshan, Taipei 116, Taiwan; e-mail: chie@aiecon.org.

with the probabilistic calculation of the odds of winning, because they often rely on heuristic strategies for handling situations. Even though the generation of the winning number is totally a random mechanism, gamblers still tend to pick nonrandom numbers; this process is called *conscious selection*. Griffiths and Wood (2002) reviewed various heuristics and biases involved in the psychology of the lottery market, such as hindsight bias, representation bias, and availability bias. It is not easy to capture such heuristics and biases by using standard rational models.

This paper is organized as follows. Section 2 briefly describes an agent-based model of the lottery market. Section 3 shows how the genetic algorithm is used. Section 4 outlines the experimental designs and Section 5 gives the results of the simulation. Finally, Section 6 provides concluding remarks.

2 AN AGENT-BASED MODEL OF THE LOTTERY MARKET

2.1 The Lottery Market and Its Design

In general, agent-based models consist of two essential parts — the *environment* and the *agent*. In this study, the environment comprises the rules or the design of a lottery game and the states of the market. Walker and Young (2001) conducted a well-known study of the design of the lottery game. Typically, the game is expressed by two parameters, x/X. In this game, gamblers pick x numbers from a total of X numbers without replacement; different prizes then are set for the various numbers that are matched on the drawing day. In a simple description of this process, let y denote the number matched. Clearly, $y = 0,1, \ldots, x$. Let S_y be the *prize pool* reserved for the winners who matched y numbers. The special term for the largest prize pool is called *Jackpot*, S_x.

A common feature of lotteries is that, if a given draw does not generate winners, the jackpot prize pool from that draw is added to the pool for the next draw; this is referred to as a *rollover*. Rollovers usually make the next draw, called the *rollover draw*, much more attractive. The prize pool is defined by the *lottery takeout rate*, τ, which is the proportion of sales that is not returned as prizes. Thus, the overall prize pool is $(1-\tau)S$, where S is sales revenue and $1-\tau$ is the *payout rate*. Therefore, a lottery game can be represented by the following x + 4-tuple vector: $L = (x, X, \tau, s_0, \ldots, s_x)$, which is shown in the control panel of our agent-based lottery software (Figure 1).

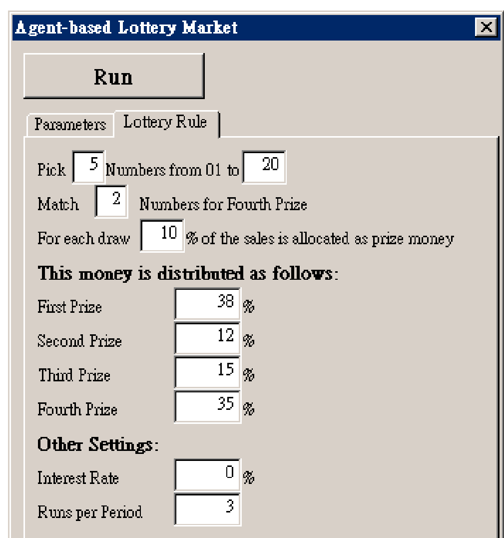

FIGURE 1 Control panel of parameter settings and lottery rules

One of the objectives for using agent-based simulation of the lottery market is to examine the effects of changes in the design L on lottery sales, and more important, on charity fund revenue. The literature shows two approaches for analyzing agents' participation in the lottery markets. In the first approach, the empirical data are used to model the principal features of the observed aggregate behavior (Farrell and Walker, 1999; Farrell, et al., 1999). In the second approach, a rational model of representative agents is used to aggregate these representative agents (Hartley and Lanot, 2000). The agent-based model is closer to the latter but does not use the attributes of rationality and homogeneity.

2.2 Agent Engineering

Since we do not know why people gamble, we do not think that a unique answer can be found to explain this issue. Therefore, many possibilities can be examined by using agent engineering. The basic principle is to ground agent engineering with theoretical and empirical observations. In this way, we minimize the degree of arbitrariness. In our agent-based model, we capture the following stylized facts of the lottery market: lottomania and the halo effect, conscious selection, and aversion to regret.

2.2.1 Lottomania and the Halo Effect

First we observe that lottery sales seem to be positively related to the size of the rollover or jackpot prize. By examining lottery market data, we find that this phenomenon is statistically significant. This phenomenon, called *halo effect* (Creigh-Tyte and Farrell, 1998; Walker and Yang, 2001), can create a bout of “lottomania,” which is propagated by the media. Therefore, we initially build the agents from a participation function, which is a measure of the participation level compared with the size of jackpot. In the standard rational analysis, the change between these two variables is in the expected value, or more generally, the expected utility, of the lottery ticket (Hartley and Lanot, 2000). However, we take a heuristic approach and assume that gamblers base their decisions on some heuristics rather than on the possibly demanding work on the computation of expectations.

The heuristic approach allows approximation of the relation by a few simple if-then rules. We represent the function of participation level by a set of fuzzy if-then rules, which are manipulated by the standard mathematical operations of the fuzzy sets as prescribed by fuzzy set theory.

2.2.2 Conscious Selection

The second important observation related to lottery markets is that gamblers are generally ignorant as to how probability operates. The phenomenon known as conscious selection refers to nonrandom selections of the combinations of numbers. Even more interesting is that there is a market for “experts,” who advise gamblers regarding which numbers to choose. To take conscious selection into account, let a vector be an X-dimensional vector, whose entities take either 0 or 1.

2.2.3 Aversion to Regret

The last feature of our model of agents is the utility function. For simplicity, most advanced-computing-environment models assume an exogenously given utility function that is homogeneous among agents. We have slightly departed from this tradition primarily because of the observation of aversion to regret. In the lottery market, regret simply refers to the utility that the decision not to gamble is based on whether there are winners. If nobody wins, gamblers do not feel regret; however, if somebody wins, they might feel regret (i.e., the prize could have been theirs if they had played the lottery).

In spirit, this consideration is in line with the regret theory proposed by Bell (1982) and Loomes and Sugden (1982). The regret theory offers explanations for numerous evident violations of the expected utility theory axioms. In regret theory, agents, after making decisions under uncertainty, may feel regret if their decisions prove to be wrong even if they seemed to be correct given the information available ex ante. This very intuitive assumption implies that an agent's utility function, among other things, should depend on the realization of alternatives not chosen and, in this sense, irrelevant.

3 GENETIC ALGORITHMS

3.1 Representation

Genetic algorithms (GAs) are motivated from the spirit of natural and are coded with the chromosomes, which is the unit of GAs. In our model, the chromosome is coded as the bit string, which is the vector ($\boldsymbol{a}_{l,t}$, $\boldsymbol{b}_{l,t}$, $\theta_{i,t}$). It fully characterizes an individual at time t. Since each component of the vector is associated with a different function, however, the coding and decoding schemes would be different. Figure 2 illustrates a fuzzy inference system with the corresponding binary string of $\boldsymbol{a}$, decoded as $\boldsymbol{a}$ = (0.2, 0.6, 0.8, 1.0) of real numbers. The input J

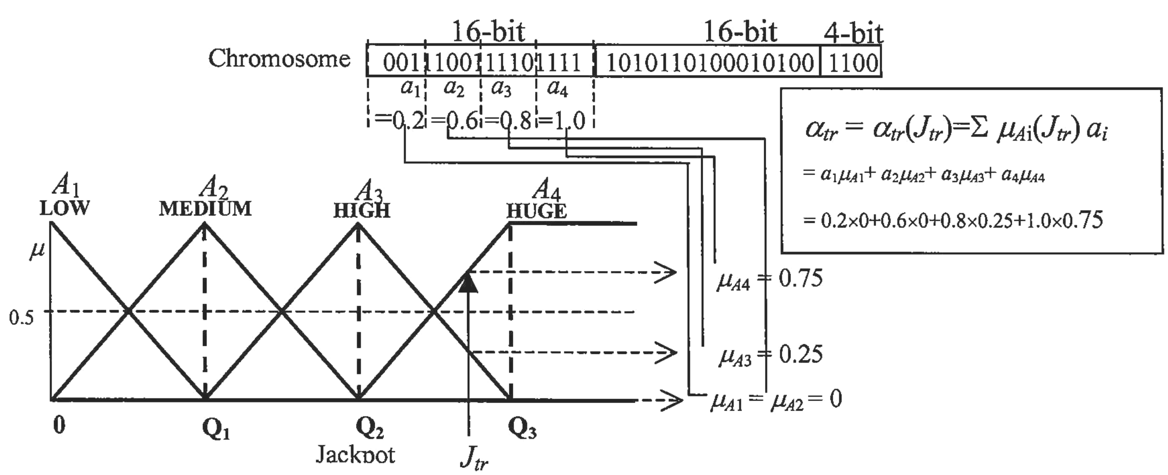

FIGURE 2 Betting heuristics based on the Sugeno fuzzy inference system

is perceived by the agent, and the membership degree of each fuzzy set is calculated as follows: $[\mu_{A1}(J), \ldots, \mu_{A4}(J)] = [0, 0, 0.75, 0.25]$. Therefore, the agent invests

$$\alpha = \sum_{i=1}^{4} \mu_{Ai}(J) a_i = 0.95$$

of his or her income to purchase lottery tickets.

It is straightforward to code $\boldsymbol{b}$, which is the number-picking vector. As mentioned earlier, $\boldsymbol{b}$ is simply an X-bit string. An example of the case $X = 20$ is shown in Figure 3.

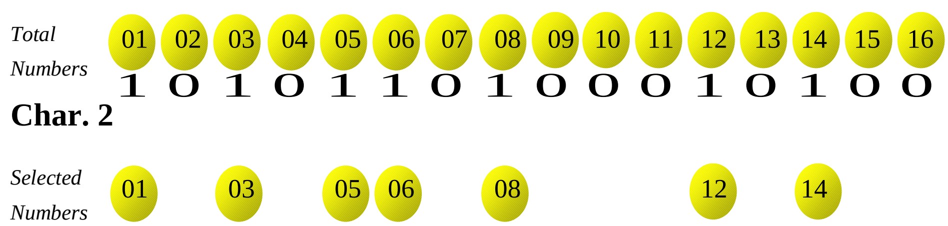

FIGURE 3 Example of numbers chosen by agents

Finally, the regret parameter θ, which also lies between 0 and 1, can be encoded in a similar fashion as binary coding by a l_θ-string bits. Therefore, the full characterization is encoded by a string with a total of $4 \times l_a + l_b + l_\theta$ bits.

3.2 Evolutionary Cycle

Genetic algorithms have two major selection schemes: *roulette-wheel selection* and *tournament selection*. Although these two selection schemes have been well studied in the GA literature, the scheme more suitable for agent-based economic modeling remains an open issue. The reason is because some of the advantages or disadvantages that are known to GA theorists may not be of relevance for social science-oriented studies. Chen (1997) argued that, for social scientists, the network behind the social dynamics is the primary criterion of the selection scheme. Generally, the roulette-wheel selection scheme implicitly assumes the existence of a well-connected global network, whereas the tournament selection scheme requires only the function of local networks. Lacking further evidence on which network assumption is appropriate, it would be beneficial to try both selection schemes to test for robustness. To narrow our focus here, we apply only tournament selection. We plan to include the other selection scheme at a later stage. The following describes the pseudo program of the evolutionary cycle:

```
begin
        Gen := 1;
        Pop := Population-Size;
        initialize(POP(Gen, Pop));
        evaluate(POP(Gen, Pop));
        while not terminate do
        begin
                for i := 1 to Pop step 2
                        Parent1 := Tournament-Select-1st(POP(Gen, Pop));
                        Parent2 := Tournament-Select-2nd(POP(Gen, Pop));
                        OffspringPOP(Gen, i) := Crossover-Mutation-1st(Parent1, Parent2);
                        OffspringPOP(Gen, i+1) := Crossover-Mutation-2nd(Parent1, Parent2);
          next i
                evaluate(OffspringPOP(Gen, Pop));
                POP(Gen+1) := OffspringPOP(Gen, Pop);
                Gen := Gen+1;
        end
end
```

4 EXPERIMENTAL DESIGNS

This paper studies the possible relation between the lottery takeout rate and the lottery sales by hypothesizing the existence of a Laffer curve and hence an optimal interior *T*. To do so, different values of *T* ranging from 10% to 90% are attempted. The remaining market parameters are treated as constants throughout the entire simulation. Figure 4 shows the parameter settings of the agent-based model of the lottery market.

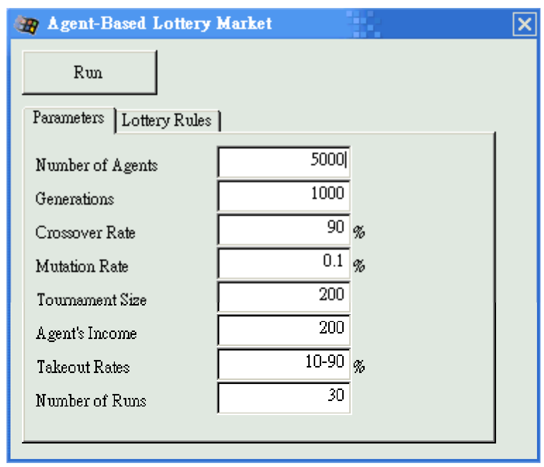

FIGURE 4 Parameter settings of agent-based lottery market

The second set of parameters concerns the control parameters of the genetic algorithm. The parameter *T* (i.e., the tournament size) is unusually large ($T = 200$), which allows for greater interaction among gamblers; this approximates the intensive attention drawn to lottery results reported by mass media. In the future, we plan to apply this agent-based lottery market to some sensitivity issues that pertain to the choice of various selection schemes, market sizes, crossover styles, etc., including their economic significance and the effect on the simulation results.

5 RESULTS

5.1 Takeout Rate and Tax Revenue

Figure 5 shows that initially, the normalized lottery revenue (effective takeout rate) increases with the lottery takeout rate τ and finally decreases with it. The highest revenue appears at τ = 40% with an effective takeout rate of 1.1%. However, the revenue curve is not unimodular; in addition to τ = 40%, it also peaks at τ = 60%. Hence, it is not a typical Laffer curve as one might suppose. The revenue does not monotonically decrease after τ = 40, and the jump at τ = 60% is not surprising. Certainly, this finding does not mean that the complex system used can have only one unique solution: τ = 0.40. Is it possible that different settings of the parameter values can lead to different results? Or are we by luck, for example, simulating a system with a set of parameters that has an optimal solution consistent with the empirical observation? This is indeed the robustness issue that must be addressed in agent-based computational modeling.

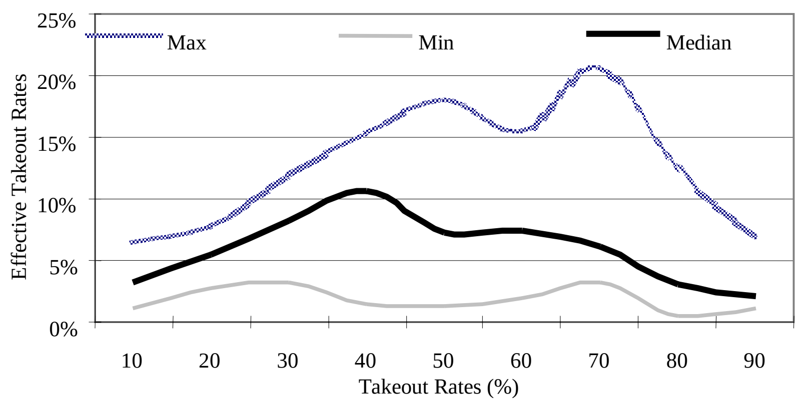

FIGURE 5 Effective tax rates statistics for 30 runs for each tax rate

5.2 Rollovers and Sales

Generally, large-sized rollovers tend to enhance the attractiveness of the lottery game. Statistics show that the mean sales that are conditional on the rollover draw are normally *higher* than those of the regular draw. For example, on the basis of the time series data for the U.K. lottery from November 19, 1994, to March 5, 2003, which comprises a total of 751 draws, the average sales are £56.0 million over the rollover draws, whereas they are £41.4 million over the regular draws. However, from a total of 112 rollover draws of the U.K. lottery, sales actually fell 25 times. On the basis of these statistics, it is interesting to see whether the patterns will be similar for our artificial lottery markets. Therefore, we use the same statistics for the simulated data.

The disappearance of the halo effect and the appearance of the anti-halo effect are certainly astonishing, especially because our agent engineering is based on the consideration of the halo effect. However, a comparison of the real data with the artificial data provides us the opportunity to reflect on something that we may take for granted. In particular, what is the essence of the phenomenon of the halo effect? Why did the agent-based system built on GA fail to deliver this feature? Is there a reasonable explanation for this?

5.3 Conscious Selection

In the real market, many "experts" who advise people on selecting numbers have analyzed the patterns of lottery numbers. In our simulations, the numbers favored by each agent are observable. The profile provides us with the opportunity to examine the behavior of conscious selection. In particular, it enables us to address the question as to whether the agents essentially believe that winning numbers are randomly selected.

5.4 Aversion to Regret

We examine the values of θ, which intensifies agents' suffering when they do not bet in the last period, and take an average from this sample. We call the average $\bar{\theta}$. We see that a culture in which people are sensitive to what others have is nursed in this lottery in this environment. The statistic nearly reaches its maximum and is independent of the takeout rate.

6 CONCLUSIONS

We introduce an agent-based model of the lottery market. This market is composed of many highly interacting agents whose decisions are inevitably interdependent. A model must allow for imitation, fashion, and contagion. In general, an agent's preference for the lottery should be adaptive and evolving rather than fixed. Agents should be modeled as a adaptive agents who, based on their past experiences, are continuously updating their anticipation of the value of lottery tickets and revising their decisions accordingly. By using GAs, we capture the decisions of lottery demand made by adaptive agents in a highly interactive environment and simulate the time series of the aggregate sales of a lottery tickets.

In this paper, the agents are primarily designed on the basis of two empirical phenomena known as the halo effect and the conscious selection of numbers. We also consider the agent's utility function. The empirical observations of the aversion to regret motivated us to find an interdependent utility function for agents. These aspects, which included unsophisticated heuristic behavior, conscious number picking, and preference, are evolving over time via the canonical genetic algorithms.

This model is a starting point for conducting some initial evaluations of the impact of the lottery takeout rate on the lottery revenue. Two observations are made in this paper. First, the Laffer curve suggests an optimal lottery takeout rate τ^*. Second, the τ^* can be sensitive to how agents are modeled. Simulations show that when the regret effect is moved from agents' preferences, the τ^* can go up. If so, the appearance of the interdependent utility function has an implication on the design of the lottery game. Empirical data from Taiwan, U.K., and South Africa national lotteries will be used to examine the performance of our agent-based model.

7 REFERENCES

Bell, D., 1982, "Regret in Decision Making under Uncertainty," *Operations Research* **30**: 961–981.

Chen, S.-H., and C.-H. Yeh, 2001, "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* **25**:363–393.

Chen, S.-H., and C.-H. Yeh, 2002, "On the Emergent Properties of Artificial Stock Markets," *Journal of Economic Behavior and Organization* **49**:217–229.

Creigh-Tyte, S., and L. Farrell, 1998, "The Economics of the National Lottery," RePEc Working Paper, Department of Economics, University of Durham, www.dur.ac.uk/Economics/Research/working_papers/9804.pdf. (RePEc is a collaborative effort to disseminate research in economics. Its home page is http://repec.org.)

Farrell, L., and I. Walker, 1999, "The Welfare Effects of Lotto: Evidence from the UK," *Journal of Public Economics* **72**:99–120.

Farrell, L., E. Morgenroth, and I. Walker, 1999, "A Time-series Analysis of UK Lottery Sales: Long and Short Run Price Elasticities," *Oxford Bulletin of Economics and Statistics* **61**:513–526.

Griffiths, M.D., and R.T.A. Wood, 2002, "Lottery Gambling and Addiction: An Overview of European Research," *European Lotteries*, Psychology Division, Nottingham Trent University.

Hartley, R., and G. Lanot, 2000, "On the Design of Lottery Games," Working Paper, Department of Economics, Keele University.

Loomes, G., and R. Sugden, 1982, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *The Economic Journal* **92**:805–824.

Paton, D., D.S. Siegel, and L. Vaughan, 2002, "The Impact of Taxation on the Demand for Gambling," *Economic Division Working Paper Series*, No. 05–2002, Nottingham University Business School.

Scoggins, J.F., 1995, "The Lotto and Expected Net Revenue," *National Tax Journal* **5**:61–70.

Walker, I., and J. Young, 2001, "An Economist's Guide to Lottery Design," *The Economic Journal* **III**:700–722, Nov.

EFFECTS OF GLOBAL INFORMATION AVAILABILITY IN NETWORKS OF SUPPLY CHAIN AGENTS

C.M. MACAL* and M.J. NORTH, Argonne National Laboratory, Argonne, IL

ABSTRACT

Supply chains are systems that touch all aspects of production, distribution, and retailing of goods and services, linking suppliers to manufacturers, to distributors, to wholesalers, to retailers, and finally to consumers. Supply chains can be thought of as collections of autonomous, interacting decision-making units (agents), such as organizations, business units, and individual decision-makers. This paper discusses the modeling of supply networks as collections of interacting social agents using agent-based simulation. A computational model is developed that captures the salient features of supply chain dynamics and system behavior. The simulation is used to study the effects of information accessibility on supply chain dynamics and the implications of agent decision rules on supply chain performance. We provide a method for quantifying the value of information from both the agent and systemic perspectives. When supply chain agents have access to only locally available information (for example, their own inventory and outstanding orders) and information visibility is minimal, supply chains may operate far from a cost-minimizing state. Attempts to control system behavior by agents who use only locally-optimizing decision rules may be highly ineffective. The value of additional information is found to increase only up to a point beyond which no further gains can be identified.

Keywords: agent-based modeling and simulation, supply chain, supply network, value of information, Beer Game

INTRODUCTION

Supply chains are everywhere in today's modern society, involving every aspect of the production and provision of goods and services. Supply chains are systems that touch all aspects of production, distribution, and retailing of goods and services, linking suppliers to manufacturers, to distributors, to wholesalers, to retailers, and finally to consumers (Figure 1). A supply chain is actually a complex network with a diverse set of members and relationships that are constantly changing over time. We use the terms supply chain and supply network interchangeably in this paper depending on the context of the discussion. A supply chain is dynamic. Goods flow down the network from primary producers to consumers in response to information, in the form of orders, which flows up the network from consumers to the production. Supply chains are an example of a system in which the complex, macro-level system behaviors can be directly traced to the behavioral decision rules of the individual agents. Even simple supply chains can exhibit complex behaviors that make them difficult to manage and control (Lee et al. 1997). The "bullwhip effect" is the name given to the tendency of a supply

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

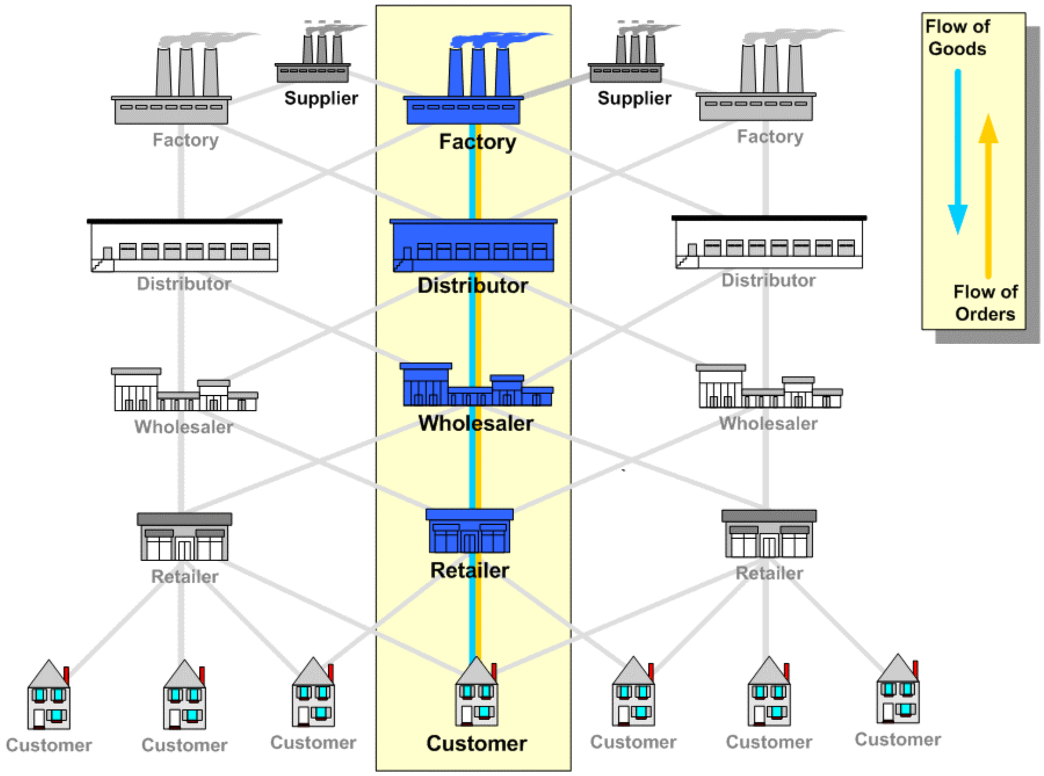

FIGURE 1 Supply network

chain to exhibit unstable behavior characterized by oscillation (orders and inventories are dominated by large amplitude fluctuations), amplification (the amplitude and variance of orders increases steadily from customers to upstream stages), and phase lag (the order rate tends to peak later as one moves upstream from the retailer to the factory stages). It is well-known that cyclic (boom and bust) behaviors and chaos can arise in even the simplest supply chains as a result of decision rules that are based on only locally available information, as is the case in most modern supply chains (Rasmussen and Mosekilde 1988, Mosekilde et al. 1991, Sterman 2000).

Supply chain agents are necessarily *social* agents, having many dimensions of interaction and exchange beyond those that are routinely modeled. They seek sources for their supply from others; they seek outlets for their products; and they negotiate on pricing, delivery, and information sharing. The value of information and the benefits of information sharing in supply networks are particularly important aspects of planning and managing dynamic supply chains. Advances in information technology allow for rapid dissemination of supply chain data. In many business relationships, data are now being shared between cooperative supply chain partners to improve overall supply chain performance and allow parties to rely less upon forecasting to make decisions. Understanding the value of information for supply chain agents and the effects of having various degrees of information on macro-level supply chain performance is an active research area and the focus of this paper. In this paper, we take an agent-based simulation approach to modeling the supply chain and investigating the value of information. Agent-based simulation is a new type of simulation modeling that focuses on modeling the diverse behaviors of individuals in a system and emphasizes the emergent properties of the system that result from the agents' interactions (Epstein and Axtell 1996, Gilbert and Troitzsch 1999, Bonabeau 2002).

Some have even made a case that computational, agent-based modeling is a third way of conducting science (Axelrod 1997) in addition to deduction and induction.

This paper is organized as follows. The second section describes the agent-based approach to modeling supply networks. The third section describes the set-up for experiments to understand the value of information to supply chain agents. An important aspect of studying information in a controlled, experimental setting concerns the effects of agent diversity and how slight variations in agents' accessibility to information can lead to larger differences in supply chain performance and emergent system effects. The fourth section presents experimental results on the role of information and diversity. The last section draws some conclusions and implications for studying real-world supply chains.

AN AGENT-BASED APPROACH TO MODELING SUPPLY NETWORKS

Supply chains lend themselves to agent-based simulation in which agents are the "decision-making members" of the supply chain. Supply chains can be thought of as collections of autonomous decision-making units (agents), such as organizations, business units, companies, and individual decision-makers. Agents are also heterogeneous, each agent having its own individual characteristics, objectives, and constraints, yet constrained in decision-making by the results of the decisions of other agents and by its environment. The essential aspect of agent simulation – the ability to represent individual agents and complex interactions between agents – allows for the possibility of developing a more complete, realistic, and recognizable model of supply chains and agent behaviors than has been previously possible.

Supply Chain Simulation Model

The starting point for the supply chain simulation is the so-called "Beer Game," (Sterman 1987, Sterman 1989, Sterman 2000), designed to illustrate the complexities of planning and operating supply chains. The "Beer Game Simulation" (BGS) is a systems dynamics simulation model (Forrester 1961) that has been used extensively to study supply chain behavior. The BGS is well-known, well-studied, and well-published. The BGS consists of a linear supply chain with five levels or stages: customer, retailer, wholesaler, distributor, and factory. One commodity is ordered and shipped between successive stages. At each time in the simulation, the customer places an order with the retailer who fills the order if the retailer's inventory allows. If the retailer cannot fill the order immediately from inventory, the order is placed in backorder, to be filled at a later time when stock is replenished. The retailer orders additional items from the wholesaler as needed to meet expected demand for the next period and to meet stock-level and "pipeline" goals. The pipeline consists of orders placed but not received. The wholesaler fills the order if the wholesaler's inventory allows. If the wholesaler cannot fill the order immediately from inventory, the order is placed in backorder to be filled when stock is replenished. The wholesaler orders additional items from the distributor, and so on, for the distributor and factory stages. If the factory cannot fill an order, it places an order for items into production. The sequence of supply chain stages from manufacturer to final customer is termed the *downstream*, and the sequence of supply chain stages from final customer to manufacturer is termed the *upstream*.

Standard assumptions include a one-period delay between orders being sent and received throughout the supply chain, a two-period shipping delay between the time items are shipped and received, a three-period production delay (equivalent to the sum of the one period ordering delay and the two-period shipping delay for the other stages) for items ordered by the factory. Other assumptions are easily incorporated into the framework, but it has been observed that the same systemic supply chain behavior results no matter what values of delay parameters one selects. A simulation consists of repeating these processes for several periods. At each time period, each agent updates its stock, backorders, inventory, and the number of items in the pipeline, according to the following flow equations:

$$\text{Stock}_t = \text{Stock}_{t\text{-}1} + \text{Receive}_t - \text{Supply}_t, \text{ for Stock}_t \geq 0$$

$$\text{Backorder}_t = \text{Backorder}_{t\text{-}1} + \text{Demand}_t - \text{Supply}_t, \text{ for Backorder}_t \geq 0$$

Inventory is defined as the difference between stock and backorder:

$$\text{Inventory}_t = \text{Stock}_t - \text{Backorder}_t$$

The backorder is always non-negative by definition, since one of the assumptions of the BGS is that an agent never supplies more than the current demand plus the backorder at any time. Stock is always non-negative by assumption in that an agent never supplies more than an amount equal to the current stock on hand plus the incoming shipment just arrived. In effect, this is the setup is for a “demand pull” supply chain.

At each time *t*, agents place an order to the upstream agent, or to manufacturing in the case of the factory. In effect, the main decision variable for each agent is how much to order from the upstream agent at each time. Ordering rules are described below. Agents also track orders in the pipeline, that is, how much they have ordered but not yet received, given by:

$$\text{Pipeline}_t = \text{Pipeline}_{t\text{-}1} + \text{Order}_t - \text{Receive}_t, \text{ for Pipeline}_t \geq 0$$

The pipeline is always nonnegative by assumption in that an agent cannot receive more than the cumulative amount they have ordered.

Cost enters into the simulation as an inventory charge for stock at $0.50 per item per period (in effect an inventory holding charge) and a charge for orders that are received but cannot be met immediately. A backorder charge is incurred of $2.00 per item per period. The exact values of the costs for stock versus backorders are of less importance in determining system behavior than the assumption that it is more costly to have an item out of stock than to carry an item in the inventory. These relative costs drive the agent goal of maintaining a positive stock level as a safety buffer.

Supply Chain Agent Model

We formulate the BGS as an agent-based simulation model, which is much different in form than the original SD model but yields mathematically equivalent results to the SD formulation. Supply chain agents vary along several dimensions, including their general characteristics, the resources they control, and the sophistication of the decision rules they employ. In general, there are several important aspects of modeling supply chain agent behavior. These include the (1) sophistication of the agents' decision models, (2) extent of information considered in making decisions, (3) the amount of information stored in memory about previous decisions and events, and (4) whether agents have internal models of how the supply chain works in the whole, including whether agents have models of other agents' decision processes and behaviors. Agents must make decisions within a limited amount of time and resources so they often use simple heuristics or "rules-of-thumb" upon which to make decisions.

Supply chain agents have a limited amount of attributes, memory, and decision-making capability. An agent carries the following minimal set of information about itself (Figure 2):

- Inventory at time t and desired inventory level (assumed to be constant),

- Orders in pipeline at t and desired pipeline level (assumed to be constant),

- Demand at t (from downstream agents) and expected demand (demand forecast) from downstream agents for the next period ($t + 1$) made at t, and

- Previous decisions made by the agent and the expected demand from downstream agents for the current period t which they made during the previous period ($t – 1$). In effect this is used by the agents to determine their demand forecasting error.

By assumption, the desired inventory and pipeline levels are decided upon and fixed for each agent through all simulation periods.

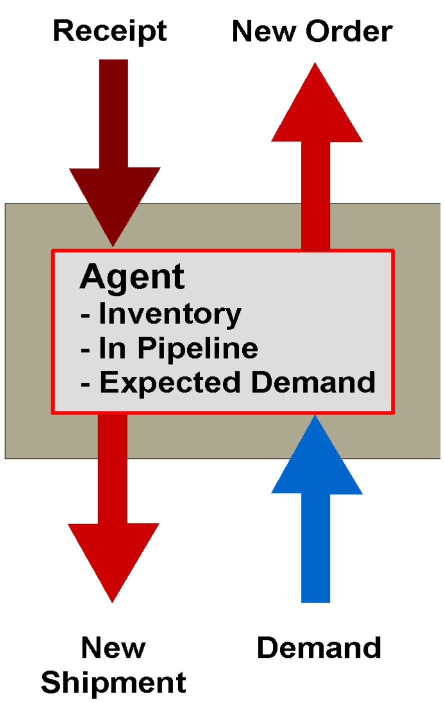

FIGURE 2 A supply chain agent

Agent Rules

Agents in the supply chain manage the flow of goods and orders for goods. Agents manage inventory (and costs) by placing orders using ordering policies in each period. Supply chain agents are faced with several difficult questions in each time period: What will be tomorrow's demand? How much should I order? How much should I ship? How should I organize these decisions? In the standard BGS, at each time period, each supply chain agent carries out the following four functions:

1. *Receive shipment:* When an agent receives a shipment from an upstream agent, it adds the shipments to its inventory.

2. *Receive demand:* An agent receives the order (demand) from the downstream agent.

3. *Decide on shipment:* In response to demand from the downstream agent and to fill any backorders that exist, an agent ships as many items as it can from its inventory to meet the current and previous demand. In effect, the shipment amount is equal to the minimum of the agent's inventory and the recent demand plus backorder. Demand that cannot be met is recorded as a backorder, or negative inventory, to be satisfied in the future. The decision made by the agent at this point is *how much* to ship to the downstream agents.

4. *Decide on order:* Finally, the agent decides how much to order from the upstream agents.

Initially, only locally available information is available to the agents for them to use in making their decisions. Information is obtained through (and created from) local interactions with other agents.

Agents have the goal of achieving desired inventory and pipeline levels. To meet these goals agents seek to close simultaneously close the gaps between the desired and actual inventory levels and the desired and actual pipeline levels. Stock adjustments to inventory are determined by:

$$\text{Stock Adjustment to Inventory}_t = \alpha_S\,(\text{Desired Inventory} - \text{Inventory}_t) \qquad (1)$$

Stock adjustments to the pipeline are determined from:

$$\text{Stock Adjustment to Pipeline}_t = \alpha_{SL}\,(\text{Desired Pipeline} - \text{Pipeline}_t) \qquad (2)$$

where the ordering parameters, α_S and α_{SL}, represent the fractions of the gap to be closed between desired and actual inventory and pipeline levels, respectively, in one time step. Ordering parameters are assumed to be in the range zero to one ($0 \leq \alpha_S \leq 1$, $0 \leq \alpha_{SL} \leq 1$). An additional ordering parameter is often defined, $\beta = (\alpha_{SL} / \alpha_S)$, which is the relative weight attached to the pipeline versus the inventory gaps. The behavioral decision rules in Equations 1 and 2 are the standard form for modeling behavior and are referred to in the literature as "anchoring and adjustment" rules (Sterman 1987). The agent decision framework consists of three types of agent rules corresponding to the four steps in the agent decision process.

Supply Rule: In the standard BGS, each agent decides on how much to supply based only on locally available information, consisting of the stock available and the incoming orders from the downstream agent. The amount supplied by an agent is equal to the lesser of the stock (if any) and the incoming order and backorder. That is:

$$\text{Supply}_t = \text{Minimum}[\text{Stock}_t, \text{Incoming Demand}_t + \text{Backorder}_t]$$

In effect, the agent supplies the complete inventory (if there is any stock) or supplies up to the order plus backorder.

Ordering Rule: Each agent makes an ordering decision based only on locally-available information. The information consists of the expected demand for the next period (demand forecast) and adjustments to the stock and the pipeline. Agents use an adaptive behavioral decision rule to determine orders

$$\text{Indicated Order}_t = \text{Expected Demand}_t + \text{Stock Adjustment to Inventory}$$

$$+ \text{Stock Adjustment to Pipeline}$$

Finally, by assumption negative orders are not allowed, and the order the agent places to its upstream agent at time t is:

$$\text{Order}_t = \text{Max}[0, \text{Indicated Order}_t]$$

This form of the ordering rule is the classic one used in the original BGS.

Demand Forecasting Rule: Agents adjust their expected demand for the next period by weighting the current demand and the previously expected demand for the current period to estimate the demand for the next period:

$$\text{Expected Demand}_t = \theta \text{ Demand}_t + (1 - \theta) \text{ Expected Demand}_{t\text{-}1},$$

where Expected Demand_t is the demand forecast for the next period that is made during the current period. The weighting factor θ is assumed to be in the range zero to one ($0 \leq \theta \leq 1$). These two rules are arguably, a good, descriptive, behavioral decision model for agent behavior in the linear supply chain (Sterman 1987, 1989).

The Network Supply Chain Model

The original BGS was formulated as a linear supply chain model. Here, to add realism to the supply chain model and to study the effects of agent diversity, we extend the linear supply chain model to a fully connected supply network. The “Network Beer Game Simulation” (NBGS) introduced here consists of the five stages in the original BGS but connected in a dense network configuration. Each agent is connected to all the agents at the next upstream stage (excluding factory agents, who do not have an upstream) and to all the agents at the next downstream stage (excluding customer agents, who do not have a downstream). Additional agent decision rules are needed for the network version of the supply chain simulation to account for additional decisions that have to be made in going from a chain to a network. These consist of rules for allocating supply to downstream agents and for allocating orders to upstream agents.

Supply Allocation Rule: In the network, an upstream agent has shipment decisions to make. The upstream agent has to decide (1) how much to supply to all downstream agents and (2) the share of the supply (allocating shipments) that should be made to each downstream agent. There are several alternatives for allocating supplies. For example, one rule is to fill the largest backorder first (LBF) until all backorders are filled or the supply runs out, whichever comes first. An alternative is to prioritize the downstream agents on the basis of smaller backorder and supply the downstream agents in order of the smallest backorder first (SBF); this scheme would seek to supply as many downstream agents as possible. For the LBF rule, the supply shipped to a downstream agent in period t is proportional to the existing backorder plus the order just received from the agent:

$$\text{newSupplies}_d = \text{newSupply} \times (\text{backorder}_d + \text{order}_d) / \textstyle\sum_d (\text{backorder}_d + \text{order}_d),$$

where newSupplies_d is the portion of the supply (newSupply) shipped to downstream agent d. If $(\text{backorder}_d + \text{supply}_d)$ is zero, no supply is allocated to the downstream agent.

Order Allocation Rule: In the network, a downstream agent has to decide how to allocate orders among upstream agents. The agent has to decide (1) how much to order in total from all upstream agents and (2) the share of the order that should be placed with the various upstream agents. There are several plausible alternatives for allocating orders. One alternative is to order less from upstream agents with relatively larger backorders and order more from upstream agents with relatively fewer backorders. Using this rule, the order placed to an upstream agent in period t is inversely proportional to the existing backorder, less the shipment just received from the agent:

$$\text{newOrders}_u = \text{newOrder} \times 1/(\text{backorder}_u - \text{supply}_u) \ / \ \textstyle\sum_u 1/(\text{backorder}_u - \text{supply}_u)$$

where newOrders_u is the portion of the order (newOrder) ordered from upstream agent u. Appropriate adjustments are made in calculating the order share if $(\text{backorder}_u - \text{supply}_u)$ is zero. Many plausible ordering and shipping rules can be postulated based on the operations management literature.

The agent-based modeling approach is particularly conducive to modeling and exploring alternatives for agent decision rules. In the NBGS experiments reported here, we assume agents use the "largest backorder first" rule for allocating supplies and the "order less from larger backorders" rule for allocating orders, as described above.

EXPERIMENTAL DESIGN FOR VALUE OF INFORMATION IN THE SUPPLY CHAIN

We study the effects of expanding the scope of the information available to an agent from the strict local domain of the agent to the broader domain extending to the entire supply network. We study the effects that access to more information can have on supply chain behavior and performance (system costs which is the total of costs incurred by all agents). We further seek to establish a framework for quantifying the value of information, or conversely, the costs associated with incomplete information.

Experimental Design

We set up two experiments to study the effect of uncertainty or imperfections in the information available to agents. Both experiments rely on the fact that an agent forecasts demand (incoming orders) for the next period as part of its ordering rule. As in the BGS chain model, in the setup for the NBGS, expected demand is based on the moving average of previous demand and current demand weighted by a user-set control parameter θ, but we replace the expected demand as computed by the weighted average formula, with the actual demand that will be incoming to the agent (Experiment 1) and with the actual final customer demand (Experiment 2). The hypothesis for both experiments is that if agents have more information, they can use the information to counter the chaotic nuances of the dynamics of the supply chain, in effect smoothing its own behavior to counter the radical dynamics of the supply chain. Specifically, we study the value to an agent of considering information beyond its immediate local neighborhood upon the agent's cost and the system behavior.

Experiment 1: Information Visibility

In Experiment 1, we study whether accurately predicting the incoming order (demand), could avoid the amplification effects of the intervening downstream stages. We have the following hypothesis:

> *Hypothesis 1:* If an agent is able to correctly forecast incoming demand each period no matter where it is located in the supply network, using this information in its ordering process will smooth the dynamics of the supply chain.

We parameterize the extent to which an agent can see into the future in terms of its incoming supplies (that is, how far it can see down the supply chain). We term the ability of an agent to see readily obtainable information (orders already in the pipeline) as *visibility*. We define visibility as the number of stages an agent is able to see beyond its local neighborhood. Visibility is a discrete parameter that ranges from 0 (local information only) to 5 (the ability to see 5 time periods up and down the supply chain). See Figure 3.

Experiment 2: Imperfect Information

In Experiment 2, we study whether if all agents were able to accurately predict final customer demand, it could avoid the amplification effects of the intervening downstream stages. In this experiment, each agent uses the actual customer demand at any time as the basis for its demand forecast for the next period, rather than basing orders on incoming demand from the downstream agents. This leads to the second hypothesis.

> *Hypothesis 2:* If an agent is able to immediately see the final customer demand each period, no matter where it is located in the supply network, using this information in its ordering process will smooth the dynamics of the supply chain.

Although in this experiment, agents may know exactly what final customer demand is, they do not know necessarily what portion of that demand will be ordered from them by the intervening

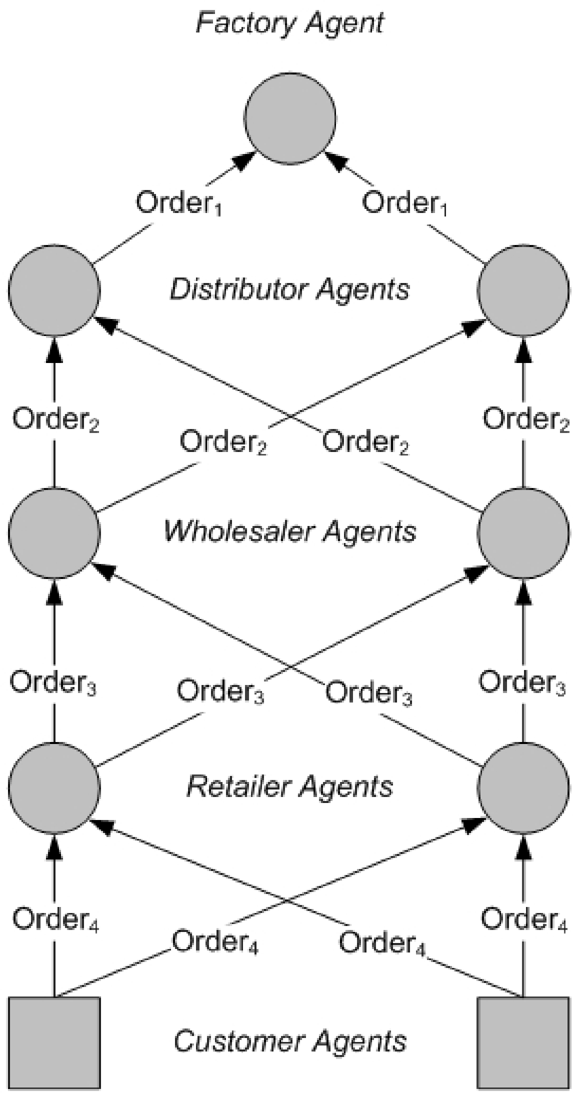

Experiment 1*: Information Accessible by Agent and Visibility Level*

Visibility	Retailer	Wholesaler	Distributor	Factory
1	$Order_4$	$Order_3$	$Order_2$	$Order_1$
2	$Order_4$	$Order_4$, $Order_3$	$Order_3$, $Order_2$	$Order_2$, $Order_1$
3	$Order_4$	$Order_4$, $Order_3$	$Order_4$, $Order_3$, $Order_2$	$Order_3$, $Order_2$, $Order_1$
4	$Order_4$	$Order_4$, $Order_3$	$Order_4$, $Order_3$, $Order_2$	$Order_4$, $Order_3$, $Order_2$, $Order_1$

Experiment 2*: Information Accessible by Agent (Final Customer Demand)*

Retailer	Wholesaler	Distributor	Factory
$Order_4$	$Order_4$	$Order_4$	$Order_4$
$Order_4$	$Order_4$	$Order_4$	$Order_4$
$Order_4$	$Order_4$	$Order_4$	$Order_4$
$Order_4$	$Order_4$	$Order_4$	$Order_4$

FIGURE 3 Setup for information visibility experiments

downstream agents that separate them from the final customers. There is necessarily some uncertainty regarding what intervening agents will decide to do, which is modeled in the following way. We introduce a parameter that indicates the degree to which the upstream agents are able to forecast the portion of the incoming demand from the downstream that is going to accrue to them. The parameter is constructed to vary between 0 and 1. Let σ_c be the portion of the customer demand that is received by agent *c* of all the agents in the same level of the supply network. Then,

$$\text{ExpectedDemand}_{at} = \sigma_c\ \text{CustomerDemand}_{ct} \,/\, (\text{Number of Customers}),$$

where:

$\text{ExpectedDemand}_{at}$ = Expected demand estimated by agent *a* at time *t* and

$\text{CustomerDemand}_{ct}$ = Customer demand for customer *c* at time *t*.

An agent is able to forecast demand correctly for the downstream stage as a whole, and assumes that the shares of orders will be equally distributed across all of the upstream agents at its level. The only difference between the actual demand and the demand the agent faces is due to the possibility that the shares of orders from each of the downstream agents will not be equally distributed across the upstream agents.

Agent Diversity

For the case of the supply network as opposed to the linear chain, the “diversity” of the agents within each stage becomes an important aspect of the study of information. For if all agents within a stage are identical in terms of their attributes, their decision rules, their interactions with upstream and downstream agents, their initial resources, the incoming supplies being received from the upstream, and the orders (demand) being received from the downstream, then all the agents within the stage will have identical behavior for each time step in the simulation, i.e., they will place exactly the same orders to upstream agents and make exactly the same shipments to downstream agents. Therefore, the study of information requires the introduction of some notion of diversity among the agents. To this end we parameterize the diversity of the agent characteristics in terms of the possible range of error that each agent could experience in predicting its portion of the current final customer demand that it will receive from downstream agents.

The amount of deviation of the shares from equality is the source of the imperfections in the information. We quantify the *degree of information deviation* in terms of the error between the expected demand as computed above and the actual demand that the agent receives. Further, we parameterize the degree of information deviation on a scale from 0 to 1, with 0 signifying that information is perfectly correct (no deviation), and 1 signifying that information is incorrect (the deviation of information from perfect ranges up to ±100%). Considering actual customer demand in place of the incoming order in the pipeline for each agent is thought to be a strategy for dampening the amplification and dynamic aspects of the supply chain (Wikner et al. 1991).

An example will illustrate the calculation. Assuming perfect information (equal shares assumed and equal shares results), distributor1 estimates expected demand to be 8 (average customer demand is 8), and distributor1's actual demand from the downstream agents, that is, wholesalers 1 through 5, is indeed 8. Here, 8 = 5 × (8 / number of customers) = 5 × 8/5 and the distribution across the wholesalers is equal: (1.6, 1.6, 1.6, 1.6, 1.6) – see Figure 4. For the case of imperfect information (equal shares assumed but variation in the shares up to ±10% result), distributor1 estimates expected demand to be 8 (distributor1 correctly estimates total demand but incorrectly assumes it will receive the average customer demand), and distributor1's actual demand coming from the downstream agents, wholesalers 1 through 5, is, for example, in one realization, (1.76, 1.50, 1.62, 1.44, 1.68), where the sum is 8, but the share components deviate from the equal shares assumption of 1.6 by ±10%.

The experimental design setup is shown in Table 1.

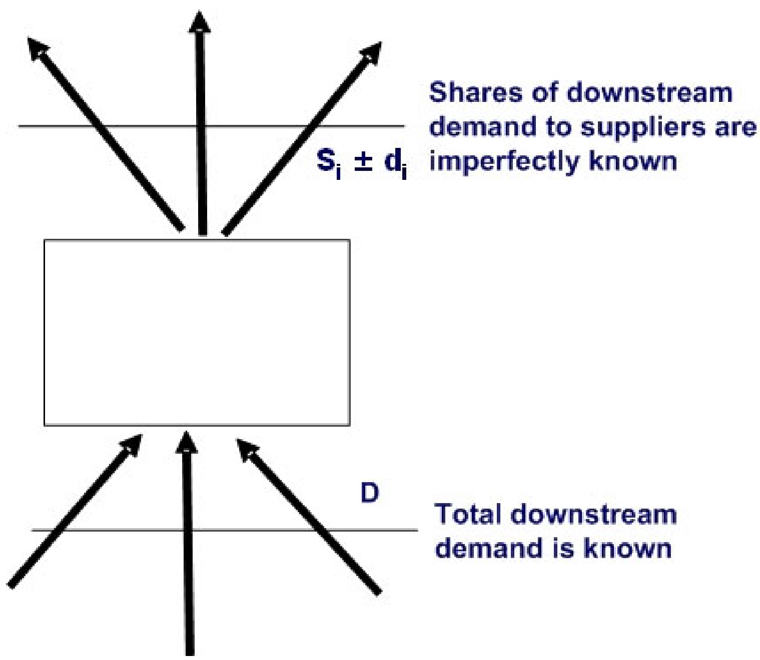

FIGURE 4 Model for imperfect information affecting agent

TABLE 1 Setup for information experiments

Experiment	Assumptions
Base Case	Expected demand based on weighted moving average
Experiment 1	Expected demand based on visibility on incoming orders Visibility extended over 1 to 7 periods
Experiment 2	Expected demand based on visibility of customer demand Variations in incoming order shares range from 0 to 100%

Parameter specifications for NBGS simulation: $\alpha_S = 0.5$, $\alpha_{SL} = 0.25$, $\theta = 1$,
$Q = 17$ (implies Desired Pipeline = Desired Inventory = 12 for all agents)

EXPERIMENTAL RESULTS: VALUE OF INFORMATION IN THE SUPPLY CHAIN

We simulate the NBGS for 1,000 periods for a network consisting of five agents at each stage. Figure 5 illustrates the agent supply network simulation at typical point in time, showing the strength of agent interactions in terms of orders and shipments. Figure 6 shows results of a single typical simulation run for the base case – inventory variation over time for agents at the distributor stage. These results are typical of the beer game simulation, showing chaotic behavior over large regions of the decision rule parameter space. Indeed, if the NBGS was reduced to a single column, from a supply network to a supply chain, the base case results would be identical to those produced by the original BGS set of simulation results.

Results of Experiment 1– Information Visibility

Figure 7 shows that the results of Experiment 1 support Hypothesis 1 but only to a limited extent. We can make the following observations:

1. As the scope of the information visibility grows larger and more information becomes available (as compared to the local information case) to the agent, there is an initial improvement as the system becomes more stable, and costs are reduced. The agent is able to buffer the incoming orders and shipments and quell the variability in its own ordering decisions.

2. The benefits of information visibility only accrue up to a point. At some point, the value of the next increment of information (marginal benefit) decreases and becomes negligible. This is due in part to the fact that as an agent looks further up- or downstream, agents intervene and transform the shipments and orders via their own ordering and shipping rules, and these transformations are not known to the other agents. For example, by the time an order is three periods away in the pipeline, that order has been processed by an agent in ways that could not have been anticipated by the agent originating the order. In effect, the transformation of the intervening agent devalues the value of the information. Several exploratory simulation runs not reported here have shown that this effect is exacerbated in supply chains that are highly dynamic to begin with, but is less problematic in stable supply chains. In stable supply chains, by definition, the transformations by agents are less volatile and more predictable. As Figure 7 illustrates, if we posit a cost for obtaining information, total information costs actually increase as more information is obtained.

We conclude that it is not enough to have access to more information for improving supply chain performance. It is a matter of how effectively the information is used. In this experiment, we used the original behavioral decision rules by simply feeding more information into them. However, it is likely that further gains could be made in supply chain performance by adjusting the decision rules or adapting the original decision rules to better exploit the additional information available.

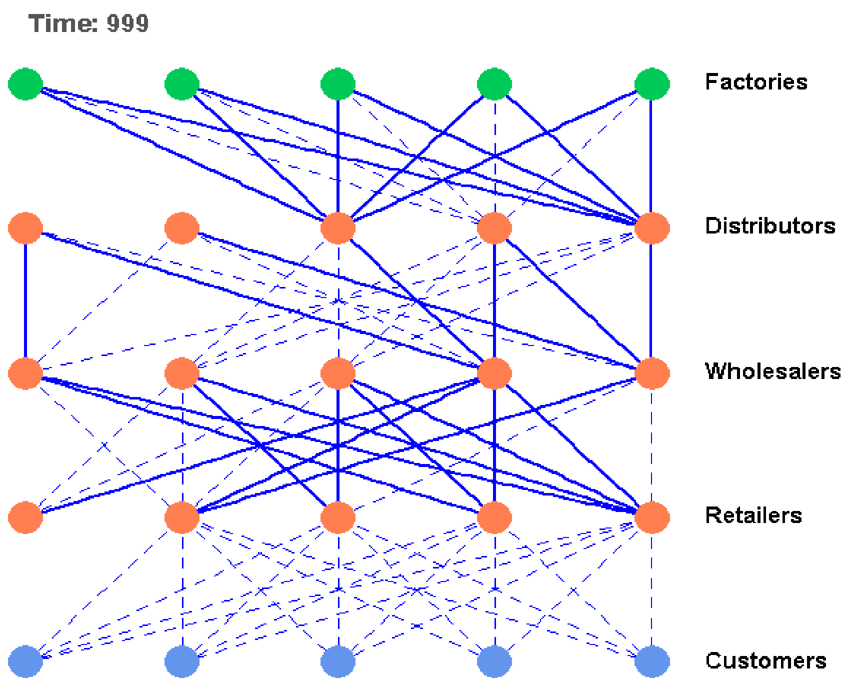

FIGURE 5 Snapshot of supply network simulation

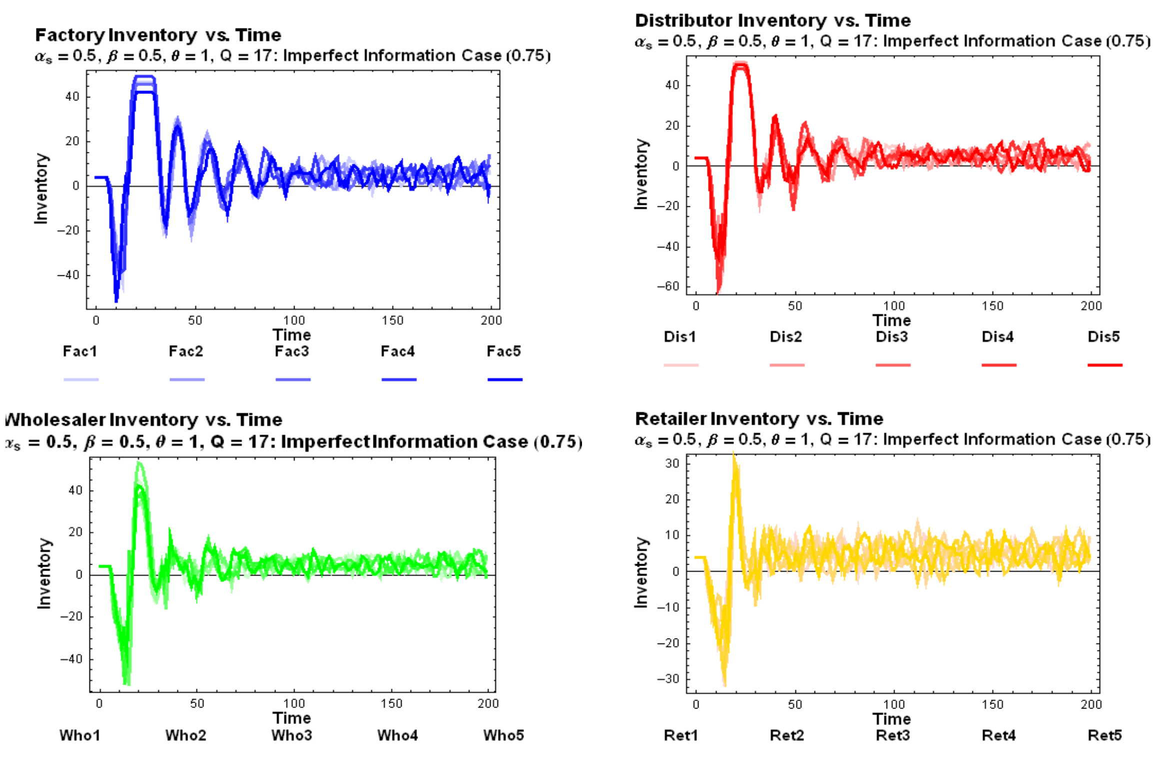

FIGURE 6 Simulation results: inventory levels over time for $\sigma_c = 0.75$

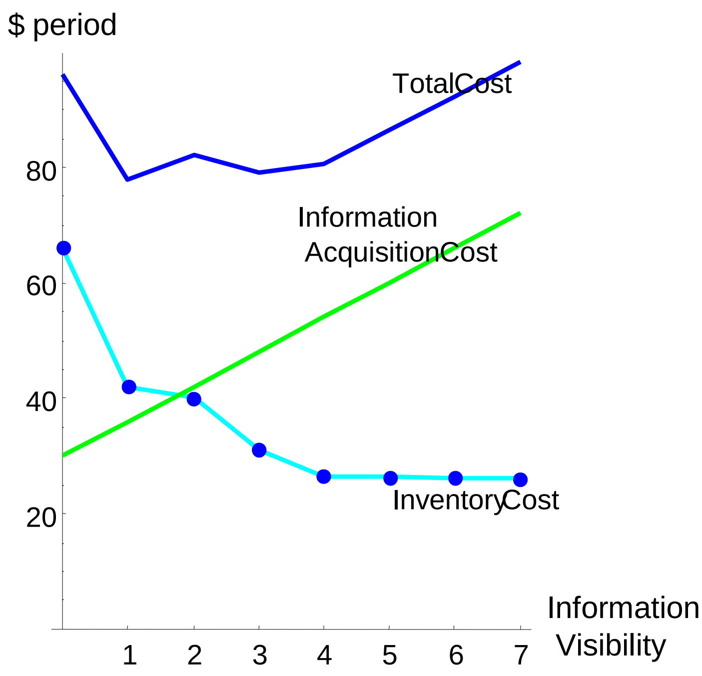

FIGURE 7 Cost vs. information reach, Experiment 1 – visibility

Results of Experiment 2– Imperfect Information

Figure 8 shows the results of Experiment 2 support Hypothesis 2 but only to a point. We make the following observations:

Perfect information, such as complete knowledge of final customer demand and the shares of customer demand accruing to each agent in a stage ($\delta = 0$), stabilizes the dynamics of the supply chain. Stabilizing the supply chain at positive inventory levels has the effect of greatly reducing costs.

1. As information becomes imperfect, that is, as $\delta \to 1$, in regions of the solution space that are close to critical behavior boundaries, supply chain costs tend to increase as information becomes more imperfect. In effect, the value of the information is overwhelmed by the larger features of system-wide behavior.

SUMMARY AND CONCLUSIONS

We have explored the value of information in the context of an agent-based model of a supply network. We relaxed the standard assumption in agent-based models that agents have access to only strictly local information. The first experiment showed that extending the “visibility” that an agent has over a broader range of the supply chain/network results in improvements but only up to the point that other agents’ decision process (transformations) intervene in the processing of the information. We conclude that it is not enough to have access

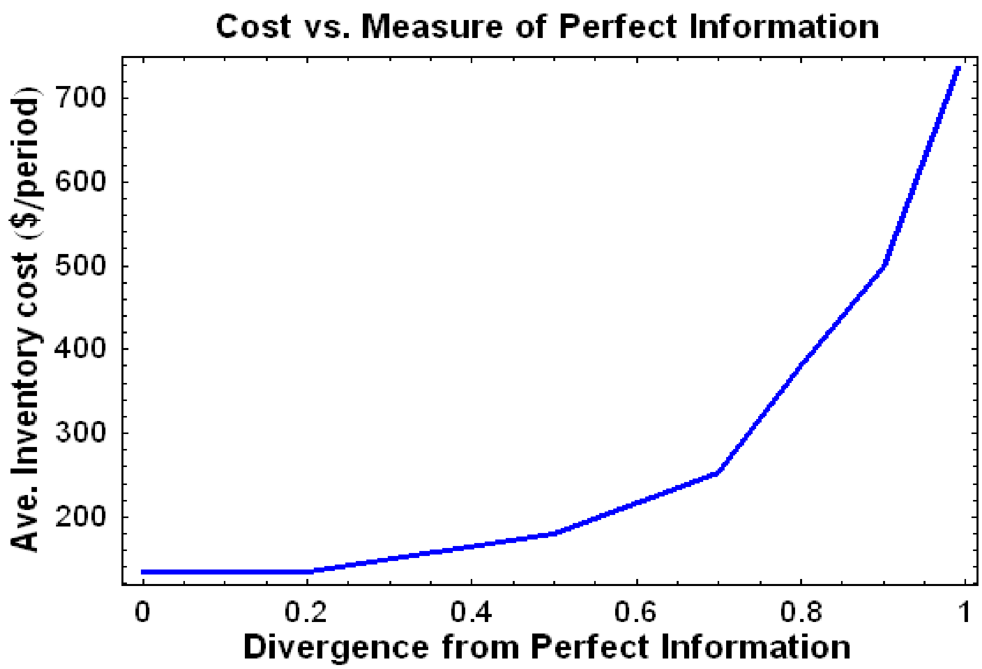

FIGURE 8 Cost vs. information accuracy, Experiment 2 – imperfect information

to more information for improving agent or supply chain performance. It is a matter of how that information is used. More information requires better decision rules that use the information most effectively. The second experiment showed that relatively small imperfections, or variations, in the information available to the agents can have drastic effects on supply chain dynamic behavior and costs. We conclude that the study of information in supply chain networks should necessarily include the study of how that information is used in the behavioral decision making processes employed for the agents. If more information becomes available, it is to be expected that this will motivate a co-evolutionary process of sorts between the adaptation of the information available and the behavioral decision rules that use this information.

ACKNOWLEDGMENT

This work is sponsored by the U.S. Department of Energy under contract W-31-109-ENG-38.

REFERENCES

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

Bonabeau, E., 2002, "Agent-based modeling: Methods and Techniques for Simulating Human Systems," *Proc. National Academy of Sciences* **99**:7280–7287.

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

Forrester, J.W., 1961, *Industrial Dynamics*, Cambridge, MA: MIT Press (Currently available from Pegasus Communications: Waltham, MA).

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

Lee, H.L., V. Padmanabhan, and S. Whang, 1997, “The Bullwhip Effect in Supply Chains,” *Sloan Management Review,* 93–102.

Mosekilde, E., E.R. Larsen, and J.D. Sterman, 1991, “Coping with Complexity: Deterministic Chaos in Human Decision Making Behavior,” in *Beyond Belief Randomness, Prediction, and Explanation in Science*, J.L. Casti and A. Karlqvist, eds., Boca Raton, FL: CRC Press, pp. 199–229.

Rasmussen, D.R., and E. Mosekilde, 1988, “Bifurcation and Chaos in a Generic Management Model,” *European Journal of Operational Research* **35**:80–88.

Sterman, J.D., 1987, “Expectation Formation in Behavioral Simulation Models,” *Behavioral Science* **32**:190–211.

Sterman, J.D., 1989, “Deterministic Chaos in an Experimental Economic System,” *Journal of Economic Behavior and Organization* **12**:1–28.

Sterman, J.D., 2000, *Business Dynamics: Systems Thinking and Modeling for a Complex World*, Boston, MA: Irwin McGraw-Hill.

Wikner, J., D.R. Towill, and M. Naim, 1991, “Smoothing Supply Chain Dynamics,” *International Journal of Production Economics* **22**:231–248.

EMCAS: AN AGENT-BASED TOOL FOR MODELING ELECTRICITY MARKETS

M.J. NORTH,* P.R. THIMMAPURAM, R. CIRILLO, C. MACAL,
G. CONZELMANN, G. BOYD, V. KORITAROV, and T. VESELKA
Argonne National Laboratory, Argonne, IL

ABSTRACT

Electricity markets are physically complex systems with interconnected pieces of equipment, such as generating units, transformers, and transmission lines. Deregulated markets, where the generation, transmission, and distribution functions are unbundled, create additional complexity. In traditional electricity markets, there is a centralized decision-making process, which will optimize the system. In deregulated markets, the market participants have conflicting interests and behave economically and strategically to enhance their objective such as profit, market share, etc. To understand and gain insight into how these markets participants behave, an agent-based modeling and simulation (ABMS) tool, the Electricity Markets Complex Adaptive System (EMCAS) has been developed. EMCAS is based on Repast, a software framework for agent-based simulation. Proper price forecasting is the key to successful implementation of the agents' strategies, which enables them to learn and adapt to changing market conditions. This paper presents the use of an agent-based modeling approach to simulate restructured power markets and demonstrate the agent learning mechanism.

Keywords: EMCAS, Repast, agent-based modeling, complex adaptive systems, electricity markets

1 INTRODUCTION

Several states in the United States and countries around the world are moving toward deregulating their electric utility systems. These electricity markets are physically complex systems with interconnected pieces of equipment, such as generating units, transformers, and transmission lines. Deregulated markets, where the generation, transmission, and distribution functions are unbundled, create additional complexity. In traditional electricity markets, there is a centralized decision-making process, which will optimize the system. In deregulated markets, the market participants have conflicting interests and behave economically and strategically to enhance their objectives, such as profit, market share, etc. To understand and gain insight into how these markets participants behave, an agent-based modeling and simulation (ABMS) tool, the Electricity Markets Complex Adaptive System (EMCAS), has been developed. EMCAS is based on Repast (Repast 2003), a software framework for agent-based simulation.

Deregulation creates several new market participants, including power brokers, marketers, and load aggregators or consolidators. As a result, a large number of entities will compete with each other. It is believed that the open markets create economic efficiency and lead towards lower consumer costs. These market participants have their own unique business

---

* *Corresponding author address:* Michael J. North, Argonne National Laboratory, 9700 South Cass Avenue, Bldg. 900, Argonne, IL, 60439; e-mail: north@anl.gov.

strategy, risk preference, and decision model. Decentralized decision-making is one of the key features of the new deregulated markets. To model these agents, EMCAS agent-based modeling techniques do not postulate a single decision-maker with a single objective for the entire system. Rather, agents are allowed to establish their own objectives and apply their own decision rules. The complex adaptive systems (CAS) modeling approach simulates agents that learn from their previous experiences and change their behavior when future opportunities arise. That is, as the simulation progresses, agents can adapt their strategies based on the success or failure of previous efforts. This paper first provides some background information on agent-based modeling. It then introduces EMCAS as a long-term deregulated market simulation tool and describes how agents are learning within EMCAS.

2 OVERVIEW OF ABMS

ABMS systems consist of a set of agents and a framework for simulating the agents' decisions and behaviors. These agents are allowed to interact with each other within the rules imposed by the modeling system under consideration. In a simple system, all agents may follow the same set of rules. In a complex system, such as electricity markets, the rules may vary by agent type. These agents have some global knowledge about the system and some private knowledge. The agents often make their decisions with incomplete and imperfect information and learn from their previous interactions with the system and other agents. Even in a simple system where all agents follow same set of rules, an aggregate system behavior can emerge. Such emergent behavior is difficult to predict with other simulation techniques such as discrete event simulation, game theory, artificial intelligence etc. An ABMS system lends itself to simulating and understanding complex systems with interacting agents. Several electricity market models have been developed (Bower and Bunn 2000; Bunn and Oliveira 2001; Petrov and Sheblé 2000). These models indicate the potential of ABMS for simulating electricity markets.

3 EMCAS APPROACH

The EMCAS framework consists of agents, interaction layers, and planning periods. The agents are market participants including generating companies (GenCos), transmission companies (TransCos), distribution companies (DistCos), demand companies (DemCos), consumers, and independent system operators (ISOs) (Figure 1).

These agents are highly specialized to perform specific tasks (Table 1). The physical equipment, such as generating units, transmission lines, and transformers, are owned and operated by these agents to maximize their own objectives with out any concern towards the overall system optimization.

The interaction layers are the several markets in which the agents interact with each other (Figure 2). Currently, EMCAS models (1) bilateral markets where GenCos directly interact with the DemCos for energy contracts and (2) day-ahead pool markets where they interact with the DemCos indirectly through an ISO by submitting bids. The pool market includes energy and ancillary services such as spinning reserve, non-spinning reserve, and replacement reserve.

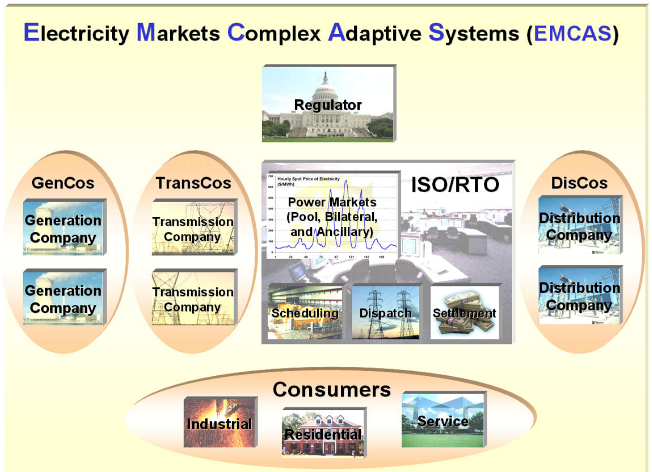

FIGURE 1 EMCAS agents

TABLE 1 Agent functions and interactions

Agent	Function
ISO	Independent system operator operates day-ahead market and dispatches the generating units
GenCo	Generating companies are unregulated owners of generating units; GenCos bid into the pool market and engage in bilateral contracts with the DemCos
DemCo	Demand companies are unregulated entities that act as load aggregators; they bid into the pool market and engage in bilateral contracts with the GenCos
TransCo	Transmission companies are passive owners of the high-voltage lines connecting the generating units and transformers
DisCo	Local distribution companies are passive owners of the low-voltage lines connected to the consumers.
Consumer	Consumers are the final recipients of electricity; they interact only with the DemCos
Regulator	The regulator sets the market rules under which the market participants operate

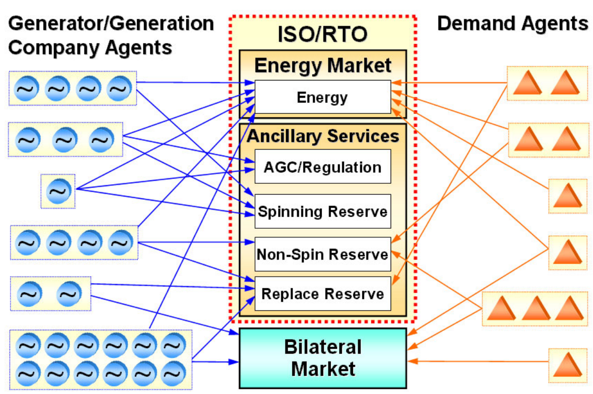

FIGURE 2 EMCAS layers (markets)

The planning periods are the range of time scales over which the agents make their decisions (Figure 3). The planning periods include hourly dispatch, day-ahead, week-ahead, month-ahead, year-ahead, and multi-year. Over longer time scales, economic behavior decisions dominate and over shorter time scales, physical laws dominate. For a detailed description of the EMCAS modeling approach, see North et al. (2002a,b) and Veselka et al. (2002).

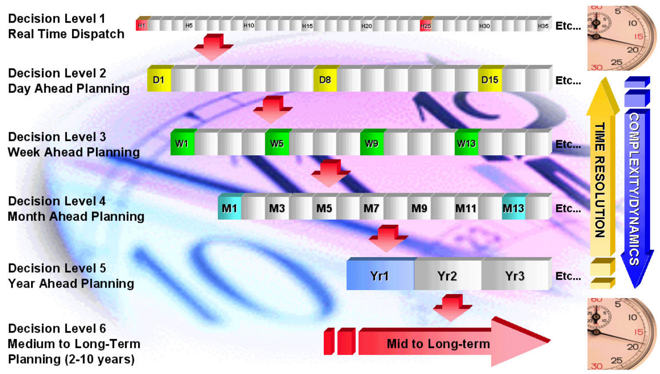

FIGURE 3 EMCAS planning periods

4 AGENT LEARNING

In electricity markets, the agent learning occurs through unit commitment; forecasting load, generation, and locational marginal price (LMP); and the bid outcome. Proper price forecasting is the key to successful implementation of the agent's strategies. The GenCos commit their units for the next day and determine the bid price based on forecasted prices. If the forecasted prices are below the marginal costs, the GenCo may withhold the generating units; if the forecasted prices are too high, the GenCo might bid too high into the market. Both actions may result in the loss of market share and profits. In EMCAS, the GenCos are provided several price-forecasting techniques. In a simple technique, agents forecast the next day's price as an average of the previous five-day LMPs (adjusted for weekday and weekend).

Recently, an adaptive regression forecasting technique has been implemented using an open-source package, OpenForecast (2003). With the adaptive regression, agents have access to several regression models, such as single variable linear regression, single variable polynomial regression, multi-variable linear regression, and moving averages. This adaptive regression enables the agents to dynamically select an appropriate model and update the regression coefficients based on the input data. In EMCAS, the agents are enabled with the 30-day rolling historical system reserve margins and bus LMPs for use with adaptive regression. The actual and forecasted prices based on adaptive regression and five-day averages for a bus are shown in Figure 4 and listed in Table 2.

As can be seen from Table 3, the higher R-square value (0.96) from adaptive regression significantly improves the learning capability of the agents. In this example, only system reserve margin is used as a predictor of prices. Additional testing is planned for using other variables as independent variable for adaptive regression price forecasting.

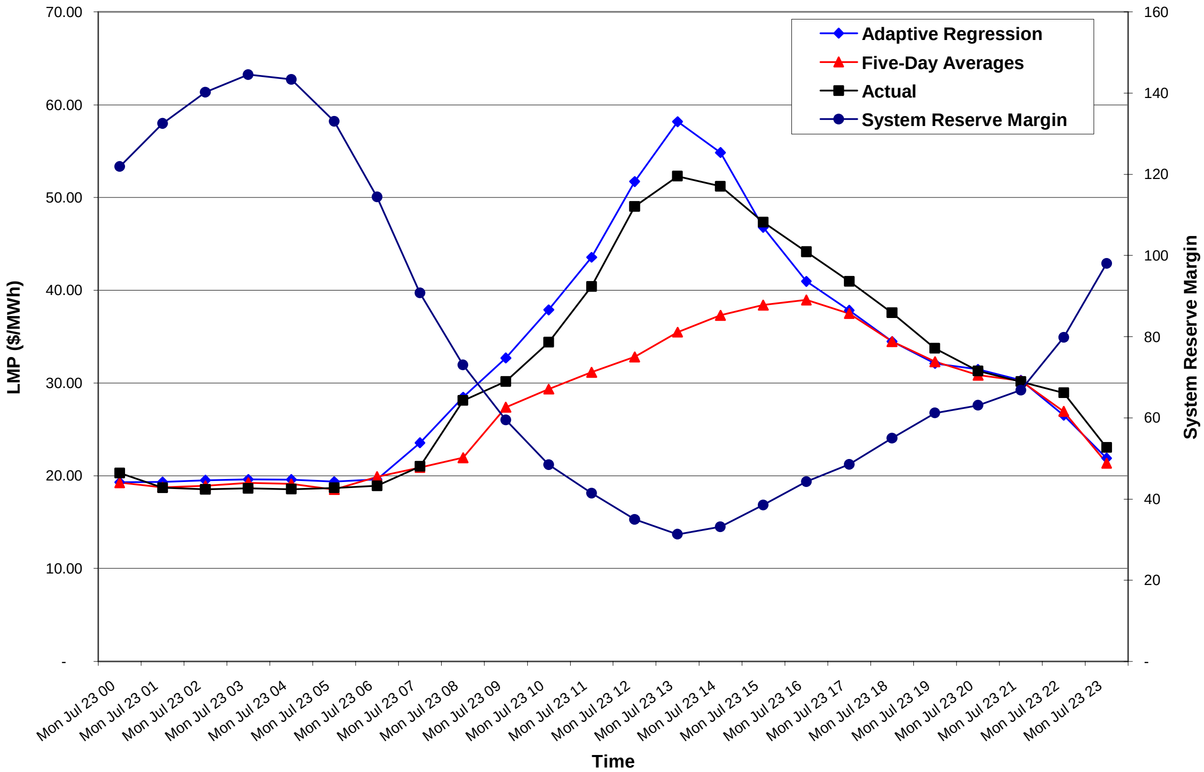

FIGURE 4 Agent learning using adaptive regression

5 REFERENCES

Bower, J., and D.W. Bunn, 2000, "A Model-based Comparison of Pool and Bilateral Market Mechanisms for Electricity Trading," *Energy Journal* **21**(3), July.

Bunn, D.W., and F.S. Oliveira, 2001, "Agent-based Simulation: An Application to the New Electricity Trading Arrangements of England and Wales," *IEEE Special Issue: Agent-Based Computational Economics*.

North, M.J., G. Conzelmann, V.S. Koritarov, C.M. Macal, P. Thimmapuram, and T.D. Veselka, 2002, "E Laboratories: Agent-based Modeling of Electricity Markets," *Proceedings of the 2002 American Power Conference*, Chicago, April.

North, M., C. Macal, R. Cirillo, G. Conzelmann, V. Koritarov, P. Thimmapuram, and T. Veselka, 2002, "Multiagent Modeling of Electricity Markets," in *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, Chicago, Oct.

OpenForecast 2003, Available at http://openforecast.sourceforge.net/.

TABLE 2 Example data used to illustrate adaptive regression price forecasting

Date	System Reserve Margin	Actual	Adaptive Regression	Five-Day Averages
Mon Jul 23 00	122	20.30	19.29	19.28
Mon Jul 23 01	133	18.72	19.34	18.77
Mon Jul 23 02	140	18.53	19.53	18.91
Mon Jul 23 03	145	18.66	19.61	19.24
Mon Jul 23 04	143	18.57	19.59	19.13
Mon Jul 23 05	133	18.69	19.35	18.51
Mon Jul 23 06	114	18.92	19.62	19.91
Mon Jul 23 07	91	21.02	23.57	20.90
Mon Jul 23 08	73	28.12	28.47	21.96
Mon Jul 23 09	59	30.17	32.71	27.38
Mon Jul 23 10	48	34.41	37.89	29.36
Mon Jul 23 11	41	40.41	43.55	31.17
Mon Jul 23 12	35	49.03	51.74	32.82
Mon Jul 23 13	31	52.29	58.17	35.48
Mon Jul 23 14	33	51.22	54.83	37.31
Mon Jul 23 15	39	47.33	46.80	38.41
Mon Jul 23 16	44	44.16	40.96	38.97
Mon Jul 23 17	49	40.97	37.85	37.49
Mon Jul 23 18	55	37.56	34.47	34.46
Mon Jul 23 19	61	33.73	32.10	32.29
Mon Jul 23 20	63	31.30	31.47	30.87
Mon Jul 23 21	67	30.14	30.31	30.23
Mon Jul 23 22	80	28.95	26.52	26.95
Mon Jul 23 23	98	23.04	21.91	21.36

TABLE 3 R-square for adaptive regression and five-day average model

Model	R-Square
Five-Day Averages	0.61
Adaptive Regression	0.96

Petrov, V., and G.B. Sheblé, 2000, "Power Auctions Bid Generation with Adaptive Agents Using Genetic Programming," *Proceedings of the 2000 North American Power Symposium, Institute of Electrical and Electronic Engineers*, Waterloo-Ontario, Canada: Oct.

Repast, 2003, Available at http://repast.sourceforge.net/.

Veselka, T.D., G. Boyd, G. Conzelmann, V.S. Koritarov, C.M. Macal, M.J. North, B. Schoepfle, and P. Thimmapuram, 2002, "Simulating the Behavior of Electricity Markets with an Agent-Based Methodology: The Electricity Market Complex Adaptive Systems (EMCAS) Model," *Proceedings of the 22nd Annual USAEE/IAEE Conference*, Vancouver, Oct.

DISCUSSION:

FINANCE AND MARKETS

(Friday, October 3, 2003, 10:00 a.m. to 12:30 p.m., Session 2)

Chair: *Charles M. Macal, Argonne National Laboratory*
Discussant: *Ed MacKerrow, Los Alamos National Laboratory*

Agent-based Modeling of Lottery Markets

Charles M. Macal: This is the session on Finance and Markets. And there is something of an applications bent to it.

Our first speaker (P.M. Beaumont) we have not seen today. We had some early warnings that his computer wasn't working or something last week and that he was having trouble getting the results out, so we're not exactly sure what happened. But, in any case, he is not here and we have a little more relaxed timing now.

There's been another slight change to the program in terms of the discussant. The discussant for this session is Ed MacKerrow, from Los Alamos National Lab. The format will be as follows: There'll be presentations by the speaker, after which there'll be perhaps time for questions, brief questions. And then after all the speakers have given their papers, we will then turn things over to the discussant, and the discussant will say a few words and perhaps cajole the speakers to say a few more words and answer a few questions, and have a more general discussion then with the audience.

So, having said that, here's our first speaker, Shu-Heng Chen, of National Chengchi University, who will be speaking on agent-based modeling of lottery markets.

[Presentation]

Macal: We have a question in the audience, from Ed. We'll get the microphone up to you.

Unidentified Speaker: It seems like as you increase the tax rate, you lower — and I'm probably making some assumptions, because I'm not a lottery expert — the total amount of money available for winnings, right?

Shu-Heng Chen: Yes.

Unidentified Speaker: What information does the market or the public have as far as what is the available pot of money? That might not be known to them. But instead what they might see is that, Steve just won $36 billion from some Lotto winning. And so what I'm wondering about is, have you thought about or looked at the integrated amount of total winnings available versus the bucketing? For example, you could just say there's going to be two prizes given out that are really big and they're really publicized, but if you add those two prizes together, it might be 10% of the total possible winnings, as if there were lots of smaller prizes. Have you thought about that effect?

Chen: No, we haven't conceded to that individual labor. But, yes, this will be a direction to think about. Thank you.

Unidentified Speaker: I have a quick question. You have a single population, am I right? There's a single population GA where all the agents are simultaneously voting rather quickly. An alternative is to have each agent have his own GA. So having a single population, could that be one of the reasons that you have rather fast learning?

Chen: Well, that's a good question, from the technical aspect. Originally, we do consider both versions; that is, single population GA and the multi-population GA. Multi-population GA really will be a little difficult to interpret. The single agent has backed up his decision with the whole population of ideas. And also, in implementation it will be difficult to evaluate all these rules behind him.

So when we run, when we choose a particular version of … multi-population GA, and we get a result quite unstable. So actually, we stopped at that point and did not investigate it further. What … signify also what has been originally found, again, in the literature; that is, a single population GA and a multi-population GA give different results.

And the reason we do not pursue this line further is because we consider how we are actually being exposed to the question how complex we are going to model our agent. Thank you for the question.

Unidentified Speaker: Have you looked into the incentive for a higher tax rate, in the sense that money is going to projects? I know it's marketed a lot that you're buying a lottery ticket; part of it is the winning, the other part of it is, "Hey, this money is going to all these nice projects." I know in Europe it's marketed that way. So is that a factor at all?

Chen: No, we didn't consider that altruistic behavior, no. The people here — it's quite simple — they participate with the hope that they can change their life.

Unidentified Speaker: I have a question, and it's going to display my ignorance of probability. Let's say you have a lottery where you could pick six numbers, say, from 1 to 50, and they're uniformly distributed. In infinite time, if you just tallied up what the winning numbers were, I think you'd have the uniform distribution; in infinite time.

So I've always wondered, what if you could take your model and have an agent just say, "Okay, you know, number 3 has not been picked ever in the last 10 years of simulation, so I'm going to bet on 3 and number 7." Have you thought about things like that, where you have these sort of, I don't know what you call them, farmer rules or whatever?

Chen: If you look at the so-called "expert," quote/unquote, on the lottery market, what it would actually be is something more fantastical than what you said, because they actually apply like data mining or machinery stuff. And actually in some magazines you can see how they recommend a number and actually correlate different lottery markets from what happened in the North Ireland to Taiwan. They can somehow get this crazy correlation.

And so actually that is what I'm introducing in the dimension of conscious selection. But my conscious selection in this moment is quite simple, okay? They have no … sophisticated machinery technique stuff. They are just randomly say, "Maybe these few numbers are good, so this time I'm going to try them." Next time, they change some numbers. But in the end, as I show in the learning curve, they are gradually being discouraged, that actually there is nothing to learn from this series of lottery numbers, to quite a large extent, but it's not entirely convinced the market to give up the learning. But it's a quite large proportion of the people will now try to learn anything from there.

Unidentified Speaker: I had one question. Have you looked at the effect of the population size, because things such as herd effect, etc., might start coming up at much larger populations, or at smaller ones you would probably just see more homogeneity and no heterogeneity emerge at all.

Chen: Right. It's a good question.

This is our agent-based simulation, and we have some natural … trend, which is the number of agents in our simulation. Over here, we take 5,000 agents. But of course, even the smallest lottery market in the country, in this world, is 5,000 people. So what we're doing here is trying to fine-tune the market size parameter; that is, the number you're going to pick, and from how many numbers.

So what we are trying to do in our presentation is to fine-tune the population size — I mean, the number of customers, and the market size so they can somehow give you a similar degree of the heartiness to win the lottery as corresponding to the real markets….

Unidentified Speakers: But if you were to look at, say, the herd effect, which is more independent of this probability of winning — rather it is a contagion sort of an effect, where you get a bunch of people — there are certain beliefs evolving into a certain direction at a certain point of time. So if you look at the rollover rates and the effects, you do see a spike in the sale, but that spike is not always proportional to the rollover. In some cases you see a much larger spike — say five times, ten times — which is rather a sudden convergence of a belief and not necessarily rational or probabilistic in some manner.

Chen: Right. I just don't have time to show the result we have, because actually we have some statistics over the six lottery markets being actually studied. Although you see that … happen in five of the markets, the degree, for example, in terms of the R squared, changes quite a lot over different markets. So I perfectly agree with what you said, yes.

Effects of Global Information Availability in Networks of Supply Chain Agents

Charles M. Macal: Okay, I'm the next speaker. My name is Charles Macal, from Argonne National Laboratory, and my co-author on this paper is Michael North…. I'm going to talk about emergent social structures arising from networks of interacting supply chain agents, or at least the presentation will be in along those lines and in that direction.

[Presentation]

Macal: Any questions?

Unidentified Speaker: Just in terms of speculating about those final graphs, it strikes me, in thinking about the real world being similar to some of the wildness in this model, to wonder whether in fact what's been going on with the just-in-time methods and so forth that have been innovated in the past decade. What these folks are doing is in fact seeking these points of very low cost by trying to have systems that are so highly efficient that they actually can hit those points, but the cost that they're bearing of doing that is, of course, that instability, that if things don't perform exactly as engineered, they can be in rough shape. Perhaps some of the things, some of the outages we've had in this or that piece of infrastructure lately, may be examples of the cost of highly efficient, but nonrobust, solutions.

Macal: Yes, I think that this kind of analysis and the prospects that it offers are highly suggestive that the real world may be poised at areas that are what we would call unstable, if we knew what the other alternative worlds around us were. And certainly the notion of zero inventory is attractive. I worked on a project once with DARPA where the whole idea of the project was to have zero staging. The trucks would come from the fort and roll out to the ship and not have that buffer area for staging. And it quickly became apparent that maybe that's not such a good idea, that buffering is good, even though it is costly.

So the only other thing I would say is that the challenge is to take a model like this and to really get it to a point where you could say something about the real world, and that people would give some credibility to it. Okay, so we have suspicions the real world does behave like that, yes.

Unidentified Speaker: Okay, I have a question about the real world. I've seen the beer model before, and so to some extent this doesn't address your talk, although I liked your talk a lot, but it's something you might be able to do with your model now that you have it, is the problem with this decision process that it's already optimized for — it seems like that step function is a very unnatural situation. So I'm wondering why, if people as you say behave so badly and go into this big chaos, you gave it a more normal sort of gradual transition in the supply. Does that actually create more chaos because it's constantly changing, or does that actually stabilize out because the function's really optimized to deal with that kind of situation?

Macal: Well, I'll give you the answer that I know to be true, and for the beer model, in particular. It will always be chaotic, no matter what the demand is, because the demand is exogenously given, and the structure of the equations is such that the chaos is in the variable part, not in the exogenous demand part. But of course we could explore how demand affects things through the simulation.

Unidentified Speaker: [Unintelligible question]

Macal: I'd actually say one, and sort of offer an extension of that answer. I'd say that on a certain sense, yes, but also a certain sense no. The actual structure of those equations was determined empirically, based on quite a few hundred business student runs. Now, whether or not you consider business students a sample of normal human beings is debatable. I should say that quietly, since we're in the business school now.

If you looked at the equations for a while and sort of see what's going on, though, the way I put it into simple terms is that people are systematically unable to estimate second derivatives properly. It seems to be a very hard thing for people to do. They get first derivatives right; they understand things are going up, they understand things are going down, but they don't understand that things are going up and slowing … or going down and slowing. They don't seem to ever quite capture the slowing part, which would be a second derivative. And so that actually seems to be a common sort of perceptual error that people do seem to have. And so I think it's grounded in a little bit more empirical observation or phenomenological observation. And so in that sense I think that you see this exact type of behavior in real systems. The real estate markets do this all the time, you know.

I can't say why people have trouble with second derivatives, but that seems to be the problem.

Jesse Voss: It seems to me that you're showing that there is an exponential effect that's associated with the linear increase in the amount of imperfection in the information. If you had an agent-based model simulation of a population environmentally situated and they were demonstrating exponential population growth that was out of synch with the record of the data that it's in, do you think that inclusion of a simple supply chain dynamic like what you have, so you've got an exponential population growth and then you've also got possibly exponentially growing effects of imperfect information that might be linear, could that linearize population growth so that it was flatter if there was a relationship between population growth and effectiveness of the supply chain?

Macal: Population growth wasn't an aspect of the model.

Voss: No, but I'm thinking about what I could use this for …

Macal: Oh, okay. Well, I guess, in terms of the imperfect information thing, I believe that — and it could probably be worked out analytically to be supported — the variances somehow are multiplicable, so the variances actually are amplified because they're being multiplied, even though everybody gets the mean right, but the tiny variance gets amplified by the time it gets reflected up and comes back down.

So if population was changing or the numbers of agents was changing on top of everything, I'm sure the system would exhibit the same kind of behavior.

Voss: Right, and I was just going to add that that's the bullwhip effect, where people are moving up the chain.

Macal: Yes, the bullwhip, right. Additional questions?

Unidentified Speaker: In this whole thing, the consumption time period and the decision-making time period are exactly identical. Consumption happens in one period, not in a staggered form. What if the retailer was to make decisions every five periods, while consumption happens every period? The distributor makes decisions on different time periods. If the decision-making is staggered out, would you now see some of this chaos go away?

Macal: I guess that remains to be investigated. I don't see that it would necessarily go away, just thinking about the mathematics of it all, even if you reduce the timeframe to, let's say, one-fifth for decision-making.

Unidentified Speaker: No, I'm talking of the reverse, where the consumption happens every period, the retailer makes decisions every five periods, and the distributor, on the other hand, makes her decision every 10 periods.

But all the information in between, up to the period before that, is available to them. So a retailer gets to see four periods before he decides what to order every fifth period, for the next five periods altogether. The distributor, on the other hand, has information for nine periods and makes a decision for the next 10 periods. So if you were to take a financial market, it's a portfolio-holding kind of a scenario, where you are trying to maximize a certain return or hedge out a risk for longer, different terms altogether. On the other hand, the production end is planning at a much, much larger scale.

Macal: Well, I don't know how to answer the question without running the simulation. But I think that it's a good lead-in to Mike North's talk on these kinds of markets, and speculations relative to near-term spot market type things versus longer-term perspectives.

So with that I will turn things over to Mike for the next presentation.

EMCAS: An Agent-Based Tool for Modeling Electricity Markets

Michael North: Thanks. I would say that probably I would think that you're right [in reference to the last comment for Macal's presentation]. They probably would have a smoothing effect. But at the same time, it's something we'd want to test, rather than simply speculate on now.

[Presentation]

North: So now I'd like to move on to questions. If people have any thoughts, comments, I'd be glad to answer them.

Unidentified Speaker: It seems like quite a complex model. I just was curious if you can give us some sense of how many lines of code, execution time, how many people to build it, some sense of its scale.

North: I would say it's actually been under construction for over three years now, I believe. Probably, tens of thousands of man-hours have gone into it. It's on the order of 150,000 lines of code, somewhere in that range, although I'd have to get a counter out to really know. It's growing by the day, because we're adding a series of new features; one interesting example is a need for something called phase shifters. It turns out that there's these moderately small devices that are in place inside the various locations, or particularly around Chicago, and we originally did a solution of this system to try to run it inside, but you can't get power through Chicago was the conclusion we were finding when we first ran the model. Then we talked to the regulators who said, "Oh yeah, we've got those phase shifters hidden inside the system." You put those in, and now the system's solvable.

And so these are the types of things we find. As we're moving on in each study, we have new things we need to work with. And so I would say, probably, that the effort is 10 person-years, something in that range.

Unidentified Speaker: You mentioned the nuclear as having different, I suppose, different physical capabilities. Do you have heterogeneous producers in that sense, in the sense of coal and various other forms of generation? What would be the impact of them, or renewables?

North: That's a good question. The answer is, yes, we definitely capture differences in generation capability. One thing that differentiates those ancillary services markets is that only certain generators are physically capable of responding to any one of those markets. And so only natural gas-fired units and a few other special types of units can get into the fastest markets. And so there's a lot of variation in those markets.

Renewables I think are a very exciting area. It's something that we're very interested in, in fact. So that's some of the discussions we're having, to model renewables. For us it's a very exciting thing, and there's a lot of potential.

Renewables are very interesting, though, because they have inconvenient physical characteristics in a lot of cases. I mean, they're good for the environment, I think they're a very positive thing, but they often are intermittent and they have other factors, which our model can capture.

Unidentified Speaker: In fact, as part of the second phase of the project, which will start up early next year, we will look at the impact that renewables might have on some of the market power that we might find under regular conditions. So there is talk about developing wind farms in certain locations in Illinois, there is biomass production that could be used, and other things that are done for the Illinois Clean Energy Association. Phase II essentially takes a look at specifically those issues.

Unidentified Speaker: Are you trying out different market rules? You didn't mention that in your presentation.

North: That's one of the things that people are very interested in. This model's designed to capture exactly that, where you can vary, say, pricing policies. And a good example would be whether or not you can get pay-as-bid. So that means if your bid is accepted, you get paid what you wrote on the contract, or a market clearing price, which means that you could get paid that amount or more. And so these are the types of things people are interested in.

It depends very much on the study we're doing, though, because in some cases people have a clear idea of what they want, you know, what the rules are likely to be, and they want to know about the physical operation. Could renewables help? In other cases, people are extremely interested in determining what types of rules would be effective.

Unidentified Speaker: But have you done that yet?

North: Oh, yes. We do have examples where we have looked at, say, market clearing price versus pay-as-bid.

Unidentified Speaker: And the results?

North: The results? Actually, I have to be careful how I answer this question, because it depends. It turns out that there are varying outcomes that you can get depending on the structure of the system. A lot of it also depends on what you look at. With pay-as-bid, sometimes you can get lower prices, but you can drive everyone out of business. And so it's complex; there's no one answer as to what's better and what's worse. It seems to depend.

Unidentified Speaker: I'm asking because I'm seeing this as the application that could put agent-based simulation on the map. It's a very big, multi-billion dollar question. It's something that is not being well addressed.

North: Oh, yes. And it's something that we're extremely interested in. In fact, this model *does* capture those and can answer those types of questions. But beyond that, the answer depends a lot on the system. The system characteristics determine the policy that's affected.

Unidentified Speaker: Is the model set up so that you can, in the future, have individual agents learn and broaden their bidding behavior to better exploit holes in the market rules?

North: Well, that's actually part of what the system does now. I think that it's always possible to improve these areas, because there's an enormous amount that can be done with machine learning and all sorts of other techniques. But even now the agents do probe to see, to try different pricing structures, to see if they're going to be effective.

And in fact we do have another example. I didn't present that here, but we had a simple market structurally. But then we allowed the agents to adjust not only their pricing amounts, but their pricing patterns. And over time they learned to evolve. And I put "learning" in quotes, because machine learning's always a debate as to how much is learning and how much is structural. But the idea is that they adjust their bidding patterns and essentially learn what we call hockey stick bidding, which is an interesting thing we hadn't really thought about until we had seen it come out of the system. And with the hockey stick bidding basically you have extremely low prices, or essentially production costs for virtually all of your capacity, except at the very end you have a very high upturn at the very edge. And so essentially what happens is you have very low risk, because virtually all of your capacity is accepted at close to the economic value, and so you won't lose money, and you may make a little bit of someone else's at a slightly high price. But in those rare times when there's an outage, you can make an enormous amount of money, because you have that extra bid, and it drives up the price, if it's a locational marginal price or you get paid the clearing price. Then you can end up with a huge payoff, kind of a lottery win.

So that's an example of something that came out, that we were able to evolve out of the model. And so the answer is, yes, we do have capabilities for that now, and we're expanding those capabilities in the future.

Do we have other questions?

Joanna Bryson: This is also more of a comment, so I'll be very quick, because I'm curious to hear his, but it came back out of the renewables.

One of the things that's really noticeable, the difference between American energy markets and some of the European ones, is that in Europe a community can decide to put up its own windmill and then supply its own energy and perhaps sell it back in. Obviously, that would be something that would be a huge thing to agent-based modeling, because you have the autonomy of new generators, which is a policy thing I know the U.S. is worrying about right now.

North: Yes. And actually many state regulators are very interested in that as well. There's an issue of microgeneration as well, and the idea is people actually are beginning to install home-scale generators out in the backyard, usually natural gas-fired. And what does that do to a system? You know, it's an interesting question. You're moving emissions, you know, from usually a relatively isolated area back into the cities now. But at the same time you're also providing relief to some of the transmission problems.

Here's a good example of how in America, the situation's reversed. Instead of deciding to install things, we're deciding *not* to install things. An example is California, where they were going to put in what I guess would have been the world's largest Internet server farm. About 50 megawatts, if I remember correctly, of power was required. They had two referendums on the ballot: "Do you want that server farm?" and "Do you want a generator to support it?" The people said yes to the server farm and no to the generator, and then wondered why the lights went out.

Brian Pijanowski: You probably addressed it throughout your talk, but one of the things I would be thinking about is modeling climates and land uses. I mean, have you thought about those kinds of exogenous factors. Or do you have surrogates for them in your model? I assume you probably get asked a lot, especially about climate change. But how are you addressing it in the model?

North: Well, I can say this isn't a climate model specifically. It's a very interesting question. In fact, we do have other obviously unrelated models now that do capture some issues of climate change. And it's something that our division has been very interested in for many years, and we've actually done quite a bit of modeling in that regard.

I'll say that the current … model does not directly look at emissions in these types of things. They can calculate what the emissions would be, and you could use that, although that's not currently something that we're focusing on, although I think with the renewables, though, that's going to become a bigger question.

And there has been some discussion — we have a couple of mesoscale meteorological models. You know, that's a small-scale weather, basically. And we have plans that we've been putting together to integrate MCAS and this mesoscale model, and so it's something that we have not done yet but we are very interested in doing in the future. I feel there's a lot of potential. So yes, definitely.

Invited Speaker:

R. Keith Sawyer

ASSESSING AGENT COMMUNICATION LANGUAGES

R.K. SAWYER,* Department of Education, Washington University

ABSTRACT

In artificial societies, macrostructure emerges from models of large numbers of agents, their connections, and their repeated communications and interchanges over time. Although interaction is a fundamental aspect of artificial societies, the communicative formalisms implemented in them have been radically simple compared with sociological theories of communicative action. Empirical studies of human communication resulted in the rejection of speech act theory by the early 1980s; yet the most widely used agent communication languages (ACLs) continue to be based in speech act theory. This paper draws on empirical studies of conversation by linguistic anthropologists and sociolinguists to show that ACLs are empirically inaccurate models of human communication, and suggests future enhancements to result in ACLs that are better able to simulate emergence in social groups.

Keywords: Agent communication language, speech act theory, artificial society, conversation analysis

1 INTRODUCTION

Within the last decade, a new social simulation technology — multi-agent systems (MAS) — has emerged from computer science. Beginning in the mid-1990s, MAS technology has been used to simulate human societies by anthropologists, sociologists, economists, ecologists, and urban planners; these simulations are called agent-based social simulations, multi-agent based simulations, or *artificial societies* (Epstein and Axtell, 1996; Sawyer, 2003a). Several sociologists have argued that this technology is not only of methodological interest, but has the potential to contribute significant theoretical insights to foundational sociological questions (Carley, 2000; Gilbert, 1999; Gilbert and Troitzsch, 1999; Macy and Willer, 2002; Sawyer, 2003a). In particular, some sociologists have argued that artificial societies can contribute to our theoretical understanding of the relation between the individual and the collective, known as the *micro-macro link* (Alexander, et al., 1987; Barnes, 2001). In MAS, the relation between the individual, or "agent," and the collective — the emergent macro behavior of the system — is the core concern of the paradigm (Conte, et al., 2001; Gilbert, 2002; Sawyer, 2001b, 2003a). Developers of MAS technology are computer scientists, not sociologists, and none of their work uses the phrase "micro-macro link." Yet they are encountering longstanding unresolved issues in sociological theory: multi-agent programmers are developing a new science of the micro-macro link.

Artificial societies focus on three distinct sociological phenomena: the model of the individual, the model of the communication language, and the observation of emergent social phenomena. These three phenomena — individuals, symbolic communication, and macro

* *Corresponding author address*: R. Keith Sawyer, Department of Education, Washington University, Campus Box 1183, St. Louis, MO 63130, e-mail: keith@keithsawyer.com.

properties of social systems — have likewise been the central concerns of sociological theory. The sociological study of the individual has focused on agency and action theory; the sociological study of symbolic communication has been the purview of microsociologists such as symbolic interactionists and conversation analysts. Most sociologists are also centrally concerned with the emergent macro behavior of social systems; structural sociologists study properties of collectives without necessarily being concerned with a model of the individual participants nor of their interaction (Blau, Mayhew, Black), and methodological individualists study how macro properties emerge from individual actions (Coleman, Homans).

Increasingly since the 1980s, sociologists have developed hybrid theories that incorporate both the micro and the macro level; this task has been called a "third phase" of postwar sociology (Alexander and Colomy, 1990, p. 43), and several theorists believe it is the central focus of contemporary sociology (Alexander and Giesen, 1987; Archer, 1995). However, these contemporary hybrid theories generally lack a theory of symbolic communication (cf., Rawls, 1987; Ritzer and Gindoff, 1992; Wiley, 1988). Rather, most micro-macro theory attempts to relate individual action and macrosocial structure without theorizing a mediating level of symbolic interaction. Theories of the micro-macro link and theories of symbolic interaction have proceeded independently and have rarely been integrated (but see, Collins, 1981; Ellis, 1999; and Ritzer and Gindoff, 1992). A great deal of sociological theory neglects symbolic interaction entirely, from the canonical syntheses of Weber and Parsons, to contemporary hybrid theories by Alexander, Giddens, and Archer.

Most sociologists assume that communication is epiphenomenal — that it has no causal consequences, either for emergent macro phenomena or for individuals. Instead, the ultimate causal forces in social life are either, for the collectivist, large macro patterns (race, class, gender, educational level, network connections) or, for the individualist, rational actions taken in the context of pairwise game-like encounters. Thus, both of these opposed camps agree in their implicit assumption that communication is of only marginal concern to the sociologist (cf. Rawls, 1987; Ritzer and Gindoff, 1992).

I argue that full explanation of the micro-macro link requires a focus on communication processes and how they contribute to micro-to-macro emergence (Sawyer, 2003b). Of course, interaction is at the center of the sociological tradition associated with Simmel, Cooley, Mead, and the Chicago School of symbolic interactionism; however, this tradition did not directly address the micro-macro link. A few contemporary sociological theorists have addressed the micro-macro link by proposing that interaction mediates between individual action and macrosocial structure (Collins, 1981; Ellis, 1999; Rawls, 1987).

In artificial societies, structures emerge from models of large numbers of agents, their connections, and their repeated communications and interchanges over time. This suggests that micro-macro theory requires an explicit theorization of interagent communication. The success of artificial societies in recreating phenomena of micro-to-macro emergence provides support to those who emphasize that micro-macro theory must consider agent interaction.

2 AGENT COMMUNICATION LANGUAGES

The communication language that agents use to communicate with each other is referred to as an *agent communication language* (ACL). Although interaction is a fundamental aspect of MAS, sociological research and theory suggests that the ACLs implemented in MAS have been radically simple. In some MAS, communications can be as simple as exchanging one bit of information. The Sugarscape artificial society uses interaction rules like the one illustrated in Example 1 (Epstein and Axtell, 1996, p. 73).

Such rules are simple versions of the forms of interaction assumed by both exchange theory and rational choice theory. In exchange theory, for example, all sociologically relevant communication is modeled as a form of exchange, as in the exchange of valuable information for status. Similarly, with Coleman's rational choice theory, social communication is modeled as a type of exchange, as in a transfer of trust from one agent to another (1990, Chap. 8).

> **Example 1. An agent interaction rule in Sugarscape**
>
> *Cultural transmission rule (tag-flipping):*
>
> For each neighbor (four orthogonally contiguous agents), a tag (one bit in an eight-bit mask) is randomly selected.
>
> If the neighbor's bit setting agrees with the agent's at that position, take no action; if they disagree, flip the neighbor's tag to agree with the agent's.

In contrast to these simple interactions, many MAS use ACLs based on the speech act theory philosophical approach. Speech act theory is the explicit theoretical foundation for the two dominant industry standard ACLs: FIPA, from the Foundation for Intelligent Physical Agents (http://www.fipa.org) and Knowledge Query and Manipulation Language (KQML) (http://www.cs.umbc.edu/kqml/). Following speech act theory, KQML messages are called *performatives*. Performatives are defined in terms of the agent's *knowledge base*, which contains two types of knowledge: *beliefs* and *goals*. The performative TELL is defined in Example 2 (following Labrou and Finin, 1997).

This model of communication is based on a theory of agency that is widely used in MAS: the *belief-desire-intention* (BDI) model (Rao and Georgieff, 1995). The italicized terms in Example 2 represent beliefs and desires of agents. Before any communicative act occurs, the agent must first have goals and beliefs about how to accomplish those goals. An "intention" captures the notion of *commitment* to a plan of action; agents communicate only after committing to a plan of action. Since 1995, theorists have developed a formal logic of BDI systems using *modal logic*, which includes logical operators such as "BELIEVE(x,y)" and "DESIRE(x,y)" (Wooldridge, 2000). Agents reason about other agent's actions, assuming that those other agents are also operating according to BDI principles. Beliefs about the BDI states of other agents play an important role in composing and interpreting messages (as in Example 2).

> **Example 2. Definition of the TELL performative in KQML**
>
> TELL (A, B, X): A states to B that A believes X to be true.
>
> Precondition for performative: This performative occurs when A *believes* X, and A *knows* that B *wants* to *know* whether X is true.
>
> Result of performative: A *knows* that B *believes* that A *believes* X. B *knows* that A *believes* X.

3 WHAT SOCIAL SCIENTISTS KNOW ABOUT HUMAN COMMUNICATION

Beginning in the 1960s, sociologists and anthropologists began to use audio and video technology to record naturally occurring interactions in dyads and groups. In the early 1960s, anthropologist Dell Hymes founded an approach known as *the ethnography of communication*, and in the early 1970s, sociologist Harvey Sacks founded *conversation analysis*. The seminal works of these two approaches led to a burst of rigorously empirical studies in the 1970s, 1980s, and 1990s, and today we have a radically better understanding of communicative processes than just a few decades ago.

These decades of empirical research were not good to speech act theory, nor to the theories of speaker intentionality that it was based on. Almost all empirical scholars of human communication reject speech act theory, and the philosophy of intentionality that drives it. That rejection is based on several discoveries about how human communication really works.

3.1 Discovery 1: The conduit metaphor is false.

ACLs assume the *conduit metaphor* of communication, where a communication represents an intention on the part of the speaker to transmit some information to the hearer (Reddy, 1979). Speaker's intentions are assumed to be formed before the act of speaking, and the addressee is seen as a passive listener whose only job is to guess what the speaker has in mind. I'll call this the individualist theory of intention and speech. However, empirical studies of conversation by ethnomethodologists, conversation analysts, and sociolinguists have shown that the conduit metaphor is an inaccurate model of human communication. There are many cases in daily life in which the meaning of a given act is not defined until the recipient has replied, thus retroactively giving the act meaning. Indirect insults among African Americans are one well-known example I'll discuss later (Fisher, 1976; Morgan, 1996). Another well-studied example is a technique used by teachers in classrooms, to rephrase a student response in terms that further the teacher's pedagogical goals — using more scientific language than the student, or creatively interpreting the response so that it more naturally furthers the day's lesson plan (O'Connor and Michaels, 1993; see Example 3).

ACLs implicitly assume a polite *turn-taking* model of interaction; they assume that only one person can speak at a time. But there are many situations where this isn't true. When your audience is interrupting, or when several people are talking at once, speech becomes collaboratively constructed. When we're talking, we're always looking at the people listening to us, checking to make sure they are still listening and still interested, looking for any signs of confusion or enthusiasm, as they quietly murmur "Hmm…" or "Uh…" during our words. Listeners communicate much of this information nonverbally, with their eyes or their posture. Conversation researchers refer

Example 3. Revoicing

Teacher:	Renee / What about you? / /
Renee:	(unintelligible)
Teacher:	We're your audience / we can't hear you / / You're doing fine / / You think it will balance because...
Renee:	(unintelligible)
Teacher:	Yes / yuh / three of them at three
Annie:	Could you speak up?
Renee:	(unintelligible)
Teacher:	So it's ten and ten?
Teacher:	So instead of / I'm going to give a little louder voice to what I think I hear Renee say / / She was saying she wasn't adding five and two / and saying it's seven / she was saying five and five / knowing that if you double something it's like adding it to itself.

to this as *backchannel communication*, and effective conversationalists can intuitively monitor these messages to calibrate their ongoing talk.

We process this feedback subconsciously, and often don't even realize how it's influencing the way that we talk. To explore these subtle cues, sociologist Chuck Goodwin (1981) videotaped conversational groups, making sure to get both the speaker and the listener on camera. He noticed that both speakers and hearers use their gaze — who they are looking at — to communicate important nonverbal information, and speakers construct their words based on where they see the listener looking. For example, speakers often stop their sentence once they get the listener's attention, and start over again.

Goodwin developed a novel way to transcribe the gaze of both the speaker and the listeners. Example 4 (Goodwin, 1981) shows a mother talking to her two children; she wants to suggest that they move closer to the teacher so they can hear better, but when she starts speaking, Brian is not looking at her. At the X, Brian turns to look at his mother, and notice how she interrupts herself and restarts her suggestion.

Example 4. Listener gaze and utterance construction

Barbara: Brian, you're gonna have — you kids'll *have* to go down *closer.*
Brian: X________________

Example 5. Listener gaze and utterance construction

Mike: Speaking of pornographic movies, I heard (---------) a while…
Carney: X________________

The horizontal line after the X indicates that Brian is looking at Barbara. Although Barbara is the only one talking, Brian's gaze contributes to the way she forms her words — when he turns to look at her, she starts again, this time addressing both children. Goodwin also noticed that speakers sometimes stop in mid-sentence, waiting for the listener to look at them; in Example 5 (Goodwin, 1981), each dash in parentheses represents a pause of one-tenth of a second.

Mike pauses because he's not sure if Carney is paying attention; he continues after Carney turns to look at him.

In both of the above examples, listeners collaborate in the construction of the utterance even though they don't say anything. These examples show that the conduit metaphor is false — it's not true that speakers formulate utterances cognitively before they speak.

An important aspect of humor, gossip, and insulting is indirection — where speech is addressed to a "mediator" but is intended to act toward an "overhearer" nearby. Indirection intentionally obscures intentionality.

This indirect insult strategy has been widely documented among African-Americans. The anthropologist Marcyliena Morgan (1996) has done some interesting work on the use of two types of indirect insults among African-American women. Perhaps the first study of this kind of indirectness was done by Lawrence Fisher in Barbados, in which this communication is called *remark dropping*. Fisher (1976) reported the example of remark dropping given in Example 6.

The apparent target of such an insult doesn't have a good way to respond. One Barbadian told Fisher "If I drop a remark to you and you challenge me, I would just say, 'Who told you I was referrin' to you? You must be hearin' things.'" It's like the Heisenberg uncertainty principle — how you perceive the utterance changes its interactional meaning.

Example 6. Remark dropping in Barbados

A woman chose to wear an overly bright shade of lipstick to a party. She overheard a woman say, "Oh, I thought your mouth was burst" to a man whose lips were in perfect order.

Just after lunch in my office, I walked across campus to get coffee at the campus food court, dressed in my suit in preparation for the afternoon's class. As I walked back to my office across the quadrangle, a group of high school boys approached me, carrying duffel bags. Just behind them, there were three adult men that were clearly their chaperones. Seeing my coffee, one of the boys said "Is the food court that way?" in what I took to be a somewhat disrespectful tone, but I brushed it off and politely said yes, it's that way, and down the stairs in the basement. As the boys turned toward the food court, the three men passed by. They had overheard our exchange. One of them asked with exaggerated deference, "Is the food court this way, Sir?"

We both knew that he had heard my directions the first time. The purpose of his question was not to request information; he was implicitly apologizing for the boy's rudeness, to acknowledge that he had also perceived the tone to be somewhat disrespectful. Even more fascinating was the way that the man spoke loudly and distinctly, so that the boy could overhear his question. With this one question, he spoke to two audiences — apologizing to me, while reprimanding the boy. And he made it clear to me that he was reprimanding the boy.

The surface meaning of the question is irrelevant — the man has already heard my directions to the food court. The apology is hidden in the words — he has not said "I apologize for my student's rudeness." The reprimand is also hidden — he has not said "How dare you be rude to this man!" This single utterance accomplished two different speech acts, each directed toward a different audience.

We've seen two problems with the accuracy of conduit metaphor: the collaborative construction of utterances, and indirectness. The third problem is even more fundamental. In his influential essay "Footing," Goffman (1981) argued that "speaker" and "hearer" were folk theoretic categories; that the analyst must break out each of these folk roles into multiple functional roles in interaction. Goffman proposed that a speaker consists (analytically) of three distinct roles: the animator, author, and principal (1981, p. 144). Levinson further elaborated these subcategories into more than 10 "producer roles" (1988, p. 172). In both cases, the attempt is to tease out the different roles that are involved in speaking. More recent theory has shifted even more radically — from an analysis of static roles to an analysis of the dynamic process whereby roles are assigned, often shifting, during an ongoing interaction (Irvine, 1996).

For example, Judy Irvine's analysis of Wolof (Senegal) ritual insult poems, performed at weddings, distinguishes seven different functional participant roles (Irvine, 1996): sponsor, formulator, speaker, co-speakers, addressee, hearers, and target. The poems are performed by low-ranking griot women, but the performer composes the poem in collaboration with the high-ranking wives married into the new husband's patrilineage. The composition process is secret, so that at the actual wedding, a particular insult can't be identified with any specific author. The griot performer can claim that she is merely the transmitter; the sponsoring women can claim that they sponsored the event only in general, and had no part in any particular poem. Thus responsibility for an insult is distributed across the transmitter, the sponsors, and the (unidentifiable) composer.

There are multiple recipients for a poem: the bride (ostensibly the addressee); the bride's kin, who are often the target of the insults; and the bride's family and friends, who are often insulted — especially those who are prominent and whose doings are of community-wide interest. Even if one of these persons is not present, word of a clever insult is sure to be relayed to them later.

The audience actually joins in and speaks the chorus of the insult poem, further distributing the speaker role, as in Example 7 (Irvine, 1996, p. 137).

Example 7. Audience collaboration in performance

Choral Couplet (introduced by the soloist and repeated as a refrain):

1. *M – G – né na, baalal ma Màka —*	M – G – said, "Forgive me, Màka —"
2. *sa xaj gi demul.*	your pilgrimage didn't work

Soloist:

3. *M – G – moo jénaxi tookër, lan*	M – G –, he is a bush-rat; whatever
4. *la mu gis jàppéwaan ni ci cop.*	he saw, he grabbed, mounting it (and spoiling it).
5. *Du ko laaj — mu dajéwoon.*	He didn't ask — he just coupled.

But Irvine's deeper point is more subtle. She notes that these participant roles don't exist only in the moment of performance itself; they only arise because the insult occurs as part of an ongoing history of interactions in the community. Everyone knows the insult was not composed on the spot, but that it emerged from prior secret conversations. And everyone knows that the insult will be repeated later by members of the audience, to nonpresent targets and their friends and relatives. "The significance of the insult … depends on this complex of implicated dialogues" (p. 139). Yet speech act theory focuses only on an isolated utterance, with no way of capturing this history.

3.2 Discovery 2: Speaker intentionality is often problematic.

> "I'm sure I didn't mean —" Alice was beginning, but the Red Queen interrupted her impatiently.
>
> "That's just what I complain of! You *should* have meant! What do you suppose is the use of a child without any meaning? Even a joke should have some meaning — and a child's more important than a joke, I hope. You couldn't deny that, even if you tried with both hands."
>
> "I don't deny things with my *hands,*" Alice objected.
>
> "Nobody said you did," said the Red Queen. "I said you couldn't if you tried."

Speech act theory is based on an overly simplistic notion of speaker intentionality; this makes it compatible with BDI agent models, but difficult to extend to more robust models of agency. For example, in many speech situations — such as a Wolof insult poem — intentionality is distributed; it is often unclear what the denotational meaning or the interactional effect of a single utterance is, apart from its performance context, because of the contingency of the ongoing interaction (Sawyer, 2001a). As I just showed, conversation research has discovered that many human communications are *cocreated*, in part with the contribution of nonverbal back channel information on the part of the listener (as in the above section). Just as stories are cocreated, intentions can also be formed and altered during the process of communication; discourse itself can be an integral part of the process of goal and belief formation. In speech act theory, meaning is identified with speaker's intentions to express certain beliefs or bring about certain changes in the world. Intentions are psychological states that exist in the speaker's mind before the act of speaking. But in many situations — in America and in other societies — intention and meaning is collectively constructed, partially determined by the addressee's response. BDI theory can't capture these real-world phenomena.

Linguistic anthropologist Alessandro Duranti studied the Samoan speech event called *fono* (1988, 1992). Duranti argued that whereas Americans think of talk as a way to communicate information from one mind to another, Samoans think of talk as a way to assign responsibility. Samoans do not talk about speaker's motives or their inner psychological states when discussing conversations. Instead, when they talk about conversation, they attempt to assign responsibility for words spoken. Unlike speech act theory, which focuses on the intentions behind words (and corresponds well to American folk theories about language, as reflected even in the legal system's focus on intention), the Samoan ethnotheory of language focuses on the consequences of words.

Duranti studied the fono in the small rural village of Falefa (Upolu) during 1978–1979. The fono is a formal meeting; Duranti focused on these special convocations of *matai*, chiefs and orators, which act as a high court and as a legislative body. Matai gathered in a fono can both make laws and decide policies related to new problems. Fonos are often antagonistic, with different powerful groups in competition, and they are public. As a result, each group attempts to control one another's actions and the public's interpretations of them. And the eventual outcome is often uncertain; as a result, it is often convenient to be "cautious, humble, and vague" early on in the meeting (1988, p. 16).

In Samoa, as in many other places in Polynesia, a special class of talented verbal artists known as *tulafale* — "orators" — has the right and the duty to represent powerful chiefs ceremonially, and to act as spokespersons and mediators in political conflicts. The tulafale often negotiate publicly and at length before the chiefs that they're speaking for speak themselves. This often puts them in the difficult position of not knowing how their chief will ultimately decide. An example of the danger facing orators is shown in Example 8 (Duranti, 1992, p. 32).

Fa'aonu'u had expressed an opinion about how a conflict between villages should be resolved, an opinion that supported a decision previously expressed by the young chief Savea (to resolve the issue in court). However, during the ensuing meeting, advocates of resolving the issue the traditional way (out of court) prevailed, and Savea changed his mind. At the conclusion of the meeting, the senior orator Moe'ono (who was an advocate of an out-of-court settlement) reprimanded Fa'aonu'u for hastily expressing an opinion, as we see in the transcript. Note that he is reprimanded even though at the time he expressed the opinion, it was that of the chief, who only later reconsidered. (In fact, one of the reasons that orators speak first on behalf of the chief is that it allows the chief to change his opinion without a loss of face. The chief's wrong decisions and changes of position are assumed publicly by the orators who speak on their behalf.)

Example 8. Samoan tulafale as scapegoat

ma: – (…) *ou ke kaukala aku fo'i Fa'aogu'u iā ke 'oe* (…)
And – … I am also talking, Fa'aonu'u, to you

mea lea e leaga ai le – le alualu i galuega
this going away (from the village) to work

sau fo'i ua– (…)
coming back to – (speak up) is bad

pei o agaleilā o le ā koe aga'i kua lo'u kāofi. (…)
as for before, my opinion is going to reach back (i.e., to what you said before)

'o le – o le makā'upu ua fikoikogu i loukou faleakua (…)
as for the – the topic that concerns your subvillage,

kaofiofi le i'u maea. (…)
moderate yourself …

ae aua le luaiga lālā mai fa'amaka o Avi'i lou kāo –
and don't show off your op(inion) like the crab that has eyes that stick out

a'o lea ua aliali gei,
now it looks like
[…]

ua fausia e Savea le – le figagalo lea e fai aku iai
Savea has agreed to say that

iga ia kākou feiloa'i ma Lufilufi. (…)
we should meet with Lufilufi …

KO'A! le fa'aukaga. (…)
HOLD (IT)! the advice …

ko'a le fa'aukaga. (…)
hold the advice …

e leai fo'i se isi Fa'aogu'u 'o oe
there is no other Fa'aonu'u but you

ga'o 'o oe ā 'o Fa'aogu'u (…)
only you, you are the Fa'aonu'u …

One strategy the tulafale can use to avoid public blame is to avoid committing to a position by being very vague. However, at various points in the fono a tulafale is forced to be more direct, as in Example 9 (Duranti, 1992, p. 29).

In Example 9, we see the spokesperson, Loa, engaging in a verbal strategy of vagueness, which has the desired result of involving the more powerful senior orator (Moe'ono), so that a difficult announcement is ultimately made jointly — with responsibility distributed.

Speakers' intentionality at the time they produce a speech act is irrelevant (Duranti, 1992, p. 34). Orators can get into trouble for public statements, even if everyone knows they were speaking on behalf of a higher-ranking chief, and retaliation can occur. Accusations and discussions of responsibility are always in terms of the practical consequences of his words; the speaker's personal motives are considered irrelevant. As Duranti wrote, "a speaker must usually deal directly with the circumstances created by his words and cannot hide behind his alleged original intentions" (1992, p. 33). This is in sharp contrast to Western contexts, where it is assumed that messengers should not be held responsible for what they say.

Based on such phenomena, Duranti argued that speech act theory is based on culturally specific Western ethnotheories of meaning — what I call the individualist theory — that "people should be held responsible only for those acts (and words) that can be clearly seen as reflecting their own individual intentions" (p. 40). In Samoa, as in many other cultures, speech acts are often collective, are often spoken on behalf of another, and are often retroactively redefined. Speaker intention is not important. Utterance meaning is not "owned" by the speaker; rather, it is a cooperative achievement: "meaning is seen as the product of an interaction (words included) and not necessarily as something that is contained in someone's mind" (p. 41). Utterance meaning (and intention) emerges from a collaborative, distributed social process.

Example 9. A tulafale's attempt to distribute speaker responsibility

(The orator Loa has just concluded the introductory speech leaving out the mention of the agenda.)

Loa: *maguia le aofia ma le fogo!*
Good luck to the assembly and the *fono*!

?: *mālō!*
Well done!

M: *'o ā makā'upu o le fogo?*
What are the topics of the *fono*?
fai mai makā'upu o le fogo.
Tell (us) the topics of the *fono*.

Loa: [*o makā'upu o le aofia ma le fogo*
The agenda of the assembly and *fono*

Moe: *fai mai* (?)
Tell us (?)

Loa: *ia e fa'akakau kogu lava i lo kākou Falelua,*
Well it's really about the two subvillages,
(CLEARS THROAT)
oga pau gā 'o makā'upu,
those are the only topics,

M: *oi!*
Oh!

Loa: *e ā?*
What?

M: *'o le isi makā'upu o Savea*
The other topic of Savea

Loa: *ia 'o le isi fo'i makā'upu e uiga i le –*
Well the other topic is about the –
le afioga iā Savea ogo'o: –
the honorable Savea 'cause –'

M: *(Fa'amakuā'igu),*
(Fa'amatuā'inu)

Loa: *'o le lā- mea fo'i ma Fa'amakuāigu,*
The – thing also with Fa'amatuā'inu,
go 'ua kukulu Savea i – i le mālō,
given that Savea has complained to – the Government,
ia'
well
go 'ua ka'ua gi fa'akosiga Fa'amakuā'igu i le paloka,
given that some illegal campaigning of Fa'amatuā'inu during the elections has been said (to occur)
ia 'oga pau gā 'o- 'o makā'upa o le aofia ma le fogo,
Well, those are the only – topics of the assembly and the *fono*,
-hh iai fo 'i gisi makā'upu
-hh (if) there are other topics,
o lo'o lē maua,
I am not getting to,
ia la'a maua i luma!
Well they will be brought to the front!

?: *mālō!*
Well done!

Moe: *mālō fekalia.*
Well done, the (honorable) speaking

?: *mālō fekalia.*
Well done, the (honorable) speaking
(…)

Moe: *ia fa'afekai aku Kafiloa. (…)*
Well, thank you Ka(o)fi(ua)iloa …
'ua 'e fa'amaga le fogo
for starting the *fono*
[…]

3.3 Discovery 3: Communication and meaning are collaboratively emergent social processes.

Linguistic anthropologist Don Brenneis (1984) studied one of the most creative styles of gossip during his research in a rural community in Fiji, a small Pacific island. Brenneis spent most of his time in Fiji with a large group of Hindi-speaking Indian immigrants, who immigrated to Fiji generations ago. The Fijian word for gossip is *talanoa* and Indian immigrants have borrowed it; it translates roughly to "idle chatter," and in Fiji, only men talk talanoa.

Talanoa stories are always told collaboratively: two or more men join together in a conversational duet. What's more, listeners are expected to jump in and contribute to the story;

sometimes a talanoa story is more like a conversational trio or quartet. Different speakers are always interrupting each other to continue the story; but no one complains about it, because it's an expected part of the collaboration.

Talanoa gossip is spoken in a rhythmic, almost poetic style of talk, with "lines" that have a metric structure like our own traditional poetry. Because of this rhythm, everyone knows when a line is going to end. And because there are clear pauses after each line, it's easy for someone to jump in and take over the story. When a man jumps in and takes over the storytelling role, he often begins by repeating the last line, and then continues the story in the same tempo and meter. Watching this, you get the impression of a connected verbal performance — a collaborative poetry.

There is a reason that *talanoa* stories are always performed collectively. In the small country villages where these Indians live, it's important not to directly insult another person, since word of it is guaranteed to get back to the person you're talking about. So if two or more people are telling the story together, neither one of them can be blamed for talking behind your back. And if the audience keeps jumping in and participating, they're equally guilty. The talanoa style of group storytelling distributes the blame for whatever is said.

In talanoa, speech acts are collectively created, and emerge from an ongoing social encounter. Linguistic anthropologist Laura Graham, in an ethnography of the Xavante Indians of Central Brazil, documented a similar phenomenon in a more public political form of discourse called *wara* (1993). Using this form of discourse, speech acts become the product of multiple selves and multiple voices, "a collage of multiple articulating voices" (p. 719). Wara discourse practice severs the link between speech and individual, and makes discourse "an emergent, intersubjectively produced social interaction" (p. 718). The function of this speech form is to promote social cohesiveness and reinforce egalitarian relations among the senior male participants. (As with the Fijian community studied by Brenneis: anonymity, and the use of metaphor, allegory, and proverbs, all make it difficult to associate intention with a specific speech act. Both are smaller, egalitarian communities, with a need to avoid direct confrontation among equals.)

The wara is a men's council that includes a morning meeting and an evening meeting, which goes on late into the night. Once it's dark it's difficult to see individual speakers; overlapping speech is common and makes it even more difficult to identify speakers.

Speakers typically address the council as a representative of their faction rather than as themselves. The features of the discourse event minimize the speaker's individual identity (p. 725):

- Many people talk at the same time.
- There are no podiums or lighting or public address systems.
- Men avoid looking at the speaker; most lie on their backs and look up at the sky.
- Many men keep up a running commentary throughout the event, resulting in a constant murmur.

- A second speaker doesn't have to wait for the first to finish before starting.

- Each speaker incorporates much of what has been said already into their delivery, resulting in a high degree of collective repetition (see Example 10).

- Skilled speakers often integrate the remarks that others are making simultaneously into their ongoing speech, revealing an ability to monitor others' speech while talking (see Example 11).

Speakers use a genre of speech that is distinct from conversational Xavante called *ihi mrèmè*, which is characterized by extensive repetition, parallelism, and a special intonation pattern. Like Brenneis' Fijian gossip, these formal features enable a more collective discursive practice.

Example 10. Collective repetition in the Wara

Eduardo	he just perfected what he saw in the dream
Jusé	yes he perfected what he saw in the dream
Eduardo	they [the ancestors] perfected themselves
Jusé	they [the ancestors] perfected themselves
Eduardo	he is remembering the story
Jusé	he is remembering the story

Finally, in interviewing the Xavante after a meeting, Graham reports that they don't claim responsibility for their own speeches, and they decline to comment on the speeches of others. No one admits to paying attention to any particular individual, and no one admits to playing a particularly prominent role in a meeting (p. 736). All of this is consistent with Graham's interpretation of the discourse event itself: speech acts are produced collectively and are distributed among the participants. "The locus of political action resides in emergent social interaction, not in any single agent" (p. 737).

Most talk occurs in groups of three or more; yet, many theories of speech are based on two-party interactions (and ACLs likewise are assumed to occur between two agents). Argument, story-telling, and family dinner-table conversation have multi-party dynamics radically different from the one-turn-at-a-time dyadic conversation assumed by speech act theory (and other linguistic theories of social action), as research by the Goodwins has so broadly documented.

Example 11. Speakers integrate others' remarks into their talk

Warodi	in the dry season…his brother… the grass's smoke rose straight up and he missed him… again…he returned…	for him to walk together	
Jusé		for him to continue walking together	
Warodi			for him to walk together
Jusé		for him [continue] walking together	
Warodi		for his brother in the dry season	
	again…he returned…for him to continue walking together		

3.4 Discovery 4: Context influences the meaning of speech.

Cultural anthropologists have found that our ways of thinking about language and about human agency are intimately linked. Theoretical attempts to explain how language works (such as speech act theory) reflect our culture's views about the nature of agentive individuals. Many other cultures hold to different theories of personal agency, and their own views of language action are correspondingly different.

Speech act theory holds that speech acts are accomplished by autonomous selves, whose deeds are not constrained by relationships and contextual expectations. But conversation researchers have discovered many speech situations in which speech action is intimately tied with social context. Anthropologist Michelle Rosaldo studied the Ilongot (in the Philippines) in 1967–1969 and again in 1974. Rosaldo found that directive speech acts (commands, requests, orders) were central to the Ilongot's cultural beliefs and system, reflecting role conceptions of men, women, children, and relative status relationships (Rosaldo, 1982).

For example, directives tend to move in lines associated with age- and sex-linked social rank; men can ask women to get something for them — and they're rarely rejected — but women rarely ask men. Women often make demands on children; and older children make demands on younger. However, in certain social contexts, it is considered appropriate for children to issue directives to parents, or for wives to issue directives to husbands.

Rosaldo argued that one cannot understand the meaning of a directive without knowing a lot about the social situation and context of utterance. In fact, she argues that relations and context are primary in interpreting the meaning of a speech act, in contrast to Western culture, which holds that speaker intentions are primary (p. 210). For Searle for example, acts of speech are not social actions, but rather the embodiments of universal goals, beliefs, and needs held by individual speakers.

For Searle, "to promise" is the paradigm speech act. "To promise" focuses on the sincerity and integrity of the speaker — it is a thing derived from inner life (unlike a greeting, which is often required conventionally by the situation). And Searle's description of the promise focuses almost entirely on the sincerity of the speaker's commitment: a promise is defined as "a sincere undertaking, by the speaker (S), of a commitment to do A, where A is something S would not ordinarily undertake, and something, furthermore, that S believes that hearer (H) desires" (Rosaldo, 1982, p. 211). Note what is missing from the definition: any invocation of context or relationship. But in fact, promises are things that we offer only to certain kinds of people at certain times. Promises to a child are typically didactic and tendentious; a promise from a candidate for political office is judged not only on sincerity but also on suspiciousness, sound bite worthiness, and grandness of vision. Promises between spouses have a different tenor from the public contractual commitments implied by Searle's definition (Rosaldo, 1982, p. 211). These are the complex social rules that surround the action of promising, all neglected by speech act theory — because they don't match the language ideology, and the assumptions about personhood and intentionality, held by the culture that developed speech act theory. Searle's definition of the promise fits with Western folk beliefs that social meaning issues from private persons.

Rosaldo's shocking observation is that the speech act of "promising" does not exist in the Ilongot speech community (1982, p. 211). Among the Ilongot, as Rosaldo claims, the directive rather than the promise is the paradigm speech act, and when it comes to the directive, it's hard to ignore the context and the relationship. Rosaldo steps through Searle's five speech act categories (assertives, directives, commissives, expressions, and declarations) and shows that the Ilongot use of each of them fundamentally presumes relationship and social context. Assertives, for example, are not meant to represent facts about the world, but rather are used to "articulate relationships and claims within the context of a history that is already known…[and] to talk about alliance and opposition in particular social groups" (1982, p. 214). The degree to which an assertion is qualified, or the way that prior acts are named, becomes the stuff of verbal duels:

> *cautious qualifying verbs*: "Well, what I will just, uncertainly, say to you"
>
> *metaphors qualifying speaker's actions*: "I'll be the one to run ahead again (and speak out) since it's the way with young dogs"; "(Let's talk until) we are filled up, contented, from hand feeding one another words"

Each of these variations is strategic and carries subtly different implications for the ensuing encounter.

Perhaps most striking is Rosaldo's claim that the Ilongot have no verbs for commissive and expressive acts, the two types that are most paradigmatic of Western folk theories of language use. There's no verb for "promise" or "apologize" or "congratulate." There are no expressive forms for "I'm sorry" or "Thank you/I appreciate that." Acts that are similar to these Western speech acts actually fall into Searle's declarative class (all the cases where saying something actually changes the world, e.g., "I marry you" or "I find you guilty") — with the result that they emphasize the consequences of the act, rather than the inner state of the speaker (1982, p. 218). As in many other cultures, among the Ilongot the eventual consequences of speech are more important than the prior intention of the speaker.

3.5 Discovery 5: Much of interaction is based on semi-scripted, ritual sequence.

Many communicative acts are socially distributed phenomena (Duranti, et al., 1991). Just as human cognition is increasingly perceived as a situated practice, and is often distributed among members of a collaborative team (as in the theoretical and empirical work of psychologists Jean Lave and Ed Hutchins), speech acts are as well (as in my examples of Samoan fono, Fijian gossip, and Xavante wara). For example, the status of an utterance as an "answer" emerges not from the utterance itself, but from its placement after a particular kind of talk — a "question" — produced by someone else. Duranti, et al. (1991, p. 4) state, "The constitution of the action as an answer is thus not situated within the intentions of a single participant, but instead emerges through a time bound process that is distributed across different participants and actions within the framework provided by the sequence of activity that they are collaboratively constructing."

In sociology, the emergent pattern of the group encounter has been called a *routine* (by conversation analysts) or an *interaction ritual chain* (Collins, 1981). Sociologists like Collins and Giddens are typically concerned with structures that emerge and perdure across repeated encounters, thus resulting in something approximating macrosocial structure.

Speech act theory focuses at the utterance level, and does not provide a theory for how social action may be found in broader interactional structures — like the sequence of social interactions implicated in the Wolof insult poem performance. When we expand the analysis to above the level of the utterance, we are analyzing interactional structures with multiple participants. MAS need to consider the effects of connected sequences of communicative behaviors, which sociological theorists refer to as *rituals* or *routines*. Sociological research suggests that MAS developers may need to explicitly model interactional routines.

The ACL standard Foundation for Intelligent Physical Agents (www.fipa.org) allows developers to create standardized sequences of communicative acts, called *protocols*; these are rigidly defined scripts. If two agents agree to engage in a protocol, it can increase their communicative efficiency, because they do not have to go through the same decision process before each communicative act. This is similar to MAS work that attempts to increase the efficiency of systems by explicitly modeling activities (Kristensen and May, 1996) or conversations (Barbuceanu and Fox, 1997). Yet there are difficult, unresolved problems associated with how agents first negotiate to engage in a protocol, and whether this negotiation might take so much energy as to offset the presumed advantages of entering the protocol. And because a protocol is a fixed script, these extensions to ACLs fail to account for opportunistic improvisation — when agents break out of the protocol, or decide to modify or embellish it.

3.6 Discovery 6: Most speech acts are implicit.

> "They gave it to me," Humpty Dumpty continued thoughtfully… " — for an un-birthday present."
>
> "I beg your pardon?" Alice said with a puzzled air.
>
> "I'm not offended," said Humpty Dumpty.
>
> "I mean, what *is* an un-birthday present?"

We all have had trouble understanding indirect speech acts.

The paradigm speech act is the explicit primary performative (EPP). But as we all know, in most cases in everyday speech, the "speech act" accomplished is not explicitly marked by the utterance's verb (recall my food court apology example). In the terms of speech act theory, most are "indirect illocutionary acts," and speech act theorists have had the most trouble accommodating these into the theory. In the context of the world's cultures, Anglo-American folk theories of language emphasize explicitness and directness; but even among us, most speech acts are indirect.

ACLs have borrowed only from speech act theory's notion of EPPs and illocutionary effects, and have ignored implicit performatives and perlocutionary effects. Here, ACL developers are drawing on the assumption of speech act theorists that the explicit performative is normative, and represents the "deep structure" of even implicit acts. In contrast, most scholars of communication have found that the normative case is implicit and indirect communication — especially once you leave the Anglo-American speech community. Implicit communications cannot be constructed nor interpreted without a profound understanding of the conversational context; a fundamental aspect of human communication is *indexicality*: the relations between communicative acts and the ongoing, co-constructed, emergent conversational context.

Of course, when people speak they often "do" things. The problem with speech act theory is that it attempts to connect speech function with speech form in a systematic way, but that way only works for EPPs, and only for certain cultures.

3.7 Discovery 7: Speech acts are less likely to result in unintended emergent effects.

Speakers use the metapragmatic function of language to "metacommunicate" about the emergent process and flow of the encounter, or about the ground rules and the communication language itself. Sawyer (2003c) has empirically demonstrated that the metapragmatic features of human communication lead to unintended emergent effects, and that these emergent effects have causal consequences for the future flow of the encounter. Yet metapragmatics have not yet been implemented in agent communication languages. In ACLs, these properties of communication are largely fixed in advance and cannot themselves be negotiated. A speech act comments on itself — saying "I promise" has a denotational meaning, but at the same time says "what I'm now saying constitutes a promise" — but this is only a small aspect of metapragmatics.

Sociological studies of improvisational groups help illustrate this point. In an improvisational theater performance, when no dialogue or plot is specified in advance, how do actors determine the variables of the interactional frame — the characters, motivations, relationships, and plot events and sequence? In Sawyer's (2003c) study of emergence in improvising theater groups, the interactional frame was shown to emerge from complex levels of metapragmatics (see Example 12).

Example 12. Emergence in improvisational theater

Dave: All the little glass figurines in my menagerie,
The store of my dreams.
Hundreds of thousands everywhere! *Turns around to admire.*

Ellen: *Slowly walks toward Dave.*

Dave: *Turns and notices Ellen.*
Yes, can I help you?

Ellen: Um, I'm looking for uh, uh, a present? *Ellen is looking down like a child, with her fingers in her mouth.*

Shared interactional understandings are created through metapragmatic communication by the collaborative efforts of the entire group. No single participant creates the frame; it emerges from the give-and-take of conversation. The interactional frame includes all of the pragmatic elements of a small group encounter: the socially recognized roles and practices enacted by each participant, the publicly shared and perceived motives of those individuals, the relationships among them, and the collective definition of the joint activity they are engaged in. The frame is constructed turn by turn; one person proposes a new development for the frame, and others respond by modifying or embellishing that proposal. Each new proposal for a development in the frame is the creative inspiration of one person, but that proposal does not become a part of the frame until it is evaluated by the others. In the subsequent flow of dialogue, the group collaborates to determine whether to accept the proposal, how to weave that proposal into the frame that has already been established, and then how to further elaborate on it.

Emergent properties are often associated with *unintended* effects of action; intended effects are not emergent by definition, because their origin can be traced to the individual motivations and advance the plans of individuals. In improvised dialogues, the actors do not have beliefs, desires, and intentions as conceived of in ACLs; rather, these are attributed to individual actions retrospectively as the dialogue evolves. In spite of their lack of plans and intentions, actors are able to coordinate their actions to generate a plausible, coherent dialogue, and stable macro patterns emerge.

I compared emergence processes in two 60-minute improvised plays, both performed by professional groups in Chicago in the early 1990s (Sawyer, 2003c). The first group, called The Family, used a format that they called “The Movie”; in The Movie, actors are allowed to step out of character and to explicitly metacommunicate about the ongoing drama, using “director talk” as if they were the director or playwright. The second group, called Jazz Freddy, did not allow their actors to step out of character at all. Both groups created their 60-minute play from a combination of two- to four-minute scenes, and the edits between scenes were emergent and collaboratively accomplished by the actors. I found that in The Movie, all scene edits were done with director talk; of course, in Jazz Freddy, all scene edits were done while remaining in character.

These differences in metapragmatic strategy resulted in the unintended emergence of two very different dramatic frames. The Jazz Freddy frame that emerged emphasized character and relationship development, but its plot was not very complex. In contrast, The Movie’s emergent frame had multiple, interwoven plot lines, but had weak characters and relationships. I used conversation analytic methods to demonstrate the step-by-step process of this emergence, showing that the metapragmatic differences were responsible for these different processes and outcomes of emergence.

There can be no conscious awareness of emergence processes in small-group improvisation, because they happen so quickly. For example, actors in The Family and in Jazz Freddy were not aware of these contrasts between their performances. As Sacks was perhaps the first to point out (1992, p. 11), during conversation people respond so fast that they could not conceivably have consciously planned and decided their action. The psychological processes underlying conversational behavior are largely preconscious. Prior research has demonstrated that speakers have great difficulty becoming aware of the metapragmatic function of their own utterances, even when they consider an interaction in retrospect (Silverstein, 1979, 1981). This is why I proposed that emergence results from the *implicit metapragmatics* of dialogue, rather than from the explicit formal features of dialogue (Sawyer, 2003c). Most offers are implicitly metapragmatic: the offer is phrased as if the proposed state of affairs already were the case (as in Turns 2 and 3 in Example 14, and as in the scene edits of Jazz Freddy). Such offers can only be explained by reference to the metapragmatic strategies used in successive turns of dialogue.

4 WHY THEN, SPEECH ACT THEORY?

If speech act theory is so wrong so much of the time, and if empirical studies of actual language use in context have shown this for several decades, why does speech act theory persist in the academy?

The linguistic anthropologist Michael Silverstein provided an explanation back in 1979: Speech act theory is popular in Western academic circles because it fits so closely with Anglo-American cultural beliefs about how language does/should work. In other words, it represents our own unexamined cultural beliefs about language, dressed up in fancy academic garb. Western societies believe in effective means-to-ends relationships, or efficiency; they believe that utterances emerge from speaker intentions; and they believe that speakers are responsible for their words, and that the meaning of the words is ultimately "owned" by the speaker. The explicit primary performative — the canonical speech act, and the one that corresponds most closely to ACL communications — describes the conventionally understood activity that speaker and addressee are engaged in at the moment of the utterance ("I promise," "I warn," etc.). Because the explicit verb describes the activity it performs, it is *metapragmatic* — it comments reflexively on the act of speaking.

Silverstein noted that Austin's tripartite division of locutionary, perlocutionary, and illocutionary results from an objectification of three different ways of reporting "what happened" in a speech event: What was said, what was done, and what the effect was. Each of these, Silverstein notes, corresponds to a distinctly English syntactic construction used to report a prior speech event (e.g., "He said X."). Silverstein concluded "it is not by chance" that speech act theory matches precisely the syntactic and semantic properties of English language (1979, p. 213). After all, "illocutionary forces are distinguished [by Austin] only insofar as distinct explicit performative formulas can be recognized" (1979, p. 213).

Silverstein concluded "There is no reason why an ideology that grows piecemeal from various metapragmatic formulations of a language should show internal consistency, nor indeed give adequate analytic insight in areas of social practice" (1979, p. 214). In other words, speech act theory fails to explain the empirical data emerging from conversation studies, but there's no reason to expect it would be successful — considering that it did not emerge from empirical scientific research, but rather from a culturally specific ideology of how language works.

Rosaldo (1982), drawing on her Ilongot data, argued that speech act theory invalidly universalizes what are culturally particular views of human social action and intention (p. 212). Summarizing how speech act theory fails to capture communicative action among the Ilongot, Rosaldo (1982) likewise concluded that "certain of our culturally shaped ideas about how human beings act have limited our grasp of speech behavior." Ultimately, Searle's work is an ethnography "of contemporary views of human personhood and action" in the West, and not a universal theory of language use (Rosaldo, 1982, p. 228).

5 CONCLUSION

The ACLs used in MAS simulate some aspects of human communication, such as explicit primary performatives that communicate preformed agent goals and beliefs, that function as isolated utterances rather than units in an ongoing routine, and that are unproblematically related to the conversational context. However, empirical study has revealed that most human communication does not meet these criteria. And in particular, communication is more likely to have unintended emergent effects when communicative actions are implicit, collaboratively distributed across three or more participants, and sequentially distributed across strips of interaction. To date, no artificial societies come close to simulating these features.

There is some promising agent work that moves beyond speech act theory. Many computer scientists have explored the kinds of communication that are necessary to sustain group action (Cohen and Levesque, 1991; Grosz and Sidner, 1990; Jennings, 1993, 1995;

Cohen, et al., 1990a; Rich, et al., 2001; Tambe, 1997). Designers of such artificial societies are particularly interested in agents that do not necessarily have identical beliefs, as is common in complex real-world environments. What would hold a team together when individual members have distinct private beliefs about the shared activity? *Collaborative systems* thus touch on key issues in the theory of intersubjectivity (Garfinkel, Matusov, Wertsch, Sawyer, Schutz), which are related to long-known issues in distributed cognition: groups can collaboratively accomplish tasks even when members hold different information, and different understandings of the task (Cohen and Levesque, 1991). The need to continually manage intersubjectivity makes communication necessary (Cohen and Levesque, 1991, p. 489). Agents in teams must be able to reason with, and communicate about, group goals and actions, in addition to the individual representations and communications supported by the ACLs of cognitive agents.

Much of this theoretical work has explored the relations between individual intentions and joint intentions; for example, if a team jointly intends to perform an action involving a sequence of steps, then the agent responsible for any step "will intend to do that part relative to the larger intention" (Cohen and Levesque, 1991, p. 505). Philosophers who have influenced this simulation work (e.g., see the essays in Cohen, et al., 1990b) have debated a version of the micro-macro issue: is collective intentionality reducible to individual intentions, or does it necessarily have a distinct ontological status? Some philosophers (Gilbert, 1989; Searle, 1990) argue that collective intentionality is not reducible. But most collaborative systems do not explicitly model collective intentions (see Tambe, 1997, pp. 113–114).

Grosz (1996) and Tambe (1997) have developed some of the most sophisticated collaborative systems. In Grosz's *SharedPlan* system, different participants have different knowledge about how to proceed to solve a problem, and they must work together as equals to solve the problem. Agents have different beliefs, intentions, and capabilities; collaboration requires that agents have the ability to reason about other agents' beliefs and intentions; and agents must be able to collaborate in both planning and acting. Planning and acting are not sequential but interleave, as the unexpected results of actions force rethinking of plans; and the plans and actions that result in such systems are emergent group phenomena, emerging not only from agent interaction but also from the unpredictability of interaction with the environment.

Grosz and Sidner (1990) argue that these emergent group plans cannot be understood as the sum of individual plans (see Grosz, 1996, p. 73), and in fact, that agents must be designed to plan differently if they are to collaborate toward emergent planning (Grosz and Kraus, 1996). For example, agents must have intentional states that are directed toward the collective group plan; Grosz and Kraus (1996) identify two such states, *intending-to* (the intention of an agent to accomplish a subtask of the group's task) and *intending-that* (the indirect intention of an agent that is depending on another agent to accomplish a subtask). "Intending-that" carries implications for the nonacting agent — not to expect other subtasks to be accomplished at the same time, and a commitment to provide help when asked. Both states involve commitment and responsibility, even though only one of the agents is directly responsible for the subtask.

In these collaborative systems, collective intentions and group plans are emergent. They emerge when agents have a mutual belief in the plan — on the overall outline of how they are going to execute the plan. There must be individual or group plans for each of the subactions. And unique to groups, there must be an "intention-that" the group will do the action; and there must be individual commitments to other agents being able to do their actions, "intentions-that" the collaborators succeed. But before this emergent state can be reached, the agents have to undergo a process of reaching at least a partial agreement on a group plan.

Teams require their agents to have explicit representations of mutual beliefs, team plans, and team goals (Tambe, 1997, p. 85). Although no single agent performs the entire team activity, but only a part of it, each agent must be able to represent and reason about the group activity. These collective representations are not supported in cognitive agent architectures nor in the ACLs that they use. Tambe (1997) noted that prior systems, particularly Grosz's SharedPlans, typically have only two or three agents. Part of Tambe's motivation for developing his STEAM system was to support teams of more agents; systems have now been developed with 8, 11, and 16 agents. A classic example is the RoboCup soccer tournament, in which different research groups develop multiagent soccer teams with 11 agents as team members (www.robocup.org). STEAM systems start with joint intentions, and then build up hierarchical structures that parallel Grosz's theory of partial SharedPlans.

But most MAS simulation work continues to use either speech act-based ACLs, or even simpler communication mechanisms (think of IPD simulations). It's not surprising that computer scientists developing agent systems have found speech act theory appealing; after all, its assumptions about individual intentions and actions fit very well with the central contemporary traits of object-oriented programming. Objects are "encapsulated"; agents are "autonomous" — corresponding quite well to the individualist theory of personhood and intentionality held by Anglo-American speech communities. The ACL adaptation of speech act theory is even more rigidly constrained by explicit primary performatives than the original theory; each communication originates in a belief-desire-intention held by an agent, and is then communicated explicitly to another agent. ACL communications are essentially like programming language instructions — rather than "ADD" or "INCREMENT 1," we see "BID" and "EXECUTE."

Such ACLs have proven quite suitable for many computer science applications, particularly distributed Internet applications such as auction markets. But my claim here is that they are woefully inadequate for modeling human social life. And in particular, we have no hope of replicating how collective properties emerge from agent interaction unless we use much more sophisticated ACLs.

6 REFERENCES

Alexander, J.C., and P. Colomy, 1990, "Neofunctionalism Today: Reconstructing a Theoretical Tradition," in *Frontiers of Social Theory: The New Syntheses*, G. Ritzer, ed., New York: Columbia University Press, pp. 33–67.

Alexander, J.C., and B. Giesen, 1987, "From Reduction to Linkage: The Long View of the Micro-Macro Link," in *The Micro-Macro Link*, J.C. Alexander, B. Giesen, R. Münch, and N.J. Smelser, eds., Berkeley, CA: University of California Press, pp. 1–42.

Alexander, J.C., B. Giesen, R. Münch, and N.J. Smelser, eds., 1987, *The Micro-Macro Link*, Berkeley, CA: University of California Press.

Archer, M.S., 1995, *Realist Social Theory: The Morphogenetic Approach*, New York: Cambridge University Press.

Barbuceanu, M., and M.S. Fox, 1997, “Integrating Communicative Action, Conversations and Decision Theory to Coordinate Agents,” in *Proceedings of the First International Conference on Autonomous Agents*, W.L. Johnson and B. Hayes-Roth, eds., New York: ACM Press, pp. 49–58.

Barnes, B., 2001, “The Macro/Micro Problem and the Problem of Structure and Agency,” in *Handbook of Social Theory*, G. Ritzer and B. Smart, eds., Thousand Oaks, CA: Sage, pp. 339–352.

Brenneis, D., 1984, “Grog and Gossip in Bhatgaon: Style and Substance in Fiji Indian Conversation,” *American Ethnologist* **11**:487–506.

Carley, K.M., 2000, “Computational Social Science: Agents, Interaction, and Dynamics,” in *Simulation of Social Agents: Architectures and Institutions*, D. Sallach and T. Wolsko, eds., Chicago, IL: Social Science Research Computation, The University of Chicago, pp. 3–27.

Cohen, P.R., and H.J. Levesque, 1991, “Teamwork,” *Noûs* **35**:487–512.

Cohen, P.R., H.J. Levesque, J.H.T. Nunes, and S.L. Oviatt, 1990a, “Task-Oriented Dialogue as a Consequence of Joint Activity,” in *The Proceedings of the Pacific Rim International Conference on Artificial Intelligence*, Nagoya, Japan: Japanese Society for Artificial Intelligence, pp. 203–208.

Cohen, P.R., J. Morgan, and M.E. Pollack, eds., 1990b, *Intentions in Communication*, Cambridge, MA: MIT Press.

Coleman, J.S., 1990, *Foundations of Social Theory*, Cambridge: Harvard University Press.

Collins, R., 1981, “On the Microfoundations of Macrosociology,” *American Journal of Sociology* **86**:984–1014.

Conte, R., B. Edmonds, S. Moss, and R.K. Sawyer, 2001, “Sociology and Social Theory in Agent Based Social Simulation: A Symposium,” *Computational and Mathematical Organization Theory* **7**:183–205.

Duranti, A., 1988, “Intentions, Language, and Social Action in a Samoan Context,” *Journal of Pragmatics* **12**:13–33.

Duranti, A., 1992, “Intentions, Self, and Responsibility: An Essay in Samoan Ethnopragmatics,” in *Responsibility and Evidence in Oral Discourse*, J.H. Hill and J.T. Irvine, eds., New York: Cambridge, pp. 24–47.

Duranti, A., C. Goodwin, and M.H. Goodwin, 1991, “Communicative Acts as Socially Distributed Phenomena,” Preliminary remarks prepared for session on “Speech Acts as Socially Distributed Phenomena,” American Anthropological Association, Chicago, Nov. 24.

Ellis, D.G, 1999, “Research on Social Interaction and the Micro-Macro Issue,” *Research on Language and Social Interaction* **32**:31–40.

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

Fisher, L., 1976, “Dropping Remarks and the Barbadian Audience,” *American Ethnologist* **3**:227–242.

Gilbert, M., 1989, *On Social Facts*, Princeton, NJ: Princeton University Press.

Gilbert, N., 1999, “Simulation: A New Way of Doing Social Science,” *American Behavioral Scientist* **42**:1485–1487.

Gilbert, N., 2002, “Varieties of Emergence in Social Simulation,” in *Agent 2002: Social Agents: Ecology, Exchange and Evolution*, Argonne, IL: Argonne National Laboratory, pp. 41–60.

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

Goffman, E., 1981, “Footing,” in *Forms of Talk*, E. Goffman, ed., Philadelphia: University of Pennsylvania Press, pp. 124–159. (Original work published in *Semiotica*, 1979).

Goodwin, C., 1981, *Conversational Organization: Interaction between Speakers and Hearers*, New York: Academic Press.

Goodwin, C., and M.H. Goodwin, 1992, “Context, Activity, and Participation,” in *The Contextualization of Language*, P. Auer and A. di Luzo, eds., Amsterdam: Benjamins, pp. 77–99.

Graham, L., 1993, “A Public Sphere in Amazonia? The Depersonalized Collaborative Construction of Discourse in Xavante,” *American Ethnologist* **20**:717–741.

Grosz, B.J., 1996, “Collaborative Systems,” *Artificial Intelligence Magazine* **17**:67–85.

Grosz, B., and S. Kraus, 1996, “Collaborative Plans for Complex Group Actions,” *Artificial Intelligence* **86**:269–357.

Grosz, B.J., and C. Sidner, 1990, “Plans for Discourse,” in *Intentions in Communication*, P.R. Cohen, J. Morgan, and M.E. Pollack, eds., Cambridge, MA: MIT Press, pp. 417–444.

Irvine, J.T., 1996, “Shadow Conversations: The Indeterminacy of Participant Roles,” in *Natural Histories of Discourse*, M. Silverstein and G. Urban, eds., Chicago: University of Chicago Press, pp. 131–159.

Jennings, N.R., 1993, “Commitments and Conventions: The Foundation of Coordination in Multi-Agent Systems,” *The Knowledge Engineering Review* **8**:223–250.

Jennings, N.R., 1995, "Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems using Joint Intentions," *Artificial Intelligence* **75**:195–240.

Kristensen, B.B., and D.C.M. May, 1996, "Activities: Abstractions for Collective Behavior," in *ECOOP '96: 10th European Conference on Object Oriented Programming*, P. Cointe, ed., Berlin: Springer, pp. 472–501.

Labrou, Y., and T. Finin, 1997, "Towards a Standard for an Agent Communication Language," in *AAAI Fall Symposium on Communicative Actions in Humans and Machines*, Cambridge, MA.

Levinson, S.C., 1988, "Putting Linguistics on a Proper Footing: Explorations in Goffman's Concepts of Participation," in *Erving Goffman: Exploring the Interaction Order*, P. Drew and A. Wootton, eds., Cambridge, UK: Polity Press, pp. 161–227.

Macy, M.W., and R. Willer, 2002, "From Factors to Actors: Computational Sociology and Agent-Based Modeling," *Annual Review of Sociology* **28**:143–166.

Morgan, M., 1996, "Conversational Signifying: Grammar and Indirectness among African American Women," in *Interaction and Grammar*, E. Ochs, E.A. Schegloff, and S.A. Thompson, eds., New York: Cambridge University Press, pp. 405–434.

O'Connor, M.C., and S. Michaels, 1993, "Aligning Academic Task and Participation Status Through Revoicing: Analysis of a Classroom Discourse Strategy," *Anthropology and Education Quarterly* **24**:318–335.

Rao, A.S., and M.P. Georgieff, 1995, "BDI Agents: From Theory to Practice," in *ICMAS-95*, V. Lesser, ed., Cambridge: MIT Press, pp. 312–319.

Rawls, AW., 1987, "The Interaction Order Sui Generis: Goffman's Contribution to Social Theory," *Sociological Theory* **5**:136–149.

Reddy, M.J., 1979, "The Conduit Metaphor: A Case of Frame Conflict in Our Language About Language," in *Metaphor and Thought*, A. Ortony, ed., New York: Cambridge University Press, pp. 284–324.

Rich, C., C.L. Sidner, and N.B. Lesh, 2001, "COLLAGEN: Applying Collaborative Discourse Theory to Human-Computer Interaction," *Artificial Intelligence Magazine* **22**:15–25.

Ritzer, G., and P. Gindoff, 1992, "Methodological Relationism: Lessons for and from Social Psychology," *Social Psychology Quarterly* **55**:128–140.

Rosaldo, M.Z., 1982, "The Things We Do with Words: Ilongot Speech Acts and Speech Act Theory in Philosophy," *Language in Society* **11**:203–237.

Sacks, H., 1992, *Lectures on Conversation,* Volume 1, Cambridge: Blackwell.

Sawyer, R.K., 2001a, *Creating Conversations: Improvisation in Everyday Discourse*, Cresskill, NJ: Hampton Press.

Sawyer, R.K., 2001b, "Emergence in Sociology: Contemporary Philosophy of Mind and Some Implications for Sociological Theory," *American Journal of Sociology* **107**:551–585.

Sawyer, R.K., 2003a, "Artificial Societies: Multi Agent Systems and the Micro-Macro Link in Sociological Theory," *Sociological Methods and Research* **31**:37–75.

Sawyer, R.K., 2003b, "Communication Matters: Agent Communication Languages and Emergence in Artificial Societies," *Unpublished manuscript*, Washington University, St. Louis, MO.

Sawyer, R.K., 2003c, *Improvised Dialogues: Emergence and Creativity in Conversation*, Westport, CT: Greenwood.

Searle, J.R., 1990, "Collective Intentionality," in *Intentions in Communication*, P.R. Cohen, J. Morgan, and M.E. Pollack, eds., Cambridge, MA: MIT Press, pp. 401–415.

Silverstein, M., 1979, "Language Structure and Linguistic Ideology," in *The Elements: A Parasession on Linguistic Units and Levels*, P.R. Clyne, ed., Chicago: Chicago Linguistic Society, pp. 193–247.

Silverstein, M., 1981, "The Limits of Awareness," Austin, TX: Southwest Educational Development Laboratory.

Tambe, M., 1997, "Towards Flexible Teamwork," *Journal of Artificial Intelligence Research* **7**:83–124.

Wiley, N., 1988, "The Micro-Macro Problem in Social Theory," *Sociological Theory* **6**:254–261.

Wooldridge, M., 2000, *Reasoning About Rational Agents*, Cambridge, MA: MIT Press.

INVITED SPEAKER:

ASSESSING AGENT COMMUNICATION LANGUAGES

(Friday, October 3, 2003, 1:45 to 2:45 p.m.)

Chair: *David Sallach, The University of Chicago*
Discussant: *Gina-Anne Levow, The University of Chicago*

David Sallach: We're pleased to have with us today as our invited speaker Keith Sawyer, from Washington University. He is embarrassingly prolific, and he does really good, rich, in-depth, qualitative social science, as I think you'll see. Yet I don't think he'll be offended if I say he's kind of in the Goffman tradition, but also has been active in thinking through what that tradition brings to, and what kinds of issues it raises for, the emerging computational social sciences, and he has published in that area as well. And that kind of bridging is what he's going to speak on today, with references to agent languages.

R. Keith Sawyer: Okay, assessing the agent communication languages. Yes, my background is a little bit too interdisciplinary perhaps, but my undergraduate degree is in computer science, and I studied artificial intelligence. This was a long time ago, even before the AI winter of the mid-1980s. But then I went back to graduate school and decided I wanted to be an empirical social scientist. So most of my research has been in blends of linguistic anthropology, sociolinguistics, and conversation analysis. And that tradition's very different from mainstream linguistics.

Mainstream linguistics has been the area of language study that has most influenced computer science, and I have some ideas about that, probably because of the roots of linguistics in philosophy as opposed to empirical social science, philosophy of language or logical formalism, whatever. But what I'm going to talk about is going to be mostly data from empirical studies of real live conversations in the world and what that says about agent communication languages that are used in multi-agent systems.

[Presentation]

Sawyer: … so I guess ultimately when I'm assessing agent communication languages, I come up with a bunch of criticisms. That doesn't mean I don't think ACLs are going to be effective for a lot of computer science applications. I think they have been effective. But if you're a member of a community that's trying to use this technology to simulate human social groups or natural social groups, then I think you're going to have to *not* use ACLs based on speech act theory.

Sallach: I'd like to introduce Gina-Anne Levow, from the Computer Science Department at The University of Chicago. She's going to be the discussant for this talk, and then we'll maybe have time for a few questions …

Gina-Anne Levow: I've got a few comments from hearing today's presentation and the copy at the top, the paper that I luckily got in advance so that I could think about this a bit. And hopefully those will bring in some outside work that we can think about and use to stimulate some further discussion.

My personal background: I'm in the Computer Science Department at the University of Chicago, and my area of research is actually in what in this context would be more called virtual or conversational agents. I work on spoken-language systems and interactive systems that use speech, natural language processing and other multi-modal-type interfaces. So it's very appealing to me to hear a talk like this, which brings together such a range of communities, both from the strong perspective of the agents, as well as the sociology perspective, sociolinguistic, socio-anthropological, and some work in previous philosophy and artificial intelligence work.

And something that the speaker didn't have time to talk about as much in the talk, though, that showed up in the full paper, is that there has been a confluence of ideas; there are a number of active research areas that are working in agent systems and in coming up with ways to start to handle some of the issues that the speaker did raise here, in particular, things like meta-pragmatic uses and conversation about conversations. And in particular, a lot of people have recognized the breakdown in this conduit metaphor; in particular, the notion of a single speaker, single hearer and sort of a single concept that's being transmitted back and forth. The importance of back channels and socially constructed utterances has become progressively more clear as our capabilities in building spoken language systems, for instance, for virtual environments have brought us into confronting the fact that we're actually having multiple agents interacting, trying to have a communication.

There's been a lot of work in the AI community, and some very recent work on multi-agent conversational systems. Today's speaker mentioned in his paper work by Gross on shared plans, which endeavors to expand both the notion that an interaction is in fact not just two speakers building their own individual plans, but that they actually collaboratively build both a plan and a communicative process, as well as some work by Tumbay on team simulation in the context of robotic soccer.

Well, these approaches have been particularly promoted in some recent work that I'd like to mention out at ISI in California; in particular some work by David Traum, Jeff Rickel, Jay Gratch and S. Marsella on negotiation over tasks in hybrid human agent teams for simulation-based training. That's quite an arcane term. But what's interesting about this work is that actually tries to generalize very strongly on speech act theory to extend it through dialogue act theory.

In human/human communication, we have a wide range of very flexible forms. At what level do we want to incorporate things like indirectness, for instance, into our agent communication language, or do we want to stay abstract past that to some other intermediate representation? If we're dealing with a multimodal dialogue, is it important to capture the modality in terms of the lower-level representation, or do we simply want to abstract the fact that somebody signaled attention in some way without necessarily specifying that it was by gaze, by gesture, or by a back-channel "mm-hmm."

Another question that I think plays into that role is what our goals are. If we're building a language where our intention is to simulate a full human society, it becomes particularly important to model all of these facets. If we're building functional agents that might be running around in the background, say, simulating auctions for electrical power, it might be less important to actually represent some of the indirectness of speech acts. And I think that's going to be an ongoing question as we develop agent communication languages for different types of activities.

One last thing that I'd like to mention (just because it's interesting) was in the notion, when you mentioned, "Well, pop in the agents and the language, and, you know, add water," well, there have been some interesting studies by a researcher at the University of Illinois at Chicago, who I think is Gmytrasiewicz.

The most recent paper I've seen was called "Toward Automated Evolution of Agent Communication Languages." And what he does is, he takes a different premise, in which he says, "There are agents, there are interactions. Let's come up using, in this case, adaptive mechanisms, in particular, optimizations of expected utility — a language that enables optimal communication, and see what we get as a language," rather than taking the question of, "Well, if we have agents and a language, what do we get as a communicational structure?"

But now I'd like to open the floor for further discussion for whatever time we might happen to have. Thanks. Comments, questions?

Unidentified Speaker: I found the paper very compelling and found myself really in agreement with the critique. It does seem to me to be posing a question, but not suggesting any answers, in the sense that, speaking as one of those who would love to come up with a way to build an agent-based model that showed some emergent social effects of the kind I think you're giving evidence of, it's pretty clear that most of the models we build really do use some version of a BDI-like framework, because the representation falls to hand. I mean, it's sitting out there … something easy to grab hold of and use to write lines of code.

And I was wondering if you had any suggestions at all about places we might look, theoretical structures that at least someone has invested some time in, that might be a place to at least begin puzzling over as a way it might be turning an algorithm or broadened model so that we can pick up some new phenomenology in our ABM constructs.

Sawyer: Yes. Well, like Gina said, I did have some more positive suggestions at the end of the paper that I left out for lack of time: collaborative agents, joint intentions. So there is some work that is trying to focus on collaborative seams that I didn't have time to talk about, that is not based on ACLs.

And I guess, yes, I would recommend, although I don't know how practical this, the use of relevant work in writings of linguistic anthropologists, some of the same people I cited. So I think it would always be a good place to start with empirical social science work first, and then work on the simulation. So, yes, there's a lot of good stuff out there. And, like I said, it goes back to the late 1970s. So it just hasn't filtered over.

Unidentified Speaker: I think your comments about our science or our agent models being based on individual perspectives is right on target. And I wholeheartedly support you. As an issue that many commentators have raised, though, it's hard to go from that criticism to a solution. And I'm not sure that you can create a sufficient glue by just noticing that by removing context you can actually come up with a better situation.

I agree with you that removing context actually creates more problems than it can solve. On the other hand, I don't think that adding context is going to provide a solution as well. It seemed to me that the thing you didn't talk about was conflict. You didn't address that as an issue. And I think until you get into conflict, where you could see resource conflicts, sexual conflicts, political conflicts, whatever they are, that you don't introduce function. And I think function is possibly — I'm not sure yet in my own mind — a glue that pulls these things together. And maybe you'd like to comment on that.

Sawyer: Sure. Yes, I think that in a model where everyone is assumed to be cooperating, then there's no conflict, right? I think that in real human small groups, that in situations of conflict and negotiation and argumentation, you're often more likely to see implicitness and indirectness and confusion about motives, and intentional confusion about motives.

So, yes, I guess I'd agree that that's important. But I would say if you want to be able to understand conflict, then my points probably become even more telling.

Unidentified Speaker: I agree with you there. What about function, social function, function in the decision … discussion process …

Sawyer: You mean the function of a speaker utterance?

Unidentified Speaker: Speaker utterance, receiver utterances. Function of the social process, whatever they're doing; why they're doing what they're doing.

Sawyer: Oh, okay. So in a sense a part of context, maybe.

Sallach: Yes…. And it turns out language is difficult, actually, and we'll probably have to draw this session to a close.

Agent Architectures

CONVERSATIONAL AGENTS

D. SYLVAN,* Graduate Institute of International Studies, Geneva

ABSTRACT

Conversations (i.e., extended verbal exchanges) are among the most salient types of social interaction. Until now, agent-based models have not focused much on this phenomenon, although in principle such models offer significant advantages as ways of exploring it. The internal and external aspects of conversational commitment are identified, and a research agenda is proposed for constructing agent-based models of such interactions in two domains: social gatherings and diplomatic negotiations.

Keywords: Agent-based models, conversation, turn-taking, third parties

INTRODUCTION

Almost a century ago, Simmel defined society as existing "where a number of individuals enter into interaction.... [Motivations] are factors in sociation only when they transform the mere aggregation of isolated individuals into specific forms of being with and for one another, forms that are subsumed under the general concept of interaction.... If, therefore, there is to be a science whose subject matter is society and nothing else, it must exclusively investigate these interactions...." (Simmel, 1908, Chap. 1). Few would disagree explicitly with this formulation even today, especially those concerned with the kind of modeling discussed at this conference. However, for all their variety and richness, the kinds of agent-based models constructed until now have tended to focus on certain types of strategic and habitual interactions, scanting arguably other important ones.

The purpose of this paper is to focus on one of those less-studied forms of interaction, namely, extended verbal exchanges (i.e., conversations). I argue first that conversations, whether among private individuals or in some kind of workplace setting, perhaps between persons serving as delegates, are of considerable importance. Second, I argue that conversations are best modeled by methodologies that take account of certain semantic and pragmatic characteristics, and that these desiderata are not satisfied by either existing agent-based approaches or certain other formal methodologies. Third, I claim, this is a pity, because a more promising tradition, namely conversation analysis, is badly in need of the potential advantages offered by agent-based modeling. Fourth, I identify internal and external aspects of conversational commitment and propose a research agenda for constructing agent-based models of these phenomena in two domains: social gatherings and diplomatic negotiations.

In making these arguments, of necessity I will be operating at a relatively high level of abstraction. The reason is not only that this paper is primarily methodological and meta-theoretical, but also, as I discuss below, because conversations have an important indexical quality; that is, they are "about" something in particular (cf. Garfinkel, 1967; Sacks, 1992,

---

* *Corresponding author address:* David Sylvan, Graduate Institute of International Studies, 11a, avenue de la Paix, CH-1202 Geneva, Switzerland; e-mail: sylvan@hei.unige.ch.

Vol. 1, Part 4; Brentano, 1874; James, 1890). This, in turn, means that models of conversation must not only vary considerably by both domain and subject matter, but that for them to be constructed as plausibly valid in the first place, they also must take into account a large number of actual examples. In this case, intuition is insufficient; model-construction must go hand-in-hand with data collection.

SIGNIFICANCE OF CONVERSATIONS

Before proceeding further, some conceptual clarification is in order. By interaction, I do not mean simply a contact between one person and another, such as might happen in an accidental collision between two persons.[1] Rather, as in the Simmel quotation above, interactions are characterized by the actions of persons, which are ways of being *with* and *for* each other, or, to build on the way in which Weber put it, actions that are "meaningfully oriented to that of others" (Weber, 1956, p. 23). Thus, if conventionally, action is distinguished from mere behavior by dint of the former's intentional quality, then, in turn, interaction is characterized by the fact that each of the parties in the interaction is intending that its actions will be relevant to the other parties (i.e., understood as such and reacted to in particular ways).

From this definition, it is clear that verbal exchanges are interactions. One need not espouse the notion that society is "socially constructed" by "discourse" to recognize that when two or more actors talk (or write) to each other, they are precisely intending to have their utterances understood by the other party and, quite often, others as well; and that certain consequences should flow from that understanding. For example, a judge who says, "I hereby sentence you to five years in prison," to a defendant in a trial is interacting with the defendant and many others, including the prison system. In this sense, verbal exchanges can be seen as being comprised of "speech acts" (Searle, 1969), even if, as is discussed below, the speech act approach misses essential features of most verbal exchanges.

Many of the most important types of interactions are comprised of such exchanges. For example, political campaigns and debates are verbal; so too are most legal proceedings; and the same can be said of negotiations. But at least as important are far more mundane types of verbal exchanges: classroom teaching, conversations among friends or co-workers, and arguments and disagreements. In short, verbal exchanges make up both a high-profile and an existentially salient quotidian part of social life. This is particularly the case for certain types of verbal exchanges, those in which interaction is extended long enough to be referred to later in that interaction or recalled in subsequent ones, thereby permitting commitments to be established and secondary interactions related to such commitments to occur (Schegloff, 1996a). A good example of this, of considerable political significance, is the way in which, during one of their telephone conversations in August 1968, the Soviet leader Brezhnev berates the Czech leader Dubček for having failed to implement political actions to which he had earlier, at length and on repeated occasions, agreed (Prague Spring, 1998). Of course, there is no *a priori* correct criterion by which momentary exchanges can be distinguished from extended ones, but the distinction is nonetheless worth trying to specify.

---

[1] An example, interestingly, given by both Simmel (1917, Chap. 1) and Weber (1956 [but drafted, 1918–1920], Chap. 1).

In the face of these observations, it may perhaps be suggested that extended verbal interactions are so much flotsam and jetsam: what counts is some kind of physical action and not the words that accompany it. This may well be true for certain types of social interactions, such as marketplace transactions. (Even there, though, the verbal accompaniment can still take on significance if it deviates from established norms. Imagine a salesperson who, on accompanying a customer to the cash register, talks to the latter in a rude and insulting way; we would expect that the sale would quickly be aborted.) But many verbal interactions have only a tenuous link, if that, to any kind of physical behavior. Take, for example, an argument in an appeals court over which precedent should be controlling in a certain case; what physical phenomenon is being pointed to here? Or consider certain types of diplomatic exchanges, such as a carefully scripted apology; this may, of course, open up future commercial transactions, but in and of itself, it is impossible to understand without reference to the words involved.

In response, it might be agreed that social interactions are not necessarily about physical behavior, but that to focus on verbal exchanges is to miss the real story, which is the overlap of preferences or lack thereof among the different parties. I address this issue in more technical terms below, but note for now that, if taken seriously, the objection leads to a *reductio ad absurdum*: actors would arrive at preferences as to what they want from other actors without taking into account what the latter are deemed to want (or, even more absurdly, making inferences about this while not attending to anyone's words about it). Hence, it is difficult to escape the conclusion that, no matter what one's views on the appropriate ways of studying extended verbal exchanges, at least certain types of them are significant social interactions.

MODELING DESIDERATA

There are several ways of thinking about extended verbal exchanges and modeling their various features. I briefly review four candidate possibilities, arguing that only one — conversation analysis — captures particularly important features of such exchanges, even if, as we see below, it is significantly limited.

The first, possibly, most prominent way of thinking about verbal exchanges is that they are collections of speech acts. Utterances may be seen as having an illocutionary force, for example, a request or an assertion on the one hand, a granting of the request or a denial of the assertion on the other. In this way, a verbal exchange can be modeled as an argument, a series of understandings, or a set of mutual commitments (Rescher, 1977; Bach and Harnish, 1979). One "output" of such types of models is a conditional proposition about what kinds of speech acts are possible for each party at different stages of an exchange (see Duffy and Goh, 2003, for a preliminary example).

Much can be said for this approach to verbal exchanges. Certainly, many kinds of policy-related exchanges are, in vital respects, arguments, and the parties to these exchanges seem to abide by many of the "rules" of arguments, including more general norms of relevance and nonrepetition. However, many conversations are not arguments, and even arguments have elements not usefully reducible to collections of speech acts. Consider, for example, the following telephone conversation between the soon-to-be assassinated South Vietnamese president, Ngo Dinh Diem, and the U.S. Ambassador Henry Cabot Lodge, whom for months had been advocating a *coup d'etat* to depose (though not necessarily kill) Diem (Diem, 1963).

1. Diem: Some units have made a rebellion, and I want to know: What is the attitude of the U.S.?
2. Lodge: I do not feel well enough informed to be able to tell you. I have heard the shooting but am not acquainted with all the facts. Also, it is 4:30 a.m. in Washington, and the U.S. government cannot possibly have a view.
3. D: But you must have some general ideas. After all, I am a Chief of State. I have tried to do my duty. I want to do now what duty and good sense require. I believe in duty above all.
4. L: You have certainly done your duty. As I told you only this morning, I admire your courage and your great contributions to your country. No one can take away from you the credit for all you have done. Now I am worried about your physical safety. I have a report that those in charge of the current activity offer you and your brother safe conduct out of the country if you resign. Had you heard this?
5. D: No. (And then after a pause) You have my telephone number.
6. L: Yes. If I can do anything for your physical safety, please call me.
7. D: I am trying to re-establish order.

This conversation operates on several levels. It is, to begin with, a request by Diem for information as to the American attitude, a request to which an answer, though explicitly avoided by Lodge, is nonetheless implicitly responded to via Lodge's change of topic from the coup as political event to Diem's safety. The conversation is also an argument: Diem tries to convince Lodge, by playing on the notion of duty, that he deserves guidance and, perhaps, a statement that the rebels are acting in an undutiful fashion. Lodge, while not agreeing to this, responds by making a counterproposal that Diem look to his own safety, perhaps by trading his post as president for exile; this, in turn, is implicitly rejected by Diem, who fails twice to respond to the possibility of physical safety. In addition, Diem attempts to have Lodge take a more proactive role by reminding him that he can call Diem if he wants; Lodge tries to escape this by offering to have Diem call him instead.

Just which speech acts are each of these persons engaging in? Is Diem making a request for information or an argument that Lodge must possess this information and is duty-bound to convey it? Is Lodge rebutting Diem's argument, or is he accepting some responsibility? In fact, it seems as if any particular utterance is performing double- or even triple-duty, in which case the same verbal exchange can be seen as giving rise to alternative speech act mappings (not very good from the standpoint of parsimony). Otherwise, the participants must be generating new, hybrid speech acts, again reducing the advantages of having models built on small numbers of concepts. Indeed, we can put this as a trade-off. If there is a pre-set list of speech acts available to conversational parties, either their utterances will fail to be mapped onto that list in a consistent fashion or the list will be so long that it will be close to indistinguishable from the utterances themselves (Sylvan, 2003).

Moreover, the same speech act — each placing responsibility on the other — appears to be engaged several times, which seems repetitive and violative of relevance maxims. In fact, neither actor is repeating himself. In turn 3, Diem poses the issue of U.S. knowledge not as a question (which Lodge had just avoided) but as a statement, which he insulates from rebuttal by emphasizing the issue of duty. Similarly, in turn 6, Lodge returns to the issue of physical safety, this time not as a piece of information (just disclaimed as knowledge by Diem) but as an offer to be triggered by Diem. Even the three-time use of the word "duty" in turn 3 is not idle repetition: it is a way of highlighting the point so that Lodge is forced to address it and thereby admit the point, implicitly rendering Diem praiseworthy. What we see here is what we can call the production of agreement by both speakers: the way in which one speaker's utterance is picked up immediately by the other and recrafted in such a way as to induce agreement in the next turn (Sacks, 1987; Schegloff, 1996a). Such agreement is easy to see at the coarser level of "brainstorming" or "problem-solving" discussions (good cases in point are the exchanges among Clemenceau, Lloyd George, and Wilson, or between Roosevelt and Stalin: Yalta, 1945; Schwab, 1994), but it also exists at a finer-grained level of consecutive utterances. Often, in fact, interlocutors seem to go out of their way to improvise agreement on issues highly local to that particular moment in the conversation.

Thus, it appears that the speech act approach to verbal exchanges is too restrictive. A similar objection can be formulated to a second approach, that which models exchanges as dialogues made up of connected topical sequences (Moeschler, 1985, 1989). The idea is that participants in dialogue move up, down, or "horizontally" within hierarchies of topics by means of different connectors such as "well," "good," "but," and so forth. This approach has the advantage of not confining itself to particular pre-set speech acts; instead, topics, which can be of considerably greater variety, are what dialogues enter, traverse, and exit. Because the connectors are standard words or phrases, the mapping problem discussed above is also avoided.

In fact, there is not much difference between dialogical modeling and speech act analysis. This can be seen clearly if we look at how topics are conceptualized and identified. A topic is understood as the "about-ness" of a proposition: a request to do X, a refusal to admit Y, a promise to enact Z, where the respective topics are X, Y, and Z. Already we can see that topics are linked to speech acts, but not only in terms of their "function" in a dialogue. For dialogical analysis to proceed, topics must be clearly identified so that a given utterance can be identified clearly and situated correctly in the hierarchy. This poses some of the mapping problems discussed earlier. Consider the following snippet of conversation in July 1958 between Eisenhower and Macmillan (Lebanon, 1958). The former has just announced that, given the day's events (a *coup d'etat* in Iraq, the murder of the pro-Western premier, and a request for intervention by the president of Lebanon), he has ordered U.S. troops to intervene in Lebanon. Macmillan tries to press for a broader intervention extending as far as Iraq; Eisenhower refuses. After several exchanges, the conversation arrives at this point:

1.	Eisenhower:	Well, now, I will tell you, of course, I would not want to go further. Today we tried to project in our discussions here and with the legislative leaders the development of the situation, and they could take many forms. If we are now planning the initiation of a big operation that could run all the way through Syria and Iraq, we are far beyond anything I have power to do constitutionally. We have had quite some trouble justifying to our own leaders what we intend to do.
2.	Macmillan:	Yes. What is your timetable?
3.	E:	Right now. ... I don't know the exact time that they will get there on account of the orders and hours, but I would not want to give any information over an open wire.
4.	M:	Of course. Now, are you going to speak to the country?
5.	E:	Yes. This is very secret. We are calling an emergency session of the Security Council for tomorrow morning. I will broadcast after that Council does something.

6. M: Tomorrow?
7. E: Well, probably.
8. M: Well, now, we have had a request from the two little chaps — the one is gone and the other is there, the king ...
9. E: We did not know what the final reports were.
10. M: I know there is little news. The second is going along for the other. We have got a sort of request from him saying that [what] we are going to do. I feel, my dear friend, that if you set off this great show, which I think is fine, you can't confine it to what you say publicly, but in fact all the trouble will blitz through on destroyers, oil fields, pipelines. Taking on Turks and getting things back. We should be ruined. I am for it. I don't want you to say that now to me, but so long as I understand we are in this together. We are doing this together.
11. E: My own idea would be this. If the situation develops where our whole national interests are abandoned and destroyed, I have to go before the Congress and ask for authority to act. We can understand and agree on that much. And that is exactly what you say, except I have to say it in my guarded terms.

Do we have here changes in topic? From one point of view, the answer is yes, since Macmillan, having been rebuffed, shifts the topic to the mechanics of the operation and its announcement, then raises another issue, a "sort of request" for intervention from the King of Jordan and (presumably) the now deposed Iraqi leader. Macmillan then returns to his idée fixe, which is that the intervention will have to be broader than Eisenhower envisages. From another point of view, though, there is not a change in topic at all: Macmillan is trying to find out Eisenhower's level of commitment and playing on the idea of an obligation to respond to requests for intervention as a way of expanding the intervention. To a limited degree, he in fact succeeds. Presumably, these various moves by Macmillan would be coded as a set of intervention subtopics, subordinate to the main topic.

The problem with this way of proceeding is that the hierarchical topic structure only emerges in the middle of turn 10, which links turns 8 and 4 into the "broad intervention" topic. But Eisenhower's utterances in 5 and 9 are only linked to that topic when he plays along in turn 11; had he simply reiterated his earlier refusal, the conversation would be made up of only juxtaposed fragments. Topics thus are cooperative achievements whose identification and presumed structure only occur in real time, in the course of the verbal exchange (Schegloff, 1996b). In spite of the fact that participants in conversations (especially diplomatic exchanges) often have an agenda of particular points they wish to make, their ability to do this is contingent on the real-time construction of the conversation. Both speech act and dialogical approaches to verbal exchanges are in this sense static and top-down.

A third approach to modeling conversations has more immediate resonance with the agent-based tradition. Agents can be seen as communicating with each other, with conversations consisting of a set of offers and requests on the one hand, and responses, on the other. These communications may be either "public," broadcast generally (e.g., by means of placing messages on a "blackboard"), or targeted at particular agents. The communicative responses, the responses to them, and so forth, can be represented as finite state automata. However, for agents engaged in multiple communications with various other agents, it is preferable to model these interactions as Petri nets (i.e., as oriented graphs linking places to transitions via augmented transition networks [respectively, Winograd and Flores, 1986, Chap. 5; Magnin and Ferber, 1993; see more generally Ferber, 1999, Chap. 6; and Kurt, 1992]). This latter approach, in particular, is quite promising because it permits modeling simultaneous verbal exchanges between multiple agents, something very much absent from the dyadic orientation of speech act and dialogical approaches.

Unfortunately, communicative models are only as good as the building blocks from which they are constructed, and, once again, it turns out that up until now, most work in communicative multi-agent systems has employed some type of speech act set of components (e.g., KQML; see Labrou and Finin, 1997, for a critique). This leads to the top-down and pre-set problems already discussed. In addition, casting verbal exchanges as communication runs the real risk of missing much of the noninformational (notably pragmatic) aspects of conversation. Consider the following example, in which Kissinger is negotiating with Brezhnev over nuclear weapons (Kissinger, 1974).

1.	Brezhnev:	We should both scrupulously observe the agreement. You are refusing to take into account forward-based systems. At whom are these aimed? Not against France, because France can't declare war on the United States.
2.	Kissinger:	But this may change if things keep up!
3.	B:	Or Holland or Belgium, or the GDR or the FRG. I can show you a map. You said the agreement should relate to American missiles that could reach the Soviet Union and Soviet missiles that could reach the United States. That is the significance of those forward-based missiles. They can reach Tashkent, or Baku.
4.	K:	The submarines?
5.	B:	Yes. And air bases. More than one-half of the European part of the Soviet Union is within range of those.
6.	K:	We have to separate the problems. First of all, if M. Jobert makes more of his speeches, we'll need some of those missiles against France.
7.	B:	You can't blame me for that! No speech ever caused destruction, only weapons have.
8.	K:	This shows submarines?
9.	B:	It shows all kinds of bases and ships.
10.	K:	So this line is the range of the submarines, and they're being counted. They are part of the agreement. They are not forward-based systems. They are counted in the Interim Agreement.
11.	Gromyko:	But they are pointed at us — whether submarines or carrier-based aircraft. The first agreement left aside strategic aviation.
12.	K:	I agree with that. That's a separate problem. These are our fighter aircraft?
13.	B:	It's not a good picture, is it? Those are European-based aircraft carrying nuclear weapons. Then nothing else remains for us but to have our aircraft carrying nuclear weapons or missiles.

We know that both Kissinger and Brezhnev went into the meeting with a pre-set, highly structured set of proposals they intended to make to the other side. But Brezhnev, in setting up his proposal, presents information aimed at putting Kissinger on the defensive, a goal which he accomplishes. Kissinger first tries to deny the relevance of the information, then to deflect its significance with a joke, and then to argue that the information has already been taken into account; finally, he is reduced to a simple recognition, giving Brezhnev an opening for his (subsequent) proposal. To call this the communication of a speech act by Brezhnev is to miss the point: it is Kissinger, not Brezhnev, who is induced (through both conversational norms of courtesy and relevance and his own desire to show that he understands the significance of the map) to specify the significance of the information. This, in turn, makes it more difficult for Kissinger later to reject Brezhnev's proposal. To model this exchange as a proposal by Brezhnev, responded to by Kissinger, will miss much of the interaction that results in their eventual agreement.[2]

What is needed, therefore, is a fourth approach to verbal exchanges that better captures the bottom-up, genuinely interactive aspect of conversation than the various speech act-based approaches do. An obvious candidate is the sociological subfield of conversation analysis; although it has significant problems, some of which I will mention below, it nonetheless satisfies the above desiderata. The reason for this has to do with the basic standpoint of conversation analysis: that conversations are locally produced social achievements. This perspective, originating in work by Schutz and Garfinkel, can be seen most clearly in the basic components, which conversation analysts see as constituting the core elements of verbal exchanges: namely, the notion of turn-taking (Sacks, et al., 1974; Clayman and Maynard, 1995).

At the heart of any conversation are "adjacency pairs": turns, i.e., utterances (understood now not in a speech act sense), which refer backward to their immediate predecessors.[3] This reference can take a number of forms, ranging from simple acknowledgment that the other party has spoken to an explicit response to the preceding turn. What matters is that there is a felt obligation of each party to do something following on a turn by the other party. In this sense, conversations are shot through with locally produced, usually improvised, coherence (Schegloff, 1984, 1990). Often, indeed quite often, a given turn may be ambiguous, and the other party then has to try and resolve the ambiguity, either by a turn that binds the preceding turn to a particular meaning or by indicating a lack of clarity. This "work" evinces a norm of consistency, a norm that operates by appealing to what we might call a default social order. A famous example of this is Sacks's (1972) discussion of a child's "story": "The baby cried. The mommy picked it up." Sacks points out that the standard interpretation of these two sentences is that the baby was picked up by its mommy, and that in making this interpretation, listeners make the sentences coherent by presuming a particular social order. In the same way, as we saw previously, Lodge connected Diem's sentences by tacitly acknowledging U.S. authority, Eisenhower connected Macmillan's sentences by tacitly acknowledging U.S. obligations to requesters, and Kissinger connected Brezhnev's map by tacitly acknowledging the threatening nature of U.S. tactical capabilities.

Note that the local and bottom-up nature of coherence as seen by conversational analysis is far more deeply social than the various top-down approaches discussed earlier. Conversations

---

2 A technical point. Negotiations (and, for that matter, ordinary conversations) frequently involve constructions of elaborate scenarios of hypothetical events, the aim being to demonstrate to the other party that a given course of action would have certain consequences, the latter being seen clearly as strongly desirable or, conversely, as highly noxious. It is difficult to imagine how these hypotheticals could all be anticipated, such that the response could be specified by the modeler in one of the transition rules.

3 Strictly speaking, an adjacency pair is a turn following on another turn; the first in the sequence may, as in the opening of a conversation, not be preceded by anything. However, if one generalizes this notion to any pair of turns, then one moves back in the direction of "topics"; it therefore is cleaner analytically to depict each turn as referring only backward.

only take place, in this sense, if the parties to it understand each other and tacitly collaborate. This does not mean they have to agree in a propositional sense, but they have to play their role in keeping the conversation going and moving it along. By contrast, the other approaches see no immanent necessity of forward motion by participants. A given speech act could at any time terminate the conversation without any problem except a possible violation of norms of courtesy (one might ask why they are even observed). For the other approaches, conversational participants operate side by side, not together, as is the case with conversation analysis.

ADVANTAGES OF AGENT-BASED MODELS

As presented above, conversation analysis is a promising approach to analyzing extended verbal exchanges. However, most studies within the field suffer from three problems:

- They tend to be pointillistic, focusing on particular mechanisms of local coherence without connecting those mechanisms into a broader picture.

- They insist, almost obsessively, on working with actual conversations, neglecting the exploration of plausible potential conversations.

- They concentrate almost exclusively on individual conversations, failing to examine how the same actor may engage in a series of conversations, with a given conversation following on earlier ones.

Each of these weaknesses can be addressed if certain features of conversations are modeled using agent-based approaches.

Pointillism

Agent-based models, as a particular class of computational models, require mechanisms to be specified precisely. This in itself facilitates explicitness as to the connection between mechanisms, but even if no such connections are laid out, the ability to run multiple, highly iterated simulations provides built-in means by which they can be abduced. What is needed is simply two or more mechanisms of local coherence. This is not a problem, as the conversational analysis literature lists a number of such mechanisms.

The real problem in modeling coherence is to find an analytical terminology with which to represent turns as adjacency pairs. It is evident from the discussion of modeling desiderata in the preceding section that an attempt to build a language from scratch is likely to run into the same problems we saw with speech act and other top-down approaches. What is needed instead is a hybrid approach (see below) in which modeling proceeds from already coded conversations; both the input and the output for each turn would therefore be indexical, semantically local utterances.

Actuality

The reason that conversation analysts work so closely on actual conversations and eschew synthetic ones is a concern that the latter are almost sure to have been "normalized" into top-down, already smoothed-out form. Hence, the near "fetishization" of transcripts made from audio and video recordings of actual conversations. Nonetheless, even if the latter should remain the starting point for conversational modeling, there is no reason that coherence mechanisms and local utterances cannot be used as the basis for generating potential conversations. With a sufficiently extensive database, it should even be possible to give a preliminary validity assessment of some of the generated hypotheticals.[4] Here, too, the fact that agent-based models are simulations makes this a tractable task.

Sequences of Conversations

The preceding two advantages have nothing in particular to do with agent-based approaches; they pertain to any type of computational model. However, when we turn to the question of multiple conversations, the features of agent-based models come to the fore. It is striking how much conversational analysis has tended to focus on self-contained exchanges between two persons. This emphasis may be due to the provenance of recorded conversations or the limitations of recording technology, but it clearly represents a reduced notion of social interaction. Most conversational analysts acknowledge this, and it is not uncommon to see brief discussions of multi-party conversations. (Multiple sequences involving at least one and perhaps both members of the original conversational dyad are even rarer.) Unfortunately, such discussions are not terribly illuminating, no doubt because of the complexity of the different conversational combinations. If, however, connections between conversations can be specified, say, by means of allusion, commitment, or other mechanisms (see below), it is possible to use agent-based means to model an entire conversational ecology.

Implication

As I am arguing the matter, there is thus an elective affinity between agent-based models and the analysis of extended verbal exchanges. Such exchanges are important forms of social interaction whose neglect has rendered our agent-based models more limited than is desirable; models of this sort seem particularly well-suited to the analysis of conversations.

RESEARCH AGENDA

Two particularly interesting conversational phenomena are what we can call the internal and external aspects of conversational commitment. In any conversation, it is possible for participants to recall what they earlier discussed, if only by alluding to it (perhaps jokingly). This phenomenon, which I call commitment, not only serves as a means of trans-local coherence but, from an interactive perspective, highlights to the participants the social link they established earlier in the conversation. This does not mean that the participants are necessarily closer to each

[4] One can imagine, for certain conversational domains, even checking certain adjacency pairs by experimental means.

other or friendlier, but it does mean that their interaction can now be referred to as a second-order phenomenon, thereby facilitating its characterization as socially (and normatively) significant.

Logically, commitments are of two sorts. One is internal to particular conversations — when participants at some point recall what they discussed in the same conversation. The other is external to particular conversations — when one or more parties to a conversation recall in a later conversation, whether between them or with third parties, what they had discussed in an earlier conversation. This second type of commitment permits conversations, as it were, to be connected, not as some sort of legally binding precedent but as a way of rendering dyadic social relations quasi-public, and hence as at least trans-temporally, if not trans-dyadically, significant.[5]

Note in passing that for now I am not attempting to model local coherence in verbal exchanges. This is not because it is unimportant, but because it seems so open-ended as to be hopeless, unless, that is, one wishes to return to the speech act approach I criticized earlier. Instead, my proposal is to take local coherence as a starting point and to try and generate later links to it. For example, we might posit a mechanism whereby acknowledgment of blame by one party leads the other party to differentiate subsequent discussion to make clear that the first party is not being blamed again; conversely, the second party may explicitly recall the acknowledgment in later conversations with certain third parties.

As even this hypothetical example makes clear, such modeling efforts require that the starting point be the early or middle stages of an actual conversation. This is due to the open-ended quality of turn-taking mentioned earlier; and so, my proposal is to piggyback on existing instances of coherence to generate certain types of subsequent internal and external commitment. In fact, the first step in a modeling effort of this sort is to amass a significant corpus of conversations in a particular domain, ideally one in which the participants both speak at length and speak subsequently with each other and third parties. Only with such a corpus is it possible to adduce various hypothesized recall mechanisms; and only with such a corpus can such mechanisms be tested for the plausibility of their generative capacities.

The examples in this paper are drawn from the domain of interstate diplomacy. This domain, which is the subject of my substantive research, is of obvious significance. Unfortunately, the availability of verbatim transcripts is somewhat less than one would wish. Partly, this is a matter of secrecy; it is not uncommon for transcripts either to be classified for decades or never to be generated (policymakers are known to discuss certain sensitive issues privately, excluding even trusted aides and secretaries). But partly it is because until recently, there was a convention for writing summaries of conversations rather than recording them by means which would later permit transcripts to be written up. For my purposes, this is problematic, because it is precisely the little, often embarrassing, interchanges which produce some of the most interesting and commitment-generating types of local coherence.

Thus, I propose supplementing the domain of diplomacy with a second, apparently quite different, domain: that of social gatherings among "ordinary" persons. Here, a different problem

5 It is tempting to see external commitment as the conversational counterpart to social networks. This may be the case, but there is no presumption here that multi-party, cross-conversational recall signifies that the parties are in some sort of interlinked network arrangement. However, many of the contact mechanisms proposed by network theorists presuppose some types of regular, extended verbal exchanges.

exists: because most of these interactions are not "problem-solving" ones, there would seem to be less incentive to engage in practices of external commitment. (This should to some extent be mitigated by close physical proximity among participants in social gatherings.) Nonetheless, this second domain appears likely to be complementary to the first.

REFERENCES

Bach, K., and R.M. Harnish, 1979, *Linguistic Communication and Speech Acts*, Cambridge, MA: MIT Press.

Brentano, F., 1874, *Psychologie vom empirischen Standpunkte*, Leipzig: Kraus.

Clayman, S.E., and D.W. Maynard, 1995, "Ethnomethodology and Conversation Analysis," in P. ten Have and G. Psathas, eds., *Situated Order: Studies in the Social Organization of Talk and Embodied Activities*, Washington, DC: International Institute for Ethnomethodology and Conversation Analysis and University Press of America, pp. 1–30.

Diem, 1963, Telephone conversation between Ngo Dinh Diem and Henry Cabot Lodge, U.S. Embassy, Saigon, to Washington 860, 1 November 1963, reprinted in *The Pentagon Papers: Senator Gravel Edition*, Boston, MA: Beacon Press, **2**:268.

Duffy, G., and E. Goh, 2003, "From Tacit to Secret Alliance: Pragmatic Analysis of Kissinger's 1973 Rapprochement Dialogues with the Chinese Leadership," presented at the *Annual Convention of the International Studies Association*, Portland, OR, February 25–March 1.

Ferber, J., 1999, *Multi-agent Systems: An Introduction to Distributed Artificial Intelligence*, Harlow, UK: Addison-Wesley.

Garfinkel, H., 1967, *Studies in Ethnomethodology*, Englewood Cliffs, NJ: Prentice-Hall.

James, W., 1890, *The Principles of Psychology*, New York City, NY: Henry Holt.

Kissinger, H., 1974, Conversation in Brezhnev's Office, March 25, 1974, 11:05 a.m. — 1:57 p.m., transcript reprinted in W. Burr, ed., *The Kissinger Transcripts: The Top Secret Talks with Beijing and Moscow*, New York City, NY: New Press, 1998, pp. 229–248.

Kurt, J., 1992, *Coloured Petri Nets: Basic Concepts, Analysis Methods and Practical Use*, Vol. 1, Berlin: Springer-Verlag.

Labrou, Y., and T. Finin, 1997, "A Proposal for a New KQML Specification," *CSEE Technical Report TR CS-97-03*, Computer Science and Electrical Engineering Department, University of Maryland, Baltimore County.

Lebanon, 1958, "Memorandum of a Telephone Conversation between President Eisenhower in Washington and Prime Minister Macmillan in London, July 14, 1958, 5:43 p.m.," reprinted in *Foreign Relations of the United States, 1958–1960*, Vol. 11, Washington, DC: U.S. Government Printing Office, 1992, Document 131.

Magnin, L., and J. Ferber, 1993, "Conception de systèmes multi-agents par composants modulaires et réseaux de Petri," in *Actes des journées du PRC-IA*, Montpellier, pp. 129–40.

Moeschler, J., 1985, *Argumentation et conversation: Eléments pour une analyse pragmatique du discours*, Paris: Hatier-Credif.

Moeschler, J., 1989, *Modélisation du dialogue: Représentation de l'inférence argumentative*, Paris: Hermès.

Prague Spring, 1998, *The Prague Spring '68*, Document 81, "Transcript of Leonid Brezhnev's Telephone Conversation with Alexander Dubček, August 13, 1968." Available at http://www.gwu.edu~nsarchiv/nsa/publications/DOC_readers/psread/doc81.htm (accessed September 1, 2003).

Rescher, N.A., 1977, *Dialectics: A Controversy-oriented Approach to the Theory of Knowledge*, Albany, NY: SUNY Press.

Sacks, H., 1972, "On the Analyzability of Stories by Children," in J.J. Gumperz and D. Hymes, eds., *Directions in Sociolinguistics: The Ethnography of Communication*, New York City, NY: Holt, Rinehart and Winston, pp. 325–45.

Sacks, H., 1987, "On the Preferences for Agreement and Contiguity in Sequences in Conversation," in G. Button and J.R.E. Lee, eds., *Talk and Social Organisation*, Cleveland and Philadelphia (U.K.): Multilingual Matters, pp. 54–69.

Sacks, H., 1992, *Lectures on Conversation*, G. Jefferson, ed., Cambridge, MA: Blackwell.

Sacks, H., E. Schegloff, and G. Jefferson, 1974, "A Simplest Systematics for the Organization of Turn-taking in Conversation," *Language* **50**:696–735.

Schegloff, E.A., 1984, "On Some Questions and Ambiguities in Conversation," in J.M. Atkinson and J. Heritage, eds., *Structures of Social Action: Studies in Conversation Analysis*, Cambridge: Cambridge University Press, pp. 28–52.

Schegloff, E.A., 1990, "On the Organization of Sequences as a Source of 'Coherence' in Talk-in-interaction," in B. Dorval, ed., *Conversational Organization and its Development*, Norwood, NJ: Ablex, pp. 51–77.

Schegloff, E.A., 1996a, "Confirming Allusions: Toward an Empirical Account of Action," *American Journal of Sociology* **102**(1):161–216.

Schegloff, E.A., 1996b, "Issues of Relevance for Discourse Analysis: Contingency in Action, Interaction and Co-Participant Context," in E.H. Hovy and D.R. Scott, eds., *Computational and Conversational Discourse: Burning Issues — An Interdisciplinary Account*, Berlin: Springer-Verlag, pp. 3–35.

Schwab, H., 1994, “The Paris Peace Conference 1919: Woodrow Wilson, Hankey, and Mantoux,” meeting of 3 June 1919, 4 p.m. Available at http://schwab-writings.com/hi/wi/1.html (accessed September 1, 2003).

Searle, J., 1969, *Speech Acts*, Cambridge, MA: Cambridge University Press.

Simmel, G., 1908, *Soziologie: Untersuchungen über die Formen der Vergesellschaftung*, trans. Chap. 1 by K. Wolff, in D. Levine, ed., *Georg Simmel: On Individuality and Social Forms*, Chicago, IL: University of Chicago Press, 1971.

Simmel, G., 1917, *Grundfragen der Soziologie*, Berlin and Leipzig: de Gruyter.

Sylvan, D., 2003, “Beyond Stipulative Semantics: Agent-based Models of ‘Thick’ Social Interaction,” presented at *Conference on Computational Modeling in the Social Sciences*, Seattle, WA: University of Washington, May 8–10.

Weber, M., 1956, *Wirtschaft und Gesellschaft: Grundriss der verstehenden Soziologie*, J. Winckelmann, ed., Tübingen: J.C.B. Mohr.

Winograd, T., and F. Flores, 1986, *Understanding Computers and Cognition: A New Foundation for Design*, Norwood, NJ: Ablex.

Yalta, 1945, “Memorandum of Conversation — Crimean Conference: Meeting of the President with Marshal Stalin, February 8, 1945.” Available at http://www.gwu.edu/~nsarchiv/coldwar/documents/episode-2/07-01.htm (accessed September 1, 2003).

ACTION SELECTION AND INDIVIDUATION IN AGENT-BASED MODELING

J.J. BRYSON,* University of Bath, Bath, UK

ABSTRACT

This paper is a tutorial on action selection for agent-based modeling. Having a clear idea of how to organize an agent's intelligence will make code cleaner and easier to maintain, and models easier to communicate to others. This paper describes four means of organizing agent action selection in increasing order of complexity: environmental determinism; finite state machines; basic reactive plans; and parallel-rooted, ordered, slip-stack hierarchical (POSH) reactive plans. Modelers should use the simplest mechanism possible. This paper describes the contexts in which more complicated mechanisms can be required and suggests coding and commenting schemes for all four systems. This paper also addresses the issue of individuated agent-based modeling (IABM), where individual agents display different behavior. It gives examples of existing IABM systems and describes how these can be moved into more mainstream agent-based modeling simulators via two relatively simple mechanisms: exploiting individual local variables or specifying different priorities within the action selection mechanism. This allows individual agents to vary in their behavior while sharing the vast majority of their code.

Keywords: Agent-based modeling, action-selection mechanisms, individuated agent-based modeling, basic reactive plans, behavior-oriented design

1 INTRODUCTION

This paper discusses action selection in the framework of agent-based modeling (ABM). Action selection is the means by which an autonomous agent solves the ongoing problem of choosing the next action. Action selection is the executive part of agent intelligence.

The remainder of agent intelligence, which is at least as important and difficult to construct, consists of the actions between which the agent chooses and the perceptions that inform the decisions and shape the acts. There is of course a trade-off in the granularity of control: the more an action is capable of, the less complexity the executive must handle, but the less control it has. For example, the action 'walk-to attended-location' represents a much larger granularity than 'extend left-knee.' If an entire agent intelligence is being developed from scratch, the best way to optimize this trade-off is to follow an iterative development process that provides heuristics to allow refactoring when either actions or action selection become too complicated (Beck, 2000; Bryson and Stein, 2001a; Bryson, 2003). For most people conducting ABM, however, the actions and perceptions are already provided, either by the simulation platform or in terms of existing libraries of behaviors from other developed simulations.

* *Corresponding author address:* Joanna J. Bryson, University of Bath, Department of Computer Science, Claverton Campus, Bath BA27AY, U.K., http://www.cs.bath.ac.uk/~jjb; e-mail: J.J.Bryson@bath.ac.uk.

Because agent-based modelers often lack experience with artificial intelligence, they frequently do not use formal action-selection mechanisms. As a result, programming the agents can become unnecessarily complicated and render simulations more difficult to understand, maintain, and extend. This paper describes a number of established idioms for action selection, each of which is useful in at least some domains. As such, it is intended to increase the clarity of their work. Part of the emphasis of this section, however, focuses on relatively complex action-selection structures for supporting *individuated* agent-based modeling (IABM), where each agent may be programmed with different behavior. Section 2 begins with a short discussion of the utility of IABM. Section 3, which makes up the bulk of the paper, describes four different action-selection idioms, which may be useful for either conventional ABM or IABM. The paper concludes with a description of the engineering issues involved in bringing the more complex forms of action selection into familiar ABM environments, such as Swarm, Repast, and NetLogo.

2 INDIVIDUATED AGENT-BASED MODELING

Much of the research exploiting ABM examines the patterns that emerge from the behavior of very simple agents. This research has shown that such simple agents can do a good job of replicating the behavior of complex real-world actors, even humans, at a high level of abstraction. This level of abstraction seems to be most useful for large numbers of agents, where individual variation or details of decision making can be treated as random noise. In recent years, a number of software tools have been developed that support this research, including Swarm, Repast, and NetLogo.

Another sort of agent-based social simulation exists, however, where the individual agents are given relatively complex intelligent controls, typically to simulate the behavior of relatively small numbers of individual actors. One of the most spectacular examples of this work is that of Tu (1999), which replicates aquatic animal behavior from swimming through fluid dynamics to mating and predation. Another example is the work by Hemelrijk (2000, 2002), which has made significant contributions to evolutionary theories of the differences in macaque social behavior. Hemelrijk’s work involves modeling colonies of primates with individual differences in initial rank. She shows that differences in social organization can emerge as a simple consequence of a single variable — the level of violence in the average antagonistic interaction (Hemelrijk, 2000).

Tyrrell (1993) built a complex simulated environment (which he creatively called the *SE*) to test a variety of action-selection mechanisms. The test agent in the *SE* is a small omnivorous animal that needs to survive and breed. This test involves six types of subproblems:

- *Finding sustenance*. This problem includes water and three forms of nutrition, which are satisfied in varying degrees by three different types of food.

- *Escaping predators*. Feline and avian predators have different perceptual capabilities and hunting strategies.

- *Avoiding hazards*. Latent dangers in the environment include wandering herds of ungulates, cliffs, poisonous food and water, temperature extremes, and periodic (nightly) darkness. The environment also provides various forms of shelter, including trees, grass, and a den.

- *Grooming*. Grooming is necessary for maintaining homeostatic temperature control and general health.

- *Sleeping*. The animal is blind at night and needs to sleep to maintain its health.

- *Reproduction*. The animal is male; thus, its reproductive task is reduced to finding, courting, and inseminating mates. Attempting to inseminate unreceptive mates is hazardous.

The success of the agent in the *SE* is counted as the number of times it mates in a lifetime. Mating is highly correlated with life length, but long life does not guarantee reproductive opportunities: these have to be actively sought. Tyrrell tested five well-known, action-selection mechanisms from psychology and artificial intelligence (AI), one of which he chose as favorite and extended. What makes Tyrrell's work noteworthy is that there are very, very few AI or Alife domains in which a single agent must meet such a diverse set of goals. These include intrinsic and extrinsic goals, homeostatic and cyclic goals as well as opportunistic event-based ones. Of course, real animals deal with such conflicting goals and desires all the time.

The work of these three researchers has two things in common:

1. Compared to most ABM, the individual agents have relatively complex individual behaviors.

2. They all have conducted their research in proprietary research environments, which they have developed or their institutions have developed for them.

Many researchers would like to be able to produce models exploring their own theories or hypotheses working at this level of complexity, but they are either unwilling or unable to accept the cost of constructing such simulations. I personally know of several theoretical biologists who want to construct models of social insects (e.g., particular species of wasps and ants) with small numbers of members per colony, at least two groups of primatologists interested in exploring and demonstrating their own hypotheses (which are at odds with those of Hemelrijk), and a cross-disciplinary group working on understanding social predator communication. It would be useful if we could enable such research within existing ABM toolkits. If we cannot, a new toolkit may be needed.

As indicated in Section 1, one way to obtain more complex behavior from existing simulators is to simplify code by using (and commenting) a good model of action selection. Section 3 goes through several such models in increasing order of complexity. We then return to examining the question of whether current toolkits can support IABM.

3 MODELS OF ACTION SELECTION

3.1 Environmental Determinism

There is no established name for the simplest way to conduct action selection, so I have called it ‘environmental determinism.’ Environmental determinism enumerates the possible salient states of the environment and states what action should be performed in each. This model assumes that:

1. There are only a limited number of salient situations in which the agent can find itself.

2. These situations are mutually exclusive.

3. Actions can easily be mapped to situations.

While these assumptions may seem unrealistic, they have been usefully applied in sufficiently abstract models. The clearest example is probably Conway’s Game of Life (Gardner, 1970), a very early ALife system, which takes place on a two-dimensional grid. If I refer to any cell of that grid rather than any live cell as an agent (which is how life is typically programmed), then Conway determined that nine environmental situations may matter, because the only thing that determines action in Conway’s system is how many neighbors one has, and one can have only 0–8 neighbors on a two-dimensional grid. Figure 1 further summarizes these into four situations. Too few neighbors (0–1) and the cell is dead, regardless of its previous state. For two neighbors, the cell holds its current state, alive or dead. Exactly three neighbors and the cell is alive regardless of its previous state, but with four or more neighbors it is again dead.

0–1	**2**	**3**	**4–8**
Die	**Stay**	**Be born**	**Die**

FIGURE 1 Environmental determinism (Example is from Conway’s Game of Life [Gardner, 1970].)

In the event that a reader is unfamiliar with the Game of Life, I strongly recommend typing “game life” into Google to view the incredible variety of emergent (higher-level) growth and action that results from this simple program. There are claims that this system is a fairly realistic model of bacterial life in a petri dish, but I leave it to the reader to determine their veracity.

Coding environmental determinism requires only a set of if-then statements. The environmental conditions should be made clear by using functions and function names to keep them clean, as well as comments. For example:

```
if (cell is dead) AND (number-of-neighbors is 3)
      then {set cell alive}; /* new cell is born */
if (cell is alive) AND (number-of-neighbors is 2 OR 3)
      then {no action}; /* Leave the cell alive */
if (cell is alive AND (number-of-neighbors is NOT 2 OR 3)
      then {set cell dead}; /*lonely or over-crowded cells die*/
if (cell is dead) AND (number-of-neighbors is NOT 3)
      then {no action}; /* leave cell dead */
```

3.2 Finite State Machines

In general, programmers prefer to think about agents rather than environments. Agents tend to be much simpler than their environment; they tend to have fewer possible behaviors than there are possible external situations. So programmers find it simpler to organize coding around actions, not events. The standard way to control many machines, notably AI game agents, is by using the abstraction of a finite state machine (FSM). An FSM enumerates the possible states the agent can be in, actions it might take, and the environmental contingencies that might make the agent change state. Programming an FSM requires two things:

1. Enumerating the states in which the agent can be, and

2. Enumerating the causes for an agent to change a state.

Again, assumptions are made that both states and transitions can be enumerated and that they are mutually exclusive.

For ALife action selection, we can think of each state being a situation in which an agent should perform a particular action. To go back to the Game of Life, we now have to think of the cell/agent as expressing one of two behaviors: looking alive or looking dead (Figure 2).

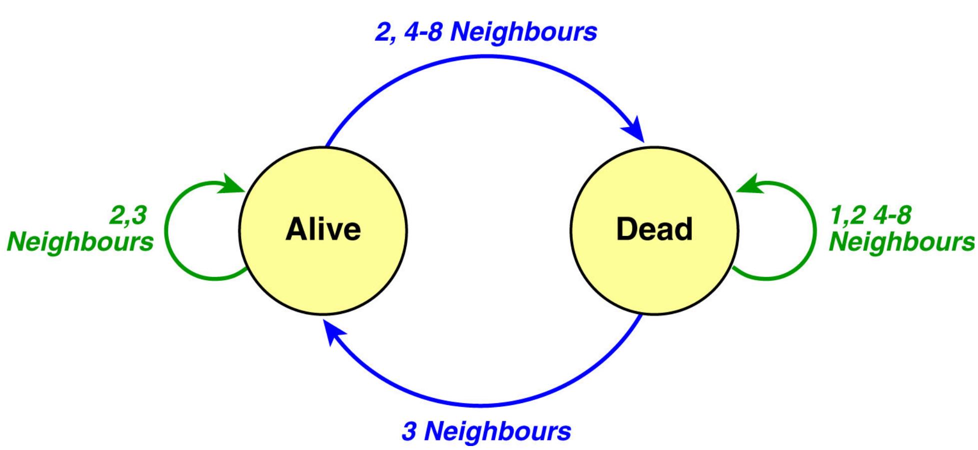

FIGURE 2 Finite state machines (Example is from Conway's Game of Life [Gardner, 1970].)

Documentation for an FSM should ideally include a diagram. The coding should be the possible transitions clustered by state:

```
/* Transitions from state DEAD */
if (self dead) AND (number-of-neighbors is 3)
   then (self be-born);
/* Transitions from state ALIVE */
if (self alive) AND (number-of-neighbors NOT in {2,3})
   then (self die)
```

To be formal, an FSM should also enumerate all possible environmental or internal events and specify the situations in which no change takes place, as described in the following paragraph.

The advantage of this approach is that one can check the code to ensure that the programmer included and coded all possible transitions. The disadvantage of this approach is that it might make the code less readable and that in most situations, such a check may be an intractably long procedure. In general, programmers should try to put in as much code and comments as required. For example, if a condition exists that the programmer had to take a long time to consider, or a condition that another person needed explained to them, such conditions deserve comments and possibly code to make the comments more explicit. For ALife modelers who run large and long simulations, it is almost never a good idea to code the cases where there is no change explicitly because they will take CPU time, though they can be coded for clarity and then commented out.

It is important to realize that in either environmental determinism or FSMs, the programmer needs to create discrete categories of both environmental events and behavioral acts. The only differences occur in the FSM, where the programmer also must enumerate states for the agent, and actions are tied to these states rather than to the environmental categories. Often the states turn out to be a useful abstraction, but that does not have to be the case. For many simulations, it may be that the environment really is the more salient actor, and the agents are simple enough that there is no reason to add theoretical entities, such as internal states for the agents.

3.3 When Enumerating Transitions Is Too Hard

Most agents, of course, have more than two possible actions. If there are limited ways of transitioning from each action to the next, FSMs are a good way to describe that behavior. However, if behavior *is* largely driven by environmental prompting, and the environment is very dynamic and unpredictable (as when it contains many other types of complex actors), there may be transitions from every state to every other state. As a result, for every new action or capability added to an agent, you need to add as many transitions to both the agent and the other states as there are other capabilities. The number of transitions (and therefore the size of the code) grows quadratically, since all N nodes must have $N - 1$ transitions in them. It would be better to have a way to code action selection that did not grow much faster than the number of possible actions.

Consider a situation where the agents approach a more humanoid intelligence than Conway's cells, for example with some arbitrary character chosen from a Jane Austen novel. The typical Austen character might be thought of as having four states: flirting, engaged, in church, and married. A first cut at an FSM for this agent might look something like Figure 3.

The problem comes when we start trying to label the transitions. When you think about it, a good Georgian English agent would not only go to church to get married but also might attend regularly at any stage of his or her life. On reading Austen further, one realizes that marriage is not in fact a terminal state; some characters continue to flirt and may even become re-engaged. In fact, the only terminal state is death, which can occur at any time (Figure 4).

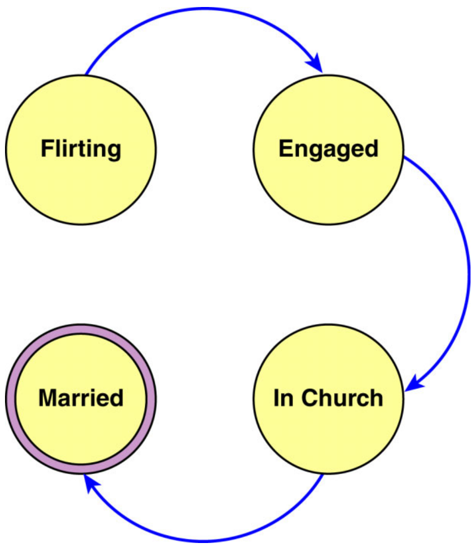

FIGURE 3 A first cut at an FSM for a Jane Austen character (The double circle on *married* indicates that it is a terminal state — the end of autonomous action selection, at least for this controller.)

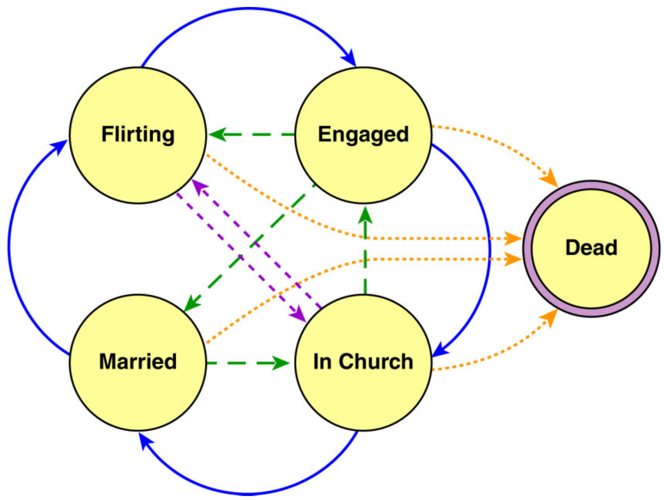

FIGURE 4 A somewhat more realistic draft FSM for an Austen character (Notice that labeling the transitions would still be fairly complex.)

The problem with the FSM is that it must represent all transitions for an agent. But if we are trying to create a working agent, *we really need to specify only the transitions that the agent makes itself*. In fact, in the AI for abstract simulations, we generally need to model only the transitions that an agent might rationally choose to take in pursuit of its own goals.

3.4 Basic Reactive Plans

The basic reactive plan (BRP) is the name given here to an idiom or pattern found in several (though not most) influential reactive planning architectures (Fikes, et al., 1972; Nilsson, 1994; Bryson and McGonigle, 1998; Bryson and Stein, 2001b). Building a BRP requires a few assumptions. For example, it requires one to assume that the agent has a goal and is capable of accomplishing a set of actions that can lead to achieving that goal. The BRP is a prioritized list of these actions. The most important action is the one that consummates the goal; the second most important action is one that enables the most important act, and so on.

If the agent can already do the most important action, it does not need to execute the other ones. A reactive plan recognizes this situation, which means it can behave *opportunistically*. Such behavior indicates that each action is paired not only with a priority, but also with a perceptual condition that allows it to know when it can execute that action. This is important not only for opportunism (skipping unnecessary steps in a conventional plan), but also for robustness (possibly repeating steps if they fail or simply need to continue to be repeated, such as digging a hole until you hit water).

Rather than explicitly coding how to recognize every possible transition between action states, a BRP programmer need only code how to recognize the situation in which each action can fire — a task made simpler by the invariant that no better (higher-priority) action can fire or the current one would not be considered. This means that for each action, the programmer need only hand-code one situations: the minimal requirements where that action might be usable. There is no reason to describe when the action should be skipped because that is encoded by its priority within the BRP.

For the Jane Austen agent, assume we are programming one of the 'pure' characters with no devious intentions. The agent's highest goal is to become married. To do so, the agent must be in a church, but there is no point in going to the church without one's fiancé, and one cannot have a fiancé without having gotten engaged. If one is not engaged, one must flirt. Flirting of course also has preconditions, but for simplicity, we stop here.

A possible BRP is shown below:

1. (fiancé present *and* in church) $\Rightarrow$ marry
2. (fiancé present) $\Rightarrow$ go to church
3. (engaged) $\Rightarrow$ go to fiancé
4. (receiving attention) $\Rightarrow$ become engaged
5. () $\Rightarrow$ flirt

The numbers assign priority, with 1 being the highest. Thus, if one is in a church but does not have a fiancé present, one does not get married; one either goes to be with the fiancé (if one exists) or flirts, unless, of course, one is receiving attention already, in which case flirting can be skipped. One can become engaged and on the next iteration become married, assuming nothing has happened to remove the new fiancé from the church between program cycles. On the other hand, a character that never receives attention can flirt indefinitely, or, if suddenly receiving notice that one has become engaged (perhaps by arrangement by parents), skip directly to going to the fiancé.

Obviously, encoding this BRP in an FSM would take at least 25 lines of code, assuming that the goal takes one line to describe and each possible transition between states could be specified in one line. Similarly, a large amount of possible environmental states have been neatly ignored as not relevant to this particular agent's pursuit of its goal.

For all its elegance, a BRP is simple to code in most programming languages. It is best coded as a switch (in Java or C) or cond (in Lisp), though it can also be coded as cascading if-then-else statements in pseudo-languages that lack this idiom. Details of building successful BRPs can be found in Bryson (2003), but essentially they are a simple derivation from a conventional sequential plan. For example:

flirt ⇒ become engaged ⇒ go to church ⇒ get married.

The BRP simply inverts the ordering and then specifies a mechanism for recognizing when each item could be activated.

3.5 Parallel-rooted, Ordered, Slip-stack Hierarchical Reactive Plans

The life-like agents described above as individuated generally have more than one goal. Further, many actions may themselves require some multiple subactions to complete, and these in turn may require a BRP to organize.

An examination of the history of the action-selection literature (for a review, see Bryson [2000a]) indicates three problems that any successful approach must address:

1. Some things must be checked at all times. For example, a loud noise will nearly always stop you from what you are doing and direct your visual attention toward the source, without any conscious processing of this attention switch.

2. Some things hardly ever need to be checked. For example, when you walk, you have a reliable pattern for controlling your legs that is characteristic of your individual gait. It would be very unusual for you to become aware of, let alone attempt to control or alter, the muscle patterns involved.

3. Some things must be checked only in particular contexts, but then in an unpredictable order. For example, if you are doing a jigsaw puzzle, you may have a set of rules about where you are piling edge pieces or pieces of a particular shade of blue. Unlike walking, you cannot predict the appearance of the next piece that you notice, so your plan cannot be ordered in advance. This situation requires a BRP.

It might seem simpler to represent everything as a sense/action pair — one enormous BRP.[1] However, it is untenable (that is, computationally intractable) to have every possible skill that requires this level of attention be equally accessible all the time. The reflexive response to the previously mentioned loud noise is one of a relatively limited number of such stimuli (some learned, some innate) that seem to be stored in a separate (and fairly small) part of the brain, the amygdalic system. This is not only a consequence of combinatorics and computational limits, but also of perception and context. In another context, that shade of blue might trigger one to follow a friend wearing a particular shirt through a large crowd or to pass a ball to a teammate rather than selecting or arranging puzzle pieces.

---

1 Or a large set of production rules, see Newell (1990).

I have developed an action-selection mechanism that supports all three of these situations with three types of representations, as well as a development methodology to determine when to apply each type (Bryson, 2001; Bryson and Stein, 2001a). The development methodology, behavior-oriented design (BOD), is an iterative process for determining not only which type of action-selection representation should be used, but also the granularity of the primitive actions. Primitives are encoded in object-like *behavior modules*, thus the name. Bryson (2003) gives a good summary of the heuristics and practicalities of using BOD.

The representational frameworks for BOD action selection are called parallel-rooted, ordered, slip-stack hierarchical (POSH) reactive plans. POSH plans contain five types of elements. Types 1 and 2 are the primitive *actions* and *senses*. There are only two differences between these types:

1. *Return values*. *A*ctions do not return a meaningful value, except in the case of radical failure, when a flag may break the system out of its current action-selection context. Senses return meaningful values that can be used in predicates for comparisons.

2. *Duration*. *S*ome actions may take some time (though usually no more than 100 ms). Senses are expected to return values very quickly, because many sensory preconditions may be checked on each cycle of the action-selection architecture, which ideally runs at least 100 Hz for real-time systems (e.g., robots) and orders of magnitude faster for simulations. Some actions (such as shifts of visual attention) also take place during sensory preconditions, but currently POSH plans have no separate type for these sense-speed acts.

Types 3–5 are *action patterns*, which are simple sequences to handle the second level of planning problem (things that almost always follow); *competences*, which are essentially BRPs and handle the third above-mentioned case (things checked in certain contexts); and *drive collections*, the things that must always be checked.

Type 5, the drive collection, is a special extension of the BRP. It serves as the root of the POSH plan hierarchy. It has several important characteristics.

- There is only one drive collection, and it is checked on every iteration of the action-selection mechanism.

- Each element of the drive collection represents a separate goal for the agent. These goals can be met in parallel, so each element of the drive collection keeps track of its immediate child, as well as what it was doing most recently — its current action selection context.

- To facilitate parallelism, the drive elements may have associated frequencies. Thus, some action (e.g., breathing, looking around) may be very high priority every few seconds, but after initiating that action, the action-selection mechanism is free to consider other, lower-priority goals over the next specified time interval.

The POSH action-selection mechanism is a sequential process, and it grants only course-grained parallelism and scheduling, because it depends on the return time of the primitives. Since the primitives are supported by independent modules (which may themselves be threaded), however, POSH agents can exhibit smooth and continuous parallel behavior. For example (from Bryson and McGonigle [1998]), consider a robot moving through a cluttered space. The robot in a dynamic environment might need to resample its sensors seven times a second, but there is no reason for it to stop moving while it does so. The primitive that creates movement can send the wheel drivers the current direction and speed for motion, with the understanding that the drivers will continue at that velocity until the next message is received.[2]

4 SUPPORTING IABM

Reflecting on the four types of action selection offered, we return to the question of IABM. Can platforms such as Swarm, NetLogo and Repast support individuated action selection? Absolutely. The main requirement is only that each individual agent be able to have its own variable state. From here, there are two possible solutions:

1. Each agent can have a copy of a common intelligence. So long as some decisions or other behaviors are made dependent on the content of the individuals' variables, their behavior will be individuated by the different values of these variables. This is effectively what Hemelrijk (2000) has done; her agents' behaviors vary only in their relative dominance ranking (their probability of success in a social context is determined by this) and by gender (males may be influenced to approach females more frequently.)

2. Each agent may have a script that describes its action-selection system as a piece of variable state. This is actually just a special case of the first solution — the common intelligence is an action-selection mechanism capable of interpreting that script.

I can demonstrate that both of these methods are plausible in existing ABM platforms. Of the three modeling platforms listed, Swarm and Repast have access to 'real' programming languages for describing agent intelligence (Objective-C and Java), while NetLogo only allows modelers access to a toy/teaching language (Logo). Recently, two Bath graduate students replicated some of the most basic Hemelrijk (2000) results in NetLogo by using Method 1 above (Muhd Fathil, 2003; Wang, 2003). Thus, even in the simplest of the three platforms, Method 1 is possible.

Method 2 has not yet been demonstrated in a major ABM platform, but Bryson (2000b) implements POSH action selection on a minor ABM platform, the *SE* developed by Tyrrell (1993). This works because of the sequential nature (described above) of the actual action-selection mechanism that exploits POSH plans. This was easy because I had a version of POSH in the native language of the *SE* available at the time. All that was required was to link the POSH code to the *SE* code and then replace the central iterator for POSH with an individual call

---

2 Good robots also have time-outs associated with their drivers, so if the action-selection mechanism is hung or crashes, the robot stops within a short time after failing to receive instructions (e.g., one second).

to a single POSH program cycle from inside the modeled agent. Code for this implementation is available from the Web page for that project, currently:

`http://www.cs.bath.ac.uk/~jjb/web/edmund.html#code.`

The main obstacle for using POSH on a major ABM platform is the absence of a version in the appropriate language. It would be impractical for NetLogo to address this problem because the overhead for the extra action-selection mechanism would be too high; the system is already significantly slower than the other two platforms mentioned. On the other hand, there is no 'in principle' reason that POSH could not be translated into Java; it currently exists in C++ (Bryson and McGonigle, 1998), Lisp (CLOS) (Bryson, 2001), and Python (Kwong, 2003). IABM by no means requires the complexity of POSH action selection.

Method 1 does not require it nor did the Hemelrijk replications. It might also be possible to make simpler systems for Method 2, though I have tried to keep POSH plan structures as parsimonious as possible. As I emphasized in the previous section, *any* sort of ABM can be improved and extended simply by making clear, intentional decisions about how the action selection will be handled, and by clearly coding and commenting the action-selection part of the agent's intelligence. In modeling as well as the rest of computer science, it's important to do the simplest, clearest thing possible for the problem at hand.

5 SUMMARY

This paper has presented two related topics: how to improve action selection for ABM and how to individuate agent-based modeling. The important steps for improving action selection are:

- Separate the problems of describing *how* the agent acts (coding its possible behaviors) and *when* it takes an action (coding its action selection).

- Use a standardized mechanism for describing an agent's intelligence. Here, 'describing' includes both coding and commenting. Four different action-selection frameworks were presented:

    1. Environmental determinism requires enumerating contexts in which the agent may find itself and then saying what it will do in each.

    2. Finite state machines (FSMs) require enumerating the things the agent might do and then describing what might make the agent switch between possible actions.

    3. Basic reactive plans require prioritizing the actions that the agent might take to achieve some goal and then describing the minimal requirements for being able to take those actions.

    4. Parallel-rooted, ordered, slip-stack hierarchical (POSH) reactive plans allow encoding of full animal-like intelligence.

The developer should choose the framework that most simply describes the minimum behaviors the agent needs.

The easiest way to individuate agent behaviors is to have all of the agents share the bulk of their behavior code (the 'how' part), then to individuate their action selection, either by having their (shared) action-selection program reference an individual agent state in its decision making or by providing different action-selection programs for different agents.

There is a danger in simplifying action-selection coding; that is, simplification can lead researchers to make more complicated models than necessary. Already, the most difficult task for modelers is to analyze and explain the group behavior that emerges from the interactions of agents — the more difficult their behavior, the harder this explanation. Having clearer code is worth this risk. The history of science is full of examples where the first successful model was not the simplest, but given that it was the first, it must have also been the most immediately intuitive. Once a good model has been built, simplifying it is part of the scientific process of analyzing how it works. Good coding techniques can make it easier to build both the first and the simplified models.

6 ACKNOWLEDGMENTS

I acknowledge that I grossly oversimplified many of Jane Austen's characters. I would like to express my gratitude to David Sallach and Tom Howe for encouraging me to write on this topic.

7 REFERENCES

Beck, K., 2000, *Extreme Programming Explained: Embrace Change*, Reading, MA: Addison-Wesley.

Bryson, J.J., 2000a, "Cross-paradigm Analysis of Autonomous Agent Architecture," *Journal of Experimental and Theoretical Artificial Intelligence* **12**(2):165–190.

Bryson, J.J., 2000b, "Hierarchy and Sequence vs. Full Parallelism in Reactive Action Selection Architectures," in *From Animals to Animats 6 (SAB00)*, Cambridge, MA: MIT Press, pp. 147–156.

Bryson, J.J., 2001, *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*, Ph.D. thesis, AI Technical Report 2001-003, Massachusetts Institute of Technology, Department of EECS.

Bryson, J.J., 2003, "The Behavior-oriented Design of Modular Agent Intelligence," in R. Kowalszyk, J.P. Müller, H. Tianfield, and R. Unland, eds., *Agent Technologies, Infrastructures, Tools, and Applications for e-Services*, Berlin: Springer, pp. 61–76.

Bryson, J.J., and B. McGonigle, 1998, "Agent Architecture as Object Oriented Design," in M.P. Singh, A.S. Rao, and M.J. Wooldridge, eds., *The Fourth International Workshop on Agent Theories, Architectures, and Languages (ATAL97)*, Berlin: Springer-Verlag, pp. 15–30.

Bryson, J.J., and L.A. Stein, 2001a, "Modularity and Design in Reactive Intelligence," in *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, Seattle, WA: Morgan Kaufmann, pp. 1115–1120.

Bryson, J.J., and L.A. Stein, 2001b, "Architectures and Idioms: Making Progress in Agent Design," in C. Castelfranchi, and Y. Lespérance, eds., *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*, Berlin: Springer.

Fikes, R.E., P.E. Hart, and N.J. Nilsson, 1972, "Learning and Executing Generalized Robot Plans," *Artificial Intelligence* **3**:251–288.

Gardner, M., 1970, "Mathematical Games: The Fantastic Combinations of John Conway's New Solitaire Game 'Life,'" *Scientific American* **223**(4):120–123.

Hemelrijk, C.K., 2000, "Towards the Integration of Social Dominance and Spatial Structure," *Animal Behaviour* **59**(5):1035–1048.

Hemelrijk, C.K., 2002, "Self-organization and Natural Selection in the Evolution of Complex Despotic Societies," *Biological Bulletin* **202**(3):283–288.

Kwong, A., 2003, *A Framework for Reactive Intelligence through Agile Component-based Behaviors*, M.S. thesis, University of Bath, Department of Computer Science.

Muhd Fathil, N.H., 2003, *The Study of the Evolution of Macaque Complex Despotic Societies*, M.S. thesis, University of Bath, Department of Computer Science.

Newell, A., 1990, *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.

Nilsson, N.J., 1994, "Teleo-reactive Programs for Agent Control," *Journal of Artificial Intelligence Research* **1**:139–158.

Tu, X., 1999, *Artificial Animals for Computer Animation: Biomechanics, Locomotion, Perception and Behavior*, Berlin: Springer.

Tyrrell, T., 1993, *Computational Mechanisms for Action Selection*, Ph.D. thesis, University of Edinburgh, Centre for Cognitive Science.

Wang, J.J., 2003, *Sexual Attraction and Inter-sexual Dominance among Virtual Agents — Replication of Hemelrijk's Domworld Model with Netlogo*, M.S. thesis, University of Bath, Department of Computer Science.

A CONCEPTUAL FRAMEWORK TO REPRESENT EMERGENT SOCIAL PHENOMENA

JORGE LOUÇÃ,* LabMAg – Laboratório de Modelação de Agentes
Faculdade de Ciências da Universidade de Lisboa, Portugal

ABSTRACT

This paper proposes a conceptual framework to represent emergent social phenomena, characterizing the reciprocity between micro and macro dimensions in an organization. The primary rationale for this research is that explicit modeling of emergent macro features might improve artificial distributed support systems. The framework represents organizational actors' beliefs and interactions and deduces emergent collective cognitive values in a multi-agent system. Cognitive mapping is used to represent actors' beliefs and collective values, and Latané's social impact theory inspires the deduction of macro-social outcomes. The paper discusses two ways of modeling collective values: a collective cognitive map, deduced from artificial agents' cognitive maps and validated *a posteriori* by direct negotiation between organizational actors, and dynamic representation of emerging social outcomes through *socio-cognitive models*. The paper suggests that these two points of view are complementary.

Keywords: Multi-agent systems, cognitive maps, emergence

1 INTRODUCTION

Human societies can be considered complex social systems because their collective properties, such as structures of collective values, emerge from interactions among individuals (social actors). Multi-agent systems have been used to represent complex social systems (Luck, et al., 2003). Nevertheless, nature and dynamics of emergent social phenomena are still very hard to identify (Axtell, 2000; Sawyer, 2003). Different approaches in social sciences — from methodological individualism to holism — have tried to explain the behavior of individuals in society. Emergentism is an alternative to both individualism and holism. The emergentist sociocultural approach in psychology proposes to look further than individual-level explanations, suggesting that we take into account a situated level, the context where social actors reason. The socioculturalist view focuses on the diversity of individual situated participations in collective activity (Sawyer, 2002). On the basis of this approach, one way of analyzing complex social phenomena is to represent high-level cognitive models that emerge from situated interactions among social actors. This paper suggests a representation of both individual beliefs and emergent collective values and adopts a concept of society composed of successive flows of microsituations (Sallach, 2003). The multi-agent paradigm is consistent with this point of view, where macro representations result from the aggregation of micro behaviors. Actually, the multi-agent model, which is composed of a set of autonomous artificial agents that operate concurrently, is particularly well suited to support interactions in a distributed organization (Louçã, 2000). Individual microsituations can be represented by using artificial agents; in a way,

* *Corresponding author address:* Prof. Jorge Louçã, ISCTE-DCTI c.296, Av.das Forças Armadas, 1649-026 Lisboa, Portugal; e-mail: Jorge.L@iscte.pt.

they help to characterize the emergence of high-level cognitive models, which can concern informal relationship nets, power structures (legal, moral, and others), and, in a general way, patterns of social life that influence interactions among social actors.

The goal of this paper is to propose a conceptual framework to represent emergent social phenomena, characterizing the reciprocity between micro and macro dimensions in an organization. The primary rationale for this research is that explicit modeling of emergent macro features of social systems might improve artificial distributed decision support in organizations. This idea requires dynamic creation of cognitive models during the run of the system, so that the macro phenomena represented would emerge from micro interactions of the organizational actors. Emergent macro patterns would then be represented in the system so that actors could act accordingly.

We consider an organization where cognitive agents support actors. Specific software tools and knowledge-based systems compose artificial agents. In this environment, interactions among actors occur through artificial agents in a multi-agent system. Figure 1 depicts this general idea, where some cognitive representation of collective values is deduced from interactions among actors.

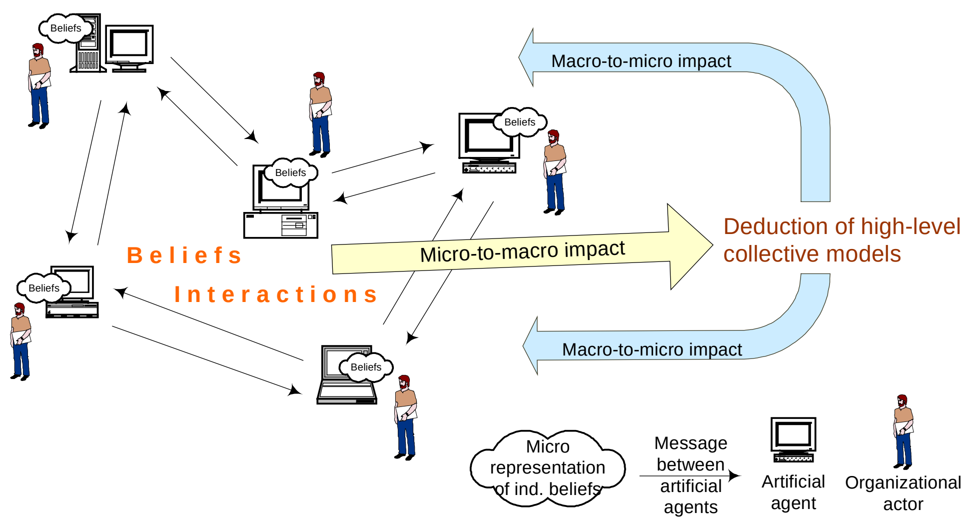

FIGURE 1 Reciprocity between micro and macro dimensions in an artificial society

The influence of actors' interactions on collective values, also called the micro-to-macro impact, is complemented by the macro-to-micro impact, where collective values constrain actors' beliefs. To represent this reciprocal influence, the framework must be able to represent actors' beliefs and interactions and deduce collective values and social regularities and structures. Cognitive mapping is used to represent actors' beliefs and cognitive interactions. On the other hand, social impact theory inspires the deduction of macro-social outcomes (Latané, 1981; Nowak and Latané, 1994).

This document is organized as follows. Section 2 presents the theoretical foundations on which this research is based, including some aspects of Latané's social impact theory, to main characteristics of multi-agent systems and to the domain of cognitive mapping. Section 3 presents STRAGENT, a multi-agent model based on multi-dimensional reasoning processes that have been proposed in previous research. Section 4 concerns the main proposition of the paper (e.g., a conceptual framework to represent emergent social phenomena, based on the reciprocity between micro and macro dimensions in an organization). Section 5 discusses the ideas proposed and references some research perspectives in the domain.

2 THEORETICAL FOUNDATIONS

The conceptual framework proposed here is based on the multi-disciplinary theoretical foundations of Latané's social impact theory to model the micro-macro link in an artificial society, on multi-agent systems to model distributed and interactive support systems, and on cognitive mapping to represent artificial agents' beliefs.

2.1 Social Impact Theory

The social impact theory proposes a general explanation as to how individuals react to social situations (Latané, 1981). According to this theory, three basic principles explain the *impact* of a group of people on an individual's behavior and beliefs (macro-to-micro influence), and vice versa (micro-to-macro influence). The first principle — *social forces* — states that impact I is caused by a combination of factors:

- Persuasiveness or strength S, that is, how important the influencing group is to the individual, acknowledging that status, age, and authority can influence strength;

- Social distance or immediacy I, that is, how close the group is to the individual in both time and space; and

- Number N of sources, that is, how many people are in the group.

The principle does not define the precise nature of the relationships among these three variables, but it presents *impact* as a function of their combination ($I = f$[sin]). Nevertheless, field research suggests that influence increases when people are important, socially close by, and numerous (Latané and Fink, 1996). The second principle — *psychosocial law* — states that impact is a power function of the scaled number of people psychologically present N, where the power function is less than 1, or $I = SNt, t < 1$. Finally, the third principle — *division of social impact* — says that the impact of an individual on others is a negative power function of strength, number, and immediacy. This principle combines the first two principles to explain how social impact acts when many individuals are close by. As the number of individuals in the group increases, the impact decreases.

Latané and his colleges used these three principles to study individual reaction to social phenomena. Some of their conclusions were that *conformity* and *imitation* are stronger and larger in four-person groups; *diffusion of responsibility* increases with the number of persons in the group; and *social loafing* (the tendency for individuals to expend less effort in groups) increases with the number, but not in direct proportion to the individuals present. An important contribution of this research is the idea of opinion clusters, which introduce the notion of social influence within a local neighborhood and among groups of actors (Nowak and Latané, 1994).

The reciprocity between micro-to-macro and macro-to-micro impacts, evidenced by the social impact theory, can be studied in light of the emergence and dynamics of collective values. This theory also considers the social context where each actor interacts (represented by the strength, immediacy, and number of the other actors), which is according to the socioculturalist approach previously introduced.

Latané's theory has been tested extensively and received empirical support. Simulations of a number of social situations have supported the development of the analytical model (Rockloff and Latané, 1996). Multi-agent systems are particularly suited for computing this kind of social simulation, that is, where cognitive artificial agents can represent social actors.

2.2 Multi-agent Systems

Multi-agent systems, composed by software agents, are distributed, with no central control. Artificial agents are autonomous and interact with each other in a proactive and asynchronous way. They are well adapted to represent complex societies, where the global behavior of the system is the result of the aggregation of agents' autonomous actions. Multi-agent systems can model both stable and dynamic interactions. For this reason, in multi-agent systems, the macro level of collective values can be explained in terms of the micro-level artificial agents. Those characteristics allow the study of what Coleman referred to as the foundations of sociology: the micro-macro relations underlying social dynamics (Sawyer, 2003).[1]

By drawing on cognitive science, artificial agents are able to support heterogeneous organizational actors through a wide range of knowledge representation and reasoning techniques, which can be based on logic, rules, frames, semantic nets, or others (Davis, et al., 1993; Luck, et al., 2003). Cognitive mapping is a knowledge representation technique recently proposed to associate agency to individual subjectivity and interpretation (Louçã, 2000, 2003a).

2.3 Cognitive Mapping

Multi-agent systems are completely distributed: the reasoning process goes on internally to each artificial agent. This feature allows the representation of heterogeneous agents, using complementary technologies and representing different cognitive models. The knowledge representation and reasoning technologies used for this purpose are chosen according to their specific features, each attending to some things and ignoring others. When choosing a given technology, we are in fact selecting a point of view about knowledge representation and reasoning. Each technology is an approach to the task of determining how well it approximates the reality we intend to represent (Sowa, 2001). For instance, logic concerns a point of view of

---

[1] A good roadmap for the next generation of agent-based computing can be found in a report by AgentLink II, a network of researchers concerned with agent-based computing (Luck, et al., 2003).

individual entities and relations among them; rule-based systems consider rules of inference; frames represent prototypical objects; and semantic nets are graphic representations of different kinds of entities through a network topology (Davis, et al., 1993). Each of these approaches has both benefits and drawbacks. In fact, the choice of a given technology is motivated by the characteristics of a given domain, as well as by some insight that indicates how people reason intelligently. On the another hand, formal technologies are problematic in practice. Recent research in multi-agent systems has searched for new technologies. These technologies should be simple and operational enough to be used in organizations, and quite powerful and adapted to hill-structured domains. According to this idea, cognitive maps have been proposed to model beliefs of cognitive agents in a multi-agent environment, as reported by Chaib-draa (2002) and Louçã (2003a).

A cognitive map is a graphic representation of the behavior of an individual or a group of individuals, concerning a particular domain. Cognitive maps can be used at a micro level to represent individual cognitive models and at an institutional level through the use of collective cognitive maps. Psychologists use cognitive maps primarily as data structures to represent knowledge. Generally, this kind of cognitive model facilitates communication inside a group, supporting discussion and negotiation between elements that have different points of view. In this way, cognitive maps can be used to detect conflicts. Several software systems are proposed to represent organizational discourse into cognitive maps, describing mental models in artificial agents (Chaib-draa, 2002; Louçã, 2000, 2003a) and allowing the use of network analysis techniques (Lewis, et al., 2001).

A cognitive map is composed of *concepts* (representing things, attitudes, actions, or ideas) and *links* between concepts. Those links can represent different connections between concepts, such as causality or influence links. Figure 2 exemplifies a cognitive map, where links can represent very positive influence (++), positive influence (+), negative influence (-), and very negative influence (--). This particular type of cognitive map is used to represent strategic thought in organizations, as reported by Louçã (2000).

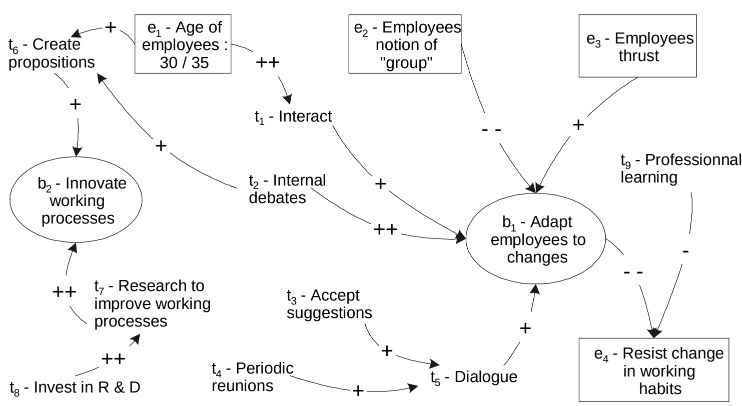

FIGURE 2 Example of a cognitive map (Louçã, 2000)

The main interest in cognitive maps is their reflexive character, allowing people to become conscious of implicit knowledge, through the visualization of direct and indirect links between concepts. Each person constructs a private *version of reality* and deals only with those constructions, which may or may not correspond to the real world (Louçã, 2000). On the other hand, organizations can be seen, at some level of abstraction, as systems of construction and interpretation of reality (Weick, 1995; Lissack and Gunz, 1999). Actually, if cognitive maps can be used to model complex systems of individual beliefs, they can also be employed to represent collective values. Collective cognitive maps can represent collective cognitive models and social values. The main advantage of such maps is to take into account shared concepts, representing a cognitive interpretation about organizational culture. Those characteristics make cognitive mapping appropriate for use in distributed software systems conceived to support interactions in an organization and to represent individual and collective values.

2.4 Previous Research

Previous research by Louçã (2000) proposed a multi-agent model called STRAGENT, which is based on multi-dimensional reasoning processes. In this multi-agent environment, individual beliefs are used to compose a collective solution to a goal through a distributed and incremental process based on agents' interactions. Cognitive maps represent the beliefs of organizational actors. They are composed, on the one hand, by concepts and causal links between those concepts in a *strictu senso* way (Weik, 1995), and on the other hand by the cognitive context of concepts (Louçã, 2003a).

Each concept is coupled to its context, which is called a *scheme* (Bougon and Komocar, 1994). More precisely, a scheme is represented by a concept and its context — a scheme represents the meaning of a given concept — and communicating schemes influence agents' beliefs. Interaction processes aim to converge to common schemes, representing the emergence of collective beliefs. This general mechanism is represented in Figure 3.

STRAGENT was tested through a prototype representing a distributed software system to support decision making in human organizations. This prototype was applied in an industrial enterprise in the domain of telecommunications and electronics, to support the collective decision-making process. Cognitive maps were designed from documents and interviews.

3 REPRESENTING THE EMERGENCE OF COLLECTIVE VALUES

The main proposition of this paper concerns the representation of collective values in an organization. This proposition is composed of two ways of modeling collective values, according to two different perspectives about the role of an artificial system in monitoring the emergence of social processes. The first is related to the composition of a collective cognitive map, deduced from artificial agents' cognitive maps and validated *a posteriori* by direct negotiation between actors. The second perspective concerns dynamic representation of emerging social outcomes. We propose that these two points of view are complementary.

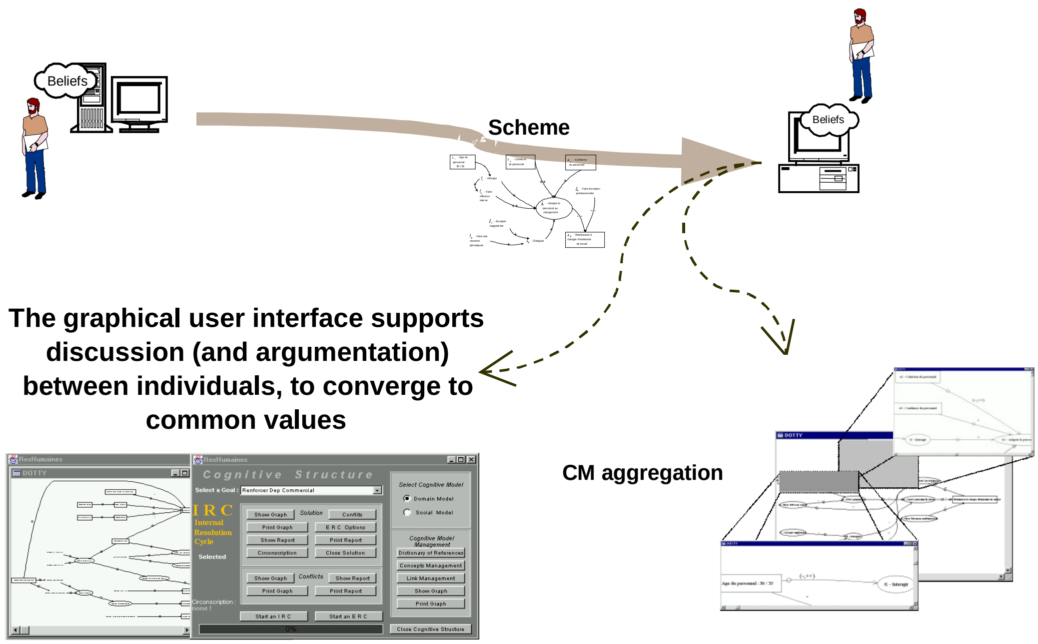

FIGURE 3 Convergence to common cognitive models (Louçã, 2003a)

3.1 Collective Cognitive Map Composition

The method of composing collective cognitive maps, through a distributed artificial system, is a development of work described above by Louçã (2000, 2003a). This method is integrated in a process of internal debate. The aim is to identify organizational goals and plans of action. First, the multi-agent system compares artificial agents' cognitive maps and automatically detects all common concepts, composing the collective cognitive map. This first version of the collective cognitive map takes into account all different points of view, including conflicts. Afterward as reported by Louçã (2000), actors discuss the collective cognitive map and negotiate until they achieve a consensual and coherent collective map with no conflicts. The overall method, including concept matching and negotiation among actors, is represented in Figure 4.

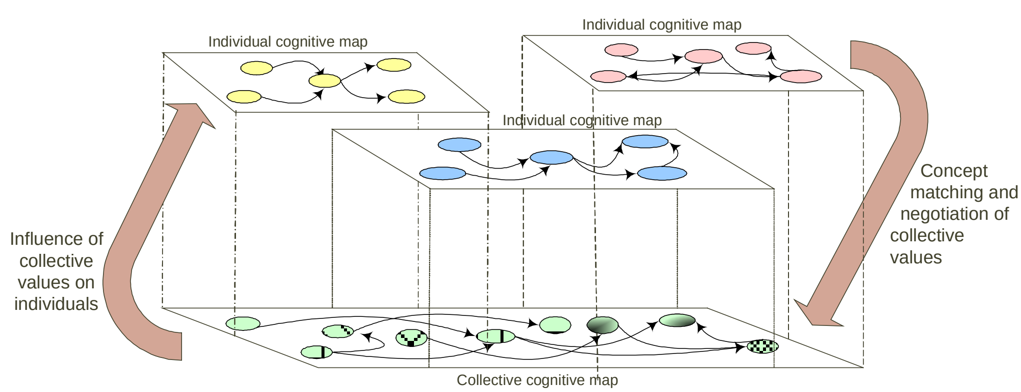

FIGURE 4 Collective cognitive map composition

Conflicts concern different points of view about the nature of the link between two concepts, for example, when one agent considers that a given link represents a positive influence (+) and another agent thinks that the link is negative (–). Conflict resolution through negotiation has two consequences: (1) conflicts are solved (the collective map becomes coherent) and (2) individual agent's cognitive maps are updated according to negotiation processes. The belief revision conducted by agents maps out all different influences to which the actor was submitted. This information is kept in each agent's cognitive map, according to the following formula:

$$y_{x,[a,b]} = \bigcup^{i=1} y_{i,[a,b]} \ ,$$

where

x = agent,

$\{1,\ldots,n\}$ = set of agents that interact with x,

a, b = concepts in x cognitive map,

y = link between those concepts, and

$\cup$ = n-tuple composed by agents' $\{1,\ldots,n\}$ opinions regarding y.

In this way, the agent's cognitive map keeps track of all conflicting opinions that influenced him.

Collective cognitive maps concern both micro-to-macro — the initial collective map composition — and macro-to-micro processes — the individual belief revision. Nevertheless, one limitation of this approach is the representation of social outcomes only at a discrete moment in time, without dynamically monitoring the system throughout interactions.

3.2 Dynamic Representation of Emerging Social Outcomes

Messages exchanged during multi-agent interaction can be used to compare concepts and match common concepts, as well as detect potential conflicts (Figure 3). This paper proposes conception of thematic subcognitive maps, called *socio-cognitive models*, concerning specific social domains, such as power relationships in the organization (legal, moral, and others), and in a general way concerning patterns of social life and collective values that influence interactions among social actors. These submaps are conceived throughout interactions composed by common concepts. On the other hand, sociocognitive models then influence agents' beliefs and behavior. This general idea is depicted in Figure 5.

Sociocognitive models are conceived dynamically, allowing the identification of different collective cognitive structures and representing diverse aspects of organizational culture. An interesting factor of sociocognitive models is their capability of explicitly representing links between common concepts and individual cognitive maps, allowing a situated cognitive perspective (i.e., taking into account the social context of each sociocognitive model). This idea

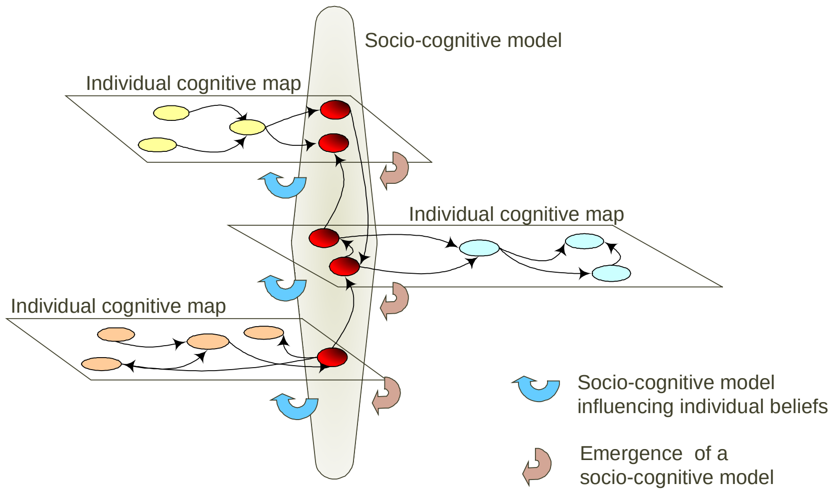

FIGURE 5 Emergence of a sociocognitive model

is according to the socioculturalist approach discussed above. We suggest considering some notions from social impact theory (Latané, 1981), identifying factors that characterize social context. In this way, the main factors of the principle of *social forces* are considered:

- Strength S, that is, how important the agent is to the group that contribute to the sociocognitive model, given that status and authority can influence strength;

- Immediacy I, that is, how close the agent is to the group; and

- Number N, that is, how many agents participate in the sociocognitive model.

The aim is to represent both the impact of individual beliefs on sociocognitive models (micro to macro), and the impact of sociocognitive models on agents' cognitive maps (macro to micro); see Figure 5. The micro-to-macro impact can be analytically represented according to the following formulas:

$$\mathrm{Imp}_{x,\S} = f(S_{x,\S}, I_{x,\S}, N_{\S}),$$

and

$$y_{[a,b],\S} = \bigcup_{n}^{i=1} (y_{[a,b],i}, \mathrm{Imp}_{x,\S}),$$

where $\mathrm{Imp}_{x,\S}$ represents the impact of the agent x on the sociocognitive model §, which is a function of strength $S_{x,\S}$, immediacy $I_{x,\S}$, and number $N_{\S}$. Therefore, the link $y_{[a,b],\S}$ between concepts a and b in § is an n-tuple, including both 1, …, n agents' opinions concerning $y_{[a,b]}$ and each agent impact $\mathrm{Imp}_{x,\S}$.

Similarly, the macro-to-micro impact can be analytically represented by the following formulas:

$$\mathrm{Imp}_{\S,x} = f(S_{\S,x}, I_{\S,x}, N_{\S}),$$

$$y_{[a,b],x} = \bigcup_{n}^{i=1} (y_{[a,b],i}, \mathrm{Imp}_{\S,i}),$$

where the measure of impact concerns the influence of 1, …, *n* sociocognitive models on agent *x* individual beliefs.

The relationship among individual cognitive maps, a collective cognitive map, and sociocognitive models is represented in Figure 6.

The process of emergence of sociocognitive models is continuous and based in concept matching during interactions. The micro-to-macro and macro-to-micro measures of impact-use notions from Latané's social impact theory allow discussion of *situated* cognitive models. The exact relation among variables *S*, *I*, and *N*, however, is still to be studied in further research. In this sense, a prototype is being developed to support interactions among artificial agents, which will allow the validation of this conceptual framework.

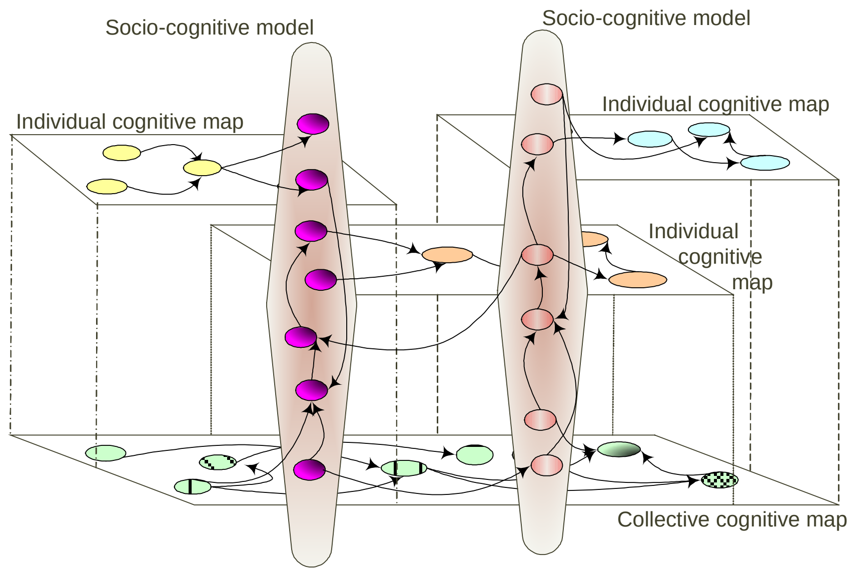

FIGURE 6 Three-dimensional cognitive representation: individual cognitive maps, collective cognitive map, and sociocognitive models

4 RELATED WORK

This research can be compared with other propositions, mainly concerning some specific aspects of the framework. The use of cognitive maps to represent knowledge can be put side by side with an artificial intelligence approach that uses a graphic notation called *semantic networks* (Sowa, 2001). Like cognitive maps, semantic networks represent knowledge through nodes connected by arcs. Nevertheless, in those networks, nodes are hierarchically typed, with derivation, according to the generality level of the nodes. Those systems are mainly used to classify or group knowledge in natural language systems. Conversely, cognitive mapping has fewer constraints and does not need particular typing; it is a general methodology. One of its strengths is precisely its ability to adapt to a large variety of domains. The same argument can be used when comparing cognitive maps with Bayesian networks. Actually, these tools have already been associated in the qualitative probabilistic networks (Wellman, 1994), a sort of cognitive mapping with causal probabilistic links, allowing Bayesian reasoning in cognitive maps. The use of the original version of cognitive maps, however, has the advantage of simplicity: cognitive maps can represent a larger domain of situations. Cognitive maps, a tool used by psychologists, allows qualitative reasoning.

The POOL2 system, proposed by Zhang, et al. (1992), composes collective maps through the aggregation of individual cognitive maps. POOL2 does not incorporate the notion of interaction among artificial agents. In A-POOL, or Agent-Oriented Open System Shell, Zhang, et al. (1994) use cognitive maps to represent artificial agents' knowledge. The communication occurs through the exchange of partial cognitive maps, and interactions are used to compose an organizational map. The most recent evolution of this system includes the propagation of numerical values (Zhang, et al., 1994). The use of quantitative inference, however, is far from the qualitative spirit of cognitive mapping. In the line of thought of A-POOL, Chaib-draa (2002) proposes a method of causal reasoning adapted to multi-agent negotiation. Chaib-draa introduces the notion of *interaction matrix* to represent different points of view. Nevertheless, the conflict detection is not dynamic throughout interactions; it is performed at a given moment. This model is not adapted to artificial agents that dynamically and continuously adjust their knowledge to a changing environment.

5 CONCLUSIONS AND FURTHER RESEARCH

Generally, in other multi-agent systems, the social macro features are preprogrammed; that is, interactions among agents modify the macro level but only within a predefined structure. On the other hand, micro-to-macro and macro-to-micro phenomena have not been modeled simultaneously (Sawyer, 2003). By drawing on cognitive science, this framework represents ill-structured emergence of social outcomes, as well as modeling both micro-to-macro and macro-to-micro phenomena. The main goal of the research, however, is not to search for an explanation of a society, but to propose some operational representation of collective values, even if this representation is a simplification of a complex reality. Here, simplification is assumed in different dimensions.

First, the use of cognitive mapping to represent agents' beliefs is obviously a simplification of agents' internal cognitive structures and interdependencies. Cognitive mapping was chosen because it concerns a well-known technology, operational enough to be used in organizations, quite powerful, and adapted to hill-structured domains. Nevertheless, the building blocks of this framework, such as the collective cognitive model and the sociocognitive models, are easily adaptable to other graphic kinds of knowledge representation technologies, like semantic networks. This would change the semantics of collective representations. Recall that a technology of knowledge representation is, among other things, a theory of intelligent reasoning and a collection of mechanisms for implementing that theory.

Another simplification hypothesis is the existence of social representations equally known by all agents. In reality, each organizational actor has its own interpretation about collective values, and interpretations can differ greatly among actors. Cognition is internal to actors, not to organizations. Social cognitive patterns emerge from interactions, but they do not really exist beside actors reasoning about their own interpretations. So, external representations, such as the collective cognitive model and the sociocognitive models, are simplifications that allow the use of collective values in multi-agent systems.

Several lines of research are opened concerning the representation of emergent social phenomena through the association of cognitive maps and multi-agent systems. One concerns the cognitive map extraction mechanism from interactions. The actual stage of the framework, however, does not consider *opinion clusters* (Rockloff and Latané, 1996). An important line of research is to study the conditions under which opinion clusters are composed. These can be based on the sociocognitive model: neighbors in a sociocognitive model can compose an opinion cluster. Also, the social influence of opinion clusters can be deduced by using some social impact formulas. The exact mechanism that relates strength, immediacy, and number, however, is not yet clear. How can we move forward and combine the impact factors? Are there any other factors? This subject is under study. Also, other theories that explain social impact are being considered, such as the social influence network theory (Friedkin, 2003).

6 REFERENCES

Axtell, R., 2000, “A Positive Theory of Emergence for Multi-agent Systems,” *Proceedings of the Conference on Computational Analysis of Social and Organizational Systems (CASOS)*, July 21–24, 2000, Pittsburgh, PA.

Bougon, M., and J. Komocar, 1994, “Façonner et diriger la stratégie. Approche holistique et dynamique,” in P. Cossette, ed., *Cartes Cognitives et Organisations*, Laval, Canada: Les presses de l’Université de Laval, pp. 57–81 (in French).

Chaib-draa, B., 1998, “A Relational Modeling of Cognitive Maps,” in R.E. Mercer and E. Neufeld, eds., *Proceedings of the Canadian Society for Computational Studies of Intelligence on Advances in Artificial Intelligence (AI-98)*, LNAI 1418:322–333, Berlin: Springer.

Chaib-draa, B., 2002, “Causal Maps: Theory, Implementation, and Practical Applications in Multiagent Environments,” *IEEE Transactions on Knowledge and Data Engineering* **14**(6):1–17.

Davis, R., H. Shrobe, and P. Szolovits, 1993, “What is a Knowledge Representation?,” *AI Magazine* **14**(1):7–33.

Firmino da Costa, A., 2001, *Sociologia*, Quimera: Lisbon (in Portuguese).

Fisher, K., and B. Chaib-draa, 1999, "A Simulation Approach Based on Negotiation and Cooperation between Agents: A Case Study," *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* **24**(9):531–545.

Friedkin, N.E., 2003, "Social Influence Network Theory: Toward a Science of Strategic Modification of Interpersonal Influence Systems," *Dynamic Social Network Analysis*, National Academy of Sciences/National Research Council Committee on Human Factors. Available at http://www.soc.ucsb.edu/faculty/friedkin/Inpress/NAS_Symp.pdf.

Latané, B., 1981, "The Psychology of Social Impact," *American Psychologist* **36**:343–356.

Latané, B., and E. Fink, 1996, "Symposium: Dynamical Social Impact Theory and Communication," *Journal of Communication* **46**:4–77.

Lawless, W.F., and J.M. Grayson, 2003, "A Quantum Perturbation Model (QPM) of Emergence," *Proceedings of the Conference of the North American Association for Computational Social and Organizational Science (NAACSOS-03)*, Carnegie Mellon University, June.

Lewis, E.T., J. Diesner, and K. Carley, 2001, "Using Automated Text Analysis to Study Self-presentation Strategies," *Proceedings of the Computational Analysis of Social and Organizational Systems (CASOS) Conference*, Pittsburgh, PA.

Lissack, M., and H. Gunz, 1999, "Managing Complexity in Organizations: A View in Many Directions," Westport, CT: Quorum Books.

Louçã, J., 2000, *Cartographie Cognitive, Réflexion Stratégique et Interaction Distribuée: une Approche Multi-agent*, Ph.D. thesis, Université Paris IX Dauphine and Faculdade de Ciências da Universidade de Lisboa (in French).

Louçã, J., 2003a, "Modeling Context-aware Distributed Knowledge," *American Association for Artificial Intelligence Spring Symposium on Agent-mediated Knowledge Management (AMKM-03)*, Stanford University, Palo Alto, CA, March 24–26.

Louçã, J., 2003b, "Representing Emergent Sociocultural Phenomena: Propositions to Support Current Emergentist Trends in Psychology and Sociology," *Conference of the North American Association for Computational Social and Organizational Science (NAACSOS-03)*, Carnegie Mellon University, Pittsburgh, PA.

Luck, M., P. McBurney, C. Preist, eds., 2003, "Agent Technology: Enabling Next Generation Computing," *Agent Link Network*. Available at http://www.agentlink.org/.

Nowak, A., and B. Latané, 1994, "Simulating the Emergence of Social Order from Individual Behavior," in N. Gilbert, and J. Doran, eds., *Simulating Societies: The Computer Simulation of Social Phenomena*, London: UCL Press, pp. 63–84.

Rockloff, M.J., and B. Latané, 1996, "Simulating the Social Context of Human Choice," in K.G. Troitzsch, U. Mueller, G.N. Gilbert, and J.E. Doran, eds., *Social Science Microsimulation*, Berlin: Springer.

Sallach, D., 2003, "The Simulation of Emergence: Mysticism, Reduction, or Social Complexity?," *Proceedings of the Conference of the North American Association for Computational Social and Organizational Science (NAACSOS-03)*, Carnegie Mellon University, Pittsburgh, PA, June 22–25.

Sawyer, R.K., 2002, "Emergence in Psychology: Lessons from the History of Non-Reductionist Science," *Human Development* **45**(1):2–28, S. Basel: Karger AG.

Sawyer, R.K., 2003, "Artificial Societies – Multiagent Systems and the Micro-Macro Link in Sociological Theory," *Sociological Methods & Research* **31**(3), Thousand Oaks, CA: Sage Publications.

Sowa, J., 2001, "Processes and Causality." Available at http://users.best.net/~sowa/ontology/causal.htm.

Weick, K., 1995, *Sensemaking in Organizations*, Thousand Oaks, CA: Sage Publications.

Wellman, M., 1994, "Some Varieties of Qualitative Probability," *Proceedings of the Fifth International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems*, pp. 437–442, July.

Zhang, W., 1990, "Cognitive Map Composition, Derivation, and Focus Generation for Distributed Group Decision Support," *Proceedings of the 23rd Hawaii International Conference Systems Science HICSS*, Vol. III.

Zhang, W., S. Chen, S., W. Wang, and R. King, 1992, "A Cognitive-map-based Approach to the Coordination of Distributed Cooperative Agents," *IEEE Transactions on Systems, Man, and Cybernetics* **22**(1), New York City, NY: IEEE.

Zhang, W., W. Wang, and R. King, 1994, "A-POOL: An Agent-oriented Open System Shell for Distributed Decision Process Modeling," *Journal of Organizational Computing* **4**(2):127–154.

INTERPRETIVE AGENTS: IDENTIFYING PRINCIPLES, DESIGNING MECHANISMS

D.L. SALLACH, The University of Chicago*

ABSTRACT

This paper takes the position that social action inherently involves meaning and that it therefore cannot be adequately modeled without representing the interpretive process among agents. However, the development of interpretive models is challenging and quickly raises issues of computational tractability. A strategy that is based on three design assumptions (agent focus, continuity reduction, and orientation fields) and three mechanisms for simulating social interaction (prototype inference, orientation accounting, and situational definition) is developed here. When the three mechanisms are incorporated into action selection, they provide a model of the social interpretation process with a higher level of verisimilitude than that achieved by many other approaches.

INTRODUCTION

In communication and action, human actors are oriented by meaning. Accordingly, they consider, discern, define, attribute, convey, question, dispute, affirm, reconsider, and evolve the meaning in every situation. Inevitably, the attribution of meaning is an indexical process: the same participants may view shared situations as having distinctive, or even conflicting, meanings. The process of attributing meaning is dynamic, often shifting rapidly as the actor's interpretation shapes and informs the subsequent flow of communications and actions.

The ability to model orientation and meaning is not a new challenge. Many significant strategies involving artificial intelligence, including semantic nets, logic-based semantics, rule-based inference, and neural networks and subsumption, have sought to address this challenge. Therefore, in a developing initiative on designing interpretive agent models, there is a responsibility to distinguish the new strategies from those already explored.

At present, agent design is bounded by the complexity exemplified by the long search for artificial intelligence and by the simplicity of reactive agents. When addressing complex environments, strategies using artificial intelligence typically encounter the limits of computational complexity. Strategies using simple agents, on the other hand, tend to require drastically simplified environments; for example, they might reduce social interactions to the random flipping of binary cultural tags (Epstein and Axtell 1996; Lustick and Miodownik 2000; Lustick, et al., 2001).

* *Corresponding author address:* David Sallach, Center for Complex Adaptive Agent Simulation Systems, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439; email: sallach@anl.gov.

The objective of the strategy developed here is to capture domain complexity without confronting computational limits by emphasizing *shared social prototypes* and the interpretations and actions through which they operate. More specifically, the goal is to design and construct three linked mechanisms that simulate the fluidity of social interaction: (1) prototype inference, (2) orientation accounting, and (3) situational definition. These mechanisms are, in turn, based on the three design assumptions, which are discussed next.

INTERPRETIVE ASSUMPTIONS

Agent Focus

One assumption of the present initiative is that agent-based modeling and simulation provide a unique and often effective research domain. The nature of this domain provides several potential advantages on which the strategy can draw. These include the ability to (1) control the complexity of the topology and artificial ecology, (2) define the action and communication capabilities available to the agent, and (3) experiment with a variety of algorithms and methods (including neural networks, genetic algorithms, and swarm algorithms). All of these opportunities provide flexibility, allowing simple, proof-of-concept models to evolve toward more complex and realistic assumptions.

Another advantage of agent models is their natural support for social processes. It is true that some agent models are only minimally social (e.g., they use individual agents as the fundamental unit of analysis, with all state and behavior defined at that level). Yet, even atomic agents that rely on simple rules (and therefore rely on limited individual intelligence) can produce a result that simulates implicitly social processes, such as local comparison or situated learning. In addition, since the models frequently assume that there are many agents distributed across space, social networks, etc., there is ample opportunity for experimentation with mechanisms that are more fully and intuitively social.[1] The very sociality of such strategies may provide a more realistic and tractable approach to the design of interpretive agents than does the direct attempt to model agent intelligence. The present discussion explores one such approach.

Reduction of Continuity

A second assumption of the present strategy is that a significant fraction of agent innovation can be represented by a translation process that goes from a continuous environment to discrete internal models that provide the basis for inference. More specifically, the present model assumes that agents are situated in a complex environment that makes available multiple, simultaneous cues (or, in a semeiotic sense, signs) that are used by the agent in defining the salient features of the current situation. Such cues include complex components of communication, such as voice tone and emphasis, facial expression, and body language. These subtle, textured situational clues and may be represented as defined on continuous domains. However, given bounded rationality, each agent "collapses" the richness of the setting into more

---

[1] As has often been argued (Minsky 1987; Sallach 1988; Gasser 1991; Axtell 2003), social algorithms and mechanisms may become increasingly important in defining the foundations of computation and information science.

discrete classifications, rules, and schemas.[2] For example, agents must determine whether a statement is a fact or a misrepresentation, a quip is a joke or an insult, or a response is indifferent or a threat, and they must analyze many other complex communications that themselves occur in a larger (ecological, technological, structural) context.

Considering the range of possible interpretations and their virtually unlimited combinations is beyond the capacity of the agent; thus, they must be reduced to a finite set of possible alternatives. However, the prospective *responses* of agents may be subtle and nuanced; therefore, they may be reasonably represented in complex and continuous forms.[3] It follows from the translation assumption that, at every step, the communicative and interpretive processes are a possible source of misunderstanding among agents.

Orientation Fields

A third assumption of this approach is that agents (dynamically) maintain an orientation field with an emotional valence for every relevant agent, object, and resource, and that this field forms the context within which inference occurs. In addition to concrete referents, an orientation field also contains typifications of various types and at diverse levels of granularity, which are the focus of agent affectivity at varying levels of intensity.

In general, there is considerable stability in affective commitments. However, events, along with associated cognitive reclassifications, can influence emotional valence (e.g., a trusted employee becomes a competitor). Accordingly, it is assumed that emotion and cognition are integrated into a co-evolving orientation field that shapes successive agent behavior.

INTERPRETIVE MECHANISMS

In the context of these three assumptions, three mechanisms can be used to at least begin to simulate the constitution of meaning. The three mechanisms are (1) prototype inference, (2) orientation accounting, and (3) situation definition. These are described now.

Prototype Inference

A primary mechanism is to use conceptual prototypes to comprehend and draw inferences about the world. In other words, human concepts take the form of prototypes rather than a set of facts, assertions, or beliefs. Agent concepts, both individual and collective, manifest a core/periphery structure with concept exemplars and with departures from that prototype, varying along the dimensions that together define the prototype concept.

Although this conceptual structure is well known and fully documented in cognitive science research (Rosch 1978, 1983; Hahn and Chater 1997), understanding the process of

---

2 The comparative advantage of particular forms of internal representations comprises an active area of research in itself.

3 For present purposes, it is not assumed that such complex responses will be expressed in natural language.

drawing inferences from a network of prototype concepts is still in an early stage. Accordingly, in the near future, methods are likely to be exploratory and perhaps domain-specific. However, the incorporation of naturally occurring data structures seems likely to add robustness and plausibility to agent models.

Orientation Accounting

This second mechanism is inspired by social pragmatism (Mead 1934; Mills 1940) and ethnomethodology (Garfinkel 1967, 2002). A simple accounting mechanism models the facts that (1) in preparing a communication or act, an agent considers the likely response of significant others, and that (2) if recent communications or acts are challenged by others, the decision must be defended, initially by invoking an anticipatory rationale. Orientation accounting locates the mechanism within the previously discussed orientation field. While there are cognitive dimensions in orientation accounting, emotional anomalies must be resolved to achieve relational stability. For example, if the challenge continues, it may be necessary to elaborate the defense and have it (minimally) accepted, accept disruption of the relationship, or acknowledge a misjudgment. Orientation accounting is implemented as a capability (set of methods) shared by all interpretive agents, and it is driven by emotional orientation.

Situational Definition

A third mechanism is situational classification. The specificity of circumstances, in conjunction with the agent response to those circumstances, provides direct input into the action selection mechanism. The ability to define the salient features of a situation has been recognized since the work of Thomas (1967) as a vital set of skills for human actors. Accordingly, it is important to model this skill set in prototype terms yet in sufficient detail to clarify domain-specific implications. Situation theory (Barwise 1989; Devlin 1991) provides the formalism with which to represent the process.

Situations are often extremely dynamic, with successive communications or actions having the potential to redefine how participants categorize and respond to the situation. As Sawyer has extensively documented, situations emerge, and participants attend to them by using their improvisational skills. The resulting experiences, which cannot be assumed to be the same, even among common participants, then feed the agent’s orientation field back into the structure.

The three mechanisms work together to shape meaning in communication and action. Communications and events create a new situation that is defined by agents in terms of existing prototype situations. Inferences are then made about causality, constraints, and probable outcomes. During the generation of communication and action alternatives, emotional accounting considers how best to justify a course of action, including the possible alteration of the course of action in order to improve the response of significant others. After a situational response, and possibly at other points in the future, prototypes are reclassified, and emotional valences are adjusted. The operation of these three mechanisms does not exhaust the components of interpretive agents, but it does provide a preliminary nucleus.

TOWARD IMPLEMENTATION

Since salient social entities and events are complex, multidimensional, variegated, and fluid, agents must have a coherent and computationally tractable means for reasoning about them. Prototypes have been explored by empirical psychology during the last several decades, but they have not yet been incorporated into agent simulation. For the purpose of this project, a prototype is defined by (1) assembling a set of dimensions[4] along which a particular entity (or related set of entities) may vary, (2) identifying clusters of core values relative to which one or more salient social entities are defined, and (3) reasoning about the location of entities (or sets of entities) and how they may influence agent orientations. These three components represent actions available to interpretive agents. They will be designed and implemented during the course of the present project.

An orientation is an emotional valence toward a salient social entity. Orientations tend to be preserved but can be called into question by events. Major or consequentially timed events can result in a wholesale reorganization of agent orientations. One constraint on the easy or frequent restructuring of orientations is the fact that they are shared with groups of other agents: groups in which the agent is a member or with whom she or he identifies. These agents, both as individuals and groups, *are* the salient social prototypes. This means that the actions constituting orientation management include (1) anticipating significant agent responses, (2) selecting and calibrating possible actions to satisfy such constraints, and (3) generating accounts in which the consequences of actions for salient individuals or groups can be justified. Together these possible actions define an aspect of agent behavior to be implemented as part of the present initiative.

Interpretive agents, complete with prototypes, orientations, and the actions with which they are managed, find themselves in situations that must be interpreted. The agents use prototypes and orientations to generate expectations and act accordingly. After the situation (itself composed of the actions of multiple agents, as well as possible exogenous developments) occurs, it is assessed, and prototypes and orientations are realigned accordingly. This process of realignment incorporates possible actions available to the agent. These actions will be designed and implemented during the course of this project.

Together, agent actions that allow the management and use of prototype inference, orientation accounting, and situational interpretation result in an *interpretative aspect* that will be available to the agents used in this project. Since prototypes and orientations are shared among groups but are individually aligned, they provide a basis of common action, but they are also a source of possible misunderstanding among agents and groups. Such misunderstandings must be negotiated or otherwise responded to in order for coordinated social action to emerge.

Thus, prototypes and orientations constitute a shared social heuristic by which coherent social behavior may be simulated without creating unbounded computational demands. The prospect of modeling interpretive agents and interpretive social processes carries the potential of a new type of social simulation that can capture the complexity inherent in meaning-oriented systems, while still remaining computationally tractable.

---

4 Mathematically and computationally, these dimensions will be represented as relational domains, with this extension: in addition to attribute values, domains may aggregate complex entities (cf., Codd 1979) and, especially, prototypes. However, every component of a complex entity must ultimately be reducible to (values defined upon) a relational domain.

DOMAIN COUPLING

While these features and capabilities are designed to be generic and thus broadly applicable, their use requires that they be mapped to particular topical domains. Each domain has its own entities, events, types, and structures that must be specifically represented in order to develop a reliable model. Thus, both designing and embedding generic mechanisms are part of the overall design process.

To illustrate the process, the generic mechanisms described above will be applied to a stylized, two-stage, electoral process. More specifically, agents will be programmed to apply prototype inference, orientation accounting, and situational definitions to the issues addressed by political actors facing elections. The interacting interpretations of diverse agents will then generate the creation and dissipation of (partially) shared political orientations relative to their political options.

The situated model containing these complex interactions is implied in the nested political games developed by Tsebelis (1990). Nested games occur when two or more games are played at different but interrelated levels, with one forming a context in which the other is conducted. Strategy choices that are apparently suboptimal may actually be a result of multilayer interaction.

For example, a two-phase election such as in the French Fifth Republic has an initial partisan phase, followed closely by a coalition phase, in which the previous competitors must close ranks within a week in order to successfully contest the general election (Tsebelis 1990, pp. 187-232). During the first round, candidates are typically motivated to attack the candidate from within their coalition (their competitor) because they seek support from the same pool of voters. During the second round, the potential for the coalition to win the office provides an incentive to support and vote for the coalition candidate.

Tsebelis's analysis of historical trends indicates that a coalition's ability to close ranks in the second phase is influenced by how competitive the coalition partners are with each other and by their prospects for winning the seat. Thus, the results of the partisan phase define a context in which the coalition phase decisions are made. Tsebelis's analysis concludes that the less competitive the parties in the coalition are, and the better their prospect of winning the office is, the more likely it is that the supporters of the losing candidate will transfer their support in the second phase. However, when this pattern occurs, it is because individuals and organizations have pre-existing orientations and interpret events in particular ways. There are also occasions when Tsebelis's generalization does not pertain, and those patterns also result from the orientations, situational definitions, and interactions of participating agents. The challenge is to construct simulation models that can capture the closing of ranks *as well as* the failure to close ranks in their situated specificity.

Tsebelis constructs a game theoretic model of these tensions. However, the resulting framework has sometimes been criticized for failing to adequately address equilibrium selection. In a broader sense, its virtues are also its faults, in the sense that the additional complexity introduced by the concept of game nesting results in a model that is analytically intractable. The nesting of layers also pushes the design of agent simulation beyond the standards of current practice, in that agents may hold views that are ambivalent and internally inconsistent. Thus, the example may usefully illustrate the prospective advantages of interpretive design as well.

For example, imagine that there are thousands of voters and dozens of issues of varying salience on which the voters may have an opinion. Multiple parties and their leading candidates attempt to attract voters who share their positions, and, in some cases, the candidates evolve toward positions that will attract key segments of the electorate. A larger number[5] of newspapers serve a particular reading audience and take positions on a variety of issues. Together, and through their interactions, these individuals and organizations define multiple orientation fields: their own and those they share. The result gives rise to a dynamic, co-evolving process.

In defining an interactive model, we can further assume that there are three small group settings in which actors influence each other: the party committee, the editorial board, and the neighborhood group. The latter group might be extended to include work groups, religious groups, etc., where the form of interaction is a sharing of orientations, as each group tries to convince other groups and/or learn the information that they have relevant to answering the question of who to support in the particular phase of the election. In each setting, individuals bring their personal orientation field, which may evolve on the basis of interactions with others, to the process of arriving at shared strategies to achieve group goals.

In sum, when the structure of a two-phase electoral process is translated into a simulation model, where heterogeneous voters from across the political spectrum are confronted with a variety of candidate strategies and first-phase outcomes as well as with the commentaries of various editorials, the contribution of situational models becomes evident. Specifically, because each agent has a potentially distinct social background and a particular position within the field of contested policies, each voter casts a ballot not only in the context of a private (subjective) state but also in the context of a unique situation.

REFERENCES

Axtell, R.L., 2003, “Economics as Distributed Computation,” pp. 3–23, in T. Terano, H. Deguchi, and K. Takadama, eds., *Meeting the Challenge of Social Problems via Agent-Based Simulation*, Springer, Tokyo, Japan.

Barwise, J., 1989, *The Situation in Logic*, CSLI, Stanford, CA.

Collins, R., 1981, “On the Microfoundations of Macrosociology,” *American Journal of Sociology* **86**(5):984–1014.

Collins, R., 2000, “Situational Stratification: A Micro-macro Theory of Inequality,” *Sociological Theory* **18**(1):17–43.

Devlin, K., 1991, *Logic and Information,* New York, NY: Cambridge University Press.

Devlin, K., 1994, “Situation Theory and Social Structure,” pp. 197–237, in M. Masuch and L. Polos, eds., *Knowledge Representation and Reasoning under Uncertainty*, Berlin, Germany: Springer-Verlag.

[5] The exact number of issues, parties and newspapers are parameters to be varied across multiple runs.

Devlin K., and D. Rosenberg, 1993, "Situation Theory and Cooperative Action," pp. 213–264, in P. Aczel, D. Israel, Y. Katagiri, and S. Peters, eds., *Situation Theory and Its Applications*, Vol. 3, CSLI, Stanford, CA.

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

Garfinkel, H., 1967, *Studies in Ethnomethodology*, Englewood Cliffs, NJ: Prentice-Hall.

Garfinkel, H., 2002, *Ethnomethodology's Program: Working out Durkheim's Aphorism*, Lanham, MD: Rowan & Littlefield.

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

Goffman, E., 1983, "The Interaction Order," *American Sociological Review* **48**(1):1–17.

Hahn, U., and N. Chater, 1997, Concepts and Similarity," pp. 43–92, in K. Lamberts and David Shanks, editors, *Knowledge Concepts and Categories*, Cambridge, MA: MIT Press.

Hilbert, R.A., 1992, *The Classical Roots of Ethnomethodology: Durkheim, Weber and Garfinkel*, Chapel Hill, NC: University of North Carolina Press.

Juarrero, A., 1999, *Dynamics in Action: Intentional Behavior as a Complex System*, Cambridge, MA: MIT Press.

Lustick, I.S., and D. Miodownik, 2000, "Deliberative Democracy and Public Discourse: The Agent-Based Argument Repertoire Model," *Complexity* **5**(4):13–30.

Lustick, I., D. Miodownik, and A.M. van der Veen, 2001, "Studying Performance and Learning with ABIR: A Research Note," in D. Sallach and T. Wolsko, eds., *Proceedings of the Workshop on Simulation of Social Agents: Architectures and Institutions, October 6–7, 2000*, Argonne National Laboratory, Argonne, IL.

Mead, G.H., 1934, *Mind, Self and Society*, Chicago, IL: University of Chicago Press.

Mills, C.W., 1940, "Situated Actions and Vocabularies of Motive," *American Sociological Review* **5**(6):904–913.

Minsky, M., 1987, *The Society of Mind*, New York, NY: Simon & Schuster.

Rawls, A.W., 1987, "The Interaction Order *Sui Generis*: Goffman's Contribution to Social Theory," *Sociological Theory* **5**(2):136–149.

Rawls, A.W., 1989, "Language, Self and Social Order: A Reformulation of Goffman and Sacks," *Human Studies* **12**:147–172.

Rawls, A.W., 2002, "Editor's Introduction," pp. 1–64, in *Ethnomethodology's Program: Working out Durkheim's Aphorism*, Lanham, MD: Rowan & Littlefield.

Rosch, E., 1978, “Principles of Categorization,” pp. 27–48, in E. Rosch and B.B. Lloyd, editors, *Cognition and Categorization*, Hillsdale, NJ: Lawrence Erlbaum.

_____, 1983, “Prototype Classification and Logical Classification,” pp. 73–86, in E.K. Scholnick, ed., *New Trends in Conceptual Representation: Challenges to Piaget’s Theory?* Hillsdale, NJ: Lawrence Erlbaum.

Sallach, D.L., 1988, “A Comparison of Parallel Architectures: Neural and Social Models of Mind,” *Proceedings of the Second ACM Symposium on Artificial Intelligence*, Norman, OK.

Thagard, P., 1999, *Coherence in Thought and Action*, Cambridge, MA: MIT Press.

Thomas, W.I., 1967, *The Unadjusted Girl*, New York, NY: Harper & Ro.

Tsebelis, G., 1990, *Nested Games: Rational Choice in Comparative Politics*, Berkeley, CA: University of California Press.

DISCUSSION:

AGENT ARCHITECTURES

(Friday, October 3, 2003, 3:00 to 5:30 p.m.)

Chair and Discussant: *Tom Howe, The University of Chicago*

Conversational Agents

Tom Howe: I think it's a fairly widely held belief at this stage in time that much of what can be achieved using the cellular automata and traditional kinds of agent-based and reactive types models has been achieved, and it's time to start investigating new ways of modeling agents, new agent architectures that can more accurately and more richly represent how societies and agents within societies actually interact.

We've switched the order around slightly in order to maintain some continuity from the previous speaker [Sawyer], and moved David Sylvan to the beginning of this session, and just bumped everyone else back. So I would at this point like to introduce David Sylvan, from the Graduate Institute of International Studies.

David Sylvan: Thank you very much.

[Presentation]

Howe: I think that was a really good description of some of the specific mechanisms that we can use to model and the problems that arise around them. And I think that that was a nice way to continue on from what Keith Sawyer said earlier about the difficulties in modeling various kinds of speech.

I have two questions. I'm going to ask one and then I'm going to open it up to the audience. You invoke Garfinkel in concepts of indexicality. It seems that most of your discussion of indexicality, though, is really focused on conversation, current and prior, but from Garfinkel's perspective there would be many, many more issues. And I'm curious how you feel is the best way to represent extended forms of indexicality as it relates to meaning within conversation.

Sylvan: It's an excellent question. I don't even have the beginnings of the answer to it, precisely because there are so many ways that one can do this. What I've done, though, just in terms again of the area that I deal with (international relations), is show how one can track some of these kinds of things exactly through a set of references. It's very rare that at least with the kind of people I'm looking at, that things are so implicit that it really is a question of a nudge and a wink and a kind of long-shared history. And you get this with certain kinds of so-called statesmen, who've interacted with each other for decades, but that doesn't happen that much, and the nice thing about, for my purposes, is that in effect the world begins again with some frequency. But you're absolutely right, for ordinary people, it's much harder. I agree with you.

Unidentified Speaker: I would agree with the last set of comments, because I can't imagine that everything is going to be found in those transcripts. And I can certainly imagine, having observed Kissinger for a number of years, that he would write a book or two about that conversation, and that a lot of the comments that he would write would be self-serving. And so it would make your job even harder to interpret and understand.

Sylvan: I agree with you. I'm looking only at the tapes.

Unidentified Speaker: I've got a question about the technology that supports the mode of communication: telephones, e-mail, chat rooms. Do the principles still hold there or is it something different, number one, and, number two, e-mail conversation seems to have the ability to actually be recorded.

Sylvan: It's a very complicated thing. The work I am doing and I started out with, is actually for a very stylized communication, for example, negotiating offers in the World Trade Organization. That doesn't occur in real time. And there you can actually use ACL-type approaches, because it's much more stylized.

Telephones, I think, when you heard they were first introduced, it was something that would be much more stultified. Already I have a lot of evidence that by the '30s, that these were becoming much more complex. But of course, there are face-to-face names, and we don't, unfortunately, have anything to access, the kind of thing you like to have, which is videotapes or things like this. That, I think, would be considered a bit much.

So it's a very good question, and we only have a very partial answer to it.

Unidentified Speaker: I just wanted to ask a sort of quick point of personal clarification in terms of what exactly you're modeling, because of your Kissinger example and also your emphasis on commitments. I'm coming from the international relations and also have serious criticisms as sort of the Copenhagen speech act approach, but I'm coming from a more materialist ontology of my criticisms, so it's not entirely clear for me, are you modeling the diplomatic process or the outcome?

Sylvan: By the process, the implicit assumption — actually *ex*plicit — is that if you can get a recall, this notion of commitment, then what that does is it get you a locally binding character … and the reason why you move toward closure or toward agreement on certain kinds of issues, or disagreement. But the point is, it gets a kind of binding nature to it, so that what happens in a very quick interaction, when you might raise your eyebrows at me and I say, "Well, too bad." That doesn't have any significance by itself, unless later on we talk further and say "Yes, we really disagreed about that." So that's what I'm really trying to get at, in a sense.

Unidentified Speaker: But the implicit assumption is that the process is in some inherent way generative of the outcome.

Sylvan: Not solely. Look, these guys go into these negotiations with all kinds of clear briefs beforehand as to what they can and can't do. But those briefs only cover approximately 5% to 10% of the topics that actually get raised. And so there's all kinds of things that happen, and if we all know about negotiators who end up having to try to sell the people back home on this …

Howe: Okay, that's all the time we have for questions for right now.

Action Selection and Individuation in Agent-based Modeling

Howe: Next up is a speech which I think is going to continue some of the themes that we've been talking about, so without further adieu, I'd like to introduce Joanna Bryson, from the University of Bath.

Joanna Bryson: I want to make a claim that there's at least three completely, almost completely, disjointed communities of people that claim that they're doing agents. And you guys are one, and then there's those guys who do KQML and that stuff. And I'm from this weird old-school group of people that used to be roboticists, and we use an "agent" to mean either a robot or an animal. And I think that what you guys are doing is very interesting, and fortunately for me at least some of you guys think what I do is very interesting. So I was asked to put it into an abstract. And I have to say, I've really enjoyed this meeting.

I'm going to talk about some of my thesis work and some of the tools that I've built, which right now are not actually integrated into the kinds of tools that you're used to, but we're going to talk about how the perspective of roboticists and people like that can help with agent-based modeling.

My title is "Action Selection and Individuation in Agent-Based Modeling," which is slightly different than what was in the program.

[Presentation]

Unidentified Speaker: I have to admit that I've been watching this architecture for probably a year and a half now, so I've been a big fan of a lot of the pieces in it and I think they're really rather exciting, particularly in the way that they create a nice, sort of discretization of action, as opposed to a continuous sense of action, and it's very difficult to decide how to actually approach these actions in a discrete way.

My question at this point would be the drive collections, the goals that you describe, particularly when you're talking about robots or animals. Those are generally predefined, exogenously defined by the modeler. What kind of support would you envision for endogenously-created drive collections and goals?

Bryson: In my papers, I've often said this is the biggest gap in my architecture. So I allow for learning, but the learning happens within the modules, and when you have to learn a new module, you have a bit of a problem.

What I'm actually really interested in doing is imitation learning, where you absorb goals, and in that case I think you may actually be forking goals. So there's nothing in my system right now that does this, but you could pretty easily imagine that you had these things lying there, and you just decided to do some simple operation on them. So you've figured out that in some context you want to do things just slightly differently, and so you just split them off and then you allow them to learn priorities.

I originally envisioned the reactive plan element as … in fact, they're dot-lap files, the little scripts for the plan part. It was supposed to be made of learnable action patterns. I eventually decided that learning was really, really hard, just like search; search and learning are the same, in fact. But I actually haven't pursued that at all, and it's for the same reasons. But I think that if you do have some information like that, you see another agent successfully do something like that, or if you've been taught a discrimination or learned a discrimination, then you want to split it. And I have some separate research, but it would take forever for four agents to be doing it, on task planning. And so in a separate thread of research, I am pursuing trying to understand how this kind of stuff would happen.

A Conceptual Framework to Represent Emergent Social Phenomena

Howe: Next up we have Jorge Louçã, from the University of Lisbon.

Jorge Louçã: The main idea of my research is that some explicit modeling of emergent macro social features in an organization might improve artificial distributed decision support in this organization. In the sense of this idea, the main goal of my research is to study micro and macro dimensions in an organization, and mainly to characterize reciprocity between these two dimensions. So my proposition is to put together a set of ideas and a set of notions in a conceptual framework, allowing them to represent emergent social phenomena, including some dynamical representation of cognitive models in the organization. This conceptual framework includes macro, representation of macro phenomena that is going to emerge from micro cognitive models and also from interactions between organizational actors.

The topics of my presentation are these: first of all, I'm going to give my point of view concerning micro and macro dimensions, and then I will present the theoretical foundations of my work, that is the sociocultural approach in psychology, social impact theory, multi-agent systems, and cognitive mapping. I'll make a brief reference to my previous research and I will explain then my proposition to describe the emergence of collective values, using some specific cognitive representation. Then, if I have the time, I will talk about related work, a conclusion and some perspectives of research.

[Presentation]

Louçã: I'd like to look to other theories that explain social impact, such as the social influence network theory or others. If anyone has a suggestion, I would be very grateful.

Howe: Thank you. I think that this represents a nice hermeneutic for defining both agent behaviors and agent values, as well as sort of group behaviors and group values. And I think it also does a nice job of providing sort of an explanation for the higher-level emergents that come out. It's not just that the collective beliefs and collective values emerge from the individual agents, but also that there's sort of a process that evolves from that as well.

I'm going to take a little bit of liberty here and ask two questions instead of just one. So my first question has to do with the concept of inner group or overlapping group, collective, cognitive maps. You talk about how individual agents' cognitive maps can be composed into a collective cognitive map, but I was wondering, what happens when you start to have overlapping groups that have influence on one another?

Louçã: Well, I haven't studied it, this problem actually. This has to be tested in practice. What I think is that we have to have some sort of mechanism to solve conflicts between different cognitive maps. And my point of view is that the best thing to do in organizations is to evidence conflicts, to show them and to propose them for people to discuss, because this is the way they accept this kind of cognitive representation. I can imagine another way to automatically choose one or another, but this is not going to be accepted, so I propose this for people to discuss.

Howe: I guess in all fairness I should open it up to the audience before I take up with the second question. So …

Unidentified Speaker: It's just a very quick question about the structure of the cognitive maps. What kind of relations are these? Would these be just some kind of simple relations between two concepts? Can you have pairs of concept, or more than two? And can you do some scenario building?

Louçã: Yes, the notion of cognitive map is very simple. It's a set of concepts and links between those concepts. Most of them concern causality links. There are works that are studying cognitive maps in a way that they can isolate a part of the map to represent scenarios or to represent some more complex idea, the notion of scheme. And, well, I think this notion can be used mainly to represent interactions and in multi-agent communication. But the initial idea is very simple.

Unidentified Speaker: Jorge, good job. One of the things that I liked about introducing Latin-A social impact theory is that you're bringing in geography or distance, and the effect that distance between people has on social influence. And I think that's an important consideration often overlooked. So thanks.

Howe: Okay, I'll go ahead and ask my second question, then, since I have a minute left and nobody else has a question right now.

As a hermeneutic for defining sort of values and behaviors, both within individuals as well as within collective groups, I was wondering how you see this sort of fitting into some of the other things that we've heard today, such as action selection, you know, language, and then, even though you can't predict what David is going to say, some of the topics of how agents make decisions.

Louçã: Well, this is a very important issue. If you want to study how agents make decisions in organizations when we use cognitive agents, the main issue is the cognitive representation, and on the other side, to reach some interaction between the artificial agent and the organizational actor that is really going to make the decisions. This is the key for the success, I think.

Interpretive Agents: Identifying Principles, Designing Mechanisms

Howe: Finally, I'd like to introduce David Sallach from The University of Chicago.

[Presentation]

[No discussion was recorded.]

Saturday, October 4, 2003

Invited Speaker:

Lars-Erik Cederman

EXPLAINING STATE SIZES: A GEOPOLITICAL MODEL

L.-E. CEDERMAN,* Swiss Federal Institute of Technology, Zurich

ABSTRACT

While international relations scholars have ignored the issue of state size, economists have explained the territorial extent of states as an optimal outcome given various constraints analogous to the theory of the firm. By focusing on full distributions rather than on average sizes, this paper adopts a systemic, generative perspective supported by agent-based modeling. Given that empirical state sizes are lognormally distributed, the goal is to reconstruct such size distributions by relying on a geopolitical model. This reconstruction task becomes possible by adding a mountainous terrain that imposes various logistical obstacles to conquest processes.

Keywords: Geopolitical modeling, international relations, agent-based modeling

INTRODUCTION

Traditionally, polarity has played a central role in international relations (IR) theory. Scholars have engaged in lengthy and inconclusive debates about whether bi- or multi-polarity is more likely to produce geopolitical stability. Despite the attention paid to the number of states, however, little has been said about their territorial extension. This fact is surprising because states are the key actors of most theories, and size is perhaps the most obvious attribute of any organization.

By contrast, economists have spent much more time accounting for state size. By drawing on the classical theory of the firm, however, their individualistic approach often downplays geopolitics. Moreover, it focuses on equilibrium outcomes at the expense of historical developments. While fitting neatly into a textbook microeconomic perspective, such scholarship is almost totally disconnected from traditional IR theory.

This paper addresses the question of state size, without losing sight of geopolitics, by adopting a systemic, distributional perspective supported by computational modeling. Size distributions contain much more information about a system than polarity, which is nothing more than a single number. Most important, distributions can be interpreted as "footprints" of the underlying mechanisms that generate them, thus helping to explain not only the states' size but their genesis and behavior in more general terms.

The starting point for this work is a new territorial dataset that shows that real-world state sizes are lognormally distributed. Because macro-historical mechanisms cannot be manipulated easily in quasi-experimental terms, a computational model is introduced, which was developed for other purposes to determine if it is possible to reconstruct the empirical distributions within

* *Corresponding author address:* Swiss Federal Institute of Technology (ETH), Center for Comparative and International Studies (CIS), Seilergraben 49, 8092 Zurich, Switzerland, +41-1-6326759; e-mail: lcederman@ethz.ch.

that framework (Cederman, 2003). The original model performs quite poorly in this respect. Adding a map of geographic obstacles, however, takes us much closer to the historical benchmark. This finding quantifies what geopolitical theorists have long suspected: rugged terrain and communication technologies have a direct impact on the scale of political organization.

EMPIRICAL STATE-SIZE DISTRIBUTIONS

The first challenge confronting a distributional perspective is to explore empirical data to determine the type of distribution that is operating. Fortunately, literature abounds on size distributions in various domains. For the present purposes, it is particularly interesting that some economists — those who choose to deviate from the microeconomic orthodoxy — have studied the systemic statistical patterns relating to firm sizes. The *locus classicus* of this scholarship can be traced back to “Gibrat’s Law,” or “the law of proportionate effect,” which states that multiplicative random walks tend to generate lognormal distributions (Sutton, 1997). More simply, this applies to processes in which an organization’s growth is proportionate to its size. Formally, lognormality implies that size S is distributed according to the following principle:

$$\log S \sim \mathrm{N}(\mu,\sigma),$$

where μ is the mean and σ is the standard deviation (Aitchison and Brown, 1957; Crow and Shimizu, 1988).

Subsequent empirical research has revealed other skew distributions that perform well as descriptive statistics of firm sizes. Following Simon and Bonini (1958), power laws are often mentioned as plausible candidates. Such distributions have a “thicker tail” that reflects a higher frequency of very large firm sizes (Axtell, 2001). Wars, measured in numbers of casualties, are power-law distributed (Richardson, 1960; Cederman, 2003).[1]

How can we distinguish lognormal distributions from power laws? The easiest way is to plot the logged cumulative frequency (c.d.f.) against logged sizes. Power laws should appear as straight lines in such frequency diagrams, whereas lognormal distributions taper off for large sizes and therefore exhibit significant bending.

Fortunately, I have been able to use a new dataset for the territorial size of states, which was generously provided by Lake and O’Mahony (2002) and Hiscox and Lake (2001). On the basis of information from several data sources, including the Correlates of War database and Polity III, their database covers the period between 1815 and 1998, excluding the world wars and the colonial empires.

Figures 1 and 2 display the frequency diagrams for the first and the last year of the sample. The diagrams depict the logged, converted c.d.f., $\log \Pr(S > s)$, as a function of logged size $\log s$. This paper uses logarithms with a base of 10. All empirical state sizes are measured in

[1] Of course, there are other skew distributions, such as the Yule distribution (see Ijiri and Simon, 1977; Reed, n.d.).

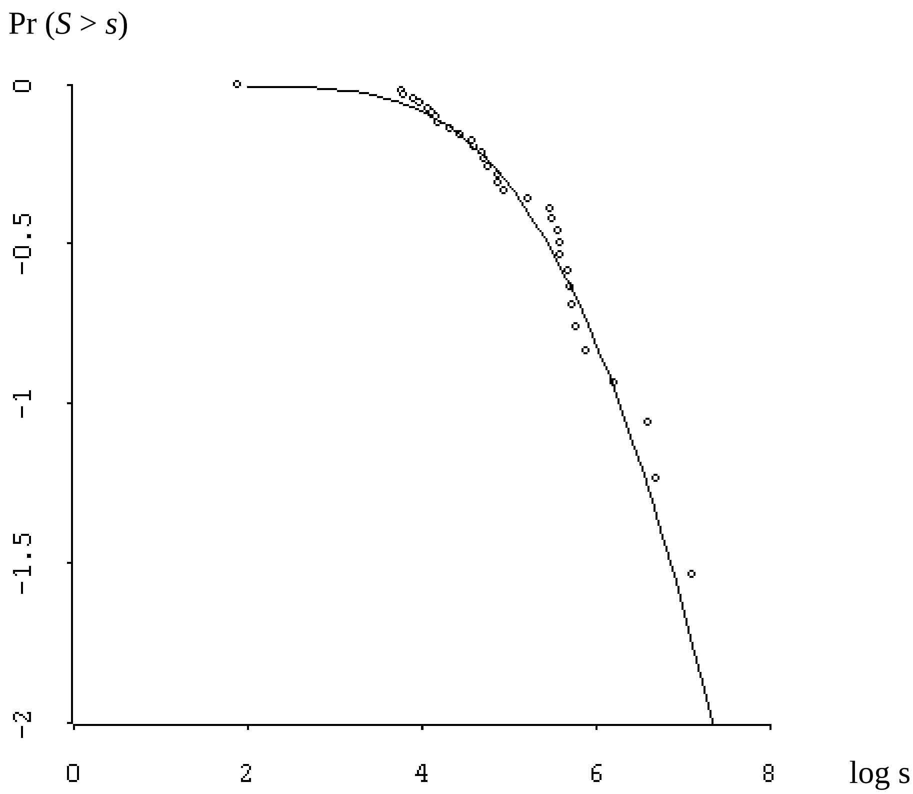

FIGURE 1 Empirical state sizes in 1815

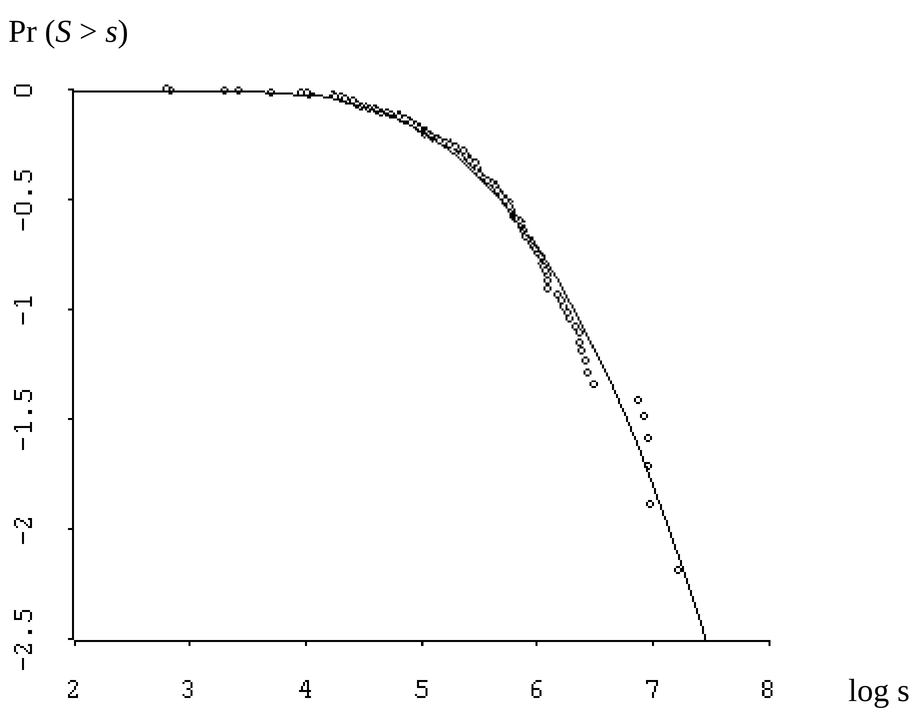

FIGURE 2 Empirical state sizes in 1998

terms of square kilometers. Intuitively, the function indicates the probability that there will be a size of an even larger size. For very small sizes, this probability is close to one, but as the size increases, it falls quickly to very small values.

Visual inspection reveals that the empirical state sizes do *not* follow a power law because of the obvious curvature of the data. To see if they are lognormally distributed, I generated such curves using maximum likelihood estimation, as indicated in Figures 1 and 2 (Aitchison and Brown, 1957). For the two years illustrated, the numerical analysis shows that a lognormal distribution captures the observations quite well.

Even in 1815, the data accurately conform to a lognormal distribution. For that year, the estimation (based on 34 observations) yields

$$\log S \sim \mathrm{N}(4.98, 1.02),$$

with a mean absolute error[2] (MAE) of 0.048. In 1998, the fit is even more impressive, as reflected by the lower MAE of 0.028 based on as many as 154 state sizes. Figure 2 demonstrates that the empirical curve approximates lognormality very well. At this point, the estimated curve approximates the following distribution:

$$\log S \sim \mathrm{N}(5.31, 0.79).$$

Further analysis, not reported here for space reasons, shows that the lognormal shape holds up for most of the intermediate years as well.

What explains the strikingly good fit? What theoretical inferences can be drawn from it? The obvious problem is the potential for an infinite number of mechanisms that are capable of producing lognormal laws (Russett, 1968, p. 315). Yet, these empirical findings are helpful because they can be used as an explanatory target. After all, far from every model is capable of generating the patterns in question. We can say, however, that any systemic theory of state size worth its salt, or possibly even any general IR macro-theory, has to reproduce this distributional footprint to claim quantitative accuracy.

Nevertheless, in enormously complex settings, such as the Westphalian state system, it is hard to match mechanisms with aggregate outcomes. For obvious reasons, counterfactual substitutions become increasingly difficult as soon as we distance ourselves from the historical path. Indeed, it is implausible that such fundamental properties as state sizes flow from superficial, short-term processes.

Computational modeling provides tools that promise to alleviate this dilemma. If we cannot conduct experiments in world history, at least we can re-create simplified, artificial worlds that lend themselves to being experimentally manipulated (Cederman, 1997, Chap. 3). Within such framework, it becomes possible to learn if specific mechanisms or conditions generate the empirically observed patterns.

---

2 The MAE offers an intuitive measure of the fit of the curves in the units of the observations. Note that the errors are calculated based on the logged axes. Other standard measures could have been used as well, such as the mean square error (MSE). The most important criterion, however, is the visual shape of the point cloud. Small MSE values may hide systemic deviations from lognormality.

GEOSIM: AN AGENT-BASED MODEL

In the past decade, I have developed a computer-based geopolitical laboratory for various analytical purposes (Cederman, 1997). Based on the agent-based tradition of Bremer and Mihalka (1977), these models share a common architecture featuring raw Hobbesian power competition among perfectly sovereign states. The current study uses a version, called GeoSim, that was developed to regenerate power-law distributed wars (Cederman, 2003). Unlike earlier versions, which were coded in Pascal, GeoSim was implemented using the Java-based toolkit Repast.

Does GeoSim produce lognormal state sizes, or is some addition needed to achieve this goal? Before attempting to answer this question, it is necessary to introduce GeoSim's main principles. Cederman (2003) describes GeoSim in great detail, so this discussion is limited to a summary; the focus then moves to the incremental changes added in this paper.

The standard initial configuration consists of a 50×50 square lattice populated by about 200 composite statelike agents interacting locally. Because of the boundary-transforming influence of conquest, the interactions among states take place in a dynamic network rather than directly in the lattice. In each time period, the actors allocate resources to each of their fronts and then choose whether or not to fight with their territorial neighbors (Figure 3). In the grid shown in Figure 3, the lines correspond to state borders, and the dots or rings, to the capitals.

All states use the same "grim trigger" strategy in their relations. Normally, they reciprocate their neighbors' actions. Should one of the adjacent actors attack, they respond relentlessly until the battle has been won by either side or ends in a draw. Unprovoked attacks can happen as soon as a state finds itself in a sufficiently superior situation vis-à-vis a neighbor. Set at a ratio of three-to-one with respect to the locally allocated resources, a stochastic threshold defines the offense-defense balance.

Because of the difficulties associated with planning an attack, actors challenge the status quo with a low probability. When the local capability balance tips decisively in favor of the stronger party, conquest results, thus implying that the victor absorbs the targeted unit. In this way, hierarchical actors form. If the target was part of another multi-province state, the latter loses its province. Successful campaigns against the capital of corporate actors lead to their complete collapse.

Territorial expansion has important consequences for the states' overall resource levels. After conquest, the capitals of conquered territories are able to "tax" the incorporated provinces, including the capital province. As shown in Figure 4, the extraction rate depends on the loss-of-strength gradient, which approaches one for the capital province but falls quickly as the distance from the center increases (Boulding, 1963; Gilpin, 1981, p. 115). This function also governs power projection for deterrence and combat. Given this formalization of logistical constraints, technological change is modeled by shifting the threshold to the right, a process that allows the capital to extract more resources and project them farther from the center. In the simulation runs reported in this paper, the transformation follows a linear process in time. (Note that in the grids, the capitals of states that have undergone at least one technological change are depicted as rings, not dots; see Figure 3.)

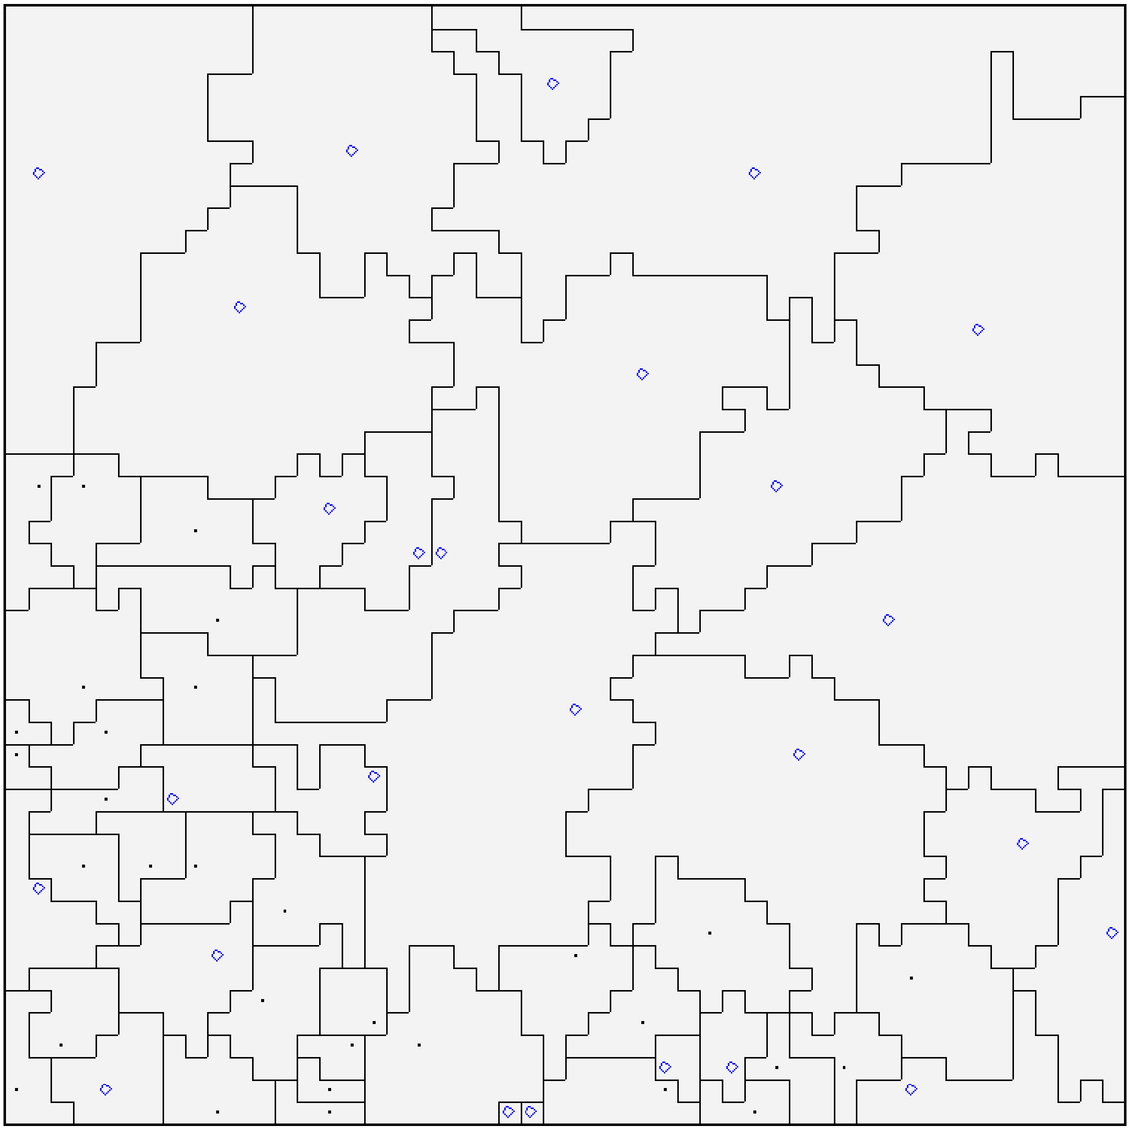

FIGURE 3 A multipolar sample system at time 5,000

Together, these rules produce four possible patterns:

1. The number of states will decrease as the power-seeking states absorb their victims.

2. As a consequence of conquest, the surviving actors will increase in territorial size.

3. Decentralized competition will create emergent boundaries around the composite actors.

4. Once both sides of a border reach a point at which no one is ready to launch an attack, a local equilibrium will materialize.

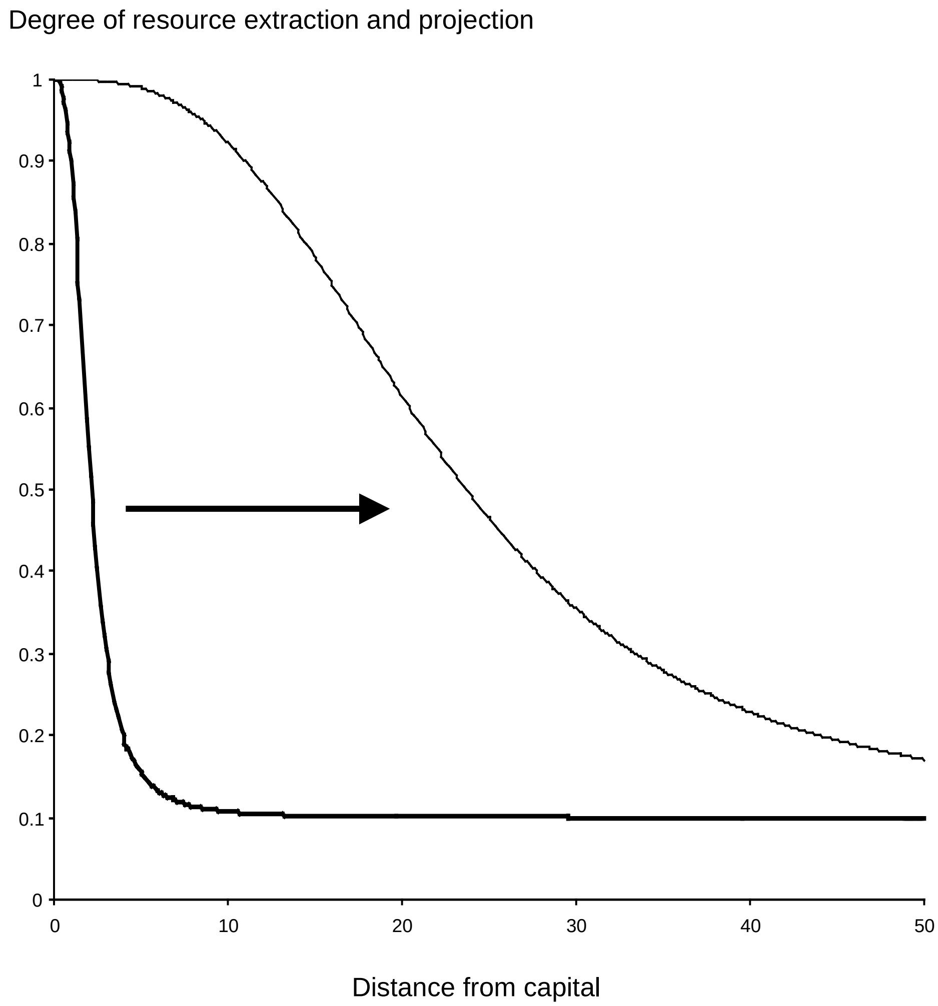

FIGURE 4 Technological change modeled as a shift of loss-of-strength gradients

STATE-SIZE DISTRIBUTIONS IN THE UNMODIFIED GEOSIM MODEL

How could the agent-based model shed light on the historical record of state sizes? It might seem that the simulation model has little to do with the empirically observed process described above. Whereas the model produces a steady fall in polarity, the dataset of Lake and colleagues points in precisely the opposite direction.

If the time perspective is lightly expanded, it becomes clear that within the Westphalian state system, polarity has decreased dramatically over the centuries. Although no precise figures exist, historians have counted about 1,000 independent political units in the Middle Ages (Jones, 1981, p. 106). Early modern Europe still had about 500 such units (Tilly, 1975). What followed was a phenomenal consolidation; by the 19th century, the number had decreased to about 25 states (i.e., half of the polarity reported for the period above).

Rather than trying to model the last 200 years of history, the goal of this paper is limited to the reproduction of state sizes as they appeared at the beginning of the empirical sample period (i.e., during the first half of the 19th century, before nationalism and democracy started to have geopolitical repercussions). This limitation seems to be reasonable, because GeoSim — in its current form — does not attempt to trace modern, participatory politics, either in terms of nationalism or democracy.[3]

Therefore, the high initial polarity characterizing the simulation runs makes sense. The main target then is to generate state-size distributions that are lognormal at a level of accuracy comparable to the estimated curves reported in the empirical section. To make a comparison possible, it is important to produce state systems with roughly the same number of states as in the 19th century (i.e., about 50).

The experimental procedure goes as follows:

1. Run a batch of 15 runs for 10,000 periods.[4]

2. Calculate the mean polarity for each time period of all the runs.[5]

3. Select a time point t* with an average of approximately 50 states.[6]

4. Estimate lognormal distributions for all the 15 runs at t*.

5. Select a representative run with the median MAE.

6. Use this value and visual inspection to evaluate the lognormal fit.

Applied to the 15 runs of the standard model in Cederman (2003), Step 3 tells us to stop at time period 5,000. At this point, the representative run with the median fit corresponds to the 55-state system shown in Figure 3.

What does the synthetic state-size distribution look like? Figure 5 suggests that it does not resemble a lognormal pattern. Rather than coinciding with the estimated c.d.f., the point cloud intersects the curve, indicating that large states are over-represented and that intermediate size states are scarce. The almost vertical drop for large state sizes suggests that these units are too similar to conform with a lognormal distribution. The relatively high MAE of 0.086 confirms the deviation from the empirical target. Compared with the distributions estimated in the 19th century, this value is at least twice as high as in most of those cases.

---

3 For attempts to extend GeoSim to such settings, see Cederman (2001, 2002).

4 Strictly speaking, all runs include an initial 500 periods before measurement starts. Thus, the total run time amounts to 10,500 periods. All run times are indicated, including the initial period.

5 To speed up the simulations, state-size distributions are computed only every 500 time steps.

6 “Atomic” states comprising only one unit are not counted because such entities reflect a lower “time resolution” in that they occur as a part of state breakups. See the rules of “structural change” described in Cederman (2003).

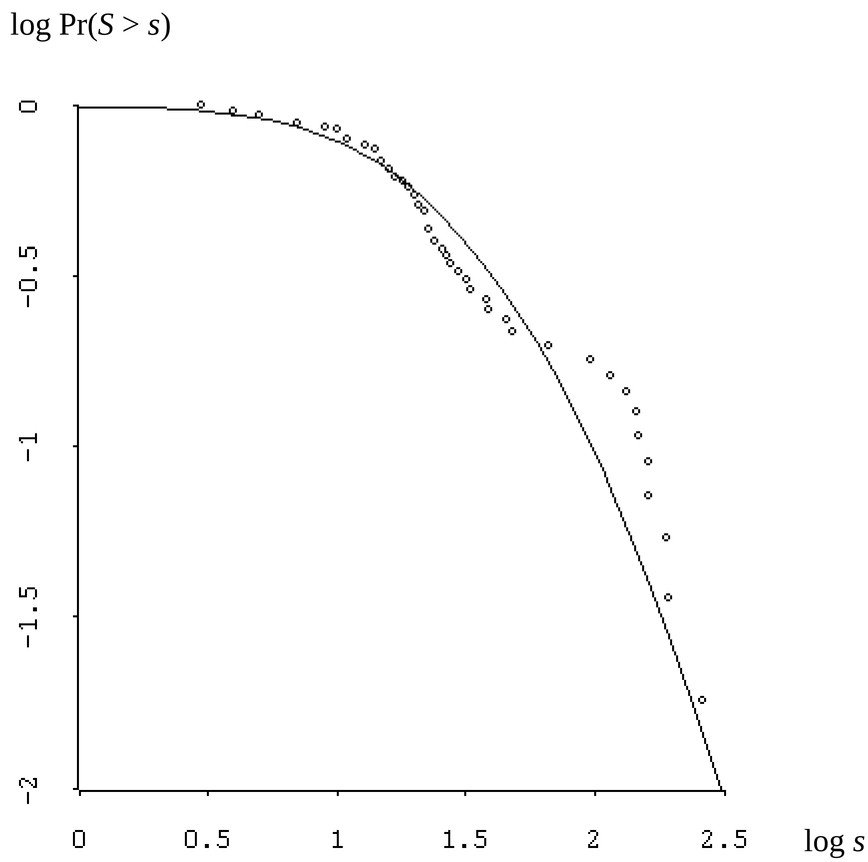

FIGURE 5 Representative size distribution at time 5,000

Careful scrutiny reveals a recurrent pattern featuring historical breaking points that allow a small number of "great powers" to outgrow the rest of the states. This phenomenon happens more or less simultaneously, thus creating historically unrealistic equality in terms of the territorial sizes. This pattern is typical for most of the runs. The histogram in Figure 6 shows that the run with the median fit is indeed representative of the entire ensemble of 15 simulations. Some error values more than triple the empirically observed rates.

In sum, GeoSim, in its standard configuration, definitely fails to generate realistically distributed state sizes. The obvious question arises: would it be possible to modify the system to bring the model's output in closer harmony with the empirical patterns?

ADDING RUGGED TERRAIN

The computational experiments described in the previous section show that a higher degree of geopolitical diversity is needed. Virtually all theories of state size postulate the presence of constraints that impose increased costs as the size of a state grows, although the identity of those obstacles varies according to the theory. In keeping with the geopolitical focus on the GeoSim framework, mountainous and inaccessible terrain was allowed to slow down

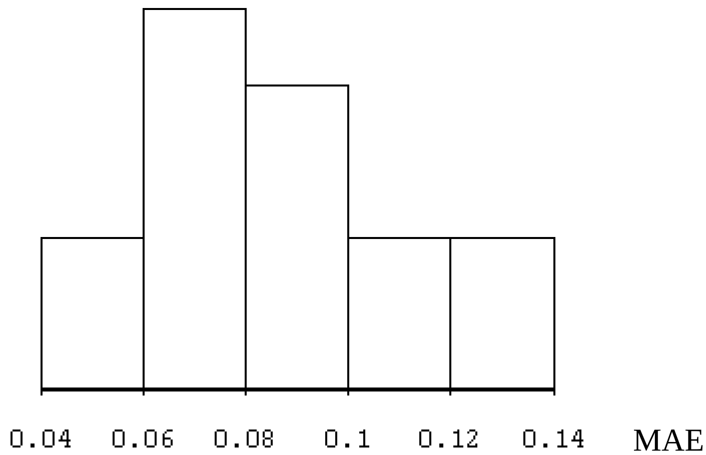

FIGURE 6 Distribution of deviations from lognormality, MAE

conquest and territorial expansion. Together with cultural differences, this factor is mentioned most often in explanations of state size (e.g., Bean, 1973). In his comprehensive study of European state formation, Jones (1981, p. 106) presents the basic logic:

> Belts of difficult terrain lying between the core-areas, and ancient ethnic and linguistic apartheid dating from early folk movements and settlement history, helped to maintain the individuality of political units. Amalgamation went so far but no farther; never to a single empire. Amalgamation costs were high. Major natural barriers protect several parcels of territory the size of modern nation-states and the more durable politics to fit the framework and there stop.

While the GeoSim framework can be adopted to feature cultural differences and nationalism (Cederman, 2002), this paper focuses exclusively on geographic obstacles. On the whole, the Jones (1981, p. 107) theoretical perspective dovetails nicely with the modeling project presented here:

> Optimal-size solutions for European states cannot be worked out as simple geometry. The spaces on the board have different values like those in the game of Monopoly and capturing and amalgamating some of them is exceedingly expensive. Where terrain did not provide much protection, units tended to disappear in takeovers by their neighbours.

Rather than trying to account for specific state sizes, Jones is content with offering a "lower-bound theory of European state formation in which other forces decide the precise outcome, but the selection of the nucleus of the rising state will be from among the richer potential cores" (p. 109). Again, this idea corresponds closely with the explanatory ambition of the current paper.

The task then is to modify GeoSim to mimic real-world geographic constraints. This change can be performed easily by creating an artificial topological map that allocates a "height" to each cell in the grid. First, a tunable number of mountain peaks is distributed randomly across the grid. A diffusion process then connects these peaks with their surrounding sites, thus creating relatively smooth mountainous terrain gently sloping down to the plains.[7]

Having described the initialization of the landscape, we need to consider the behavioral implications of the rugged terrain. The key to these modifications resides in the notion of effective, rather than geographic, distance. Whereas the base model uses Cartesian distances as input to the logistical curve presented in Figure 4, the modified version replaces this measure with one that takes the difficulty of the local terrain into account. The additional obstacle plus the distance per unit depend on the "altitude" of the mountains, where the peaks represent the maximum logistical penalty. The simulation below assumes that the initial parameter is three. Each time a province is conquered, the effective distance from the capital is calculated on the basis of the effective distance of the conquering province, adding the terrain-corrected value of the conquered area. If a mountain peak is conquered, three effective distance units (rather than one) are added to the accumulated distance from the capital. Apart from this change, the logistical distance curve shown in Figure 4 remains unchanged. It should be stressed again that capitals must cope with these constraints, both when allocating and when projecting resources.[8]

Figure 7 displays a grid with a geographic setting of this type. The gray shades correspond to the altitude of the virtual mountains. This particular snapshot describes the situation in time period 7,000. State formation has already generated a number of larger, compound states. As expected by Jones (1981), the largest states are located in the more easily accessible basins (which are depicted as the brighter areas), whereas the smaller units can be found in the mountains and the periphery of the system. Moreover, as a rule, state borders tend to coincide with the mountain ranges.

What differences do geographic obstacles make in terms of state-size distributions? Following the same experimental steps as in the previous section, I concluded that the 15 replications generated about 50 states at time period 7,000. It is not surprising that t^* is somewhat higher with the mountainous terrain than without because the additional obstacles can be expected to slow down state formation. Again, the run associated with the MAE was selected and the distribution estimated. Figure 8 presents the results of this procedure.

It is readily apparent that this distribution comes much closer to the empirical benchmark. While some deviations occur for intermediate-size states, on the whole the observations conform roughly with the estimated curve. This is reflected in a much lower mean MAE of 0.050 than the

---

[7] To be more precise, the terrain module allocates heights to every cell in the 50×50 grid. The algorithm starts by creating a random selection of mountain summits constituting a fraction `propMountains` = 0.05 of all unitary cells. A recursive algorithm is then run repeatedly `timesSmooth` = 20 times, smoothing the height of the surrounding cells. This "brush" continues to a cell in the von Neumann neighborhood with a probability of `prVisit` = 0.8. For each "visit," the algorithm sets the neighboring call to a weighted average (`propSmooth` = 0.7) of the initial cell and its previous value. From an intuitive standpoint, this algorithm is similar to water flowing from the mountain peaks, dragging with it soil that leads to a smoother landscape around the summits.

[8] This algorithm calculates distances based on the square grid, as if the system were a big Manhattan, rather than as the crow flies.

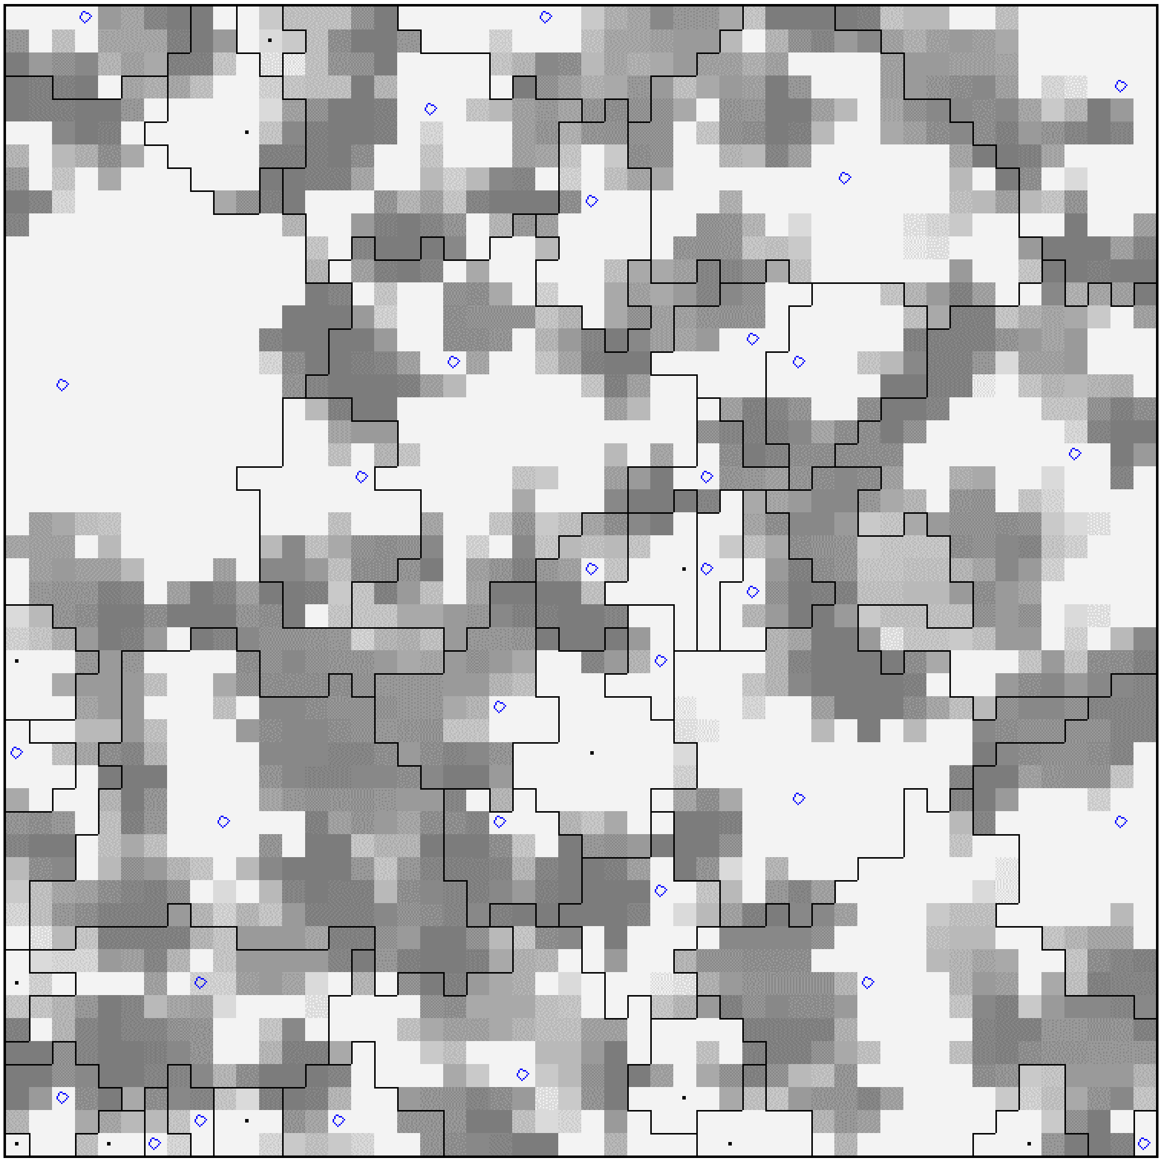

FIGURE 7 Snapshot from system with difficult terrain at time period 7,000

one obtained with the base model. The histogram shown in Figure 9 helps to gauge the fit in all 15 runs. A comparison with Figure 6 reveals that the fit for the geography-dependent runs approximates lognormal much better than those of the base model.

This experimental procedure prescribes the selection of a particular polarity level. Although the choice of about 50 states reflects a specific historical benchmark, the finding that terrain helps generate empirically realistic state sizes would be potentially fragile if it hinged on the number of states in the system. Therefore, I compiled an additional graph that plots the average of the MAEs in the 15 runs over time (Figure 10). This step allows a dynamic comparison of the two experimental configurations, with and without terrain.

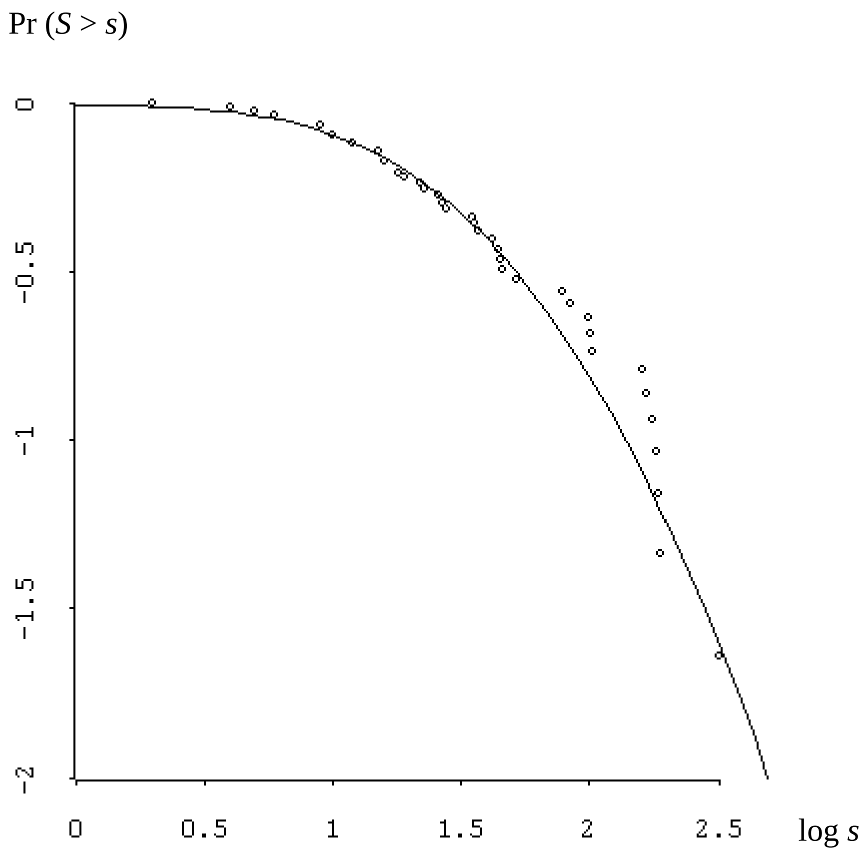

FIGURE 8 Representative size distribution for system with terrain at time 7,000

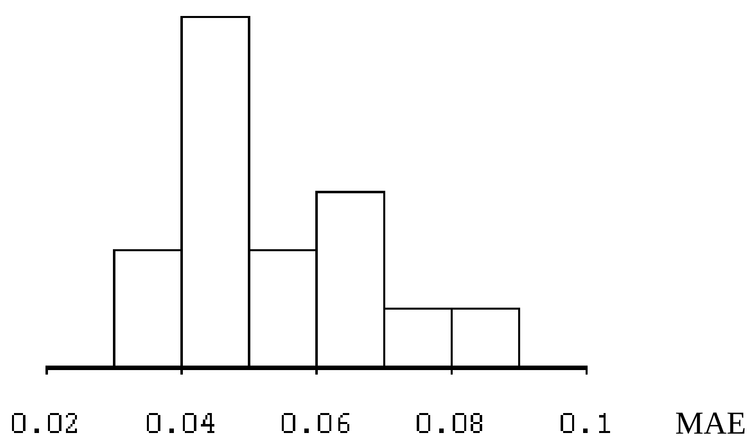

FIGURE 9 Distribution of lognormal fits in system with terrain at time 7,000

Mean MAE in 15 runs

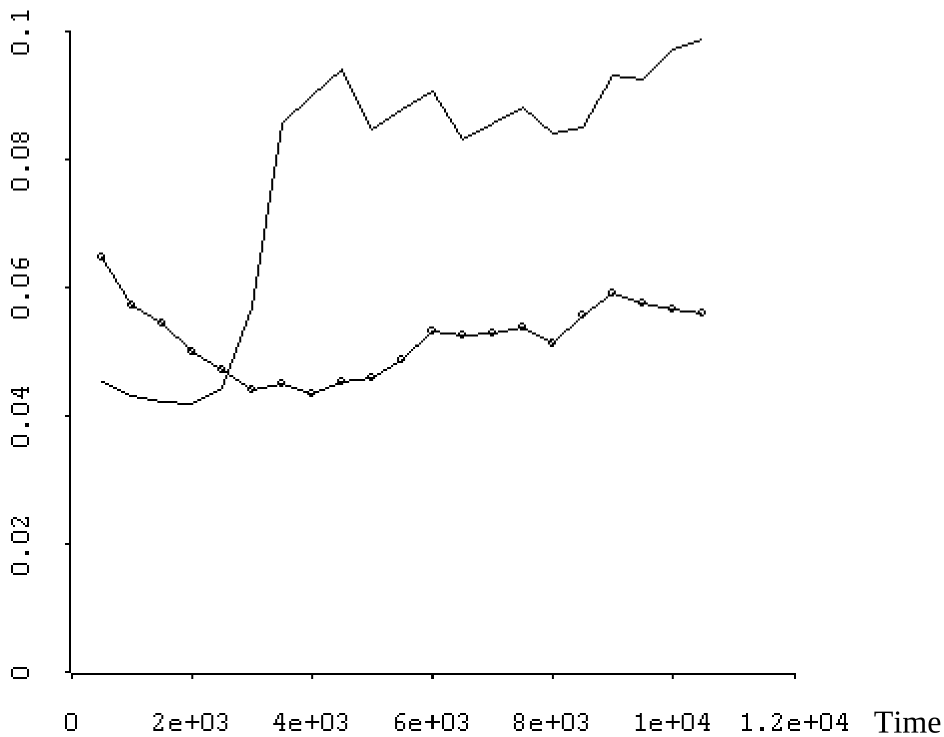

Legend:
Straight line = Systems without terrain
Dotted line = Systems with terrain (height three)

FIGURE 10 Mean value of loglinear fits over time with and without terrain

On the basis of the findings shown in Figure 7, it can be concluded that the runs with terrain (marked with dots) produce a persistently tighter fit than without it — nearly at the level of real-world data. Moreover, on average, the discrepancies remain steady throughout the simulations. In the base model, however, the runs shift decisively away from the lognormal curve between 3,000 and 4,000 time steps into the simulations. This effect is precisely what was captured in the snapshot in Figure 5.

It is obviously premature, however, to draw any far-reaching conclusions about the influence of terrain on state size. All that has been done so far is to show that in a particular artificial world, the finding can be upheld. Because most of the parameters have been intuitively tuned for the model to behave realistically, sensitivity testing and further empirical calibration are still urgently needed.

PRELIMINARY SENSITIVITY ANALYSIS

This paper offers very limited insights into the robustness of the computational findings. This section explores three theoretical dimensions:

- Consideration of different levels of terrain obstacles,
- Process of technological change, and
- Reconstruction of a size distribution that approximates the current state of the international system.

To facilitate analytical comparisons, Table 1 provides an overview of these experiments. Corresponding to a set of 15 runs, each line contains information about representative distributions selected according to the six-step procedure, and statistics about the entire set of runs. Shaded lines 1 and 3 refer to the two configurations that have been discussed. Line 1 is associated with the base model, whereas Line 3 introduces rugged terrain in the standard configuration with an altitude of three.

Focusing on the difficulty of the terrain, Lines 2, 3, and 4 together tell us that the other levels of obstacles produce similar, if not equally impressive, results. Yet, in all cases, the lognormal fit clearly improves over the base runs in Line 1. Thus, other things being equal, the main finding appears to hold reasonably robustly.

The second dimension highlights technological change. Tuned to provide power laws over a large size range in Cederman (2003), the standard rate shifts the loss-of-strength gradient by as much as 20 units from an initial 2 units during each simulation (Figure 4). To explore a less

TABLE 1 Results from the Sensitivity Analysis[a]

Parameters[b]				Representative Distribution					Output from all 15 runs				
#	dim	tech	terr	t^*	#states	μ	σ	MAE	mean #states	mean μ	mean σ	min MAE	max MAE
1	50	20	0	5,000	57	1.38	0.48	0.085	49	1.54	0.47	0.046	0.126
2	50	20	2	7,000	62	1.51	0.35	0.055	50	1.51	0.47	0.039	0.085
3	50	20	3	7,000	43	1.47	0.53	0.050	53	1.45	0.48	0.035	0.088
4	50	20	4	7,000	60	1.37	0.48	0.059	52	1.42	0.52	0.045	0.089
5	50	10	0	8,500	50	1.52	0.36	0.086	52	1.49	0.41	0.058	0.137
6	50	10	2	10,500	53	1.48	0.43	0.049	54	1.50	0.43	0.036	0.084
7	50	10	3	10,500	51	1.48	0.46	0.043	59	1.45	0.43	0.031	0.062
8	50	10	4	10,500	54	1.39	0.52	0.048	54	1.46	0.46	0.030	0.076
9	75	10	0	7,000	172	1.34	0.40	0.092	140	1.40	0.42	0.047	0.117
10	75	10	3	9,000	146	1.41	0.44	0.032	154	1.38	0.45	0.020	0.040

[a] The shaded lines refer to the runs discussed before the sensitivity analysis.

[b] dim = dimension of grid; tech = rate of technological change; and terr = max level of terrain obstacles (mountain altitude).

dramatic geopolitical transformation, the next group of runs features half that rate (i.e., a shift by 10 units). While the runs without terrain, shown in Line 5, differ very little from the base runs in Line 1, geographic constraints contribute even more strongly to generate realistic output. Lines 6, 7, and 8 reveal that for each level of ruggedness, the lognormal fit improves compared to the first set of runs with a high rate of technological change. In fact, for mountain heights of three, the MAE value of 0.043 of the representative run is one-half of the run without terrain.

Still, these results fall short of reproducing state sizes comparable to those in today's international system. Although this is not the prime scenario because of the model's limited suitability in such scenarios, it is interesting to see how far the model can be pushed. Because the real-world distribution in 1998 has as many as 154 states, it seems reasonable to generate around 150 rather than 50 states at time t^*. To produce polarity levels of this magnitude, however, it is necessary to increase the dimensions of the grid from 50×50 to 75×75. Lines 9 and 10 report the results. Here, the difference in terrain is even more marked than in the smaller grids. With a median MAE of 0.032, the topologically modified runs come very close to the empirical fit (which has a MAE of 0.027). Despite the low MAE value, the shape does not live up to loglinearity because the points exhibit a slightly exaggerated curvature, especially for large states. Whether this depends on the absence of nationalism or a fragility in the present model requires further research, which is beyond the scope of this study.

Before considering such extensions, however, it is very important to check whether the power-law result in terms of war sizes of Cederman (2003) remains robust despite the topological alterations introduced in this paper. Recall that Line 1 produces exactly the same war behavior as found in the earlier paper. A quick check of the standard model in Line 2 suggests that scaling survives the changes without problems. Inspection of all the 15 log-log plots (not shown here) confirms that linearity is upheld across the board. Numerically, the result is indistinguishable from that reported in Cederman (2003) because the median $\boldsymbol{R}^2$ is as high as 0.994 as opposed to 0.991 for Line 1, and the range of $\boldsymbol{R}^2$ values has a higher minimum: [0.983, 0.997]. As might be expected, the slopes become steeper (a median of –0.69 rather than –0.55).[9] Thus, it can be concluded that the geographic modification generates realistic sizes of both wars and states.

Having investigated the findings' robustness along a few selected dimensions, I am still not in a position to draw any firm conclusions about the independent effect of logistical constraints on state sizes. Further sensitivity analysis and empirical calibration are clearly needed. A promising avenue of analysis would build on geographic information system tools and other sources of geographic data to derive empirical measures of terrain and communications. Such data could then be used to calibrate the model (Lake and O'Mahony, 2002). For example, investigations of the entire international system appear seriously incomplete without paying more attention to naval warfare and sea communications (e.g., Rasler and Thompson, 1989).

This paper has suggested that the decrease in average state size throughout the 20th century, together with the explosive proliferation of independent states, is closely associated with nationalism and participatory politics: "These nineteenth-century developments differed fundamentally from other histories of territorial consolidation in Europe" (Rokkan, 1999,

---

9 For shock levels at 10, the improvement is equally significant (see Line 7). In this case, the median $\boldsymbol{R}^2$ value is 0.991 with a range [0.980, 0.996] and the median slope –0.72.

p. 263). To capture such effects, the model needs to be explicitly extended, which is fortunately within the realm of possibilities. Such an extension requires both a mechanism for secession and a representation of culture and identity (Cederman, 2002). In addition, it would be desirable to develop empirical measures of the colonial empires' territorial scope.

Beyond this important challenge, the long-term goal would be to join the economists and other scholars who attempt to explain the size of particular types of states; however, the current systemic approach provides a more solid theoretical method because it allows rejection of models that fail to produce the distributional benchmarks. Moreover, it appears very likely that state size is to a large extent an inherently systemic attribute.

Finally, elaborations of this type represent one possible strategy. Equally interesting, however, is the idea of radically simplifying the model. Such a project might say more about the logic of the underlying mechanisms, the detailed operation of which is very hard to trace in the GeoSim framework (e.g., Stanley, et al., 1996).

CONCLUSION

This study demonstrates that realistic, lognormal state-size distributions can be "grown" artificially. To my knowledge, no one has proposed a model with this capability. The main finding confirms the intuition of systemic theories that stress topological constraints as a course of geopolitical diversity. Furthermore, the model generates power-law distributed war sizes, thus hitting two important empirical targets simultaneously.

Of course, this paper is not the definitive word on the topic of state size. On the contrary, I hope that it will inspire others to take on the challenge of modeling macro-historical processes with computational tools, for this approach has many obvious advantages. First, as a contribution to IR theory, it goes well beyond the traditional debate about polarity, with its unclear definitions and failure to endogenize the number of states. As a complement to traditional quantitative studies and rationalistic model building, computational models of this kind can serve an important purpose in checking the empirical plausibility and internal consistency of systemic theorizing in IR.

Second, agent-based modeling also constitutes a useful alternative to individualist theories of state size, especially to those that regard this property as an equilibrium outcome. By offering an explicit representation of the dynamic mechanisms constituting nonequilibrium processes, the generative approach sheds more light on the sources of the empirical patterns than does comparative statics analysis. Since longitudinal data are available, it would be worthwhile to take advantage of this information in theory building.

Third, by introducing a new way to experiment with what is profoundly unalterable in the real world, the current computational approach will hopefully revive long-neglected geopolitical scholarship and put macro-sociological analysis on firmer ground. Without the "accounting mechanism" of agent-based modeling, it is hard to assure that intuitively compelling arguments are connected with observed macro patterns.

ACKNOWLEDGMENTS

This paper was prepared for the Agent 2003 Conference on “Challenges in Social Simulation,” October 2–4. I would like to thank David Lake and his colleagues for permission to use their new dataset.

REFERENCES

Aitchison, J., and J.A.C. Brown, 1957, *The Lognormal Distribution*, Cambridge, U.K.: Cambridge University Press.

Axtell, R.L., 2001, “Zipf Distribution of U.S. Firm Sizes,” *Science* **293**:1818–1820.

Bean, R., 1973, “War and the Birth of the Nation State,” *Journal of Economic History* **33**:203–221.

Boulding, K.E., 1963, *Conflict and Defense*, New York City, NY: Harper and Row.

Bremer, S.A., and M. Mihalka, 1977, “Machiavelli in Machina: Or Politics among Hexagons,” in K.W. Deutsch, ed., *Problems of World Modeling*, Boston, MA: Ballinger.

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

Cederman, L.-E., 2001, “Modeling the Democratic Peace as a Kantian Selection Process,” *Journal of Conflict Resolution* **45**:470–502, Aug.

Cederman, L.-E., 2002, “Endogenizing Geopolitical Boundaries with Agent-based Modeling,” *Proceedings of the National Academy U.S.A.* **99** (Suppl. 3):7296–7303, May 14.

Cederman, L.-E., 2003, “Modeling the Size of Wars: From Billiard Balls to Sandpiles,” *American Political Science Review* **1**(97):135–150.

Crow, E.L., and K. Shimizu, 1988, *Lognormal Distributions: Theory and Applications*, New York City, NY: Marcel Dekker.

Gilpin, R., 1981, *War and Change in World Politics*, Cambridge, U.K.: Cambridge University Press.

Hiscox, M.J., and D.A. Lake, 2001, *Democracy and the Size of States*, Institute for International, Comparative, and Area Studies, Department of Political Science, University of California, San Diego.

Ijiri, Y., and H.A. Simon, 1977, *Skew Distributions and the Sizes of Business Firms*, New York City, NY: North-Holland.

Jones, E.L., 1981, *The European Miracle: Environments, Economies and Geopolitics in the History of Europe and Asia*, 2nd Ed., Cambridge, U.K.: Cambridge University Press.

Lake, D.A., and A. O'Mahony, 2002, *The Incredible Shrinking State: Explaining Change in the Territorial Size of Nations*, Institute for International, Comparative, and Area Studies, University of California, San Diego.

Rasler, K.A., and W.R. Thompson, 1989, *War and State Making: The Shaping of Global Powers*, Boston, MA: Unwin Hyman.

Reed, W.J., n.d., *The Pareto, Zipf and Other Power Laws*, Department of Mathematics and Statistics, University of Victoria, B.C., Canada.

Richardson, L.F., 1960, *Statistics of Deadly Quarrels*, Chicago, IL: Quadrangle Books.

Rokkan, S., 1999, *State Formation, Nation-building, and Mass Politics in Europe: The Theory of Stein Rokkan*, P. Flora, S. Kuhnle, and D. Urwin, eds., Oxford, U.K.: Oxford University Press.

Russett, B.M., 1968, "Is There a Long-run Trend toward Concentration in the International System?," in B.M. Russett, ed., *Economic Theories of International Politics*, Chicago, IL: Markham.

Simon, H.A., and C.P. Bonini, 1958, "The Size Distribution of Business Firms," *American Economic Review* **48**(4):601–617.

Stanley, M.H.R., et al., 1996, "Scaling Behavior in the Growth of Companies," *Nature* **379**:804–806.

Sutton, J., 1997, "Gibrat's Legacy," *Journal of Economic Literature* **35**:40–59.

Tilly, C., 1975, "Reflections on the History of European State-making," in C. Tilly, ed., *The Formation of National States in Western Europe*, Princeton, NJ: Princeton University Press.

INVITED SPEAKER:

EXPLAINING STATE SIZES: A GEOPOLITICAL MODEL

(Saturday, October 4, 2003, 8:30 a.m.)

Lars-Erik Cederman, Swiss Federal Institute of Technology
Chair and Discussant: *John Padgett, The University of Chicago*

John Padgett: Lars-Erik, as you probably realize, is the most well-recognized agent-based modeler in political science. It's not just my opinion, but the field's opinion has done this. He's had a career, an Axelrod student in Michigan, gone to UCLA, gone to Harvard University, and now he's at the Swiss Federal Institute of Technology, and his work has been actually quite influential in international relations. And indeed he is a sign, you might say, that the field of political science, which, like most disciplines, is not that risk-taking, is actually starting slowly to accept this whole agent-based modeling paradigm, and Lars-Erik is the leader in that.

Lars-Erik Cederman: Well, thank you, John, for these very kind words.

Okay, so let me see if we can get going here. This is the outline on my presentation, and afterwards I'll say a few things about the theoretical background. I will show you a few pictures of the real-world state-size distributions. And then we're going computational, and I'm going to tell you a bit about how I was trying to reconstruct these distributions. Before closing, I'll also say a couple of words about the robustness.

[Presentation]

Unidentified Speaker: Could you briefly define a state?

Cederman: The sovereign entities in the international system. You can't just have a little ethnic group in the mountains; you must have a sovereign unit. I don't have the exact operational definition that Lake and company used.

Unidentified Speaker: For several years in World War II, there was occupation of a number of states in Europe. Those states still exist.

Cederman: Yes, but that's a detail compared to the grand array that I'm interested in here. I'm interested in the fact that some process is coming down from several hundreds of years, leveling out before something new happens. I'm modeling *this*. I want to land in my simulations *here* and, at the same time, have not just 50 states, but those I want to have in a long, normal row.

That's the challenge here. We'll come back to the underlying reasons. I postulate that this is about nationalism and the breakup of empires. Decolonization is a huge factor, but the breakup of the European empires will be important too. And my model in its current form doesn't really cover that. So let's focus instead on what the model does.

This work is based on my own geopolitical framework that I created in the early 1990s and has been gradually improved and reimplemented. Now it has found its optimal form due to Repast, entirely in Java. You have a dynamic network that's overlain on a territorial grid. Some people think that this is a CA basically. But it isn't; it's a dynamic network that happens to have some kind of spatial representation that is grid-based.

[Presentation Continues]

Unidentified Speaker: What does this mean for the local interaction between agents?

Cederman: There's no difference in the rules; one state can attack another, as before. But the weights put on the calculations when these states are evaluating the benefits of attacking will be affected by the distance. The hope is that this is going to lead to a rather realistic situation, where large states reside in the plains where there is fertile soil. The larger units seem to build up around these fertile bases, and that's exactly what Jones anticipates.

Unidentified Speaker: When you start the simulation, are all the states the same size?

Cederman: The simulations all start with a system seeded with 200 states, almost equally sized. During those first 500 periods, I try to let the systems settle and use the local interaction to form a more spontaneous distribution. There are no wraparound effects, for the simple reason that in the real world, in geopolitics, unless you look at the whole globe, there aren't. It's not realistic.

[Presentation Concludes]

Commentary

Padgett: Last year, Lars-Erik gave a social-theory style talk, and a distinguished social theorist, Alex Wendt, commented. This year he's giving a more data-oriented talk, so my comments will be data- and methods-oriented. I want to reflect on model validation issues, which have been very central issues for a long time. Let me frame things in the following way. Lars-Erik represents a small but growing number of agent-based modelers who are looking into distributional forms as an approach to model validation. He mentioned Doyne Farmer on stock market issues and a number of people who have processual dynamic modeling style inclinations and are turning to distributional forms as an approach to model validation. I support that move. I think this is a good occasion to reflect on how to do that and on what the strengths and weaknesses in that approach are.

Why is this attractive? Why do people want to look at distributional form rather than just mean variance type of regression approaches? The answer is very straightforward. Processual distributional form is a signature of process. Many different distributions and textbooks in the stochastic process literature show particular detailed processes of cumulation, multiplication, subtraction, division, and so forth. Sometimes they result in normal, or chi squared, or power law distributions. There's a wide range of distributions, and all of them have different processual roots.

In my early days, I did work mostly in stochastic processes. I was attracted to this work for the same reasons that people are attracted to computer simulations; namely, that it's dynamic; it revels in heterogeneity; it's not a homogeneous distribution; it's processual. There's a resemblance between the old stochastic process literature and the new computational literature that is being brought to the surface by this emphasis on distributional form. So although these two types of literature developed independently, the more people have turned to distributional forms, the more the connection between stochastic process literature and computational literature have been highlighted. I applaud this move, because the intuitions and sensitivities of these two literatures — heterogeneity, interactions, distributional outcomes — are similar.

They differ with regard to one huge variable, though. The stochastic process is overwhelmingly KISS (keep it simple, stupid). It wants really stripped-down, minimalist types of models. Computational models typically go a little bit in the other direction. Although the approaches differ a lot with regard to this one dimension, they are very similar with regard to every other dimension. Think about possible deeper and more explicit connections between computational modeling and the old stochastic process approach.

Let me comment about data before I return to stochastic processes. There were some questions early on, such as "What is a state?" The historian thinks that's a very important question, not only for its own sake. If you look at the actual trends over time and the "blips" (e.g., the World War II blip and the big states in the middle blip), the rise and collapse of an empire is apparent. It's not obvious what a state is. And the unit (of a state) is not necessarily consistent over this long time scale. Units are typically not consistent over long time scales. So historians are picky with regard to the stochastic process, computer simulation, Lake's data set, this whole enterprise. Fundamentally these are radically changing units. So doing a single model across time seems crazy.

In response, the modeler will say, "Look at those lognormal distributions. They're amazingly good." They're really good at the beginning. Lars-Erik didn't show us the middle, which I would like to see, because the middle represents the big empire. I would be much encouraged if it goes down in the middle period. But let's assume it doesn't. Then you have the following problem. The historians and common sense indicate that the really fundamental issues about changing definitions of states have been, to some extent, "swept under the rug." Yet the models are doing powerfully well at explaining this, in spite of these complications.

So how are we going to achieve goodness of fit between simple models and historical richness? Computational modeling is attractive because it at least offers the promise (delivery is another issue, of course) of filling in the continuum. You'd like to have some "knobs" that tune the world down to a very simple Gibrat's Law and then tune it up to actual history, showing the rise and fall of empires and so forth. Lars-Erik doesn't claim he's there. Yet this is the promise, and it is why we're here in this room: we believe that you could make some disciplined, principled judgments on this continuum. Don't dismiss the historians but respond to them. Determine when their issues matter with regard to lognormal distribution and when they don't. That's my philosophy on the general picture of what we are trying to do in this enterprise.

Now I'll be more specific. When you turn to what Lars-Erik is doing — fitting lognormal distributions of city and nation sizes — a question arises with regard to all these distributional sorts of models. Lars-Erik has the data on lognormal distribution, and he has a model that fits that data. He makes a move, and you learn something from that move, specifically about terrain.

He used the data as a way of pushing forward his own logic. Both his own logic and general theorizing are improved by this exercise, but that doesn't really prove that the model is valid. It means that he has learned something and he's making progress. That's an internal test. But in terms of the external test, is this what's going on in the real world? We cannot say yes or no because a lot of processes can lead to lognormals. Gibrat's original processes led to lognormals, and it was a multiplicative, randomly distributed error term. Instead of having normally distributed error plus, normally distributed error plus, normally distributed error plus, you have normally distributed error times, normally distributed error times. It is a large number of multiplicative terms that are added or multiplied together. It shows and generates Gibrat's law, and the original application in firm size distributions. This means random profit rates. For example, say there is a random distribution of profit rates, +5%, +6%, +7%, -5%, 10%, and so forth. If you are looking at the distribution of assets at the end of the day, you multiply percentages, you don't add percentages. So when you multiply this, a random walk on the profit rate generates a lognormal distribution of firm sizes. That's a classic stochastic process, minimalist, KISS type of thing that's almost a nihilist result, in the sense that a really random distribution of profits fits the data. This is a nihilist-type conclusion, yet it generates lognormal distributions.

If you look in the literature, you'll see many other processes that generate lognormal distributions. For example, I just saw one about splitting. You take a unit line and then randomly split, split, split. As N goes to infinity, you get a lognormal distribution of interval lengths. The point is that for this distribution, like any other distribution, a number of processes underlie it. Does that mean that the real world is just Gibrat's law or this "splitting thing"? No, it just means that we have to think. Now that we've thought about distributions and so forth, we have to think about external validity in addition to internal validity, and what we can do.

There are three things you can do. The first is what I asked Lars-Erik about, which is on the diagnostic level. It's great that you could show good fits, but it would be even better if you could show where the model expects bad fits. The middle period, which is the empire world, is not really well captured by the model. If the data are also not well captured by the model, that's a point in your favor. But if the thing fits the data just as well in the middle period as it did in the front and end period, that's a point against you. That's a little internal diagnostic I would emphasize.

Stick with the distribution. Work with the conditions under which you expect the distribution to work or not work. See if the rank order of fit is consistent with your self-critique of your own model. That's an internal test you can do.

The second thing you can do is what Lars-Erik did magnificently — that thing about war-size distribution. A standard critique of any distribution fitting is that it is just curve-fitting. You could take a model and try to get lognormal. In response to that, Lars-Erik could say, "We don't have to look at just one distribution. Let's look at other distributions." Say we have done curve-fitting on the target distribution. Can we use another distribution, like war-size distribution, that's not based on curve-fitting to evaluate the model? In his case, that worked very well. And the more alternative multiple distributions you have, the more confidence you can have in your results.

As a footnote, I would like to add that historians have a very different approach to validating models or interpretations than do social scientists. Social scientists take a given data set and try to fit it. Historians say a model or interpretation is good or bad, depending on how many different archives they have visited. In other words, how many heterogeneous unconnected data sources does your one interpretation fit? It may fit any one of these data sources poorly, actually. But if it fits multiple frameworks sort of coherently, that's what historians are impressed by — not goodness of fit of one particular type of data set. What I've just said about multiple distributions is a stochastic process way of making that same point: look at multiple distributions and many different data sets, and rely on that as a mechanism.

The power of both computational modeling and the stochastic process is that they generate predictions about multiple distributions. One should focus on that.

The third thing you can do is what Lars-Erik had thrown away at the end. He's sympathetic to this, but I would like him to take it seriously: distributions of change. Not just distributions of final outcome but distributions of year-by-year change. That gets much closer to the actual mechanism underneath these marginals, which is really what you're talking about: marginal one, marginal two. What does the structure of the interior of a cell look like? These distributions of change are crucial. In my work on stochastic processes, they are what I found to be most revealing. Distributions of outcomes are a good first step; they help you make sure you're "in the right ballpark." But then you can have a problem: too many models fit the same outcome. If you look at the detailed change distributions instead, you'll find radically stronger divergences among these sorts of things.

I'll conclude with a pitch. The earlier literature on the stochastic process was motivated by some of the same concerns as those held by current-world computational modelers. But basically the two worlds aren't meeting. There are no conferences attended by both stochastic process people and agent-based modelers. The two aren't connecting even though they should. The benefit of connecting agent-based models to the stochastic process is that it gives you not only well-defined literature but also more powerful statistics.

I agree strongly with what Lars-Erik said about "eyeball stuff." However, I don't go all the way in that direction, because I would like to know whether the mean absolute errors are statistically significant or not. For example, you showed some plots. I don't know if there was a good or bad fit. I have the same sorts of questions any statistician would have. I don't think eyeballing is quite good enough. But the stochastic process literature has quite a lot about sampling from lognormal distributions. You could do a lot better than what you did in terms of actual statistical significance tests if you examined that literature.

My final point is that in order build a bridge between the stochastic process and computer simulation (which will strengthen model validations and speak to historians), agent-based modelers need to have a certain style. This style relates to old debates between KISS and reality testing. You must have some "knobs" that turn your model into real KISS — that strip it down into real baseline, low-level, stochastic process stuff — so that you actually have a computer implementation of some of the standard stochastic process methods. Then as you tune it up, you can see the degree to which some parameters don't alter the basic KISS lognormal outcome. Some parameters do alter it. And you can have control over this sort of exercise. This is an argument for KISS, but I'm also saying don't stop there. I'm saying you need some parameters to tune it down so you can actually solve it in a simplified version. If we modelers did that exercise, I think we would be in a much better position to connect to the whole body of stochastic process literature, which would really strengthen our statistical and model validation sides.

Questions and Answers

Greg Madey: With regard to the distributions, you observe a lognormal on the data and suggest that lognormals are often generated by multiplicative processes. Do you have a hypothesis about what's going on in the real world that's multiplicative that would be generating the lognormal? Then when you observe the lognormal in your simulation, do you know what multiplicative process is generating it? Are they the same thing? Also, what kinds of things would change if your data were to go from lognormal to power law? Like your size of wars: what's different about the size of wars relative to the size of states? What changed? Why isn't the size of wars lognormal, or why aren't the sizes of states power law, for example?

Cederman: I'll start with the last question first. One of the most fundamental differences has to do with the fact that wars are not limited in size, as states are spatially limited. So you have a territorial system that is the upper limit. Wars, however, can drag out along a temporal axis. Also, they are partly spatial. This may be the main reason why you can't get the "fat tails" or "parlor tails" in terms of state sizes, because the world has an upper limit; the planets cannot grow.

When it comes to the first question, I'm still searching for possible mechanisms. But this is an internal search for now. It's too early for me to draw very tight links between the processes in the real world and the stylized processes I'm plugging into my models. I completely agree with John on his third recommendation. Without starting to calibrate those mechanisms using real-world data, I cannot make stronger claims. So it's still at the intuitive level. There is something multiplicative about it; that's how that plays out. It would probably require a lot of research. Even if it may seem to be an internal search, there are external repercussions. You can certainly say certain things about models that fail to produce this kind of pattern. It's like a filter explanation.

There is probably a huge number of models that can generate these outcomes. But there are even more models that cannot. In that sense, I can claim to be on the right track. For example, if Alassina and company make the assumption that all states are equally large, that assumption doesn't even begin to answer this question. There is some heuristic value.

Padgett: Do you have insight on where in your model there's a multiplicative process or something similar that's producing the observed lognormal? One of the benefits of this type of modeling is that that if you stumble across that process, that's at least a conjecture as to what's going on in the real world.

Cederman: Yes. The answer is the loss of strength gradient. If I have no terrain, then I get just a bunch of almost exactly equally sized states. So with the distance dependence, but without terrain, I'm somewhat closer to something resembling lognormal distribution. If you diversify it even more, you get even closer. My instinct is that it's really about the spatial distribution of state behavior, and how you can extract and project resources over space. That's the crucial dimension.

Unidentified Speaker: I want to take issue with the discussant rather than the speaker. I liked your three recommendations about what to do to elaborate on the results. But I want to contest this idea of the individual search. What about the philosophy of science? How long do you keep working on a model before you publish it? Nobody in science can come up with absolute certain proof. People don't prove things with absolute first principles. They just have the most likely current explanation. You can't say whether one individual in isolation has done anything good. You have to ask, "Is this currently the best explanation? Is this the simplest explanation? Is it the most powerful explanation for various reasons?" So if you come up with two things that are equally simple and powerful, you have to examine them more and find a way that they are different and separate them. Yet at some labs, people sit on results for years, and at other labs, people get results out and allow other people to play with them and attack them. I favor the latter.

Claudio Cioffi-Revilla: Lars-Erik, there's no question that GeoSim today is *the* standard for doing computational work on political research. But what would you consider to be the few challenges in model design that would be significantly innovative with respect to the GeoSim standard today or that you foresee in the near future?

Cederman: I make no claims of presenting a standard. I can point to a couple of places where I'm going with this research. I'm moving toward [tying] in nationalism, and I've presented a couple of papers on that. That's one area where I want to infuse the landscape with cultural entities. More fundamentally, when it comes to improving the whole framework, the greatest challenge may actually be to loosen up some very drastic assumptions that I've made from the beginning. I made them to make life somewhat simpler, but they may actually take away some of the real fun.

I'm thinking of the sovereignty assumption: that these entities have very sharp borders, and that when they attack an area, they may get it and incorporate that province. It would be very exciting to start to loosen up these borders, because empires just don't look like this. Most historical empires have loose fringes: frontiers rather than borders. That is a huge challenge that is on my agenda, but I don't know how many years are going to elapse before I take that one on. It has to do with the interaction topology. The progress that has been made within Repast with this new package that Tom Howe and others have presented could be very helpful. Perhaps at some point I would re-engineer the whole system to include not just civil wars and all sorts of violations against sovereignty but also even the emergence of sovereignty. How did these entities in early modern Europe emerge in the first place?

The ultimate problem with that is that you need to come up with some kind of organizational code. The entities themselves would have to have some kind of a simple representation for their organizational forms. Also I'd like to breed sovereignty from first principles. That's a tall order, to say the least. These are a couple of things on the agenda.

David Sylvan: I wonder if you might be able to commission an alternate modeling exercise, or, as a way to incorporate some other kinds of things, to try to generate something that would also get at some of these same distributions. For example, if you had a more social model, you could bring in things, not just attacking, but marrying, inheriting, going bankrupt (which was true for a lot of these princes), and coexisting. The really interesting question is less about the ability to generate these things, because I think the points that have been made about that are correct, and more about the particular kinds of absorptions. *Who* is absorbing *whom*? And the particular kind of microdynamics about the typical paths by which a particular unit ends up absorbing parts of another one.

Cederman: These are all excellent suggestions. In fact, I've been dreaming of the moment when some historian would be willing to approach this kind of research. I'd like to do teamwork, because there's no way I'm going to be able to maintain and continue to develop this type of system while at the same time being a historian who visits a number of archives. Such a meeting of minds or cultures could result in very serious scientific progress.

There are obvious reasons why it's never happened. John did a good job in his comments to explain exactly why that is so. You're right when it comes to the dynastic marriages. For instance, the Hapsburgs had a saying, "Others may fight, but you, lucky Austrians, marry" or something like that. In reality, the social networks among the sovereigns — that whole game — was actually very small, and a very small number of figures participated in it. And it's well documented, if not quantitatively, at least historically and qualitatively. The most exciting thing would be to link that game to the territorial equation and see how both interact over time. But I would certainly need help with that. That's not a one-man show.

Brian Pijanowski: I was fascinated by that historical trend line. Could you introduce into your model a test of how technology — at least technology of transportation — would change the assumptions of the interactions between the states over time, and could you vary and relax it? Have you done that?

Cederman: In fact, I'm doing that. I'm changing the communication technologies, in the sense that the loss of strength gradient that I show as this curve is actually shifting outward. So I have a very simple model of technological change. The reason is that I was unable to produce power loss without that. And the base model of GeoSim, or the predecessor, didn't have that process. That's something I added along the line. It's so fundamental that if you want to study centuries, you must have some kind of simple model of that process.

Joe Jeffrey: I was fascinated with the idea of the impact of technology on distance power projection. What do you think would happen if you had a more heterogeneous set of laws for different actors? Certainly there are lots of cases when some sovereign state has some unique technology (this includes a general notion of technology, where it's not just science or engineering, but something like management techniques for organizing armies and so forth). You might get a lot of interesting effects on your power law distributions. What's your take on it?

Cederman: The very simple process I built in is capturing some of what you're talking about, because this sliding curve represents the state of the art in the system. It doesn't mean that all states update to that curve. In fact, I have a jerky update; in all probability, the state may acquire the newest technology. So because of that function, I have differences that are absolutely crucial. It would be even nicer if I had endogenous updating, where, for instance, fighting would spread the weapons technology. Right now, it's stochastic updating that's happening.

Your point about organizational innovation is also well-taken. This is abstract enough to capture that, too. But what you're saying is right. I don't draw any strict line between technological and administrative improvements.

Social and Cultural Dynamics

INCLUDING RELATIONSHIPS, MOTIVATION, AND ACTOR-OBSERVER-CRITIC IN THE HUMAN COMMUNITY FRAMEWORK

H. JOEL JEFFREY,* Department of Computer Science
Northern Illinois University

ABSTRACT

The Human Community Model (HCM) is designed specifically for simulating human action in the context of a social system. Actions are modeled as agents engaging in the practices of a community, as opposed to the simpler input-process-output model, which is suitable for simulating machines or animals. The HCM can be used in the construction of formal models that capture a significantly wider range of the factors involved in group human behaviors, including identity, individual and group preferences, and the complex logical structure of behaviors. This paper extends previous work in three ways. First, the HCM incorporates the logical relationship between an agent's motivations and his/her relationships, providing a more complete and sensitive representation of the fluid or "smooth" changes in moment-by-moment actions. Second, the model is expanded by explicit representation of three logically distinct types of agent functioning: (1) actor, in which the agent engages in a social practice of a community; (2) observer, in which the agent simply observes the facts of a situation; and (3) critic, in which evaluations of the situation supply motivations to engage or not engage in various actions. Finally, the paper discusses the implications of the HCM for simulating emergent characteristics and features of social systems.

Keywords: Agent-based simulation, human community model, human behavior, social systems, social practices, social practice descriptions

INTRODUCTION

The Human Community Model (HCM) is intended to prove explicit representation of the logic of human action in social systems. Social systems are defined as organized groups of humans (rather than merely biological or physical agents). In the HCM, human actions are modeled as agents engaging in the practices of a community, as opposed to the simpler input-process-output model, which is suitable for simulating machines and animals. Both communities and behavior have formal representations that can capture a wide range of facets, thus allowing construction of formal, computer-implementable models that secure a significantly wider range of factors involved in group human behaviors, including identity, individual and group preferences, and the complex logical structure of behaviors (Jeffrey, 2003).

In the HCM, rather than simply being a way of thinking, all aspects, including communities, practices, relationships, and motivation, have formal articulations, as do the logical relationships among these factors. Having a formal articulation of these concepts allows us to

* *Corresponding author address:* H. Joel Jeffrey, Northern Illinois University, Department of Computer Science, DeKalb, IL 60115; e-mail: Jeffrey@cs.niu.edu.

represent, in machine-usable form, a much wider range of the aspects of behavior. The goal is to have more complete and realistic simulations of social systems.

Formal representation of social practices has been used to build software systems that model the practices of several different communities, including some question-answering systems with knowledge bases of 50 to 150 practice descriptions, and a system to actually carry out practices in a large multi-office bank (Jeffrey, et al., 1981; Jeffrey, 1983). One of the more interesting aspects of previous work is the application to practices not usually considered as formally describable, such as analyzing a person's intrinsic motivations, or a manager collaborating with a colleague.

Previous work applying the HCM to agent-based simulation focused on two aspects: the hierarchical structure of human social practices and the logical relationships between behavior and the actor's place (formally, their status) in their community, and their communities (Jeffrey, 2003). Briefly, the model presented in Jeffrey (2003) includes aspects discussed in the following section.

DESCRIPTION OF THE HUMAN COMMUNITY MODEL

The HCM is based on the central concept that human action, as contrasted with animal or machine processes, is most appropriately and usefully seen and modeled as the actor engaging in the social practices of a community. Both social practices and communities are formally modeled, which gives a significantly richer representation of the facts of social action. This paper extends the HCM by including interagent relationships and the motivational significance of those relationships for behavioral choice by the agent, and the representation of three "modes" of agent behavior: actor, observer, and critic.

A *community* is a group of persons, or "members," characterized by specific social practices, choice principles, statuses, languages, concepts, and world. Formally, a community is defined as (Putnam, 1981):

$$\text{Community} = \langle M, P, Cp, S, C, L, W \rangle,$$

where

M = members,

P = practices,

Cp = choice principles (govern choices in practices),

S = statuses,

C = concepts (distinctions),

L = language, and

W = world (of, for example, art, accounting, baseball, Lakota).

Every member occupies one or more statuses in any community in which he/she is a member. Each status has one or more (typically very few) social practices that are intrinsic to that status (i.e., practices engaged in by members that position simply by virtue of being in that position). Chess players play chess; mountain climbers climb mountains; and agent-based modelers create agent-based models of systems. Pragmatically, the intrinsic practices provide the "fuel" for a system of human actors; that is, the actions that do not need any other "cause." Rather, they are performed by an actor in that status simply when he/she has an opportunity. Every action in which a person engages or everything a person does is a social practice of a community (of which the person is a member). Thus, all actions are instances of engaging in one or more social practices of a community.

The choice principles of the actor's community function as priorities on the person's actions. "Preferences" are not an undifferentiated set (ordered or partially ordered) but include both personal preferences (in the sense of the word) and priorities among actions reflecting community values. As a number of authors have pointed out (e.g., Lustick, 2000), these values are quite significant and very different from one community to another, leading to markedly different actions. Because a person typically is a member of more than one community and has more than one position within a community, the actions chosen reflect a complex and often conflicting set of preferences and priorities.

Each practice is a hierarchically structured set of subpractices, sub-subpractices, and so on. A practice is not "atomic" but is carried out by engaging in one of its known versions. A version is a set of smaller practices such that completing the set is a case of completing the larger practice.

Social Practice Descriptions

Formally, a practice is specified by a *social practice description:*

$$\text{Social Practice Description} = <S, D, G, C, V, Sk, Kn>,$$

where

S = stages/options of the practice, the other, smaller, practices necessarily or optionally involved in carrying out this practice;

E = elements, or logical roles of the practice;

G = eligibilities of individuals for each role, specified by lists of individuals or formal rules;

C = constraints on the occurrence of combinations of stages/options, either on other stages/options or on specified states of affairs being the case;

V = versions, the sets of stages/options that are considered by members of this community to be valid instances of this practice;

Sk = skills, acquired by relevant history of engaging in practices, required to carry out this practice; and

Kn = facts, concepts, or perspectives required for the practice.

The social practice description is a formalism originally devised by Putman (1982) to extend the Ossorio (1978) basic process unit. Engaging in practices by engaging in versions, coupled with the complexities of membership and status, provides a great deal of sophistication in addressing issues of choice in action and representation of individual behaviors. Simulation proceeds by selecting the practice with the highest priority and engaging in a version of that practice, recursively carrying out Stage A by carrying out a version of Practice A, and so on.

Several researchers, including the above-mentioned formal work, have shown the value of using hierarchically structured sets of social practices in communities to capture several aspects of human action in social groups. Such work does not capture certain other aspects, however, such as actor-observer-critic functioning; the logical role of motivationally significant relationships: and the smooth, seamless way in which persons switch from one behavior to another in response to changing circumstances. The difficulty is not in devising an algorithm to model the fact that an agent's behavior changes in response to new situations, but rather in the *prima facie* verisimilitude of the simulation.

Because agent-based simulation with the HCM proceeds with each agent carrying out a set of steps that comprise a version of a practice, any reasonable simulation must address the assessment of current facts, including the state of affairs that is the outcome (by definition) of the step just completed. Certain characteristics of this assessment are important in devising the logical architecture of the system that includes it. First, in the real world, circumstances often change in ways that have nothing to do with the practice currently being enacted by an agent. It seems indisputable that on many occasions agents make observations unrelated to the practice being carried out. Persons (agents) often stop in the middle of carrying out some action, even though there is no lower-level, more detailed breakdown of that action into smaller behaviors, in light of new circumstances, for example, stopping with the forkful of food on the way to one's mouth. Modeling this fact by an a overt observation step engaged in by an actor, while possible technically, seems a poor fit with the facts — a technical "kluge."

Second, behavior, especially human behavior, is often highly sensitive to facts involving relationships between the agent and other agents or the agent and other states of affairs, including desired objects. The intensity with which agents act to obtain a valued state of affairs changes as the agent's relationship to the state of affairs changes (physical proximity being the most obvious of these relationship changes). Agents change their behavior when a practice is discovered to be harmful or would endanger another agent with whom the agent has a positive relationship, etc. The "smoothness" or "fluidity" of this change in behavior is difficult to model with only the logic of hierarchically structured social practices and community factors, as fruitful as those appear to be for capturing certain aspects of human behavior.[1]

---

1 The author is indebted to David Sallach (2003) for observations on the fluidity of human behavior in response to changed circumstances, a cogent reminder of a criticism that has been made of all frame systems and "chunking" approaches for a number of years (Dreyfus, 1992).

This paper addresses these two considerations by:

1. Incorporating the logical relationship between an agent's motivations and relationships, providing a more complete and sensitive representation of the fluid, smooth changes in moment-by-moment actions that is not represented in the formal structure of social practices *per se*.

2. Expanding the agent model to include explicit representation of three logically distinct types of agent functioning:

    a. Actor – the agent engages in a social practice of a community;

    b. Observer – the agent simply observes the facts of a situation; and

    c. Critic – evaluations of the situation supply motivations to engage or not engage in various actions.

3. Discussing the implications of this model for studying emergent phenomena in social systems, in particular creating new communities and subcommunities with distinct values and practices.

Functions of Actor-Observer-Critic

When an *agent* carries out a practice, three kinds of functioning are required. These can be thought of as either logical roles (Ossario, 1981, p. 58; 1998, p. 25) or distinct modes of functioning. The actor engages in social practices of the community, chooses which community to act as a member during a conflict, and chooses what status to act on. Acting as a member of a community in the chosen status provides a reason to engage in the practice or practices intrinsic to that status.

The *observer*, implemented as a separate process, observes all facts and "posts" them to the list of known facts.

The *critic* appraises each known fact for motivational significance. This appraisal automatically gives the agent a reason to engage in behaviors — behaviors that the agent knows provide an opportunity for achieving the change in state of affairs indicated by the critic.

Actor-Observer-Critic for Each Status

Actor, observer, and critic are types of functioning, not statuses in communities. These functions are required for carrying out the practices associated with any status because they are simply a formal representation of three logical requirements for successful functioning in a situation in which success of a behavior is not guaranteed: the agent must carry out (part of) the practice; he/she must observe the results of the action; and he/she must correct its functioning. (These fundamental facts could be phrased in terms of feedback loops, but for our purposes, rephrasing of these basic behavioral facts in the language of electronic circuits seems to have little value.)

Because each kind of functioning is associated with a status, an agent's identity repertoire (the statuses he/she occupies in each community of which he/she is a member) corresponds to an actor-observer-critic repertoire, representing the three kinds of functioning of the agent in each status.

Relationships

It is commonplace that persons (and some animals) have relationships with other agents and other states of affairs that affect their behavior. These affairs can be either simple, such as having food or avoiding danger, or much more complex, such as engaging in an intrinsic practice of a religious community.

The relationships between the agent (person or animal) and other agents or states of affairs are logically related to motivation for the agent to engage in practices: having relationship *R* (to another agent, situation, object) gives an agent a reason (motivation) to engage in certain practices and not to engage in others. If *A* has the relationship "friend" to *B*, he/she is motivated to engage in certain kinds practices with *B*, to refrain from engaging in others, and so forth.

Because relationships and behaviors have this fundamental motivational character, which plays a large role in the behavior chosen by an agent, the HCM has been expanded to include this aspect of agent functioning, based on the articulation of this range of facts in Ossorio (1998, pp. 87–91).

Each relationship is associated with a set of behaviors practiced in the community. Having relationship R_i gives the agent reason (i.e., motivation) to engage in practices $P_{i,1}$ ... $P_{i,n}$. Intuitively, reasons can be stronger or weaker, which is modeled by associating with reason R_i to engage in practice P_j, $-1.0 \leq S_{i,j} \leq +1.0$. The agent's total motivation to engage in P_j is the sum of the strengths of the motivations to engage in P_j.

Agent Functioning

Each agent engages in the practice P_j for which the total motivation is highest. The agent engages in P_j by becoming involved in a version of it via the actor. The fact that the agent is participating in the practice, a version of the practice, and the particular stage (or substage, sub-substage, etc.) currently being engaged in is recorded in the representation of the current state of the agent's world — the facts known to be the case at that time.

Independently (in separate threads), the observer of each status occupied by the agent watches facts. Each fact, when noted, is recorded in the representation of currently known facts. In a parallel set of threads, the critic of each such status appraises the known facts, adding those appraisals to the known facts (since an appraisal is a special kind of fact, one with motivational weight). The actor thread, carrying out P_j, is interrupted when the practice called for differs from the practice being engaged in (a fact reported by the observer and appraised by the critic).

When a practice is interrupted, it does not acquire a special status, such as being placed on a push-down stack. Rather, it is handled as any other fact: the fact that it is "in progress" is observed and appraised, which automatically gives the agent a certain amount of motivation to return to it.

Examples of the HCM

Everyday Example

The elaboration of the HCM, and its operation, can be illustrated by an example used in earlier work to highlight the enhancements to the model. In this example, a couple goes out to dinner to celebrate their wedding anniversary. A partially complete formal description of the practice involved is given to illustrate the correspondence between the formalism and the ordinary-language description. At dinner, each spouse has a glass of wine. The husband sips the wine. The couple engages in several practices, beginning with sipping the wine and proceeding to larger and more significant practices, all of which add up to celebrating an anniversary. This last practice is intrinsic to the status of husband (and wife). The practices followed are listed below:

- The husband sips the wine.
- The husband has a meal.
- The husband is dining.
- The couple has a meal together.
- The couple dines at a nice restaurant.
- The couple celebrates an anniversary.

(In each case, the verb phrase names the practice. Note that the phrase is a formal name. Thus, "dines at a nice restaurant" is the formal rather than informal name for a practice.)

A partial description of the top levels of this hierarchy of practices is shown in Table 1. In celebrating their anniversary, a couple engages in Paradigm 2, Version 1, a paradigm case that consists of Stages 1a, 2, 3, 4, 5, 6, 7, 8, and 9. Stage 7 takes place by engaging in the paradigm case version 7a, 7b, and 7c. Stages 7a and 7b take place by engaging in the paradigm case Version 7ai through 7av.

In a straightforward, unproblematic case, this practice begins and is carried through to completion. A number of authors have pointed out, however, that human complexities surface in the cases that are not straightforward (i.e., ones that are either realistic or not).

Problematic Example

In light of the everyday example, let us consider a problematic case. In this case, the husband (or wife) drinks a glass of wine when a diner at another table drops a glass, breaking it,

TABLE 1 Top levels of hierarchy of practice (partial description)

Example 1: HCM elaboration and operation:

Couple celebrates a wedding anniversary

- *Paradigm 1*: Couple buys gifts for each other
- *Paradigm 2*: Couple dines at a nice restaurant
- *Paradigm 3*: Couple goes on a cruise

Couple dines at a nice restaurant

1. Couple goes the restaurant via one of three options:
    a. By car
    b. By train
    c. By walking
2. Couple is seated and engages in one or more options (items 3–5)
3. Couple examines menu
4. Couple examines wine menu
5. Couple orders wine
6. Couple orders food
7. Couple eats meal together
    a. Husband eats meal
        i. Person eats salad (optional)
        ii. Person eats soup (optional)
        iii. Person eats main course
        iv. Person eats dessert (optional)
        v. Person drinks wine (optional)
    b. Wife eats meal
        i. Person eats salad (optional)
        ii. Person eats soup (optional)
        iii. Person eats main course
        iv. Person eats dessert (optional)
        v. Person drinks wine (optional)
    c. Husband and wife converse
8. Couple pays
9. Couple departs the restaurant

and cuts his/her hand on the broken glass. The other diner is bleeding. Table 2 provides details of the HCM and highlights its functionality by giving details of the problematic example.

Advantages of the HCM

An advantage of the HCM is that it provides an explicit framework for integrating research into replication of what is sometimes called “cognitive” functioning (i.e., algorithms based on representations of facts, including relational ones, to derive new facts, including both observations and appraisals). It formally represents the relationship between cognition and

TABLE 2 HCM problematic case

Example 2: HCM elaboration and operation for a problematic case

1. Couple dines at a restaurant.
2. Guest at adjacent table cuts hand on broken water glass — a fact with no place in the practice of dining.
3. Couple observes that the other diner is bleeding. This fact is posted to the known facts for each agent as is the state of affairs that the relationship to the injured diner is fellow member (the “default” relationship between members of a community).
4. The critic function of each agent (more precisely, the critic function of the status of member in good standing of the community) appraises the new fact: injury to fellow member of community. Action called for: see that the member gets aid.
5. Husband and wife suspend dining; both observe whether activities involved in helping an injured person (subpractices of emergency help practices known to the husband and wife) are occurring.
6. Husband and wife observe help arrive.
7. Husband and wife observe that fellow member has now received aid.
8. The practice with the highest motivational value is chosen. (Dining has previous value plus the fact that it is in progress.) The fact that one is engaging in an intrinsic practice carries automatic reason to continue it.

action, and therefore for assessment of the implications for agent functioning of theories, algorithms, or empirical findings in cognitive modeling and/or artificial intelligence.

SIMULATION OF EMERGENT PHENOMENA

The HCM does not include any explicit theory on how choice principles develop or evolve. The reason for this is that although evolution of a facet of a community can be modeled as a process, that process is not *per se* a social practice of the community (although a community may, and some do, have practices in which the outcome is to change specific practices). The fundamental goal has been to provide formal representation of all of the logic of actions in a community, thereby providing the means for simulating the life of a community.

However, there is a closer relationship between this goal and that of understanding emergent phenomena than is initially apparent. The HCM provides three mechanisms for modeling emergence. First, the basic “unit” of simulation in the HCM is the action of the individual agent (keeping in mind that one of the central features of the HCM is that individual actions reflect several kinds of individual preferences *and* community principles). The laws that govern individual behavior can result in the emergence of global phenomena and complex states, which is perhaps the central concept in the discipline of chaos and complex systems, as discussed by Gilbert (2003), Page (2003), and other researchers in the field. The elements of the HCM provide a significantly expanded set of logical primitives involved in human behavior in social groups, in comparison with the usual primitives for describing human action. The richer set of primitives provides a logically more complex system in which this kind of emergent-from-individual-behaviors phenomena can occur. (The situation is perhaps analogous to having more dimensions in a physical system.)

Second, many animal and human practices have an associated latency (i.e., a time between episodes of a practice). This fact holds even for practices that are intrinsic, that is, for which the agent needs no reason, only an opportunity. Thus, for example, if we model a chess community, we know that playing a game is intrinsic to chess enthusiasts and is (tautologically) engaged in whenever an opportunity arises. Realistically, however, actual human chess players delay playing their next game for some period (which varies from agent to agent and clearly must be modeled statistically in the case of many chess players). This latency results in a cyclical pattern often found in practices, including eating (or for humans dining), sexual activity, sanitary activity, growing of crops, hunting, building, and so forth. In addition, the heat and light of the sun create 24-hour patterns found in several areas of human life. If we consider such biological processes as plant growth, we can see that these cyclical patterns can extend over several seasons or years (e.g., tree growth, reforestation, or even millenia-long weather patterns). In summary, these cyclical nonbehavioral conditions create cyclical emergent *behavioral* phenomena, at all time scales, modeled by including these conditions in the constraints of the social practice descriptions of an HCM of the community of interest. (The cause of the latency may or may not be biological. Of interest here is the set of cycles, at various time scales, and their impact on practices.) This seems particularly interesting when coupled with the observer and critic functioning, as it appears to be a model of the subjective phenomenon of a need or desire to do some entirely unrelated action arising in the midst of one's activity. (Has anyone *not* experienced the need for food or coffee in the middle of an intense and interesting intellectual discussion?)

The third connection to emergent phenomena and change, the one we personally find most interesting, is that human agents act deliberately to create new communities and choice principles. That is, humans commonly engage in social practices that specifically result in a new community. The simplest example of community creation is perhaps one child asking, "Will you be my friend?" If the answer is affirmative (the second action in the practice), the result is a new two-member human community. The same logic can be observed in the deliberate formation of alliances between groups and nation states — creation of a new community by deliberate action.

Consider how this seemingly simple addition changes the dynamics of a basic predator-prey community. The basics of this community consist of two statuses — predator and prey, each with practices and choice principles:

- Practices
  1. Predators: prey, fight, mate, raise young
  2. Prey: flee predators, grow food, mate, raise young

- Choice principles
  1. Predators: high priority for prey
  2. Prey: high priority for fleeing

These practices are modeled via a practice description, each of which fundamentally combines task analysis (with constraints on allowable sets of tasks), role specification, and versions. Thus, the practice to "grow food" might include "sow seed, tend crop, harvest," and so on.

To this simple model, let us add the status of pacifist and the following practices:

3. Prey: convert predator to pacifist
4. Prey: convert predator to defender
5. Pacifist: convince predator to spare prey

Here, we see the basis for modeling situations, which for humans, is more realistically termed "ally" and "enemy," rather than predator and prey; allies can sometimes be convinced to change sides or become peacemakers.

Next, consider a community in which pacifists have following additional practice:

6. Pacifist: convert predator or prey to pacifist

Clearly, we now have the basis for a simulation in which the group of pacifists in the community is likely to increase because pacifists create more pacifists.

A second form of community creation is one that can be observed empirically with some regularity — the deliberate creation of a community *with a specific community value (choice principle).* If Agents A_1 and A_2 both have the individual value V, V may or may not be a community choice principle. But A_1 and A_2 can engage in a different deliberate action — forming community C in which value V is part of the *definition* of C. This phenomenon occurs when a group of people, for example, create a community service club or a society to promote a cause. In this case, V is now a choice principle of C. ("God wants us to spread this religion" is recognizable as verbal behavior indicative of this principle of community creation.)

Finally, one of the most common forms of community creation, and also one of the most far-reaching, is perhaps the creation of a community with a single overarching, social practice, the practice for which that community exists, that is, an *organization* (Putnam, 1983). Thus, "Agent A creates an organization" is a third kind of community-creation practice, arguably one the most significant for simulating human societies in light of the impact of organizations (including businesses). Once created, organizations carry out the practice for which they were created: building houses, growing corn, conquering territory, creating theories, publishing newspapers, and so on, almost *ad infinitum*.

CONCLUSION

The central concept of the Human Community Model is that human action, as contrasted with animal or machine processes, is most appropriately and usefully seen and modeled as the actor engaging in social practices of a community. Both social practices and communities are formally modeled. The formalism allows explicit representation of a significantly wider range of facts of human action than does the usual input-output-process–type model. This paper extends the HCM by including interagent relationships and the motivational significance of those relationships for behavioral choice by the agent, and the representation of three modes of agent behavior, namely, actor-observer-critic. The purpose of the extensions is to provide greater accuracy and verisimilitude when simulating societies, including both the daily life of the society and certain kinds of social evolution and change.

REFERENCES

Dreyfus, H., 1992, *What Computers Still Can't Do: A Critique of Artificial Reason*, Cambridge, MA: MIT Press (revised from 1972 and 1979 editions).

Gilbert, N., 2003, "Varieties of Emergence," in C. Macal and D. Sallach, eds., *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, Argonne National Laboratory, Argonne, IL, pp. 41–50.

Jeffrey, H.J., 1983, "The MENTOR Project: Replicating the Functioning of an Organization," in K.E. Davis and R.M. Bergner, eds., *Advances in Descriptive Psychology, Vol. III*, Greenwich, CT: JAI Press.

Jeffrey, H.J., 2003, "Behavior as Social Practice in Community as the Basis for Agent-based Simulation," presented at the *Second Lake Arrowhead Conference on Human Complex Systems*, Lake Arrowhead, CA, March 19–23.

Jeffrey, H.J., H.P. Zeiger, T. Schmid, and A.O. Putnam, 1989, "LDS/UCC: Intelligent Control of the Loan Documentation Process," in *Proceedings of the Second International Conference on Industrial & Engineering Applications of Artificial Intelligence and Expert Systems*, New York City, NY: ACM Press, pp. 577–588.

Lustick, I.S., 2000, "An Agent-based Modelling of Collective Identity: Testing Constructivist Theory," *Journal of Artificial Societies and Social Simulation* **3**(1). Available at http://www.soc.surrey.ac.uk/JASSS/3/1/1.html.

Ossorio, P.G., 1978, *What Actually Happens*, Columbia, SC: University of South Carolina Press, pp. 52–53.

Ossorio, P.G., 1981, "Outline of Descriptive Psychology," in K.E. Davis, ed., *Advances in Descriptive Psychology, Vol. 1*, Greenwich, CT: JAI Press.

Ossorio, P.G., 1998, *Place*, Ann Arbor, MI: Descriptive Psychology Press; originally published as LRI Report 30a, Boulder, CO: Linguistic Research Institute.

Page, S.E., 2003,"The Interplay of Differences," in C. Macal and D. Sallach, eds., *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, Argonne National Laboratory, Argonne, IL, pp. 311–327.

Putman, A.O., 1981, "Communities," in K.E. Davis, ed., *Advances in Descriptive Psychology, Vol. 1*, Greenwich, CT: JAI Press.

Putnam, A.O., 1982, unpublished manuscript.

Putman, A.O., 1983, "Organizations," in K.E. Davis and R.M. Bergner, eds., *Advances in Descriptive Psychology, Vol. III*, Greenwich, CT: JAI Press.

Sallach, D., 2003, Personal Comment, Network Studies Session, presented at the *Second Lake Arrowhead Conference on Human Complex Systems*, Lake Arrowhead, CA, March 22.

SOCIOECONOMIC SIMULATION TO ANALYZE ROOT-CAUSE MOTIVATIONS OF MIDDLE EASTERN TERRORISTS*

EDWARD MACKERROW,[1] Los Alamos National Laboratory

ABSTRACT

The rise of Islamic terrorism is the result of many complex and interrelated issues associated with globalization and cultural penetration of the West into predominantly Muslim regions. Many perceived causes of this social unrest have been well-stated by Islamic fundamentalists as obvious determinants of social conflict. Other, less obvious causal factors and long-term conditions might be better understood by applying methods from agent-based approaches. An agent-based simulation framework has been developed to examine the usefulness of this approach in understanding the "why" behind Islamic terrorism. It also aims to support decisions in pre- and post-conflict analyses. The design of the simulation framework allows a "modular" approach to various micro-social models used for agent interaction, cultural transmission, and social network dynamics, so the model is not "hard-coded" to particular social theories and allows for a research test-bed for various micro-social models applied to terrorism. The goal of this framework is to provide policy makers with decision support based on socioeconomic computer experiments: scenario generation representing known militant and terrorist groups, ethnic and culturally defined groups of agents, Western and Eastern regimes, and their interrelated political economies. Since this effort is designed to deliver a specific desktop computer social simulation framework to the analyst community, it requires very specific data and features relevant to real-world socioeconomic conditions. This paper discusses the trade-offs and challenges used to deliver the simulation, demonstrates the simulation, presents results, discusses the pros and cons of this research, discusses actionable decisions that could be supported by this simulation, and suggests future improvements.

---

* At the time of publication, the full paper for this presentation had not been received.

1 *Corresponding author address:* Edward MacKerrow, Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM 87545; e-mail: mackerrow@lanl.gov

THE EFFECT OF RECIPROCAL EXCHANGE ON THE RESILIENCE OF SOCIAL NETWORKS: AN EXAMPLE USING THE MESA VERDE PRE-HISPANIC PUEBLO CULTURE

R.G. REYNOLDS,* Wayne State University
Z. KOBTI, University of Windsor
T. KOHLER, Washington State University

ABSTRACT

The initial version of the model used in this study, Village 1.0, was implemented by Tim Kohler and a team of developers from Washington State University. Village 1.0 addressed environmental constraints, but not social ones. Recently, we used cultural algorithms as a framework for adding the social aspects. This paper extends our previous model by adding the ability for agents to perform symmetric reciprocal exchange. We developed a state model for agents' knowledge and the various responses of given agents based on this knowledge. Experiments showed that the system without reciprocity was not only the most resilient, but also the least complex. Because we allowed agents more opportunities to exchange resources, we produced more complex social structures and larger populations. Furthermore, allowing the agents to buffer their requests reduced the variability inherent in these larger systems but did not remove it. Introducing reciprocity to be triggered by both requestors and donors produced the largest number of successful donations. This work represents the synergy produced by using the information from two complementary situations within the network. Thus, the network has more information with which it can work.

Keywords: Cultural algorithm, multi-agent, network resilience, reciprocity, small world networks

1 INTRODUCTION

The initial version of the model used in this study, Village 1.0, was implemented by Tim Kohler and a team of developers from Washington State University. The simulation relives the settlement and farming practices of the Pueblo Indians of the Mesa Verde region of Southern Colorado; it is based on archeological, geological, tree-ring, and other data sources. The model presents an approach developed to help to understand the behavior of the region's inhabitants and the reasons that led to their eventual disappearance from the region, given modern archaeological knowledge of that area (Kohler, 2000). The reasons for the disappearance of this social system — one that had occupied the region since A.D. 600 — is one of the key issues in the archaeology of the Americas. Many theories have been proposed to solve this "mystery." Kohler's approach started with the most fundamental of these propositions — climatic change.

The Little Ice Age (LIA) has been invoked as one of the reasons for the depopulation of the Northern San Juan region in the thirteenth century. The term *LIA* refers to "an interval within

* *Corresponding author address:* Robert G. Reynolds, Wayne State University, Department of Computer Science, Detroit, MI 48202; e-mail: reynolds@wayne.edu.

the last millennium when alpine glaciers in many parts of the world advanced and when climatic conditions were significantly cooler than today" (Van West and Dean, 2000, p. 27). In his original simulation, Kohler tested the hypothesis that changes in high-frequency precipitation patterns would accompany the LIA, affecting the distribution of arable land in the region, and lead to depopulation in the region. Using tree-ring data to estimate precipitation in the region, he generated a Swarm-based, multi-agent simulation model for land use in the area between A.D. 900 and 1300.

Kohler's original simulation suggested that something other than these climatic factors affected the social history of the valley, since the model did not generate depopulation of the region using only climatic factors. The approach of the original model concentrated on farming practices and the distribution of land suitable for farming. The social interaction of individual households (agents) was purposely omitted to see how much of an explanation for the systems behavior could come directly from the rainfall and land distribution alone.

We extended Kohler's original model by establishing a framework using cultural algorithms, in which to embed a population of social agents, and to document the impact that their social interaction and cultural learning have on system performance. In particular, we are interested in the nature of the social networks produced by various combinations of social parameters. While the cultural algorithm provides a framework in which agents can learn to select various combinations of co-adapted parameters, our goal is to observe how varying these parameters causes certain aspects of the social networks to appear or dissipate.

Specifically, we focus on a single parameter — search or move radius. This parameter determines how far away from the original family a new family would settle, or if a family decides to leave its current location how far away it could look for a new location. We view it as a surrogate measure of propensity for social aggregation. Thus, when a new agent household is formed or relocated, its position is governed by the range over which it can search for another position. Links are then made between the household and the households of its relatives. We assume that these links represent communication links between agents, and that they are necessary to coordinate activities. We do not, however, consider the details of these activities at this time. We then observe the social structures that emerge from the interaction of the agents and their environment.

Our previous work looked at the trade-offs between the tendency for social aggregation on the one hand and the environmental constraints that tend to force dispersal on the other (Kobti, et al., 2003). In the simulation, the agents effectively tried to strike a "balance" or "find a centering" between the social forces for aggregation and the environmental forces in the region. The notion of finding a balance or center is a common theme in the study of southwestern Indian cultures (Thompson, 2002). We demonstrated that agents tend to form "small-world networks" based on kinship, and that these network structures can be differentially affected by environmental perturbations.

However, no explicit movement of information, or resources within the network, was considered. This paper extends our model of the agent household to produce a preliminary state model of each household agent. This state model is the basis upon which various models of reciprocal exchange among the individuals within the network are produced. This capability allows resources to move between individuals that have different states, e.g., between those with excess productivity and those in need. We provide three different frameworks for such flows:

1. Needy households ask their relatives for assistance.

2. Households with excess productivity poll nearby relatives to see if they need assistance.

3. Both items 2 and 3 within the same network are combined. It is suggested that the combined approach synergistically enhances the resiliency of the network.

Section 2 briefly gives an overview of the cultural algorithm framework in which the model is embedded. Section 3 describes how the social networks are modeled. Section 4 focuses on describing the state model of an agent's behavior. Section 5 presents the basic simulation system used here. Section 6 gives the results of our experiments; Section 7 gives our conclusions.

2 THE CULTURAL ALGORITHM FRAMEWORK

Holland (1975) developed a formal framework for any generic adaptive system. This adaptation framework involves a system that can alter its structure or behavior based on experience in some set of performance environments (Reynolds, 1979). Adaptability is the capacity to (1) function in an uncertain or unknown environment and (2) use information to evolve and learn (Conrad, 1983). Adaptation can take place at three different levels: population, individual, and component (Angeline, 1995).

Cultural algorithms consist of a social population and a belief space (Reynolds, 1979), as shown in Figure 1. Selected individuals from the population space contribute to the cultural knowledge by means of the acceptance function. The knowledge resides in the belief space where it is stored and manipulated on the basis of individual experiences and their successes or failures. In turn, the knowledge controls the evolution of the population by means of an influence function. A cultural algorithm thus provides a framework in which to accumulate and communicate knowledge so as to allow self-adaptation in an evolving model.

Five basic categories of knowledge are important in the belief space of a cultural evolution model: situational, normative, topographic, historical or temporal, and domain knowledge (Reynolds and Kobti, 2003a). All of these knowledge sources are present in our cultural model. For example, in our current model, we assume that agents have access to knowledge regarding the distribution of agricultural land (topographic knowledge), the distribution of rainfall as it occurs over time (history or temporal knowledge), and agricultural planting and harvesting techniques (domain knowledge). These three knowledge sources are fixed at this time.

We concentrate on the acquisition or learning of just two types of knowledge by agents at this point — situational and normative. *Situational knowledge* is a "snapshot" of the state of the world. The world can be viewed as a sequence of situations linked by social behaviors (Russell and Norvig, 1995). Examples of specific individual experiences correspond to a set of situational knowledge or relationships between individuals and their physical and social environments. *Normative knowledge*, on the other hand, describes how a rational agent should act in terms of ranges of acceptable behavior (Russell and Norvig, 1995). In other words, normative knowledge defines a standard or ideal that can be used to judge which behavior is desirable or undesirable (Valente and Breuker, 1994). For example, in our scenario, agents can learn in general which kin relations are most likely to respond to requests for resources and which are most likely to require aid. This knowledge can reduce the search effort of individuals and improve the flow of resources within the system.

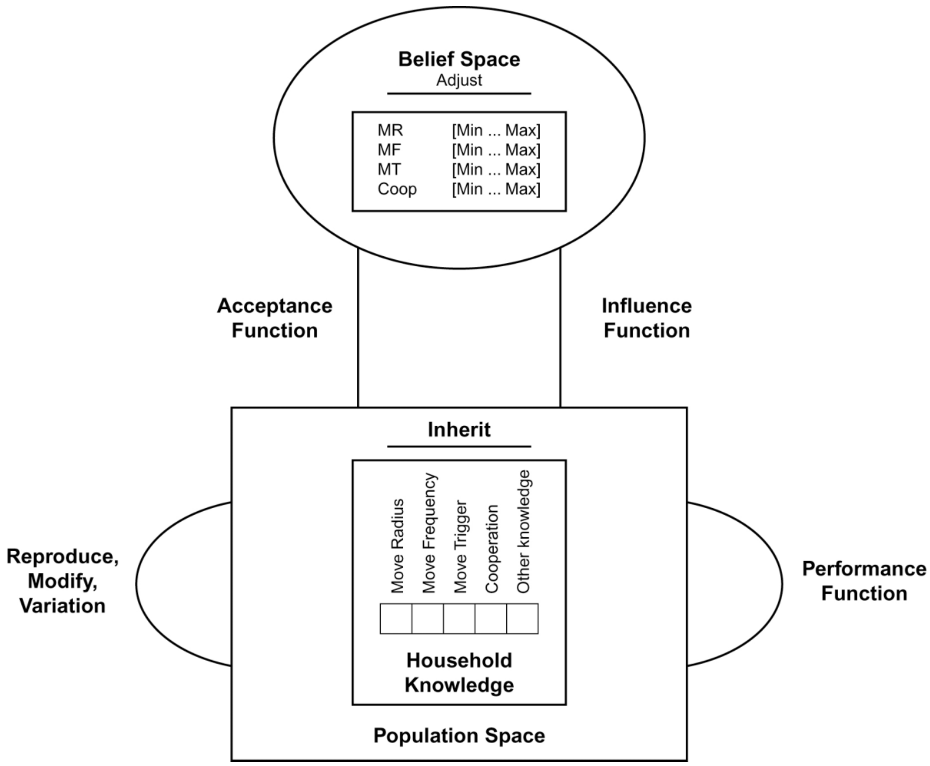

FIGURE 1 The cultural algorithm framework

3 MODELING THE SOCIAL NETWORKS

The first step in establishing social relationships between households is to develop the social network through which these relationships can be expressed. In this model, the basic relationship is based on kinship, and the strength of the relationship is affected by the distance between agents that share a kinship relationship. Each agent is a household that is composed of a husband, wife, and their children. Household members live together in the same location, share their agricultural production, and are affected by the same environmental conditions in the region. Children grow up, marry, and move out to form their own households; however, they maintain their connection to their parent households and siblings. Likewise, the parents maintain ties to their children. When one parent in a household dies, the other can form a new household with an available single agent. The initial structure of the social network supports the notions of parents, siblings, and grandparents. Members of a given household know who their parental households are on both sides of the family.

Table 1 shows the layout of the social network from the perspective of a household.

TABLE 1 Connected nodes identified by the kinship social network

Tag Identifier	Description
ParentHHTagA	Link to the parent from the mother's side
ParentHHTagB	Link to the parent from the father's side
ChildHHTag	One link to each child that moves away from this household and forms its own household
RelativeHHTag	One link to each extended family member

Each household (agent) is identified by a unique tag identifier in the system. This tag is basically the household's name that uniquely identifies it. In the initial step, agents are formed without any links (or relations). In subsequent annual steps, an agent updates its links based on a set of simple rules, as discussed below.

If a new child marries and moves away to form his or her own household, the parent household has a link to it. The child household in turn maintains a link to the parent household. This effort is duplicated because each child comes from a different household. So cumulatively, a child household has, at a minimum, links to two parent households: one link to the parent household from the wife's side and another link to the parent household from the husband's side. A child household can also remember its siblings. They are immediate relatives and hence the child has links to them. A household member can remarry if the spouse dies. In that case, the household keeps a link to the initial household, except now as a relative, and updates a new parent link to the new spouse's household. Kinship relation rules can be extended to other family members. Other concepts, such as friendships and neighbors, could also be modeled; these are necessarily based on kinship relations but could be based on a household's physical location and proximity. The overall social network is maintained dynamically, as it is updated every time step.

The set of kinship relations between agents cumulatively form a directed graph. At each time step, the current graph is stored in the form of an adjacency list. We can then plot the graph and examine the distribution of structural properties for the vertices (agents) and their edges (relations) as shown in Figure 2.

This kinship model can set the stage for the flow of materials between agents. In particular, we assume the custom of generalized reciprocity as practiced among kin. Such a practice is common in many societies (Flannery, et al., 1989). Thus, we investigate how such reciprocity shapes the population and social networks within the region and in turn how those networks are subsequently affected by the results of climatic change, such as a drought. Section 4 discusses how reciprocity is modeled here.

Social Link = {< Parent Household from Wife's side>
<Parent Household from Husband's side>
<Child Household>$_1$, <Child Household>$_2$, …, <Child Household>$_c$
<Relative Household>$_1$, <Relative Household>$_2$, …, <Relative Household>$_r$}
Graph$_{\text{Social network}}$ = {<Social Link>$_1$, <Social Link>$_2$, …, <Social Link>$_s$}

The format of the output file containing the adjacency list generated for one time step (each year) is:

Tag	Agent tag whose links we are describing
X, Y	Position of this agent in the world
ParentTagA, ParentTagB	Links to each parent's tag (–1 means no link)
ChildHHTag$_1$, …, ChildHHTag$_c$	Link to each child
RelativeHHTag$_1$, …, RelativeHHTag$_2$	Link to each relative

FIGURE 2 Social network structure as defined by the kinship relations

4 COOPERATION FRAMEWORK

Three strategies for symmetrical reciprocal aid were explored along with the case in which no goods are exchanged between agents. Table 2 lists the methods of aid used and briefly describes each. Reciprocal exchange is defined in terms of the exchange of maize between one agent related through kinship to another agent. Unlike trade between agents, the model of generalized reciprocity used here does not keep a record of debts owed. Modeled after the compassionate and human response of social beings, agents can ask their relatives for food in a time of need, while others donate their surplus to their relatives during prosperous times. In other words, the exchange is ritually activated by the requestor, the donor, or both. Each version is potentially reciprocal; the only difference is who provides the information that triggers the exchange.

TABLE 2 Description of the different cooperation methods at the kinship level

Cooperation Method	Description
0	No cooperation. No reciprocity of food between households.
1	When an agent requires food, it is allowed to select and request food from within its kinship network in order to survive.
2	When an agent has excess food, above a predetermined threshold, it is allowed to select an individual(s) from its kinship network and donate some of its excess.
3	Both methods 1 and 2 are enabled simultaneously.

4.1 Individual Agent Strategy

Each agent's basic strategy is to farm, harvest, and store enough maize to survive. In a time of drought or on lands of relatively low yields, the household may not have enough food to sustain it. Starvation in a noncooperative model directly triggers the death of the entire household. To avert this situation, the starving household can seek the help of another household known to it from its kin network. The selection method can be one of several types (random, fitness proportional, or round robin), but the model is flexible enough to allow other strategies to be implemented in the future.

If the selected agent is in the appropriate state, resources are exchanged. If the counterpart is not in the appropriate state, the requesting agent can ask another available agent. The number of attempts that can be made by an agent in a year is another model parameter. On the other hand, an agent may be in the appropriate state but unable to fulfill the entire request. In the case of partial fulfillment, the requesting agent can make another attempt to request the needed amount of food to survive, if desired.

Figure 3 shows the basic states of agents. An agent can have more food stored than needed and can therefore be a donor. An agent can have more than the minimum but less than what is needed to be a donor. Likewise, an agent can be taking in somewhat less than needed; the hungry state or the food level for an agent can be less than the starvation point. They can ask for or receive food from a donor. Three cases of exchange are based on the states of the individuals. In the first, an individual who is in a hungry or critical state can ask for resources from a kin-related agent who is in a donor state. In the second method, an individual agent who is in the donor state can provide resources to kin who are in the hungry or critical state. In the third case, reciprocal aid can be triggered either by agents who are hungry or by agents who are in the donor state.

In any case of symmetric reciprocal exchange, a limiting constraint is the physical distance between the two agents. For instance, an agent cannot request food from another agent

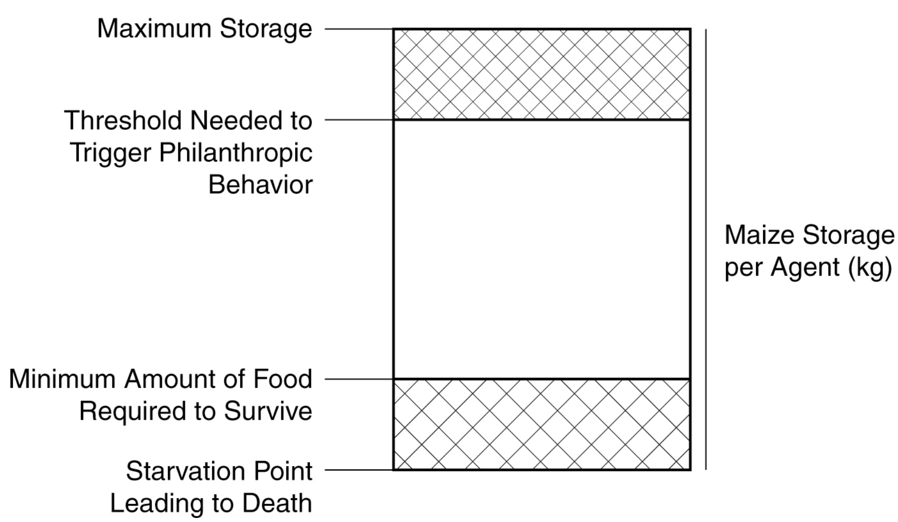

FIGURE 3 Actual maize amount in storage determines the state of an agent

who is beyond its search area. A similar case applies for an agent giving away its surplus. A radius distance parameter is implemented in each of the two cooperation strategies.

4.2 Selection Methods

A number of possible household selection methods are implemented in the system. An agent can be allowed to adopt any approach to select another agent to cooperate with it. Three basic selection methods are potentially present in the current model: random, roulette wheel or fitness proportional selection, and round-robin selection. Table 3 gives the three different choices. In this paper, the selection scheme is random. An agent can randomly select the agent it wants to cooperate with as long as this agent is within its cooperation radius (which is set to 20 here) and is in the appropriate state. If the requestor initiates cooperation, the requestor must be in a state of need and must ask for a donation from an agent in a state of excess. Likewise, if the donor initiates cooperation, the donor must be in a state of excess and must select an agent in a state of need. Figure 4 describes these and other state-base agent interactions supported within the model.

TABLE 3 List of the implemented selection methods

Method	Description
Random	An agent is randomly chosen from the kinship network within a given range.
Roulette	An agent is randomly selected (equal initial wheel portion) from the kinship network, then rewarded or penalized based on the size of the portion depending on whether the agents cooperated with or declined the request.
Round robin	Each agent that is kin with the agents and in the given cooperative radius is systematically given a turn.

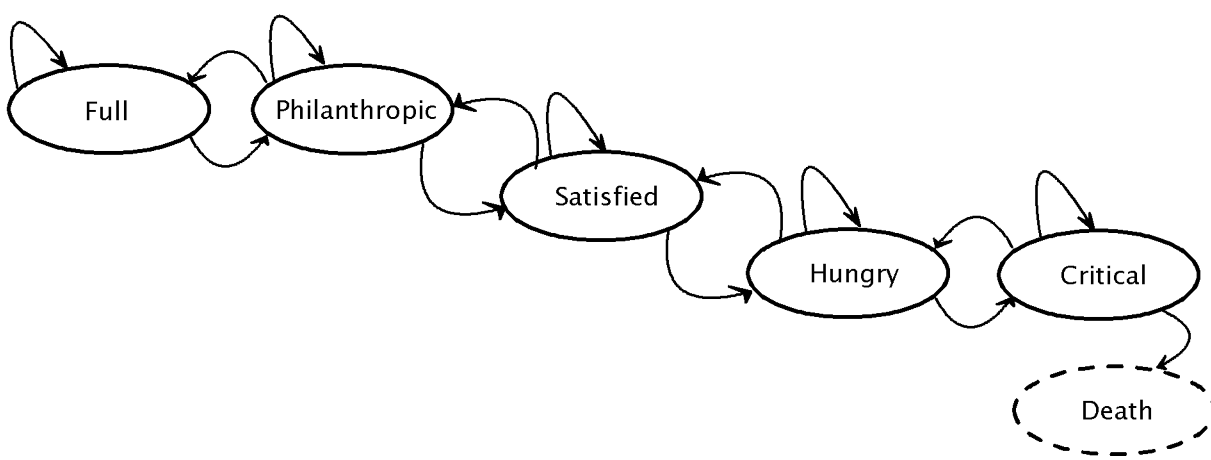

FIGURE 4 Agent state transition diagram (Note that additional states transitions are possible directly between F, P, S, H, and C states.)

4.3 A State-based Model of Agent Behavior

The approach used to implement our symmetric reciprocal strategy is based on an agent's set of states. The state of an agent can change dynamically each year on the basis of the food (maize) produced by the agent. An agent normally acquires its food from the harvest of its plots of land. The yields directly depend on the environmental factors affecting the farming area as well as on the farming practices (currently fixed in the model) and the number of labor force a household can spare.

How agents participate in the symmetric reciprocity process depends on their current state. For instance, agents doing well may either choose to donate some of their excess food or keep it for themselves in storage. At the other end of the spectrum, an agent straining under an unfortunate event of nature and bad crop yields is doomed to starve and die. With a cooperative social network, the agent has the opportunity to tap into the wealth of its nearby relatives and ask for food to survive. Essentially, the system is now formed by reactive agents in a dynamic environment.

Figure 4 gives the state transition diagram for each agent. Arcs between nodes represent state changes that can take place as the result of reciprocal aid. The arc that points from the state back to itself represents the fact that the agent can change its state or remain in that state from year to year. The states are described as follows:

- *Satisfied (S)* – An agent is in a satisfied state when it has sufficient food in storage to feed the entire family in the household.

- *Philanthropic (P)* – An agent becomes a philanthropist when it has a surplus of food in storage, defined in terms of stored maize in excess of a given threshold. For instance, an agent that stocks 90% or more of its storage capacity would be able to donate its food to others.

- *Hungry (H)* – A safety buffer zone is implemented as the level below which the agent should consider asking for additional food. When the agent has its last food ration, it enters a "hungry" state that triggers precautionary requests for food to avoid starvation.

- *Critical (C)* – An agent with insufficient or no food has no choice but to ask for food or face starvation and imminent death. If even after moving the agent does not receive its ration to feed the entire family, it will die.

- *Death (D)* – An agent is marked for immediate removal from the system. An interaction diagram shown in Figure 5 allows us to examine the cooperative relationships that are possible between two individual agents based on their respective states.

Agent B ⇨				
	Philanthropic	**Satisfied**	**Hungry**	**Critical**
Agent A ⇩				
Philanthropic, P	N/E	P donates S	P donates H	P donates C
Satisfied, S	N/E	N/E	N/E	N/E
Hungry, H	H requests from P	H requests from S	H requests from H	H requests from C
Critical, C	C requests from P	C requests from S	C requests from H	C requests from C

FIGURE 5 Matrix of possible interactions between agents A and B (agent A, depending on its current state, interacts with agent B differently according to its own current state. N/E = no exchange.)

5 EXPERIMENTAL FRAMEWORK

The cultural algorithm framework discussed in this paper has been implemented in the Swarm simulation environment (Langton, et al., 1995). This environment provides a framework for facilitating the development and experimentation with a large number of agents in a dynamic environment. Currently, the system is written entirely in Objective-C and uses the Swarm 2.1.1 libraries. The model is a graphic multi-agent simulation that allows us to probe individual agents in a dynamic environment. Agents reside in a cellular space that corresponds to the basic geographic region. At every model step in the simulation, the environmental state of each cell is updated by using the database of collected environmental data. Geographic data for the region were compiled from a number of sources and rainfall estimates were produced from tree-ring data (Van West, 1994). Agents can be observed either by means of probes that examine their internal properties or by accumulating output files for later examination. A dynamic viewer has been developed in Visual Basic to show the results of the run over time.

The simulation was run five times over the period from A.D. 900 to 1281 for each propensity/constraint combination. For each time step, the network is generated and stored into a file named “links<YEAR>.out.” This file is used to examine the properties of the emerging network and any characteristic found relevant to network resiliency. Visualization of the network is written in a separate package, either MatLab or Visual Basic (shown in Figure 6). Our program allows plotting of the graph and close examination of the distribution and densities of the links between agents. As the graph becomes more and more dense, especially as more agents develop more social links over time, we can visually filter out weak links and display only those edges attached to a node with a certain associated out-degree. In other words, we can identify the agent with the highest connectivity in the social network in terms of the number of associated links that it maintains.

In the experiments we conducted (see Section 6), the search radius that the agent uses when looking for a new plot to move into was varied from 5 to 30 pixels, and movement was possibly further constrained by existing land uses (e.g., two households cannot farm the same plot). For the different values of the move radius tested, we generated the network volume and the number of links over time. The network volume is the sum of all of the links associated with each individual household (agent) present. The idea is that with a larger move radius the agents must move farther from their previous location. Because relocation occurs frequently when a group fissions as a result of population growth, a larger move radius means that they relocate

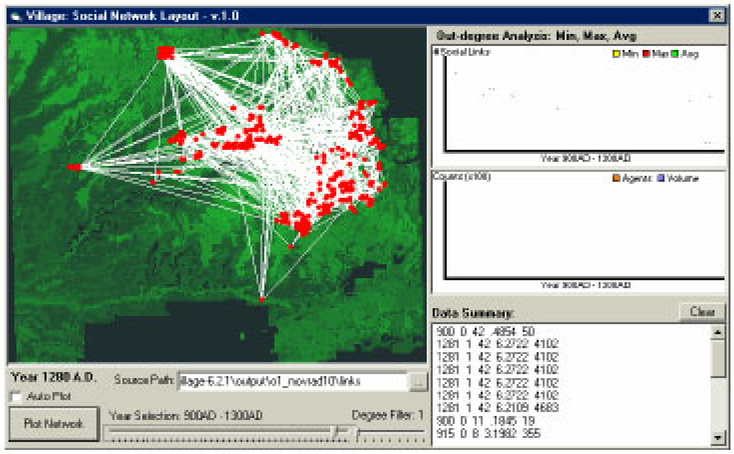

FIGURE 6 Using a search radius of 10 (MOV_RAD 10) the network is plotted for 1280 A.D. showing highly dense nodes

farther from a previously successful area. The success of the strategy of increasing search radius depends on the distribution of productive land in the region.

6 RESULTS

To test the various effects of cooperation under changing environmental conditions built into the model, fixed rules for reciprocity are implemented to guide the strategic behavior of the agent. In turn, the outcome of such cooperation techniques is examined to evaluate its effect on the agents in terms of their survival rates. Of specific interest are known drought years that have devastating effects on crop yields unknown to the agents. Experiments are set up to examine the agents' response over time in terms of (1) the number of requests and fulfilled requests and (2) the number of donations available and accepted. With each of these two measures, we tracked the amounts of food that changed hands each year. The social network is examined in terms of its hub node densities and any observable collapse of major nodes. The number of surviving agents is recorded yearly to examine the effect of environmental variability on the population counts in the presence (and absence) of cooperation in its different forms. Three possible cooperative strategies were examined: the first involves an agent requesting food when it reaches its critical state, hungry state, or either state; the second involves an agent donating food when it reaches a full state, philanthropic state, or either state; and the third combines the effect of the first two cooperation strategies.

To make comparative observations, the selection of the reciprocal aid strategy required symmetry in the state. For instance, when we tested an agent who could ask for food only when it reached its critical state, we also tested an agent that would donate only when it reached its full state. Similarly, we tested an agent who could reach the hungry or critical state along with agents that could reach the philanthropic or full states. In addition, we fixed the target agent state to F, P, S and C, H, S, reflecting that an agent could donate to another agent who is in its satisfied, hungry, or critical state, and an agent could request from another who is in a full, philanthropic, or satisfied state. Of course, many other combinations could be tested, but these selected combinations reflect the scenarios most likely, in our judgment, to be encountered in human social systems (Table 4).

An important aspect of the early cooperative strategies examined is to allow the agents to maintain a safety buffer in their food storage levels. For instance, an agent does not need to wait until it reaches the critical state to ask for food. If an agent is in a critical state and does not receive sufficient food to survive from another agent, the former agent will terminate. A safety buffer is therefore established based on a measure that an agent computes to determine if it is about to eat its last food ration. This measure depends primarily on the makeup of the household members and their food requirements. Currently, the system defines the hungry state as that time when an agent has only one food ration left. After it eats, it falls into the critical state where it has no food and must seek food from others. One set of experiments tested the scenarios with the presence of such buffering in anticipation of allowing the agent to foresee and bail itself out of starvation, while another set tested the scenarios without the presence of buffering, where agents procrastinate and gamble for their life (Tables 4 and 5, respectively).

A fourth set of experiments focused on the control scenario when there is no reciprocity (Table 6). These experiments allow us to establish a base level against which the impact of various reciprocity scenarios can be assessed.

6.1 The Impact of Reciprocity on Network Resiliency

We performed runs for each combination of states for three different move radii; however, in this section, we give only the results for move radius 30. Figure 7 consists of three figures that summarize five runs of the simulation in which no cooperation occurs between the agents. The agents produce kinship networks over time, but the networks are not used explicitly to provide symmetric reciprocal aid to kin. Figure 7a gives the maximum, minimum, and average number of links between the agents. The average number of links is six, which is very characteristic of a small-world network. The average and minimum number of links are not affected by change in precipitation over time, but the maximum number of links per node is. The nodes with large numbers of links are the hub nodes, which provide the network connectivity required for an individual agent to search the network. Figures 7b and 7c give the volume of the social network over time, while Figure 7c gives the number of households. The household numbers decrease later in the period, as drought conditions begin to emerge.

When reciprocity is allowed (see Figure 8), the numbers of social agents, link volume, and the maximum number of links increase. Figures 8a–8c give the results when agents in critical need (Figure 8c) request from agents who are full. Figures 8d–8f give the results when buffering is allowed in that agents who are hungry, but not yet in critical need, can ask as well. This represents a buffering situation in which the system is less constrained to produce complete help for each request.

The addition of reciprocity increases the complexity of the system by producing larger, more complex networks, but these networks are much more variable in response to the same environmental input. For example, the maximum number of links shown in Figure 8a varies markedly throughout the run, but in particular near the end, as do social network volume and the total number of households. It is interesting that while both the buffered and unbuffered cases

TABLE 4 Experimental setup with buffering[a]

Coop Type	Move Radius	Donor	Requestor	Donate to	Request from
1	10	F, P	C, H	C, H, S	F, P, S
2	10	F, P	C, H	C, H, S	F, P, S
3	10	F, P	C, H	C, H, S	F, P, S
1	20	F, P	C, H	C, H, S	F, P, S
2	20	F, P	C, H	C, H, S	F, P, S
3	20	F, P	C, H	C, H, S	F, P, S
1	30	F, P	C, H	C, H, S	F, P, S

[a] F = full; P = philanthropist; S = satisfied; H = hungry; and C = critical.

TABLE 5 Experimental setup without buffering

Coop Type	Move Radius	Donor	Requestor	Donate to	Request from
1	10	F	C	C, H, S	F, P, S
2	10	F	C	C, H, S	F, P, S
3	10	F	C	C, H, S	F, P, S
1	20	F	C	C, H, S	F, P, S
2	20	F	C	C, H, S	F, P, S
3	20	F	C	C, H, S	F, P, S
1	30	F	C	C, H, S	F, P, S
2	30	F	C	C, H, S	F, P, S
3	30	F	C	C, H, S	F, P, S

[a] F = full; P = philanthropist; S = satisfied; H = hungry; and C = critical.

TABLE 6 Experimental baseline control: no cooperation

Coop Type	Move Radius	Donor	Requestor	Donate to	Request from
0	10	–	–	–	–
0	20	–	–	–	–
0	30	–	–	–	–

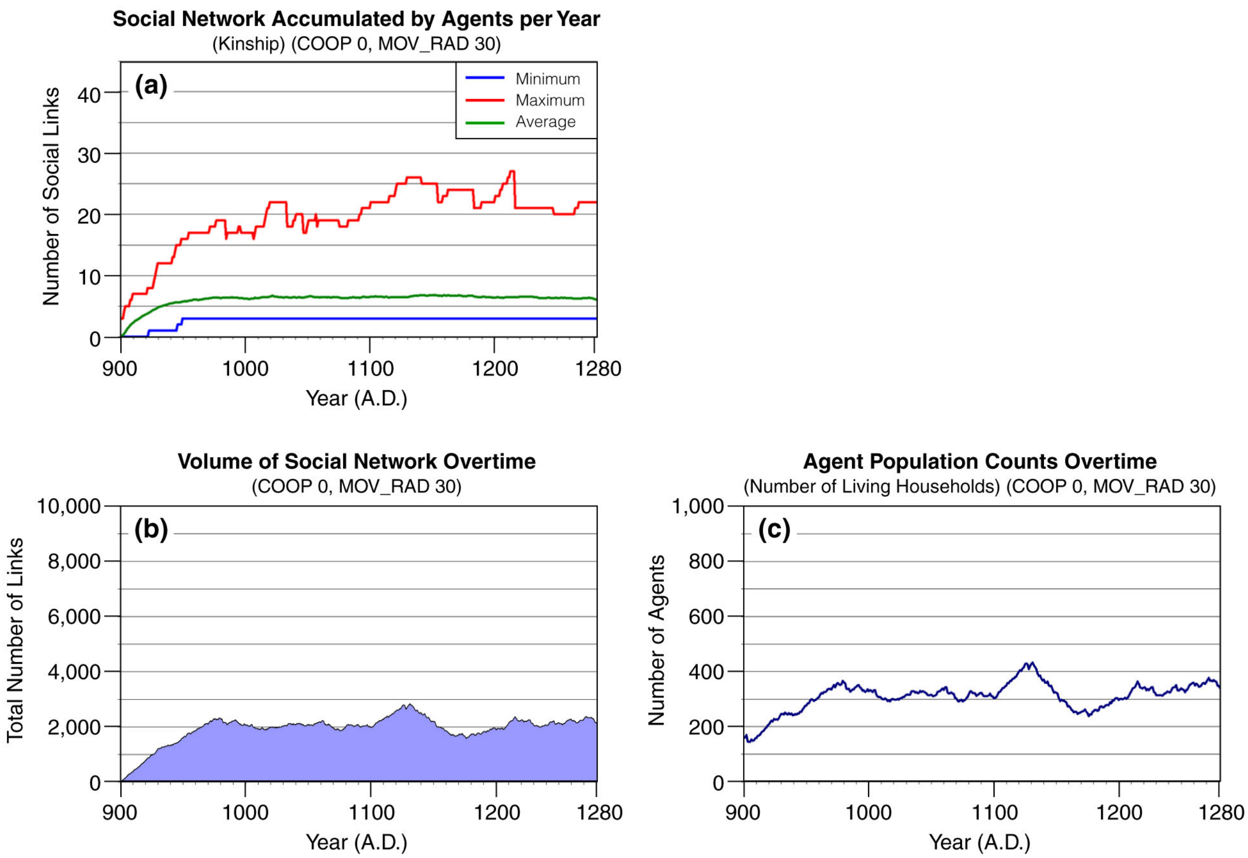

FIGURE 7 No reciprocity across the network (Coop 0)

of reciprocity are more complex than the no exchange case, the buffered version has more uniform variability and slightly larger peak values for household number, link volume, and maximum link number. Thus, adding in the buffering capability dampens the fluctuations while maintaining an increased complexity over the no exchange case.

Figures 9a through 9f give the case in which reciprocity is triggered by the donor and not the requestor. As before, Figures 9a–9c concern the unbuffered case where only agents in a critical state ask for aid from their kin, while Figures 9d–9f concern the buffered case. The complexity of the emergent system is not as great as the networks produced when the receiver initiates the exchange. Also, the complexity of the buffered case and its variability are reduced compared with those of the unbuffered case. Thus, allowing the donor to initiate the exchange identifies where in the network the excess is but may not allow sufficient time to get to those in critical need.

Several things are notable:

1. Buffering reduces variability in household number, maximum link size, and social volume in both the donor- and the requestor-initiated exchange.

2. The requestor-initiated exchange tends to produce more complex networks than the donor-initiated exchange.

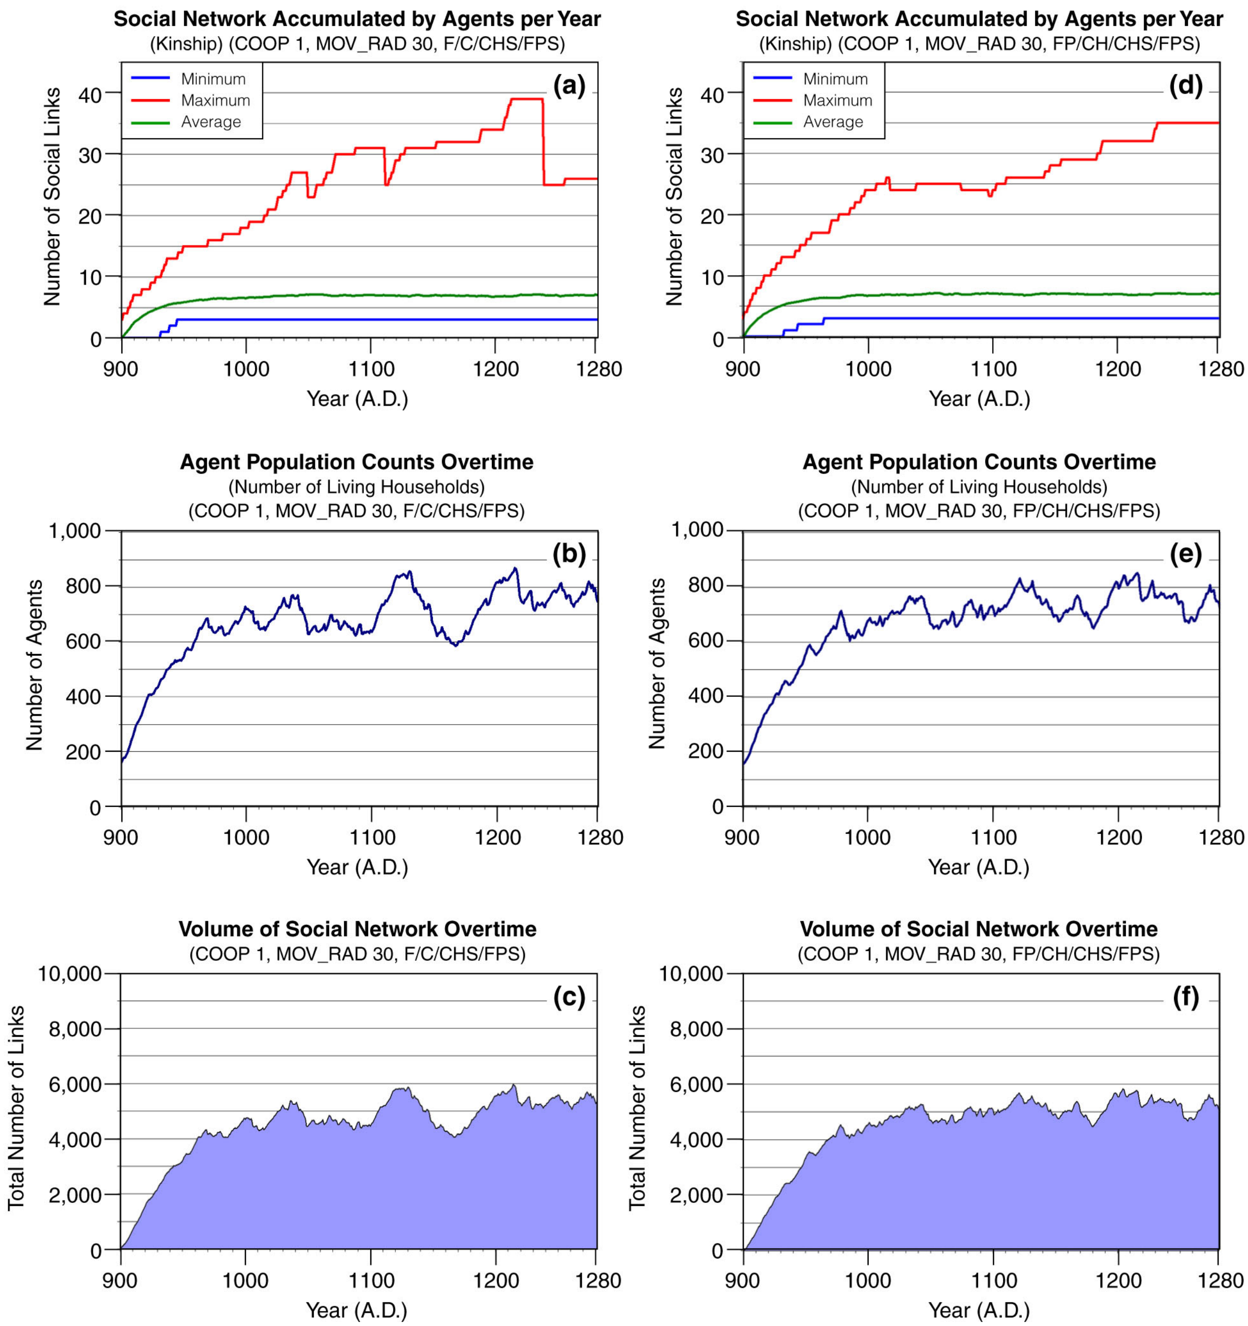

FIGURE 8 Coop 1 Network volume and agent counts over time

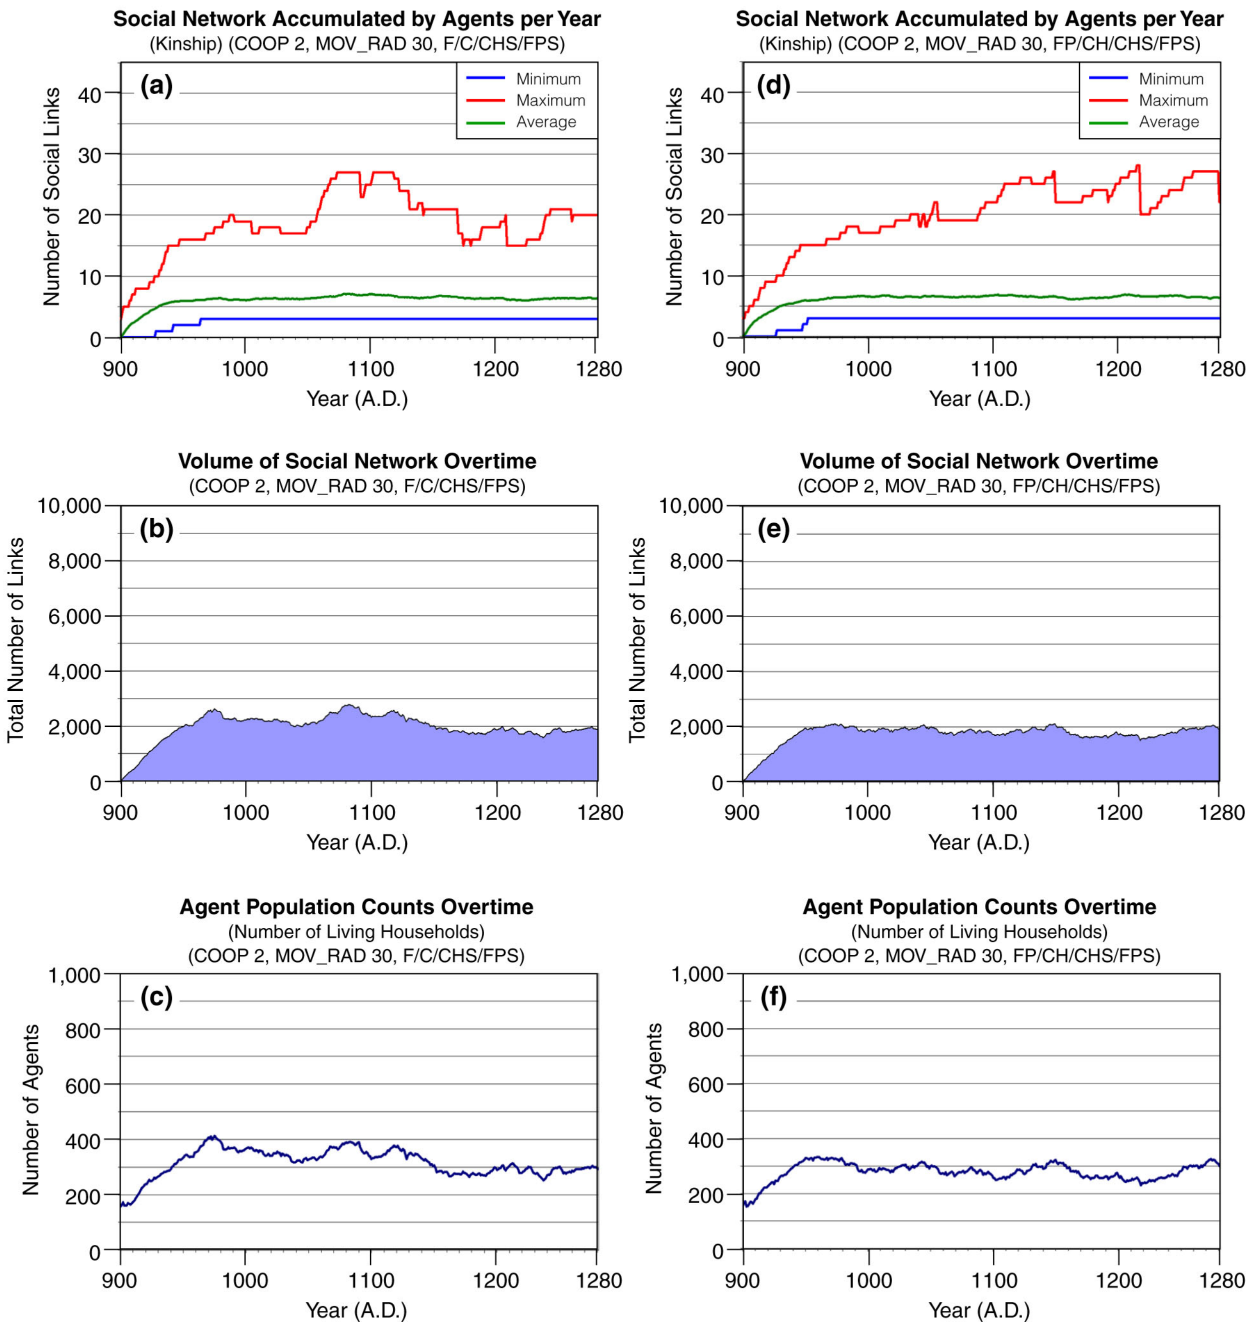

FIGURE 9 Coop 2 Network volume and agent counts over time

3. However, the donor-initiated exchange produces networks that are less sensitive to environmental perturbation than the requestor-based approaches.

There is a potential for the two approaches that initiate the exchange to be complementary. The requestor approaches produced larger populations in both the critical and buffered case, but at the cost of increased variability in the wake of environmental change.

Figure 10 shows the synergy produced by combining both the requestor and the donor modes of reciprocal exchange. In both the buffered and unbuffered cases, the system that is produced is more complicated and larger than the networks produced by either one alone. The unbuffered system, however, exhibits large variability in social volume, household numbers, and maximum links. Adding the buffering capability to the system generates a slight decrease in complexity and produces a more stable system.

An interesting result is that by adding reciprocity and buffering to the system, we effectively produce more complicated social structures with larger population sizes, network volumes, and hub complexity. The downside of this result is that these systems exhibit the most variation of the various configurations tested. The presence of buffering in the combined case makes the variation more predictable, but it is still substantial. Reduced resilience may be a price that one pays for social complexity.

Figures 11 and 12 demonstrate why the combined scheme produces the largest and most complex social networks. Figure 11 gives the number of successful donations produced by the four different exchange combinations for each of the three different search radii tested. The figures show that as the radius increases, the volume of the donations decreases. Within any given radius, the buffered solutions outperformed the unbuffered ones. Figure 12 shows that the number of requests for each exchange configuration decreases as the move radius increases. The fact that there is little difference between the number of requests in each configuration reflects the fact that the same environmental perturbations are being presented in each case. The fact that the configurations differ in terms of the number of donations made, however, reflects the ability of the system to opportunistically deal with shortfalls.

7 CONCLUSIONS AND FUTURE WORK

This paper extends our previous model by adding the ability for agents to perform symmetric reciprocal exchange. We developed a state model for agents' knowledge and, given agents' different responses based on this knowledge, we arrived at some general findings:

1. The system without reciprocity was the most resilient, but least complex.

2. As we allowed agents more opportunities to exchange resources, we produced more complex social structures and larger populations.

3. Allowing the agents to buffer their requests reduced the variability inherent in these larger systems but did not remove it.

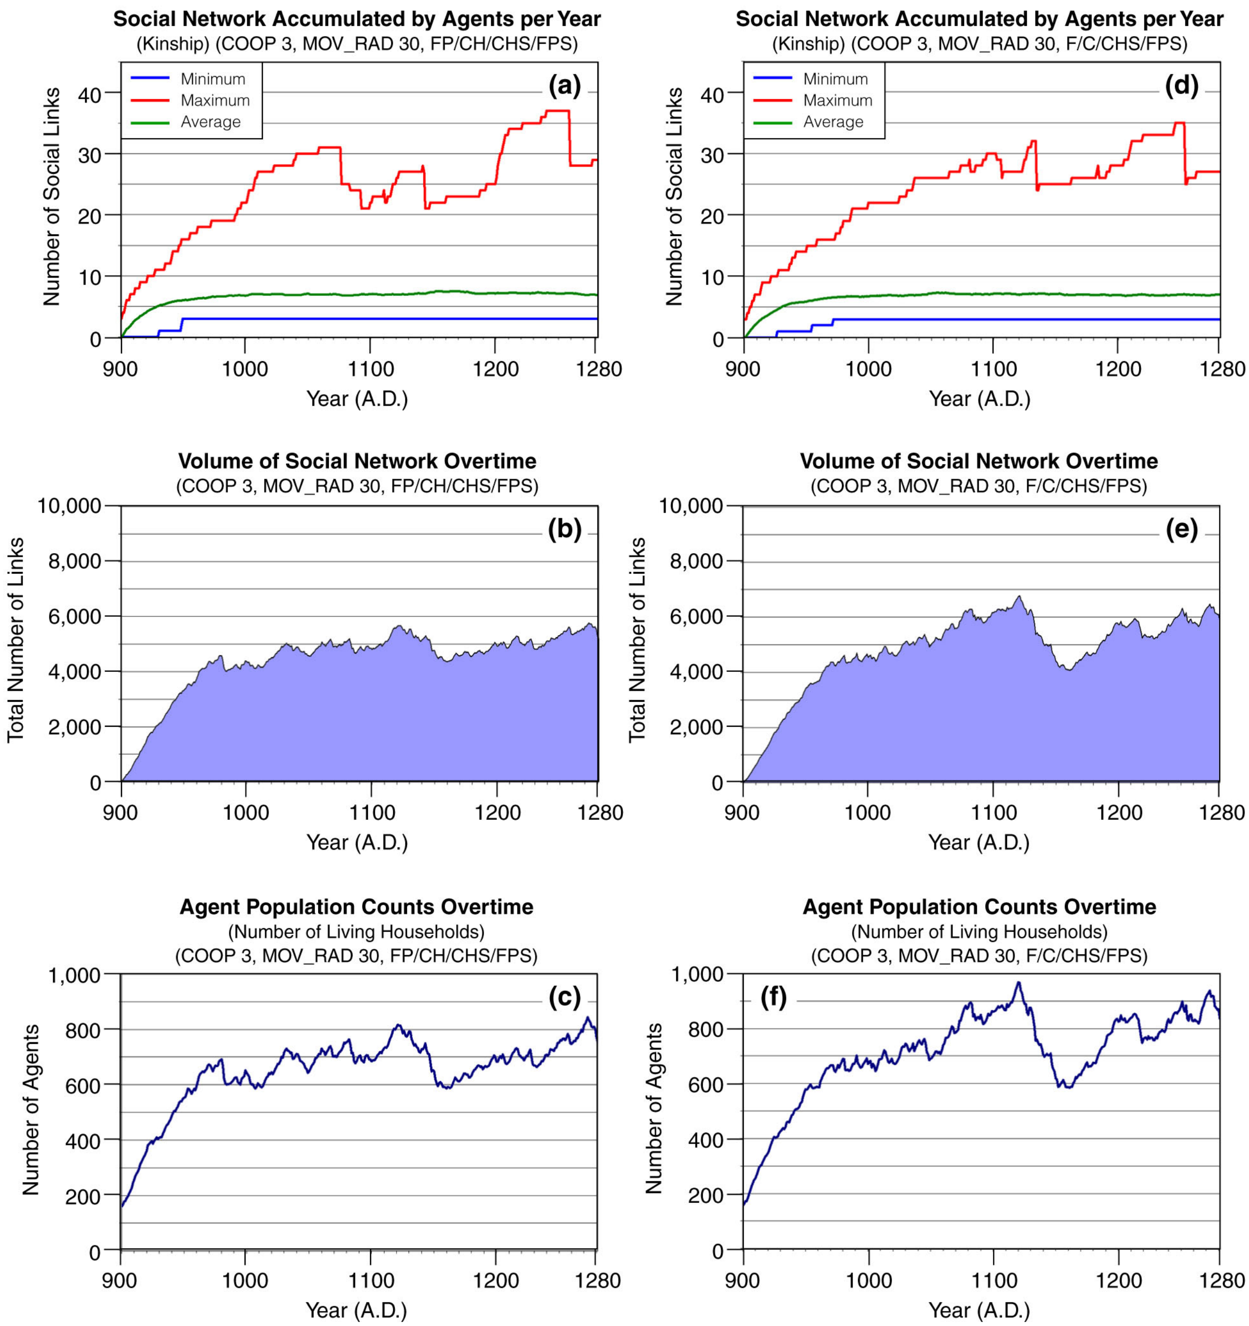

FIGURE 10 Coop 3 Network volume and agent counts over time

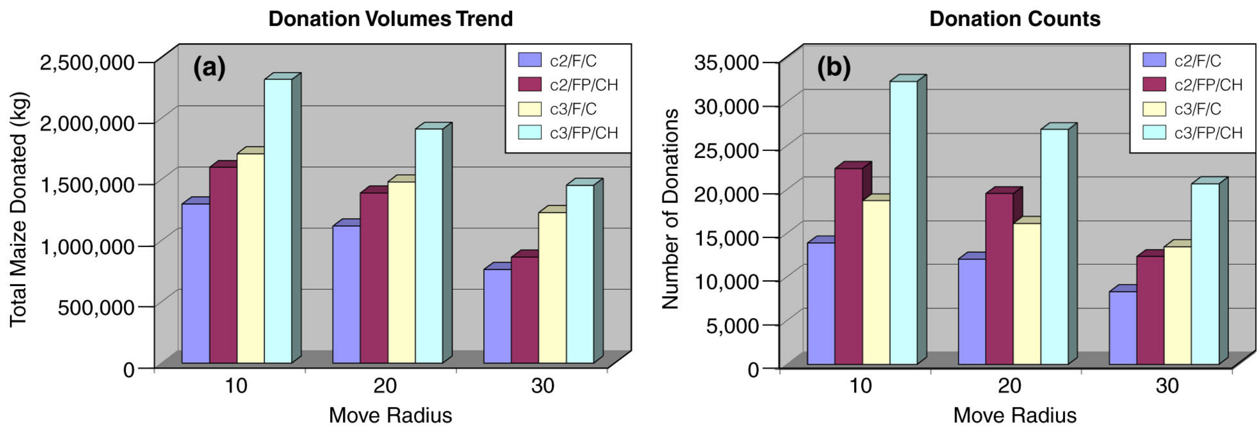

FIGURE 11 Donation trends: as the radius increases, the volume of the donations decreases

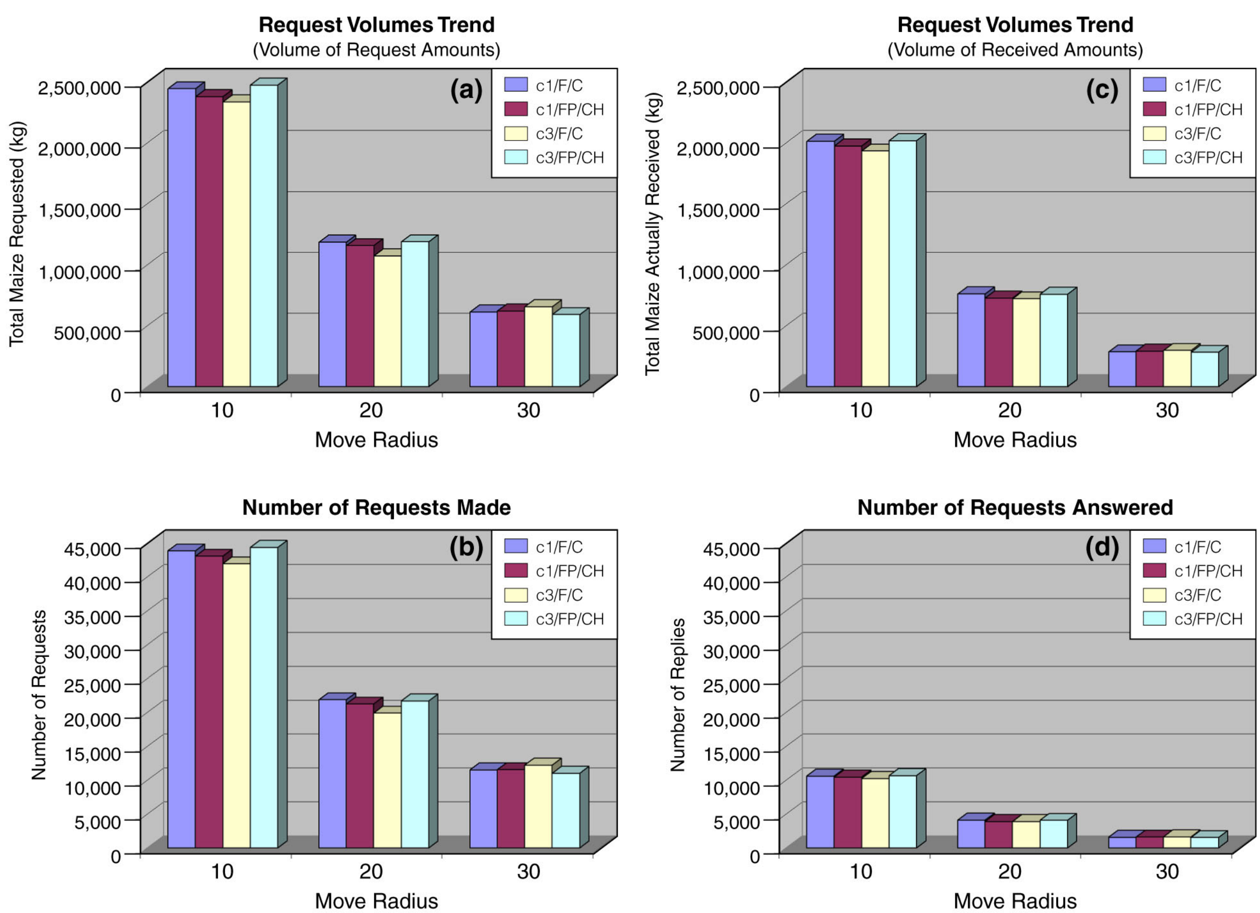

FIGURE 12 Donation trends: the number of requests for each exchange configuration decreases as the move radius increases

4. Allowing reciprocity to be triggered by both requestors and donors produced the largest number of successful donations. This represents the synergy produced by using the information from two complementary situations within the network. Thus, the network has more information with which it can work.

In future work, we plan to extend the state model so that agents can include other activities, including trade and warfare.

8 ACKNOWLEDGMENT

This work was funded by National Science Foundation grant #BCS-0119981.

9 REFERENCES

Angeline, P.A., 1995, "Adaptive and Self-Adaptive Evolutionary Computation," in M. Palaniswami, et al., eds., *Communication Intelligence*, New York City, NY: IEEE Press, pp. 152–163.

Chung, C.-J., 1997, *Knowledge-based Approaches to Self-adaptation in Cultural Algorithms*, Ph.D. dissertation, Wayne State University.

Conrad, M., 1983, *Adaptability: The Significance of Variability from Molecule to Ecosystem*, Plenum Press, New York.

Flannery, K.J., J. Marcus, and R. Reynolds, 1989, *The Flocks of the Wamani*, New York City, NY: Academic Press.

Holland, J.H., 1975, *Adaptation in Natural and Artificial Systems*, First edition, Ann Arbor: The University of Michigan Press, 1975; First MIT Press edition, 1992.

Kobti, Z., R. Reynolds, and T. Kohler, 2003, "A Multi-agent Simulation Using Cultural Algorithms: The Effect of Culture on the Resilience of Social Systems," *Proceedings of the IEEE Congress on Evolutionary Computation*, December 5–12, Canberra, Australia (to be published).

Kohler, T., 2000, "The Final 400 Years of Pre-Hispanic Agricultural Society in the Mesa Verde Region," *Kiva* **66**:191–264.

Langton, C., et al., 1995, "The Swarm Simulation System A Tool for Studying Complex Systems," Santa Fe Institute.

Reynolds, R.G., 1979, *An Adaptive Computer Model of the Evolution of Agriculture for Huntergatherers in the Valley of Oaxaca, Mexico*, Ph.D. dissertation, University of Michigan, Ann Arbor.

Reynolds, R.G., 1999, "An Overview of Cultural Algorithms," in *Advances in Evolutionary Computation*, McGraw Hill Press.

Reynolds, R., and Z. Kobti, 2003a, "A Multi-agent Simulation Using Cultural Algorithms: The Effect of Culture on the Resilience of Social Systems," *Santa Fe Institute Workshop on Human Environmental Interaction*, May 16–18.

Reynolds, R., and Z. Kobti, 2003b, "The Effect of Environmental Variability on the Resilience of Social Networks: An Example Using the Mesa Verde Pueblo Culture," presented at the *68th Annual Meeting of the Society for American Archeology*, April 9–13.

Reynolds, R., T. Kohler, and Z. Kobti, 2003, "Robustness in Coupled Human/Natural Systems in the Northern Prehispanic Southwest," presented at workshop on *Robustness of Coupled Natural and Human Systems*, Santa Fe Institute, May.

Reynolds, R., J. Brewster, and R. Iacoban, 2003, "Cultural Swarms: The Impact of Culture on Social Interaction and Problem Solving," *Proceedings of IEEE Swarm Intelligence Symposium*, Indianapolis, IN, April 24.

Russell, S., and P. Norvig, 1995, *Artificial Intelligence — A Modern Approach*, Englewood Cliffs, NJ: Prentice Hall.

Thompson, I., 2002, "Native American Perspectives on Sand Canyon Pueblo and Other Ancestral Sites," in *Seeking the Center Place: Archaeology and Ancient Communities in the Mesa Verde Region*, M. Varien and R.H. Wilshusen, eds., University of Utah Press.

Valente, A., and J. Breuker, 1994, "A Commonsense Formalization of Normative Systems," *Proceedings of the ECAI'94 Workshop on Artificial Normative Reasoning*.

Van West, C., 1994, *Modeling Prehistoric Agricultural Productivity in Southwestern Colorado: A GIS Approach*, Department of Anthropology Reports of Investigations 67, Pullman, WA: Washington State University.

Van West, C., and J. Dean, 2000, "Environmental Characteristics of the A.D. 900–1300 Period in the Central Mesa Verde Region," *Kiva* **66**:19-44.

White, D., and M. Houseman, 2002, "The Navigability of Strong Ties: Small Worlds, Tie Strength, and Network Topology," *Complexity* **8**(1):72–81, Sept./Oct.

DISCUSSION:

SOCIAL AND CULTURAL DYNAMICS

(Saturday October 4, 2003, 9:45 a.m., Session 1)

Chair and Discussant: *Claudio Cioffi-Revilla, George Mason University*

Including Relationships, Motivation, and Actor-observer-critic in the Human Community Framework

H.J. Jeffrey: The interesting thing to look at here is the subtitle, not the formal title, which is always sort of boring. It's a semi-tongue-in-cheek reminder that human beings, the kind of agents that many of us here are particularly interested in modeling, aren't primarily biological organisms. Obviously we have biological characteristics, but people, all of us at each moment, exist *within* a fabric of a human society, and these biological things like eating and so forth take place as ways of engaging in human, specifically human, social practices.

[Presentation]

Claudio Cioffi-Revilla: Thank you. Any questions on this presentation?

Konstantinos Alexandridis: I understand what you're saying on an individual level, but sometimes as cultures and as groups we have also a collective perception. And when we do that, we face options like the choice sets. It appears to have different distribution among groups. Or we face different choice sets.

So if we go about to assess those choice sets, it's a very data-intensive process, and I don't know if it's feasible to go through and interview an entire population in order to reveal those things. There are some economic methods of addressing our revealed preferences or some kind of contingent valuation methods, but when we're talking with the trade-offs between the individual and the collective behavior or choices, we face problems, like what's going on with the freedom of choice where, as we reduce our choice sets and things like that.

H.J. Jeffrey: The way that I see that one is, it's a question at what level you want to model things at. Clearly, an approach to simulation that says, "Well, go out and build a model of each individual human beings in the culture" — and there's, what, 250 million Americans or something like that? Or 12 million, say 12 million Swedes. Clearly, if you have to build a 12-million, many-element data set, it's infeasible. So it's a question of what level of granularity you want to simulate at.

My notion is that you pick the level that is the usual compromise between what you need and what you can afford to do, and you model it statistically. You make guesses using the data sources that give you indications of how many people tend to make which choice, tend to have which values, tend to have which priorities, and model it using that. In other words, you say, well, so many people will make this choice, so many people will make that choice. You don't drive that by interviewing all these people. You don't go out and build it bottom up from individual choices, but you use exactly the tools you'd expect. So in other words, take in the data that you can and say, "Here are the choices statistically."

For example, you might have, well, on the abortion issue, in a given community X% of the people pick this, and Y% of the people pick this, and Z% of the people pick this other. And that's what you use directly in your simulation; not trying to interview all the, as I say, 12 million Swedes or 250 million Americans or whatever it happens to be. So you're driven by the data you can use, in other words.

Does that answer the question?

Alexandridis: Well, sometimes you have to build a model that moves across scales and how you merge those scales together is a challenge sometimes, and that's what I was trying to point out.

Jeffrey: Yes. The usual situation is that you build what you can, and then you see how it matches the data.

Socioeconomic Simulation to Analyze Root-cause Motivations of Middle Eastern Terrorists

Claudio Cioffi-Revilla: So now this is followed by Ed MacKerrow from Los Alamos on his terrorist network agent-based model.

Edward MacKerrow: Thanks, Claudio,

[Presentation]

The Effect of Reciprocal Exchange on the Resilience of Social Networks: An Example Using the Mesa Verde Region Pre-Hispanic Pueblo Culture

Claudio Cioffi-Revilla: Okay. Bob Reynolds will present on behalf of his whole team a model of the Mesa Verde Southwestern Region, near the Four Corners.

Robert G. Reynolds: This is truly as multimedia as time will permit. Basically, the study area is here in the Four Corners: Utah, Arizona, Colorado and New Mexico. The group under study is the Anisazi, the ancient Pueblo peoples….

[Presentation]

Reynolds: The simulation is going to crash. Here goes it. Wham! Okay, and it takes a little time for the system to kind of recover from it, but all of a sudden these are going to go away.

Unidentified Speaker: What causes the crash?

R.G. Reynolds: The crash has to do with the fact that the distribution of rainfall has been reduced, and the impact of the distribution is felt differentially over the landscape. It turns out that those individuals who are hubs are hubs because they have been productive, and they have generated lots of links to other people in the network. And it turns out that the rainfall differentially affects those areas where the hubs are.

MacKerrow: Bob, I have a question having to do with the bark beetle problem in the Rocky Mountains right now in the Southwest. Do you think that effect is the same as the drought, or is it different?

Reynolds: That's an interesting question. I think that the bark beetle problem relates to the drought, because basically what happens is that when you have sufficient water, it means that there's more sap in the trees, and the sap keeps the bark beetles out. And so when you have less water, basically there's less sap in the trees and more room for the bark beetles to get in.

MacKerrow: Yes, but what I'm asking is, in your model, do you think the effects of something like a bark beetle would play out differently in a drought? You know, you could take a forest fire percolation model, that will topographically have the same effect? Do you see where I'm going?

Reynolds: Yes. The interesting thing is that a drought is a regional thing. Bark beetles and forest fires can be very localized phenomena. But in fact it's very possible if they're localized relative to a habitation center, the impact of the drought could be amplified. And so, in other words, the drought by itself has an impact regionwide. However, if you throw in other factors that are going to effectively be exacerbated by the drought and allow them to provide localized perturbations, then you're in big trouble. Excellent point.

Cioffi-Revilla: One more?

Jesse Voss: My question has to do with how you operationalize the lineage controls. I was interested in how you expressed the kin relations. To what degree of granularity to you have the clan lineage household connections, or is it more limited? I mean, how is that, how far do you break it down?

Reynolds: Well, we have it basically down to the lowest level. We keep all of that information at this point. In fact, that's one of the issues in making the model initially run as slow as it did. We've expedited it by coming up with data structures that make searching these networks an easier thing to do.

In other words, these hub nodes are a lot like routers on the Internet. In fact, one of the big issues is getting these routers. You have rules that basically govern how messages come through the Internet. One of the key issues is making sure that these rules can be easily checked to route as many mail messages as possible, especially when these rules are changing all the time. And in fact the situation here is very similar. And so we have to be checking — these connections are rules that determine where things flow, and we have to check these things all the time. And so the data structure problem is a key to make this. It ran here in front of your very eyes, but we have run it for, initially for days. So that's a key question.

Unidentified Speaker: I'm just wondering, have you compared this to a simple pooling of resources? And could part of the volatility be due to the fact that by sharing you're increasing survivability, but when you cross a threshold, the problem becomes too severe, and then you kind of revert to a more normal level.

So I'm just wondering, have you compared this to a case where all the resources in the community were pooled and what the distinction is.

Reynolds: Well, actually, the way we have it set up with moved radius, if we just set that equal to, let's say, 10, that's basically what we're doing. We're pooling resources, but within kin. Now we haven't pooled resources in the community outside of kin-based relations. But in fact that would be one of the next steps. So we have other things that we can add in, of course: trade, exchange, and that sort of thing. But at this point we wanted to see as far as we could get with kin.

But you're right. I think certainly one of the things by just pooling in localized areas, well, it wasn't as successful as having a larger range. That's because you're only sharing with your kin. If, as you say, you can share with non-kin, then you certainly will have more resources and you don't have to look as far in your network. We haven't done that yet, but that's the next layer. But each of these layers in and of itself is, it's easy to describe, but it's not necessarily easy to implement.

Cioffi-Revilla: Thank you, Bob.

The organizers asked me to both chair and discuss this session. So I'm not going to provide an extensive discussion of these papers, because we're really out of time, but I just want to raise a few points.

On Jeffrey's paper there are, I thought, some really amazing parallels or quasi-parallels between the approach you take and what has been done in political analysis, in political science, especially in comparative politics and international relations. There we have a subfield called "event data analysis" where the idea is to code in a systematic way the longitudinal patterns of behavior of what countries do to each other, for example, or what different groups do domestically.

There's extensive literature on this, and today this is being carried out through computational tools. The leader in this field is Phil Shrout, of the University of Kansas, who has been using Holland classifiers to machine code, for example, the UPI and AP chronicle of what countries do to each other internationally, because human coding of that data has become impossible because of the sheer flow there.

Something I would point to is that there's a great deal of need for a systematic notation that could be at least as efficient as the notational system that exists in music, ballet, and other areas. The notational music for ballet is very interesting, and it's very tough to learn. It's not a weekend job. It takes quite a bit of work. But it's very efficient, and so much so that you can take people that have never danced together in a formal piece, you show them this graphic script, and they know how to execute it exactly as was intended by their choreographer. A similar thing, of course, happens in music, so for describing behavioral events in social science we've never had this sort of notational system. And I think that that would be a good thing to have, because it would also facilitate a comparison between empirical records of behavior and synthetic ones generated by agent-based modeling.

Now I'll comment on MacKerrow's model of the terrorist network. I think there are some great ideas in this project. I would just point out a few of them. First of all, the dry-grass notion related to the concept of potential energy. I would say that rather than apologize for its introduction from physics, we should be very happy of that, because social scientists in many different areas of social science have had an intuitive idea about this in the past. But there's been nothing anywhere near as elegant as a theory of the potential in physics.

For example, international relations scholars in political science often define a crisis as a situation in which the probability, not the certainty, but the probability of war is significantly increased. Well, in terms of a bifurcation, this means entering a bifurcation set and creating a metastable state, which is a modification of the potential function in the dynamics. So the intuitive idea is there, it just hasn't been formalized. And some help along those lines would certainly be welcome.

There are other parts of this that I think are very interesting and really meritorious, but I would just like to point out that in terms of Atran's article in *Science*, there's a follow-up to that which appeared in the September or October issue of the *American Political Science Review* on suicide bombing. And this is the first systematic database of suicidal terrorism that has been put together and published. It's not a sample, it's a population. And so it has tremendous value. And I believe that the analysis presented in this APSR article is minimal. And so the minimal analyzable paper you published? There's a lot more that can be squeezed out of that data set. And for validation purposes, it should be very helpful.

There is, by the way, Middle Eastern agent-based modeling simulation coming.

Because of lack of time I'd like to say one thing only about Bob Reynolds' project. This is really great proof of the fertility and the power of a great initial modeling framework like that of the cultural algorithms. When Bob began developing this a number of years ago, it was really impossible to do — I guess he had a vision and he could see some future directions. But, of course, rolling all that out, one modeling stage at a time, for me it's been fascinating to see over the years the application of this framework to the rise of Oajaca civilization in the Southwest. Now with the Mesa Verde case, I really see a great deal of future for this framework. And similar to what I mentioned to Lars-Erik earlier this morning, since this really defines quite a standard in this area, it's always important to keep in mind where do we jump ahead, to maintain the frontier as vibrant and as exciting as possible? But this is a really very exciting new development.

One surprise I had is that maybe looking at the ideas of Steve Lexon in his model of the Chacoan meridian may have some bearing on the ideas that are being modeled in this project. Lexon's idea is of a polity in the Southwest that initiates with a capital in Chaco Canyon, which then moves on the same meridian to the modern-day city of Aztec, and then drops all the way down into the Chihuahua desert across the Mexican border to Paquime, which is today called Casas Grandes. And so these people had this moving capital, and Mesa Verde comes into the picture from as being a northern sort of a, not quite a frontier, but a northern community of this larger Southwestern polity with the initial capital in Chaco.

Network Complexities

EMPIRICAL ESTIMATION AND MULTI-AGENT BASED SIMULATION OF A DISCRETE CHOICE MODEL WITH NETWORK INTERACTION EFFECTS: AN EXPLORATION OF THE ROLE OF GLOBAL VERSUS LOCAL AND SOCIAL VERSUS SPATIAL NETWORKS IN TRANSPORTATION MODE CHOICE BEHAVIOR IN THE NETHERLANDS

E. DUGUNDJI,* University of Amsterdam, Amsterdam, Netherlands
L. GULYAS, Hungarian Academy of Sciences, Budapest, Hungary

ABSTRACT

An outstanding challenge in econometric discrete choice analysis is the treatment of the interdependence of decision makers' choices. Furthermore, in an activity-based transportation demand modeling framework, both social and spatial networks may be relevant. This paper illustrates theoretical issues in the estimation of a discrete choice model with global and local interactions. A multi-agent-based simulation model is presented to highlight some main hypothesized interaction effects with two broad classes of abstract networks: Erdős-Rényi and Watts-Strogatz graphs. Initial results suggest that when a network representing the interactions between an agent and the aggregate behavior of other (local) agents has the small-world property, the system behaves in the long run as that with global mean field information. Testing for the small-world property may be an alternative to collecting data on the precise details of a social network. Limitations in the present work are summarized and suggestions for future research efforts are outlined.

Keywords: Discrete choice analysis, activity-based transportation demand modeling, network interaction effects, multi-agent-based social simulation, Erdős-Rényi and Watts-Strogatz graphs

INTRODUCTION

A wide spectrum of policy measures have been put forward over the past decade to try to address the infamous rush hour road congestion in the "Randstad," the western region of the Netherlands marked by the ring of the cities Amsterdam–Utrecht–The Hague–Rotterdam. These measures range from flexible work hours to congestion pricing to light rail to facilitation of park-and-ride to road construction. The research reported here is a small part of a larger work aimed at understanding, measuring, and modeling the combined residential choice and multi-modal transportation choice behavior of households residing in the north wing of the Randstad, that is, the Amsterdam–Utrecht greater region, focusing particularly on multi-modality as a land use transportation planning policy instrument for reducing road congestion. Here is understood both the promotion and facilitation of carpooling (slow mode or single-driver private vehicle to multiple passenger private vehicle) as well as the use of so-called park-and-ride "transferia" (slow mode or single-driver private vehicle to multiple passenger transit vehicle). A central aspect of the research approach is the intended treatment of social dynamics.

---

* *Corresponding author address:* Elenna Dugundji, University of Amsterdam, Department of Geography and Planning, Nieuwe Prinsengracht 130, 1018 VZ Amsterdam, Netherlands; e-mail: E.R.Dugundji@uva.nl.

Pioneered in the domain of multidimensional choice of location, housing, automobile ownership, and mode to work by Lerman (1975), discrete choice analysis has become an industry standard in land use and transportation planning models. Some subsequent elegant and elaborate operational examples of the development of this methodology are Wegener and Spiekermann (1996), Waddell and Nourzad (unpublished), and Hensher and Ton (2001), to cite just a few. Meanwhile, the field itself has flourished in the past 30 years, ultimately extending the basic random utility model to incorporate cognitive and behavioral processes, flexible error structures, and different types of data in so-called hybrid choice models (Ben-Akiva et al., unpublished). However, as discrete choice theory is fundamentally grounded in individual choice, an outstanding challenge remains in the treatment of the interdependence of various decision-makers' choices, be that via global or local interactions. The formulation of the nature of the interaction in turn raises the issues of networks and network evolution. However when considering the problem domain of the relation between residential choice behavior and given multi-modal transportation planning and policy, not only social networks but also spatial land use and transportation networks may be relevant (Dugundji et al., 2001).
A framework for the mechanisms of interaction is proposed as follows:

(i) Interactions among individuals within households — for example, joint residential location choice in a dual-income earner household (Timmermans et al., 1992); coordinating activity schedules and travel patterns within a household.

(ii) Interactions between identifiable households proximally situated in a spatial network — for example, both nuisance from neighbors and (conversely) satisfaction with neighbors are very strong factors in the inclination to relocate, both for under 55 and over 55 age groups in the Netherlands (Hooimeijer and van Ham, 2000); coordinating carpooling with neighbors or co-workers.

(iii) Interactions between identifiable households tangentially situated in a spatial network — for example, coordinating carpooling via a carpool facility.

(iv) Interactions between identifiable households associated in a social network, not necessarily proximally or tangentially situated in a spatial network — for example, attraction to a particular municipality in choice of residential location because friends or family live there; *awareness about availability of certain alternatives in the choice set generation process through information transmission in the social network* via friends, family, neighbors, and/or co-workers, be that the suitability of a particular neighborhood in residential location choice, the suitability of using a park-and-ride transferium for a commute, or the existence of a carpool facility.

(v) Interactions between a household and the aggregate actions of other households proximally situated in a spatial network — for example, high volatility or (conversely) stagnancy of turnover in housing stock in a particular neighborhood, affecting the general desirability of a neighborhood or the possibility to move there; social pressure to own a car because other neighbors or other co-workers on average do, regardless of whether there is any direct social contact with these persons; improved feasibility for higher level of public transit service associated with higher volume of public transit ridership in a particular region.

(vi) Interactions between a household and the aggregate actions of other households tangentially situated in a spatial network — for example, many other households passing through a given household's (prospective) neighborhood on their commute trip may lead to negative traffic externalities affecting that household's evaluation of that neighborhood, such as conditions unsafe for young children, noise and air pollution, etc.

(vii) Interactions between a household and the aggregate actions of other households associated in a social network, not necessarily proximally or tangentially situated in a spatial network — for example, preference for a particular *type* of housing situation [as opposed to preference for a specific municipality, see (iv)]; social acceptance of cycling or public transit because friends, family, neighbors and/or co-workers also cycle or use public transit.

(viii) Interactions between a household and the aggregate actions of other households in a subpopulation, not necessarily associated in a social network nor proximally or tangentially situated in a spatial network — for example, as in (vii) above, not because of a household being influenced by others in a social network, but rather because of a more general trend or societal bandwagon effect.

Furthermore, an important distinction can be understood in this particular problem domain among (social and/or spatial) network interactions impacting choices, such as transport mode choice, which do not necessarily endogenously affect the household's reference position in a network (e.g., whether a household chooses carpool versus transit in a multi-modal trip, or chooses a unimodal trip, will not spatially affect the fact of who the household's neighbors or co-workers are), as opposed to network interactions affecting "sorting" type choices, such as residential location choice, which obviously endogenously impacts the household's reference position in a spatial network and potentially also within a social network (e.g., in moving to a new neighborhood, a household by definition acquires new neighbors).

In short, a distinction is hypothesized between social network interactions versus spatial network interactions, identifiable versus aggregate interactions, proximal versus tangential versus global interactions, and exogenous versus endogenous interactions. The research reported here explores mechanisms (v), (vii), and (viii) and to some extent interaction mechanisms (ii) and (iv), for the exogenous network case in the given problem domain, that is, transportation mode choice (see Figure 1). Technically, however, interactions of types (i) and (iii) may also be modeled as non-directed graphs, and thus results reported here may prove to be useful in those areas as well. The authors are currently exploring representations for the endogenous case, that is, residential choice. As similarly proposed by Brock and Durlauf (2002), the nested logit model is seen as a promising direction for coupling the exogenous network case (transportation mode choice) and the endogenous network case (residential choice).

Some examples	*Interactions between…*	Identifiable households	Aggregate households
Spatial network	Proximal	Coordinating carpooling with neighbors	**Feasibility of high level of public transit service**
	Tangential	Coordinating carpooling via a carpool facility	Traffic unsafety; noise and air pollution
Social network		Awareness about mode choice alternatives	**Social acceptance of cycling/transit**

FIGURE 1 Interaction mechanism framework (This paper considers [1] interactions between a household and the aggregate actions of other households proximally situated in a spatial network and [2] interactions between a household and the aggregate actions of other households associated in a social network. Interactions between a household and the aggregate actions of other households in the population (i.e., bandwagon effects, are also addressed as the special limiting case of a *fully connected* network].)

DISCRETE CHOICE WITH AGGREGATE GLOBAL INFORMATION

Discrete choice theory allows prediction based on computed individual choice probabilities for heterogeneous agents' evaluation of alternatives. In accordance with the notation and convention in Ben-Akiva and Lerman (1985) and Ben-Akiva and Bierlaire (1999), the so-called *multinomial logit model* well known in econometrics and discrete choice theory is specified as follows. Assume a population of N decision-making entities indexed (1,..., n, ..., N) each faced with a choice among J_n alternatives indexed (1,..., j, ..., J_n) in subset C_n of some universal choice set C. Let V_{in} be the deterministic (to the modeler) or so-called "systematic" utility that a given decision-making entity n is presumed to associate with a particular alternative i in its choice set C_n. Then the probability that the individual decision-making entity n chooses alternative i within the choice set C_n is given by:

$$P_{in} \equiv P(i \mid C_n) = \frac{e^{\mu V_{in}}}{\sum_{\forall j \in C_n} e^{\mu V_{jin}}}, \tag{1}$$

where μ is a strictly positive scale parameter that we generally normalize to 1.

For the binary case of choice with the universal set containing only two alternatives, $C = \{i,j\}$, say car versus public transit, we have the simplification:

$$P_{in} = \frac{e^{\mu V_{in}}}{e^{\mu V_{in}} + e^{\mu V_{jn}}} = \frac{e^{\mu (V_{in} - V_{jn})}}{e^{\mu (V_{in} - V_{jn})} + 1}. \tag{2}$$

We now follow the spirit of Aoki (1995) and introduce social dynamics by allowing the term $V_{in} - V_{jn}$ to be a linear-in-parameter β function of the proportions x_i and $x_j = (1 - x_i)$ of decision-making entities[1] who have made each choice:

$$\begin{aligned} &V_{in} - V_{jn} = \beta f(x) \\ &x \equiv x_i - x_j = x_i - (1 - x_i) = 2x_i - 1 \end{aligned}. \tag{3}$$

The variable x is termed a "field variable," as motivated by Aoki: "Knowledge of a field variable relieves agents (at least partially) of the need for detailed information on interaction patterns. Any macroeconomic variable that serves this decentralizing function is called a field variable." Such an approach can be particularly useful if simplifying assumptions, such as having constant interactions among all possible pairs of microeconomic agents or interactions only with other agents in neighborhoods in the sense of Markov random fields, are not appropriate.

Substituting Equation 3 into Equation 2 and normalizing the scale parameter $\mu = 1$, we have:

$$P_{in}(x) = \frac{e^{\beta f(x)}}{e^{\beta f(x)} + 1}. \tag{4}$$

Note that the parameter β indicates the level of certainty in the model. If it is fairly certain that the utility of alternative i is greater than the utility of alternative j, then $\beta >> 0$, and we have effectively deterministic choice:

$$P_{in} = \frac{e^{\beta f(x)}}{e^{\beta f(x)} + 1} \approx 1, \text{ for } \beta >> 0. \tag{5}$$

If there is uncertainty as to which technology is more profitable, then $\beta \sim 0$, and we have effectively a "fair coin toss" between the two alternatives:

$$P_{in} = \frac{e^{\beta f(x)}}{e^{\beta f(x)} + 1} \approx \frac{e^0}{e^0 + 1} = 1/2, \text{ for } \beta \approx 0. \tag{6}$$

Aoki (1995) models decision-making entities "as jump Markov processes, and the dynamics of interactions (among these entities) as birth-and-death stochastic processes where they switch their choices randomly and asynchronously with transition rates that are functions of the aggregate situations summarized by the proportion of decision-making entities who have taken the same choices." As remarked in Blume and Durlauf (2002) in an analogous continuous time Markov process changing state in discrete jumps with a uniform global interactions model

---

[1] To be clear about our notation in that it differs slightly from Aoki (1995), we define the proportion $x_i = N_i/N$ and $x_j = (1 - x_i) = 1 - (N_i/N)$, where N_i in the number of decision-making entities have chosen alternative i at time t.

(that is, a model with interactions of constant strength between all pairs of decision-making entities), “Implicit... is the fact that players are myopic in (stochastically) best-responding to the current play of the population rather than some forecast of future paths of play.” In this paper, we will accept this myopic assumption for our exogenous network case with transportation mode choice; however, particularly for the endogenous network case with residential choice, this assumption may be worth revisiting.

Let us formalize these assumptions by considering the aggregate behavior of the population of N decision-making entities instead of the behavior of an individual decision-making entity. Let $P(N_i, t)$ denote the probability that N_i number of decision-making entities have chosen alternative i at time t. The total number of possible states of the population of N decision-making entities is $N + 1$, since the number of decision-making entities choosing alternative i can range from 0 to N, and the number of decision-making entities choosing alternative j is fully determined given the number choosing alternative i, for our binary choice case. Let $W_{Ni,Ni'}$ denote the transition rate between the states of the population with N_i and $N_{i'}$ number of decision-making entities choosing alternative i, and let $W_{Ni,Ni}$ be the rate of the inverse transition. Aoki (1995) uses the backward Chapman-Kolmogorov equation, or so-called “master equation,” to govern the time evolution of the probability density. The master equation is fully specified once the transition rates are given between the states.

$$\begin{aligned}\frac{\partial P(N_i,t)}{\partial t} &= \sum_{\forall N_i'\neq N_i} P(N_i',t)W_{N_i',N_i} - \sum_{\forall N_i'\neq N_i} P(N_i,t)W_{N_i,N_i'} \\ &= \sum_{\forall N_i'\neq N_i}\left\{P(N_i',t)W_{N_i',N_i} - P(N_i,t)W_{N_i,N_i'}\right\}\end{aligned} \quad . \tag{7}$$

As remarked in Reif (1965), “Note that all terms... are real and that the time t enters linearly in the first derivative. Hence the master equation does not remain invariant as the sign of the time t is reversed from t to $-t$. This equation describes, therefore, the irreversible behavior of a system.” Nonetheless, as motivated by Reif, there is assumed to be a symmetry property relating a transition to its inverse:

$$W_{N_i',N_i} = W_{N_i,N_i'} \ . \tag{8}$$

In general, we find that the probability density $P(N_i, t)$ tends to increase with time because the population transitions from other states to the given state with N_i number of decision-making entities choosing alternative i, and the probability density tends to decrease with time because the population in the given state transitions to other states. For the earlier mentioned assumption of asynchronous choices of the decision-making entities, however, we have a convenient simplification, since the only states to which the population in the given state with N_i number of decision-making entities choosing alternative i can possibly transition to, are the states with $N_i' = N_i + 1$ and/or $N_i' = N_i - 1$ number of decision-making entities choosing alternative i. In short, as remarked by Aoki (1995), in the birth-and-death processes of this paper, the transition rates are non-zero only for N_i', which is either $N_i + 1$ (a so-called “birth”) or $N_i - 1$ (a so-called “death”). We can thus simplify the master equation for this continuous- time discrete state Markov process:

$$\begin{aligned}\frac{\partial P\left(N_i,t\right)}{\partial t} &= \left\{P\left(N_i+1,t\right)W_{N_i+1,N_i} - P\left(N_i,t\right)W_{N_i,N_i+1}\right\} \\ &\quad + \left\{P\left(N_i-1,t\right)W_{N_i-1,N_i} - P\left(N_i,t\right)W_{N_i,N_i-1}\right\} \\ &= \left\{P\left(N_i+1,t\right)-P\left(N_i,t\right)\right\}W_{N_i,N_i+1} + \left\{P\left(N_i-1,t\right)-P\left(N_i,t\right)\right\}W_{N_i,N_i-1}\end{aligned} \quad . \tag{9}$$

The assumption that only one decision-making entity revises its choice per unit time may be reasonable for analytical purposes if we consider an arbitrarily small time unit. In practical situations, however, we can also imagine that there can be non-negligible time-lag in the spread of information in the population, whereby multiple decision-making entities may revise their choices per unit time interval, before the knowledge about changes in the system is disseminated. In the multi-agent simulation implementation of the model we relax this assumption of asynchronous choices, allowing explicitly for revisions by multiple decision-making entities per unit time via an external parameter that can be set by the researcher.

In the simplest birth-and-death processes, the transition rates $W_{Ni,Ni+1}$ and $W_{Ni,Ni-1}$ are given by:

$$\begin{aligned}W_{N_i,N_i+1} &= \kappa\left(N-N_i\right) = N\kappa\left(1-\frac{N_i}{N}\right) \\ W_{N_i,N_i-1} &= \lambda N = N\lambda\frac{N_i}{N}\end{aligned} \quad . \tag{10}$$

More generally we can express the transition rates $W_{Ni,Ni+1}$ and $W_{Ni,Ni-1}$ as expansions in powers of (1/N):

$$\begin{aligned}W_{N_i,N_i+1} &= g\left(N\right)\left[\gamma_0\left(N_i/N\right)+\left(1/N\right)\gamma_1\left(N_i/N\right)+O\left(N^{-2}\right)\right] \\ W_{N_i,N_i-1} &= g\left(N\right)\left[\rho_0\left(N_i/N\right)+\left(1/N\right)\rho_1\left(N_i/N\right)+O\left(N^{-2}\right)\right]\end{aligned} \quad . \tag{11}$$

Dropping all terms of order N^{-1} and higher, and making the simplifying assumptions that:

$$g\left(N\right) = N \quad , \tag{12}$$

and that the “birth” transition rate $W_{Ni,Ni+1}$ is linear in the individual choice probability P in that alternative i is superior to alternative j, and the “death” transition rate $W_{Ni,Ni-1}$ is linear in the individual choice probability P_{jn} that alternative j is superior to alternative i, we have:

$$\begin{aligned}W_{N_i,N_i+1} &= N\gamma_0\left(N_i/N\right) = N\kappa\left(1-\frac{N_i}{N}\right)P_{in}\left(x\right) = N\kappa\frac{1-x}{2}P_{in}\left(x\right) \\ W_{N_i,N_i-1} &= N\rho_0\left(N_i/N\right) = N\lambda\left(\frac{N_i}{N}\right)P_{jn}\left(x\right) = N\lambda\frac{1+x}{2}P_{jn}\left(x\right)\end{aligned} \quad . \tag{13}$$

Aoki (1995) shows that the mean φ of the field variable x is governed by the deterministic differential equation:

$$\frac{d\varphi}{dt} = \gamma_0(\varphi) - \rho_0(\varphi) = \kappa\frac{1-\varphi}{2}P_{in}(\varphi) - \lambda\frac{1+\varphi}{2}P_{jn}(\varphi)\ . \tag{14}$$

Normalizing κ=1 and λ=1, we have:

$$\begin{aligned}\frac{d\varphi}{dt} &= \frac{1-\varphi}{2}\left(\frac{e^{\beta f(\varphi)}}{e^{\beta f(\varphi)}+1}\right) - \frac{1+\varphi}{2}\left(\frac{1}{e^{\beta f(\varphi)}+1}\right) = \frac{e^{\beta f(\varphi)} - \varphi e^{\beta f(\varphi)} - 1 - \varphi}{2\left(e^{\beta f(\varphi)}+1\right)} \\ &= \frac{1}{2}\frac{e^{\beta f(\varphi)}-1}{e^{\beta f(\varphi)}+1} - \frac{\varphi}{2}\frac{e^{\beta f(\varphi)}+1}{e^{\beta f(\varphi)}+1} \\ &= \frac{1}{2}\frac{e^{\frac{1}{2}\beta f(\varphi)} - e^{-\frac{1}{2}\beta f(\varphi)}}{e^{\frac{1}{2}\beta f(\varphi)} + e^{-\frac{1}{2}\beta f(\varphi)}} - \frac{\varphi}{2} \\ &= \frac{1}{2}\left(\tanh\frac{1}{2}\beta f(\varphi)\right) - \frac{\varphi}{2}\end{aligned}\ . \tag{15}$$

Stationary points are zeros of $d\varphi/dt$. Thus the key equation to determine local equilibria is:

$$\frac{d\varphi}{dt} = 0: \quad \varphi = \tanh\frac{1}{2}\beta f(\varphi)\ . \tag{16}$$

This equation can be solved conveniently graphically, by plotting the left side and the right side on a graph, and finding their intersection (see Figure 2). Depending on the specification of $f(\varphi)$, this equation may have more than one solution. Equivalently, from the first line of Equation 15, we could instead write:

$$\begin{aligned}\frac{d\varphi}{dt} = 0: \quad (1-\varphi)e^{\beta f(\varphi)} &= 1+\varphi \\ e^{\beta f(\varphi)} &= \frac{1+\varphi}{1-\varphi} \\ f(\varphi) &= \frac{1}{\beta}\ln\left(\frac{1+\varphi}{1-\varphi}\right)\end{aligned}\ . \tag{17}$$

A stationary point of the mean φ for the field variable x is locally stable in perturbations of the mean if the derivative $d^2\varphi/d\varphi\, dt$ is negative:

$$\frac{d}{d\varphi}\left(\frac{d\varphi}{dt}\right)\bigg|_{\varphi=\tanh\frac{1}{2}\beta f(\varphi)} = \frac{d}{d\varphi}\left(\frac{1}{2}\left(\tanh\frac{1}{2}\beta f(\varphi)\right)-\frac{\varphi}{2}\right)\bigg|_{\varphi=\tanh\frac{1}{2}\beta f(\varphi)}$$

$$= \frac{1}{2}\left(1-\tanh^2\frac{1}{2}\beta f(\varphi)\right)\left(\frac{1}{2}\beta f'(\varphi)\right)-\frac{1}{2}\left(\frac{d\varphi}{d\varphi}\right)\bigg|_{\varphi=\tanh\frac{1}{2}\beta f(\varphi)} \quad . \tag{18}$$

$$= \frac{1}{4}\left(1-\varphi^2\right)\beta f'(\varphi)-\frac{1}{2}$$

Thus we have the condition for local stability:

$$\frac{d}{d\varphi}\left(\frac{d\varphi}{dt}\right)\bigg|_{\varphi=\tanh\frac{1}{2}\beta f(\varphi)} \leq 0 : \quad \left(1-\varphi^2\right)\beta f'(\varphi) \leq 2 \ . \tag{19}$$

If the derivative $f'(\varphi)$ is nonpositive with β nonnegative, this local stability condition is always satisfied, since we have defined $x = x_i - x_j = x_i - (1 - x_j) = 2x_i - 1$ on the interval [-1,1], and thus $\varphi = E(x) \leq 1$. When β is large and if $f'(\varphi)$ is positive, this inequality may be violated, with the equilibrium becoming unstable (see Figure 2).

Using the Repast agent-based modeling platform, we created a computational version of this model and replicated Aoki's original work. Example results for $f(x) = x$ are shown in Figure 3. With a low value of the parameter β, we obtain precisely the case of one stable equilibrium at $x = 0$. Conversely, using a high value of the parameter β, we obtain precisely the unstable equilibrium at $x = 0$, with the system being driven away with equal probability to either the extreme $x = +1$ or the extreme $x = -1$. Cumulative histograms for 500 runs (plotting the value of x after 2000 iterations) are shown in Figure 4.

An example of the Repast graphical user interface (GUI) is shown in Figure 5.

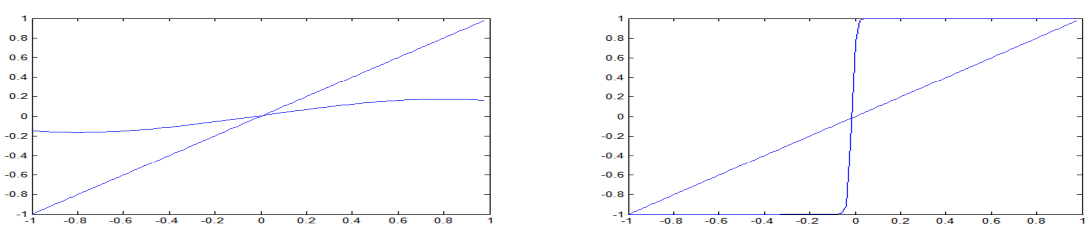

(a) $\beta \sim 0$ φ (b) $\beta >> 0$

FIGURE 2 Plotted graphs of $y = \varphi$, and $y = \tanh(1/2)\beta f(\varphi)$ with $f'(\varphi)$ positive, showing (a) one stable equilibrium at $\varphi = 0$ and (b) one unstable equilibrium at $\varphi = 0$

(a) $\beta = 0.03$, Random seed = 1

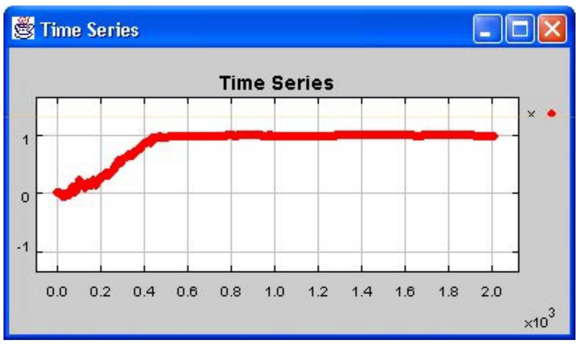

(b) $\beta = 5$, Random seed = 1

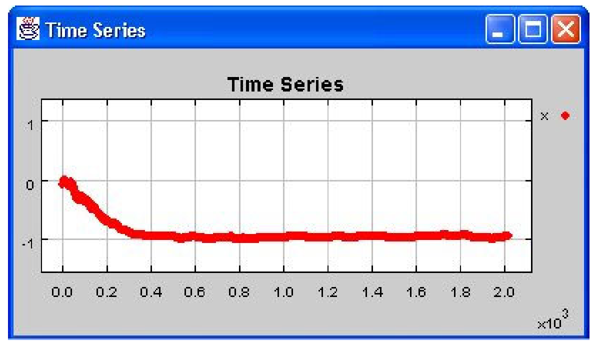

(c) $\beta = 5$, Random seed = 3

FIGURE 3 Example time series for 100 agents with $f(x) = x$ for (a) low certainty (see also Figure 5) and (b) and (c) high certainty with two distinct random seeds

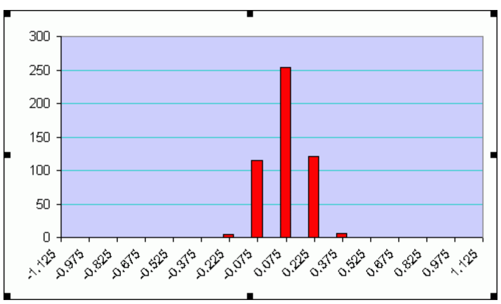

(a) $\beta = 0.03$

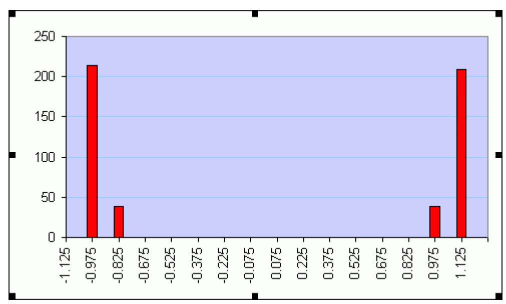

(b) $\beta = 5$

FIGURE 4 Computational results for 100 agents with (a) low and (b) high certainty; cumulative histograms of $f(x) = x$ over 500 runs (i.e., 500 distinct random seeds)

(a) Control panel

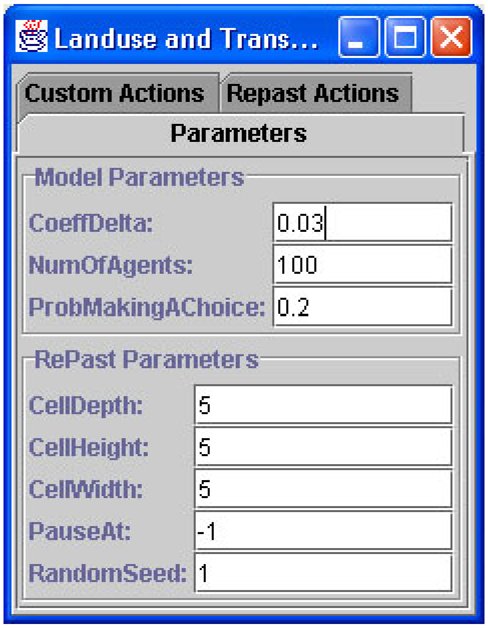

(b) Simulation settings

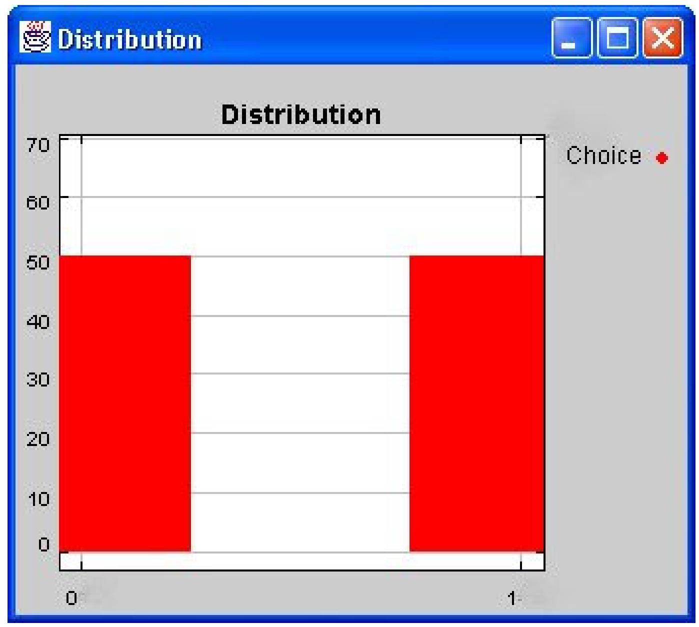

(c) Distribution of decision-making entities having chosen each alternative at the indicated tick count

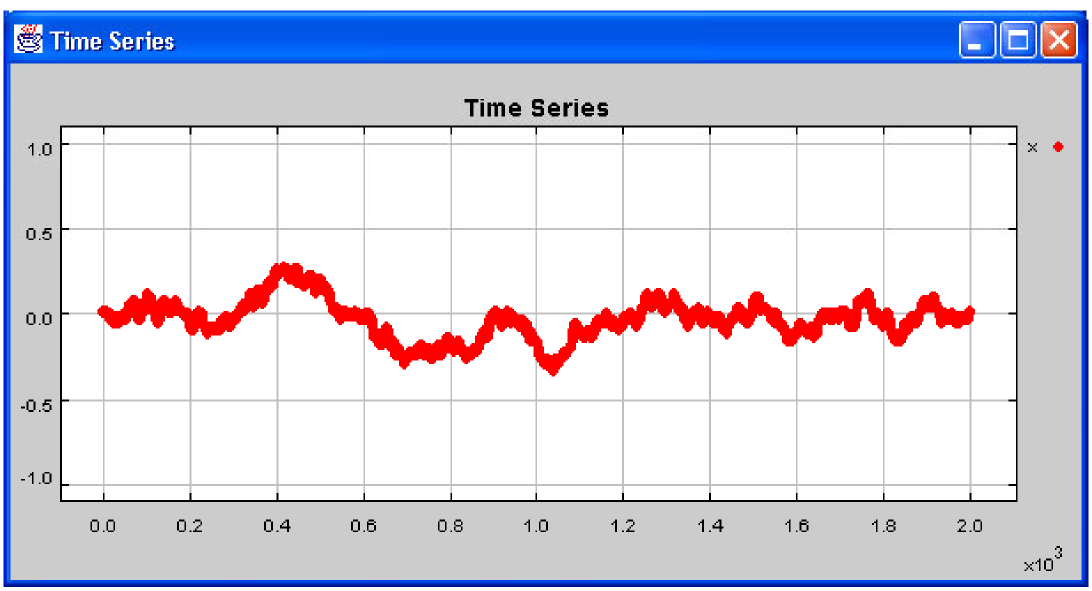

(d) Time series for the field variable *x*

FIGURE 5 Repast graphical user interface for an example run (The setting “CoeffDelta” in [b] refers to the parameter β in the text indicating the level of certainty in the model.)

AN AGENT-BASED APPROACH: HETEROGENEOUS LOCAL INFORMATION

Aoki's model described in the previous section assumes uniform, global and perfect information access, unlike in the real world. The very fact that certain influences are transferred via social interactions, and thus via social networks, implies heterogeneous local information. Therefore, the model is extended to explicitly model interaction networks. The perceived information is still perfect; however, noise and errors are not modeled explicitly.

In our model, each decision-making entity n is assigned a set of "reference" decision-making entities influencing its choice. At each time step, the decision-making entities look at the choices their particular reference entities made in the previous round, plus their own choice, and calculate *localized* values of the difference in systematic utility between the alternatives.

Obviously, the "reference" relationships introduced above define a graph or network. Let us denote this graph by G = (**N**, L), where **N** is the set of nodes (vertices) and L is the set of edges (links) between them. In our case, each decision-making entity is a node, i.e., **N** = {1,...,n,...,N} and a decision-maker's "reference" entities are defined by its links L(n). We assume that edges are symmetric and that the graph contains no loops (i.e., no node can have a link to itself).

It is hypothesized that different network structures yield different system behavior. Empirically however, it can be in practice difficult to reveal the exact details of the relevant network(s) of reference entities influencing the choice of each decision-making entity. Moreover, the actual reference entities for a given decision-making entity may not be among those in the data sample. Therefore, we have turned our attention towards studying abstract classes of networks in the hope of identifying classes of networks that yield similar results.

Erdős-Rényi Networks

The first class of interaction networks investigated is the Erdős-Rényi graph or random network. A random network consists of a number of nodes and set of random edges between them, such that the probability of the existence of a given link is uniform across all possible edges. The actual number of the links is determined by the density p of the network, which is usually perceived as a parameter of the Erdős-Rényi graph. Here network density p is defined as the ratio of the number of existing links versus the number of all possible links, that is:

$$p = \#\{\text{actual links}\} \,/\, \#\{\text{all possible links}\}. \tag{20}$$

One advantage of studying random networks is that they are perhaps the simplest possible networks that are general enough to describe a wide range of graphs, from unconnected nodes to a fully connected network (i.e., a graph that contains all possible links, as in our initial replication of the Aoki model). In addition, they accomplish this without introducing any explicit bias into the structure of the network. Moreover, results are known about important properties, such as at approximately what value of p will the network become connected (i.e., when each node is "reachable" along the edges from any other node), or when a so-called "giant component" will emerge. Finally, an important feature of random networks that is observed in real-life social networks is the so-called "small-world" property: the average path length l (the average number of “hops” between an arbitrary pair of nodes) is less than or of the order ln(N), where N is the number of nodes.

With the discrete choice model on a random network, both analytical results and computational replications confirm that low certainty values effectively yield a “fair coin toss” between the alternatives. This is due to the fact that the relative importance of others’ choices becomes low. Therefore, there is no reason to believe that the structure of the underlying interaction network would play any significant role in determining the outcome. The following experiments were therefore carried out with a high certainty value, $\beta = 5$.

As a base consistency check, we first tested our model with $p = 1.0$, which, should yield a model equivalent to the one discussed in the previous section. Simulations[2] confirmed this expectation. Experiments were then carried out with density values ranging from 0 to 0.025 (p = 0, 0.005, 0.01, 0.015, 0.02, and 0.025). The results show that low densities yield behavior similar to that of low certainty, while higher p values display tendencies toward the high-certainty outcome of the global information model (see Figure 6).

The actual density values tested were designed to embrace the critical point at which a giant component emerges, i.e., when, in practical terms, the graph becomes connected. It is known (Molloy and Reed, 1998) that this occurs around $p = 1/N + \varepsilon$, where N is the number of nodes and $\varepsilon > 0$ is a small value. In case of 100 agents, this formula gives $p = 0.01$ as the critical point. Indeed, simulation results show “random outcomes” for subcritical densities. Also, a significant change in behavior occurs around a density of 0.02. The range in between yields more ambiguous outcomes, which may warrant further study. Nonetheless, the overall picture suggests

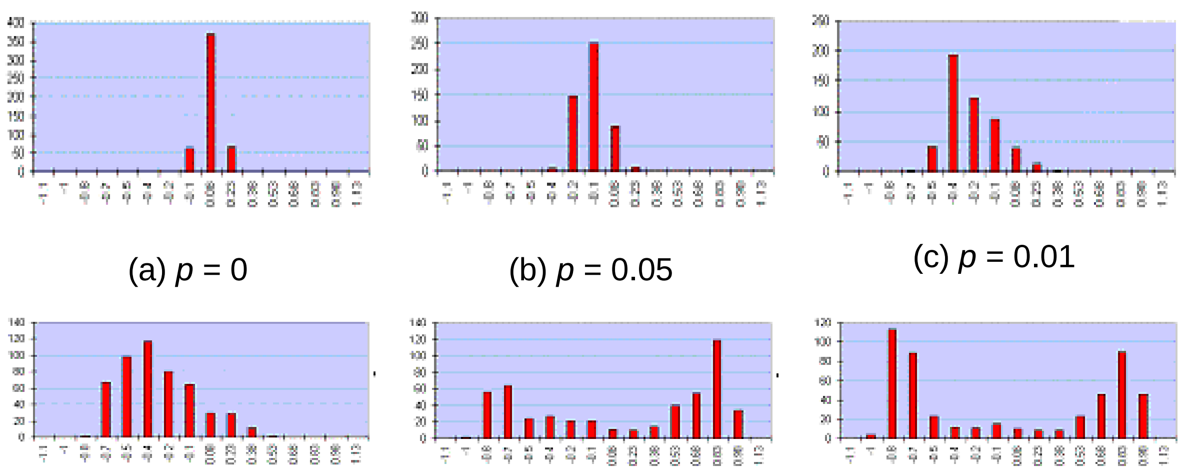

(a) $p = 0$ (b) $p = 0.05$ (c) $p = 0.01$

(d) $p = 0.015$ (e) $p = 0.02$ (f) $p = 0.025$

FIGURE 6 Cumulative histograms of $f(x) = x$ over 500 runs for different network densities

2 All simulations reported here consisted of 100 agents and had 20 different networks for each density value. With each individual network, 500 independent runs were carried out. Each run was stopped after 2,000 iterations.

that it does not really matter what structure the interaction network has; as long as it is connected, it starts yielding outcomes similar to the fully connected graph.

Watts-Strogatz Networks

In the previous section, we found that random networks yield behavior similar to that based on global aggregate information, given connectedness, or more precisely, the existence of a giant component. This result may be counter-intuitive as it appears relatively easy to craft examples where it is, at least, unlikely that one alternative will eventually reach total dominance. Therefore, we re-visit the unbiased nature of the Erdős-Rényi (1959) graph. An important property that is observed in real-life social networks, but not embodied in a random network, is "clustering": i.e., two friends of a certain person are more likely to be mutual friends themselves than an arbitrary pair of individuals.

The Watts-Strogatz (1998) model starts from an ordered network, or lattice, which contrary to the random network, has high clustering, but long average path length, or in effect, no small-world property. The dimensionality of the lattice is a parameter, although only 1D and 2D models are commonly discussed (see Figure 7). The extent to which the neighbors are connected is also a parameter of the Watts-Strogatz model. In Figure 7, nodes are linked to their immediate neighbors, that is, with extent equal to 1. To avoid artificial boundary effects, thorical lattices are considered that are "wrapped around," that is, nodes on the boundary of the system link to nodes at the opposite boundary.

"Shortcuts" are then introduced into these systems to create the Watts-Strogatz model, by randomly rewiring a few links. The controlling parameter is the probability of rewiring w, which gives the probability that each original link in the system is replaced by a random connection. Only a very few shortcuts (i.e., a fairly low w) are needed to achieve the small-world property.

Experimenting with Watts-Strogatz networks for our discrete choice model, the first thing to consider is the density of the generated graphs. In our experiments this value was $p = 0.04$ (1D with extent = 2), a density that is sufficiently high to be in the two-equilibrium regime of the random network model. Indeed, experiments with $w = 1.0$, which renders the

(a) 1D thorical lattice

(b) 2D thorical lattice

FIGURE 7 Examples of ordered networks (with extent = 1)

Watts-Strogatz model equivalent to that of Erdős and Rényi, confirm this.[3] Varying w from 0 to 0.3 shows that as the average path length l falls below the small-world threshold, the outcome converges to the two certain choices behavior (see Figure 8).

CONCLUSIONS AND FUTURE WORK

Our initial results with both Erdős-Rényi and Watts-Strogatz graphs suggest that when a network representing the interactions between a decision-making entity and the aggregate behavior of other (local) reference entities has the small-world property, the system behaves in the *long run* as Aoki's original model with global mean field information (Aoki 1995). If we are only interested in long-run behavior and not how long the system takes to transition there, testing for the small-world property may be an empirically advantageous alternative to collecting data on the precise details of a social network.

Another important and widely studied property of social networks is the degree distribution, that is, in our case, the distribution of the number of "references" the decision-making entities have. Both the Erdős-Rényi graph and Watts-Strogatz graph yield a Poisson degree distribution, which is unlike various real-life cases. Further studies with models yielding different degree distributions may be fruitful. Another interesting question is what class of networks may be needed to address interactions between *identifiable* decision-making entities,

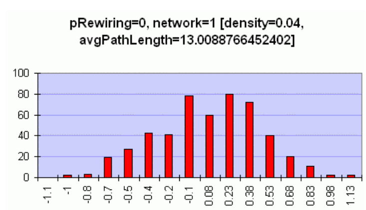

(a) w = 0, l ~ 13

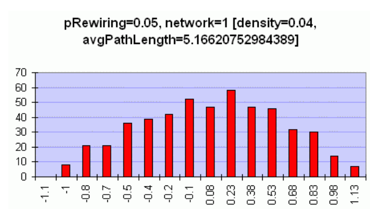

(b) w = 0.04, l ~ 5.2

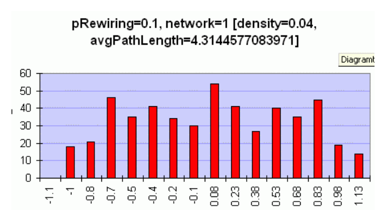

(c) w = 0.1, l ~ 4.3

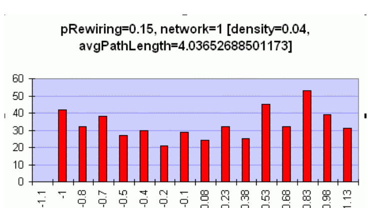

(d) w = 0.15, l = 4.0

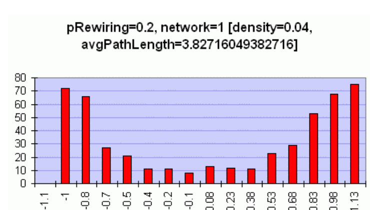

(e) w = 0.2, l = 3.8

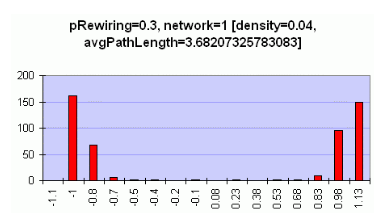

(f) w = 0.3, l ~3.7

FIGURE 8 The effect of rewiring w and decreasing average path length l. The "small-world threshold" is ln(100) ~ 4.6

---

3 The results reported here are based on experiments with 20 different networks for each parameter set. Each network was tested with 500 independent runs, and each run lasted 5,000 iterations.

as technically they may also be well-modeled by interaction graphs, but in particular, as the current results suggest that more *sparse* networks are more dependent on the actual reference structure.

ACKNOWLEDGMENTS

This research is made possible in part through a generous grant by the Dutch National Science Organization (NWO) for Project II of the AMADEUS program (Assessing the time-varying effects of Multi-modal transportation networks on Activity and Destination Choices in Urban Systems), directed by Harry Timmermans and coordinated by Theo Arentze. Additional financial and institutional support was provided by the University of Amsterdam. Elenna Dugundji would like to thank Frank le Clercq, Loek Kapoen, and Moshe Ben-Akiva for primary guidance, as well as Masanao Aoki, Axel Leijonhufvud, Peter Howitt, Robert Axtell, and other lecturers and participants at the CEEL Third Trento Summer School on Adaptive Economic Processes. Laszlo Gulyas would like to thank George Kampis, József Váncza, and András Márkus for fruitful discussions.

REFERENCES

Aoki, M., 1995, "Economic Fluctuations with Interactive Agents: Dynamic and Stochastic Externalities," *Japanese Economic Review* **46**(2):148–165.

Ben-Akiva, M., and M. Bierlaire, 1999, "Discrete choice methods and their applications to short-term travel decisions," in R. Hall, ed., *Handbook of Transportation Science*, International Series in Operations Research and Management Science, Kluwer, Dordrecht.

Ben-Akiva, M. and S.R. Lerman, 1985, *Discrete Choice Analysis: Theory and Application to Travel Demand*, Cambridge, MA: The MIT Press.

Ben-Akiva, M., D. McFadden, K. Train, J. Walker, C. Bhat, M. Bierlaire, D. Bolduc, A. Boersch-Supan, D. Brownstone, D.S. Bunch, A. Daly, A. de Palma, D. Gopinath, A. Karlstrom, and M.A. Munizaga, "Hybrid Choice Models: Progress and Challenges," *Journal of Marketing Research* (forthcoming).

Blume, L.S., and Durlauf, 2002, "Equilibrium Concepts for Social Interaction Models," *SSRI Working Papers*, Social Systems Research Institute, University of Wisconsin, Madison.

Brock, W., and S. Durlauf, 2002, "A Multinomial Choice Model of Neighborhood Effects, *American Economic Review* **92**:298–303.

Dugundji, E., L.L. Kapoen, F. le Clercq, T.A. Arentze, H.J.P. Timmermans, and K.J. Veldhuisen 2001, "The Long-term Effects of Multi-modal Transportation Networks: The Residential Choice Behavior of Households," *Proceedings of the 9th World Conference on Transportation Research*, Seoul, Korea, July.

Erdős, P.A., and Rényi, 1959, "On Random Graphs," *Publicationes Mathematicae* **6:**290–297.

Hensher, D.A., and T. Ton, 2001, "TRESIS: A Transportation, Land Use and Environmental Strategy Impact Simulator for Urban Areas," *Proceedings of the 9th World Conference on Transportation Research*, Seoul, Korea, July.

Hooimeijer, P., and M. van Ham, 2000, "Leefbaarheid en Buurtkwaliteit (Livability and Neighborhood Quality)," in W. Relou, ed., *Perpectief op Wonen: Rapportage van het WoningBehoefte Onderzoek 1998*, Netherlands Ministry of Housing, Spatial Planning and the Environment (VROM), The Hague.

Lerman, S.R., 1975, *A Disaggregate Behavioral Model of Urban Mobility Decisions*, Ph.D. dissertation, Department of Civil Engineering, MIT, Cambridge, MA.

Molloy, M., and B. Reed, 1998, "The Size of the Largest Component of a Random Graph on a Fixed Degree Sequence," *Combinatorics, Probability and Computing* **7**:295–306.

Reif, F., 1965, *Fundamentals of Statistical and Thermal Physics*, New York, NY: McGraw-Hill.

Timmermans, H.J.P., A.W.J. Borgers, J. van Dijk, and H. Oppewal, 1992, "Residential Choice Behavior of Dual-Earner Households: A Decompositional Joint Choice Model," *Environment and Planning A* **24**:517–533.

Waddell, P., and F. Nourzad, "Incorporating Non-motorized Mode and Neighborhood Accessibility in an Integrated Land Use and Transportation Model System," in *Transportation Research Record* (forthcoming).

Watts, D.J., and S.H. Strogatz, 1998, "Collective Dynamics of 'Small-World' networks," *Nature* **393**:440–442.

Wegener, M., and K. Spiekermann, 1996, "The Potential of Microsimulation for Urban Models," in G.P. Clarke, ed., *Microsimulation for Urban and Regional Policy Analysis*, European Research in Regional Science, Vol. 6, Pion, London.

EMERGENCE OF TRADING NETWORKS: A HYPERCYCLE APPROACH*

JOHN E. PADGETT,† DOOWAN LEE, and NICK COLLIER, The University of Chicago

ABSTRACT

On the basis of the hypercycle model of the origins of chemical life on earth, this paper develops an autocatalytic model of the co-evolution of economic production and economic firms, represented as skills. Production and distribution of goods by firms are only a part of what is accomplished in markets. Through a learning process, firms are produced and transformed as goods pass through them. By using both agent-based and analytic modeling, we establish three principles of social organization: structured topology, altruistic learning, and stigmergy. These principles provide sufficient foundations for the unconscious evolution of technological complexity.

Keywords: Hypercycle, Repast, agent-based modeling, economics production.

INTRODUCTION

The production and distribution of goods by firms represent only half of what happens in economic markets. The production and transformation of the firms themselves also occur as a result of the goods passing through them. This transformation is not just a matter of an increase or decrease in profits. Skills and the core competencies that define firms are developed and maintained through "learning by doing" and other learning processes that are triggered by exchanges among firms. In periods of decentralization and outsourcing, like today, it is more evident than ever that linked chains of skills are distributed across firms. In this context especially, the learning and evolution of distributed skill sets reverberate directly in the reconstitution of firms. Evolving links among firms, in turn, guide and shape the recombinant new-product possibilities latent in distributed skill sets.

The duality of this co-evolution between product and organization is often ignored by analysts as they "assume away" one side of the dynamics in order to focus attention on the other. One place to find analytic inspiration is the field of chemistry. From the chemical perspective, life is an interacting ensemble of chemicals that reproduces itself through time, in the face of turnover of its parts.[1] Biological organisms are not fixed entities; they are autocatalytic networks of chemical transformations that continually reconstruct both themselves and their physical containers. The origin-of-life problem, under this view, is how such an ensemble can self-organize from a "soup" of random chemicals that are interacting and in flux.

---

* This paper updates a related work, titled "Economic Production as Chemistry," that was published in 2003 in *Industrial and Corporate Change*, Vol. 12, pp. 843-877.

† *Corresponding author address*: John E. Padgett, Department of Political Science, The University of Chicago, Chicago, IL; e-mail: jpadgett@uchicago.edu.

1 From the physics and biological points of view, additional criteria are sometimes added for defining "life." Physicists (e.g., Prigogine and Glansdorff 1971) sometimes add the criterion of far-from-equilibrium throughput of energy. Biologists (e.g., Maturana and Varela 1980) sometimes add the criterion of permeable encapsulation.

This chemical perspective can be applied to the analysis of the co-evolution of products and firms through the following analogy. Skills, like chemical reactions, are rules that transform products into other products. Products, like chemicals, are transformed by skills. Firms, like organisms, are containers of skills that transform products. Trade, like food, passes transformed products around through exchange networks, renewing skills and thereby firms in the process. In the macroeconomic aggregate, product inputs flow into this trading network of firms and skills, and outputs flow out of it.

In this view, firms are sites through which a distributed production process flows, akin to a chemical reaction. At a minimum, firms can be considered to merely be collection bins for diverse skills. Trading among firms regulates both the activation and the evolution of the skill sets that are distributed across firms. The composition of the skills within firms evolves through learning-by-doing, among other methods: the more a skill is used, the more the skill is reinforced. Skills that are not used are forgotten. These two processes of learning and forgetting impose selection pressure on an evolving network-of-skills-through-firms production system. The origin-of-life problem for markets, then, is to discover how a randomly distributed set of skills across firms can self-organize, through exchange, into a coherent product-transformation network,[2] which then reproduces itself through time and "grows" a set of firms to sustain it.

Inspired by the literature in chemistry on hypercycles, we first develop a family of economic production models that operationalizes this co-evolutionary perspective on markets. We then discuss extensions beyond the hypercycle framework at the end of this paper.

The "hypercycle" is a specific model of the chemical origin of life pioneered by Eigen (1971) and Eigen and Schuster (1979) and extended by others (e.g., Hofbauer and Sigmund 1988; Kauffman 1986, 1993; Fontana and Buss 1994; full literature reviewed in Stadler and Stadler 2002). From random distributions of chemicals, the hypercycle model seeks to find and to grow sets of chemical transformations that include self-reinforcing loops: $\{1 \rightarrow 2, 2 \rightarrow 3, 3 \rightarrow 4, \ldots, n \rightarrow 1\}$. Chemical cycles are crucial to the issue of life because they are the motors behind the self-reproduction of metabolic networks, in the face of continuous turnover in component chemicals. Without cycles, there is no positive feedback for growth; without them, any chemical reaction left to itself will stop or "die." Eigen and Schuster, Hofbauer and Sigmund, and others have explored how variations in the reaction rates, density, and number of components can affect the dynamic stability or "survivability" of various classes of hypercyclic chemical reactions within a well-stirred liquid reaction tank. Boerlijst and Hogeweg (1991) and Padgett (1997) extended the investigation beyond the original liquid context to a spatial topology of interaction.

Viewing economics as chemistry entails extraordinarily minimalist assumptions about economic production: firms become nothing more than bins of transformation rules; products randomly flow in and through these bins, without purpose; rules reproduce or die only as functions of use. There is no guiding intelligence, either at the level of the market or at the level of the firm.[3] In such a minimalist setup, the analytic question is this: Can any coherent and self-reproducing systems of production (that is, coevolved sets of products and firms) emerge? And if they can, what mechanisms affect the likelihood of such emergence? A priori, one might expect that not much complex economic organization could result from randomly iterated rules. Yet the

---

2 Such a network could be called a "metabolism" or a "technology," depending on the application context.

3 This is not only bounded rationality, this is the absence of consciousness altogether.

history of chemical and biological life on earth suggests that minimalist systems can generate astounding complexity under the right circumstances. Intelligence, we speculate, may not have been necessary for markets to emerge.[4] We are not arguing thereby that humans are no more complicated than chemicals. We are arguing that a surprising amount of social and economic organization does not depend on humans being complicated.

METHODOLOGICAL CONSIDERATIONS

Here we describe our hypercycle model of economic production in pseudo-algorithmic fashion, since we have implemented it in the form of an agent-based simulation.[5] We used the Repast simulation development package developed at the Social Science Research Computer Center at The University of Chicago. The method we used to verify the specification validity of our model is a "scaling technique." Scaling is essentially a comparison of analytic results with simulation outcomes while the experimental complexity of the model is gradually increased. Here we start with a simple setup of parameters and gradually increase the complexity of the experimental variations with subsequent simulations. Then we formulate probability functions of the model's basic setup and analytically derive the probability of each parameter's permutations by solving the mathematical expressions. Finally, we compare the analytic solutions to the simulation outcomes obtained by the same setup of parameters to show that the model is behaving as expected. This way, we know with certainty that the simulation outcomes are not some artifact of the source code or accidental convergence. We believe that this is a reliable way of testing the specification validity of the model when there are no empirical data for verification.

HYPERCYCLE MODEL OF ECONOMIC PRODUCTION

First we describe our core models of production and learning. These illustrate the logic behind our basic "dependent variable": hypercycle emergence. Then we describe experimental variations in our core model: number of products, interaction topology, mode of learning, input environment, and input search. These are the "independent variables" that may affect the likelihood of hypercycle emergence. The simplest versions of our spatial hypercycle model can be solved analytically. We present such solutions below in order to both verify and aid in the interpretation of the simulation results.[6]

---

[4] Hayek (1948) made an argument about the "self-organizing" operation of markets that is similar to the one we make here about the emergence of markets.

[5] Our agent-based model is publicly available for both demonstration and modification. It is available at http://repast.sourceforge.net under the application module HYPERCYCLE. Repast is a comprehensive software framework and library for creating agent-based simulations, built in the Java language.

[6] In economics, though not in physics, there is frequently a fruitless methodological debate about agent-based modeling versus analytic modeling. Our position is that one can and should do both: namely, solve simple settings analytically and then scale up through computer modeling. Analytic solutions are more transparent than computer simulations, but they frequently require the imposition of highly restrictive and unrealistic homogeneity assumptions. Computers can numerically solve highly nonlinear models with heterogeneous agents in nonhomogeneous topologies, and there is no reason not to let them do so as long as one can understand the results.

Core Model of Production

1. There are three components in the model: rules ("skills"), balls ("products"), and bins ("firms").

2. Rules/skills transform balls/products into other balls/products. For example, if balls/products are indexed by i = 1,2,3,…,n, then the set of transformation rules obeying a cyclic structure[7] would be represented as {1→2, 2→3, 3→4, …, n→1}. The set of transformation rules we call a "technology"; n indexes the "complexity" of the technology.

3. Rules/skills are contained in bins/firms. At the beginning of each simulation run, skills are just randomly distributed across available firms, without any logic. The number of firms initially is large.

4. Bins/firms are arrayed on a spatial grid, with wraparound boundaries. Each firm has eight possible nearest-neighbor trading partners.

5. At each asynchronous iteration of the model, a random rule is chosen as "looking for action." The firm containing that rule/skill reaches into the input environment (modeled as an urn) and draws an input ball/product. If the input ball/product selected is compatible with that rule, then the ball/product is transformed according to that rule. (For example, if a firm possessed an activated 1→2 skill, and it drew a 1 as input from the urn environment, then it would transform the input 1 into the output 2.) If the ball/product selected could not be processed by the activated rule, then the input ball/product would pass through the firm into the output environment (also modeled as an urn) unchanged.

6. Products successfully transformed within the firm are passed randomly to one of the firm's eight possible trading partners. If that trading partner possesses a compatible skill, then it transforms the product further, and passes the transformed product along in a random direction. (For example, if the second firm possessed a 2→3, then, after receiving the output 2 from the first firm, it would transform the 2 into a 3, and then pass that on to a third firm or possibly back to the first.) In this way, transformed products pass through sequences or chains of skills.

7. Bins/firms continue passing around transformed products among themselves until the product lands on a firm that does not possess a compatible skill to transform it further. At that point, the product is ejected into the output environment, and a new input ball is selected to restart the iterative process.

---

7 Rule sets that are more general than the cyclic structure are quite possible to set up and explore. In this paper, however, we restrict ourselves to the cyclic structure in order to root ourselves in the preexisting literature on hypercycles. In future papers, we intend to explore other rule sets. The HYPERCYCLE code has already been written with this extension in mind.

Overall, the production process is as follows: Input balls/products come in from an input environment, then pass in random directions through randomly distributed production chains of skills. They are transformed en route until they pass back into an output environment. For this random production process to self-organize into coherence, there must be some sort of a feedback mechanism. In our case, this is learning through trade.

Core Model of Learning

1. "Learning by doing" is modeled in a chemical fashion as follows. If one skill transforms a product and then passes it on to another transforming skill, the skill is reproduced. We call such a sequence a "successful transaction," since both sides transform products.[8] Which of the two skills is reproduced in a successful transaction is an experimental variation within the model, as discussed below.

2. "Forgetting" is modeled in a chemical fashion as follows. Whenever one skill is reproduced anywhere in the system, another skill, randomly chosen from the overall population of skills, is killed. The total population volume of skills in the population is thereby held constant.[9]

3. Once a firm loses all its skills, it "goes bankrupt" or "dies," never to recover any skills.

For firms, learning is thus equivalent to reproducing skills. We argue that learning and the reproduction of skills are the same process; they may simply be called by a different name at different levels of analysis.

Our focus is analogous to taking a germ's eye view of disease. Instead of focusing on the organism getting sick, we focus instead on the reproduction and spread of germs. Firms learn and adapt in our model, but the underlying mechanism is not conscious reasoning. Rather it is the reproduction of the inherited skills through use.[10] Firms are kept alive or are killed solely on the basis of the skills that operate through them like chemical reactions.

This combination of learning, forgetting, and dying imposes selection pressure on the production system of skills. In the face of inexorable forgetting, skills must reproduce in order to survive. In the harsh conservation-of-skills setup employed here, the success of the rules in one place in the system imposes sharply competitive selection pressure on the rules elsewhere in the system. Heavily used subsets of the distributed skill set reproduce, and rarely used subsets of the distributed skill set disappear. The death of a firm is an absorbing state that permanently

---

8 Final consumption is the output urn.

9 This conservation-of-skills assumption mimics the conservation-of-mass assumption in chemistry. While perhaps too harsh an assumption for many human populations, this constraint is one chemistry-style way to model competition among firms.

10 In future extensions of this model, we intend to add diffusion of skills among trading firms, in order to mimic "collaborative dialogue."

eliminates its unsuccessful skills.[11] As the composition of the skills within rules within firms thereby evolves, surviving firms cluster into mutually reinforcing trading groups, reminiscent of Marshallian industrial districts. Production chains of compatibly sequenced rules self-organize their way through these spatially contiguous groups of firms.

A conscious desire to cooperate (indeed, consciousness itself) is not necessary for mutually reinforcing clusters of trading firms to emerge or survive. In this model, the minimal requirement for the long-term survival of both firms and clusters is to participate in at least one spatially distributed production chain that closes in on itself to form a loop. Not all production chains within a trading cluster need to be closed into loops. And more than one loop within a cluster is possible, in which case, there may be a dense hypercyclic network of spatially distributed production rules. However, loops within distributed chains of skill are crucial, not for production itself but for the competitive reproduction of skills. Loops set up positive feedbacks of growth in skills; these give the firms that participate in the loops the reproductive ability to outproduce firms that do not participate. Put another way, clusters of firms can produce products with or without hypercycles, but firms whose skill sets participate in production chains that include loops have the further capacity to keep renewing each other through time. This is the chemical definition of life.

From a chemical perspective, therefore, the secret to understanding the competitive success of both firms and industrial districts is to find the conditions that foster the spontaneous self-organization of skills into self-reinforcing hypercyclic production chains that wend their way through firms, knitting them together in trade and helping them to foster each other's reproduction through continuous learning.

Experimental Variations

Five independent variables (i.e., experimental treatments in the simulation model) affect the likelihood of finding and sustaining self-organized hypercycles of skills. The first three are discussed here.

1. *Complexity*. A parametrically fixed volume of rules or skills is scattered randomly around the space of firms at the beginning of each run. In this paper, 200 rules are scattered, and the composition, or complexity, of the rule set is varied. In a cyclic structure of rules, complexity is indexed by n. We vary n from 2 to 9; that is, we explore two-skill hypercycles, three-skill hypercycles, and so forth, up to nine-skill hypercycles.

2. *Interaction topology*. The basic spatial topology for trading explored in this paper is the 10×10 wraparound grid. At the beginning of each run, there are 100 firms (one firm per cell in the grid), each of which can trade products with its eight nearest neighbors. This is the so-called Moore-neighborhood topology.[12]

---

[11] Allowing the entry of new firms is another obvious extension to our model that we do not explore here.

[12] In future work, we plan to investigate additional topologies as well. Padgett (1997) used four-neighbor (von Neumann) neighborhoods. The impact of social networks of various kinds, such as cliques and small worlds, is an especially important avenue to explore.

As an experimental variation, we compare the hypercycle behavior of this spatial topology to that of the nonspatial “well-stirred liquid reactor” topology, which is more traditional in chemistry. In nonspatial or random topology, every rule is equally likely to pass a product to any other surviving rule, irrespective of the firm’s spatial location.

A major finding in the existing hypercycle literature (Hofbauer and Sigmund 1988, page 96) is that nonspatial hypercycles are dynamically stable up to four elements but not beyond. In other words, in nonspatial interaction when hypercyclic sets consist of five elements or more, one or more of the component chemicals is always driven to zero during the reaction process, thereby breaking the reproductive loop and causing the hypercycle to “crash.” This is a “complexity barrier” that self-organizing hypercycles, and hence “life,” cannot penetrate when the chemical interaction is nonspatial or random. Padgett (1997) has shown that in spatial interaction topologies, dynamically stable hypercycles with a complexity of five elements and more can be grown, although at increasingly lower frequencies at higher levels of complexity. Spatial interaction, in other words, can break the complexity barrier. Presumably this is one reason why complicated chemical life is embodied. We reconfirm both the nonspatial findings of Hofbauer and Sigmund (1988) and the spatial findings of Padgett (1997) here in a new context.[13]

3. *Learning/reproduction.* In the spatial topology setting, two variants of learning by doing are explored[14]:

    a. “Source reproduction” occurs when the originating rule in a successful transaction is reproduced.

    b. “Target reproduction” occurs when the receiving rule in a successful transaction is reproduced.

    For example, if $1 \rightarrow 2$ receives a 1 from the input environment, transforms it into a 2, and then successfully passes that 2 onto a neighboring $2 \rightarrow 3$ that transforms it again, then source reproduction occurs where the initiator $1 \rightarrow 2$ reproduces, and target reproduction occurs where the recipient $2 \rightarrow 3$ reproduces.[15] The variation in the mode of reproduction thus defines who benefits from the transaction.

---

[13] The models in this paper are extensions of the model presented in Padgett (1997). The main extension is the addition of explicit products that are being transformed. In the earlier paper, there were action-reaction chains of “play,” but nothing was actually produced or accomplished. We believe that the setup in Padgett (1997) was appropriate for modeling the emergence of informal organization among people within a firm, whereas the setup here is more appropriate for modeling trading among firms in an economy.

[14] In nonspatial interaction, these two reproduction modes behave identically. Space is what separates target from source. In Padgett (1997), a third mode was also explored: “joint reproduction,” where both rules in a successful transaction reproduce. Because two rules are reproduced in this hybrid, two offsetting skills need to be killed to preserve conservation-of-mass.

[15] Of course the recipient $2 \rightarrow 3$ could easily turn into an initiator the next time, if a neighboring $3 \rightarrow 4$ is subsequently found.

Source reproduction is considered “selfish learning” because the initiator of the successful transaction (like a teacher) reaps the reward. Target reproduction is considered “altruistic learning” because the recipient of the successful transaction (like a student) reaps the reward. Although “selfish” and “altruistic” accurately characterize who benefits, they are suggestive labels, and one should avoid importing motivational connotations. In the minimalist models developed here, there are no motivations — just actions and reactions, like in chemistry.

Padgett (1997) demonstrated that, in contrast to source reproduction, target reproduction dramatically increases the likelihood of growing stable hypercycles. It also increases the spatial extensiveness and complexity of the firm cluster that is produced by the hypercycles. Both of these findings are reconfirmed here.

In addition to these three experimental manipulations, two more experiments vary the input environment in which hypercycles grow. Such additional experiments were not possible in Padgett (1997), because previously there was no explicit modeling of products or of product environments.

4. *Input environment.* Input environments of resources or products can be conceived of as being fixed or variable and rich or poor.

   Among fixed resource environments are two types. Rich input environments are modeled by letting the input urn of resources contain all possible inputs, never to be depleted, even as products/resources are withdrawn. Poor input environments are modeled by letting the input urn of resources contain only one possible input (by convention, it is called “1”), which is not depleted, even as products/resources are withdrawn.

   Among variable resource environments is the endogenous input environment. This environment is modeled by letting the input urn be constructed over time by the outputs of the production system. Under the endogenous environment variant, our model withdraws one input product, transforms it into other products through distributed production chains, and then places the final output back into the original input urn.

   Presumably, rich input environments are more congenial to hypercycle emergence than are poor environments. What is less clear is where endogenous environments rank. Given that we have defined “rich” as being virtually nirvana (i.e., all possible inputs are available all the time, never to be depleted), our expectation is that nothing can outperform that. However, modelers of social insect behavior (e.g., Camazine et al. 2001) have discovered that “stigmergy” (i.e., the ability of social insects to transform their physical environments into nests, mounds, paths, and the like) can sometimes provide surprisingly powerful feedback that affects the development of the social organization itself. The open question, therefore, is whether the social organization that is achievable in an endogenous environment is superior in any way to that achievable in a rich environment.

5. *Input search.* The final experimental manipulation varies the precision of the search through the environment. A random search occurs when an activated rule reaches into the input environmental urn and chooses inputs randomly, in proportion to what is there. A selective search occurs when an activated rule reaches into the input environmental urn and selects the exact input it needs to transform, if it is there. A random search is like literal chemistry. (Metaphorically, we think of this as the intelligence of an atom, bouncing around.) A selective search is more like animal behavior. (Metaphorically, we think of this as the intelligence of a cow, looking for grass.) This is the only place in the model where we vary degree of intelligence. We expect the more intelligent selective search procedure to outperform the random search procedure in finding and nurturing production hypercycles.

FINDINGS AND DISCUSSION

The basic findings from our agent-based simulation model of hypercycle economies are presented in Figures 1 and 2. The y-axis represents the dependent variable: the long-term[16] probability of hypercycle survival. The x-axis represents the varying degrees of complexity in the simulated economies' technologies: simple two-skill technologies, slightly more complicated three-skill technologies, and so forth, up to the most complex nine-skill technologies. Different lines within the graphs present the results of our various experimental manipulations: interaction topology, mode of reproduction/learning, and input environments. Figures 1a and 1b show comparisons for selective search; Figures 2a and 2b show them for random search. The findings are discussed here.

1. Unstructured interaction topologies are not conducive to the emergence of complex technologies. Without help through embodiment, long sequences of skills cannot dynamically regulate their own stable reproduction. "Structured topology" does not have to mean spatial, as it does here (Cohen et al. 2001). But constraints on interaction are necessary for two reasons: (1) to break the symmetry of full mixing and induce localized heterogeneity and (2) to allow positive reproductive feedback to turn that raw heterogeneity into path-dependent memory of past successes. This is the chemistry answer to why

[16] Our operational definition of "long term" came inductively from observing a great many runs and seeing how long even the slowest among them took to converge to equilibrium. We finally chose the following liberal stopping points for our simulations: 30,000 ticks for two-element hypercycles, 40,000 ticks for three-element hypercycles, 60,000 ticks for four-element hypercycles, 80,000 ticks for five-element hypercycles, 120,000 ticks for six-element hypercycles, 180,000 ticks for seven-element hypercycles, 250,000 ticks for eight-element hypercycles, and 300,000 ticks for nine-element hypercycles. For the more complex of these hypercycles, much computing time was required. Each point in these graphs represents the average of 30 simulation runs.

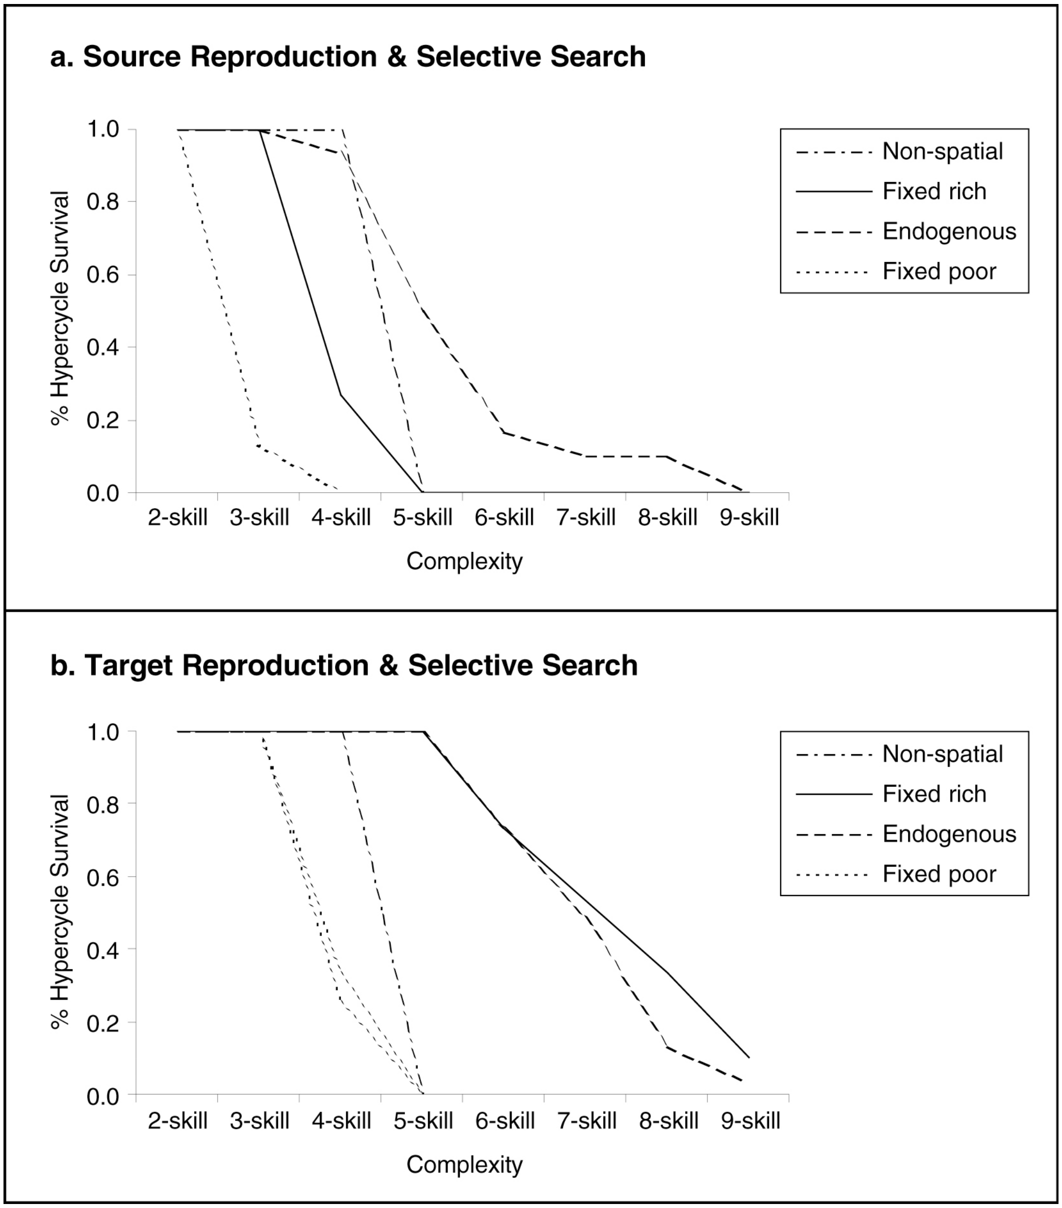

FIGURE 1 (a) Source reproduction and selective search and (b) target reproduction and selective search (Each point is an average of 30 simulation runs.)

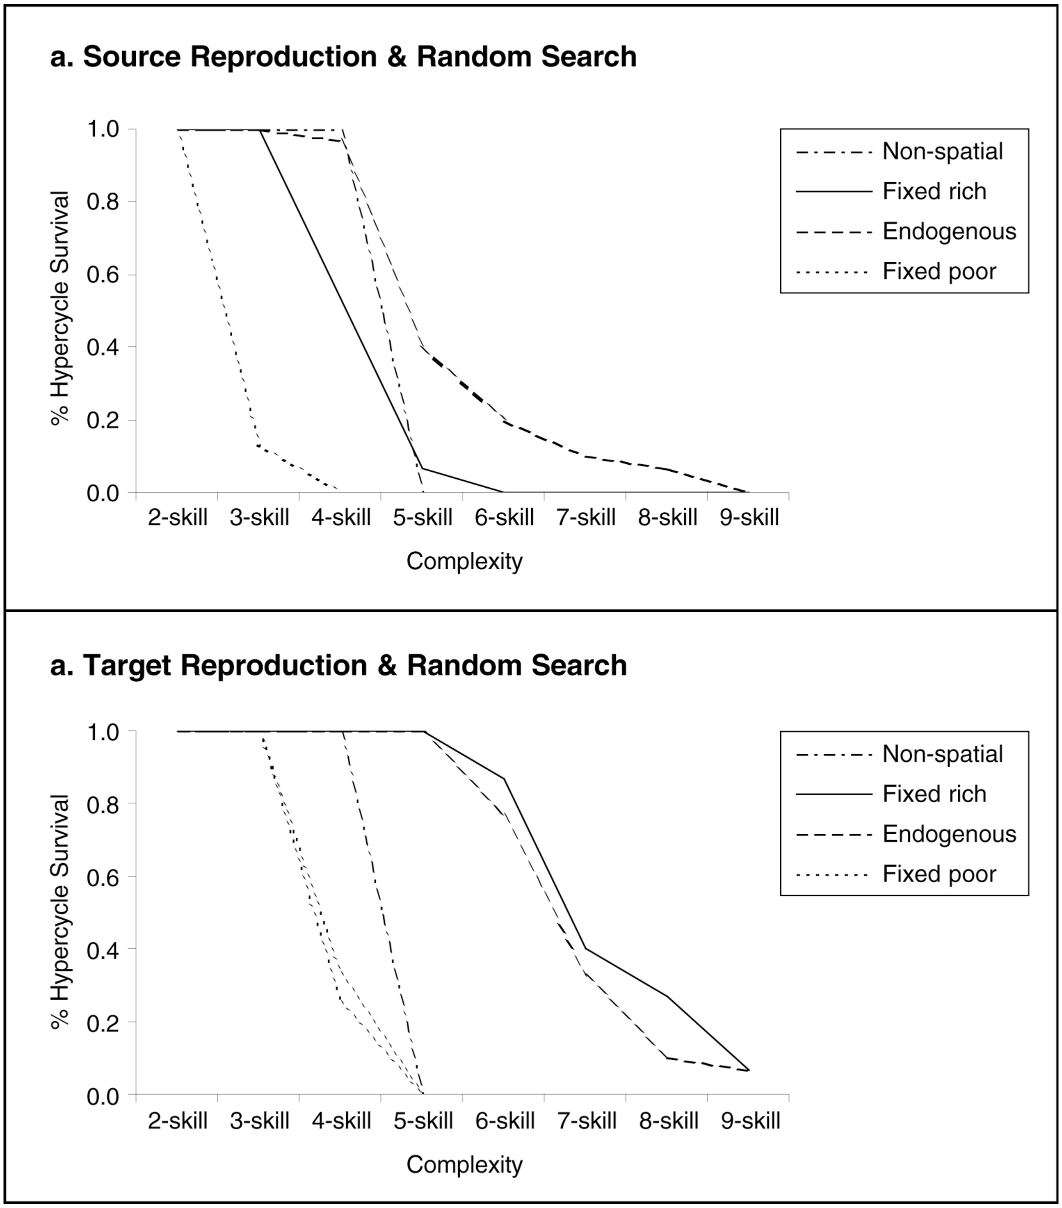

FIGURE 2 (a) Source reproduction and random search and (b) target reproduction and random search (Each point is an average of 30 simulation runs.)

firms exist:[17] dynamic barriers of technological complexity can be transcended, once “global” is transformed into the concatenation of “locals.” This finding is evident in Figures 1 and 2.

Classic Marshallian industrial districts receive the benefits of physical space naturally. However, in this era of globalization, densely interconnected firms may or may not be so fortunate. What our model implies is that trading within these new “virtual industrial districts” will have to become interactionally constrained in order for technological progress, and not instability, to be the consequence of increased connectivity (cf. May 1974; Davidow 2000).

2. The potential benefits of localized embodiment are more easily reaped through altruistic learning than through selfish learning. When the recipients rather than the initiators of transactions receive the reproductive rewards, complex technologies are more readily nurtured and repaired.[18] “Free-riding” does occur, but it does not threaten system stability. This repair mechanism can also be shown analytically for one special case of our model. We derived differential equations of skill growth for the extremely simple two-skill-hypercycle setting of a single dyad: two adjacent firms trading only with each other. Table 1 collates the differential equation results for ease of inspection. In our agent-based simulations, interlinked dyads proliferated across the entire grid, generating interaction effects not captured in the stripped-down dyadic setting. Simplification does, however, permit analytical solutions not otherwise possible. Such solutions are useful, both to increase transparency and to double-check computer code.

   The analytic contrast between target reproduction and source reproduction is most sharply illustrated in the setting of a fixed rich environment. There, in both of the target reproduction equations, $E(n_{12,t+1})$ always goes up when $n_{12,t}$ is less than $n_{21,t}$, and it always goes down when $n_{12,t}$ is more than $n_{21,t}$. The converses are true for $E(n_{21,t+1})$. In other words, target reproduction generates a consistent tendency toward homeostatic stability over the entire range of n_{12}. In sharp contrast, in both of the corresponding source reproduction equations, both $E(n_{12,t+1})$ and $E(n_{21,t+1})$ equal zero. Source reproduction exhibits no built-in tendency toward homeostatic stability: n_{12}

[17] Padgett (1997, pages 119-200) discusses why the traditional explanations for the firm given in neoclassical economics (namely, transaction-cost economics and principal-agent theory) are inadequate from a biological perspective: “Such a transposition of ‘the firm’ down into a series of dyadic negotiations overlooks the institutionalized autonomy of all stable organizations. In organisms, social or biological, rules of action and patterns of interaction persist and reproduce even in the face of constant turnover in component parts, be these cells, molecules, principals, or agents. In the constant flow of people through organizations, the collectivity typically is not renegotiated anew. Rather, within constraints, component parts are transformed and molded into the ongoing flow of action.”

[18] Sabel (1994) recommends squeezing the temporal distance between the two sides of an iterated transaction until this distinction is effaced. Such relational constraints are consistent with our first conclusion. Regarding our second conclusion, however, Padgett (1997) demonstrated that joint reproduction, the closest analog in chemistry to this recommendation, does not succeed in breaking complexity barriers.

TABLE 1 Growth-of-skill equations for the two-element hypercycle in a trading dyad

---

Transformation Structure

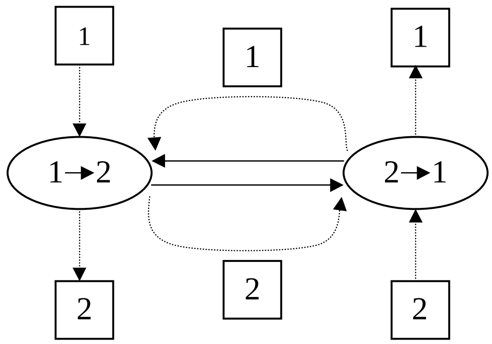

Spatial Layout

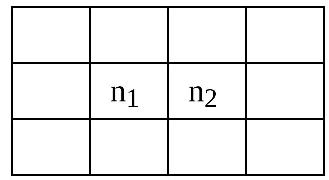

Volume of Rules/Skills: $n_{12} + n_{21} = N$
Volume of Balls/Products: $b_1 + b_2 = B$

Growth-of-Skill Equations

1. Nonspatial topology: Unchained

$$\frac{d}{dt}E(n_{12}) = \left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

2. Spatial topology: Fixed rich environment, with selective search

   (a) Source reproduction of rules

$$\frac{d}{dt}E(n_{12}) = 0$$

   (b) Target reproduction of rules

$$\frac{d}{dt}E(n_{12}) = \left(\frac{9}{64}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

3. Spatial topology: Fixed rich environment, with random search

   (a) Source reproduction of rules

$$\frac{d}{dt}E(n_{12}) = 0$$

   (b) Target reproduction of rules

$$\frac{d}{dt}E(n_{12}) = \left(\frac{17}{256}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

TABLE 1 (Cont.)

4. Spatial topology: Fixed poor environment, with either selective or random search

(a) Source reproduction of rules

$$\frac{d}{dt}E(n_{12}) = \left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[8+\left(\frac{n_{12}}{N}\right)\right]$$

(b) Target reproduction of rules

$$\frac{d}{dt}E(n_{12}) = -\left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)^2\left[8+\left(\frac{n_{12}}{N}\right)\right]$$

5. Spatial topology: Endogenous environment, with selective search

(a) Source reproduction of rules

$\frac{d}{dt}E(n_{12}) = 0$, as long as $b_1 > 0$ and $b_2 > 0$

(b) Target reproduction of rules

$\frac{d}{dt}E(n_{12}) = \left(\frac{9}{64}\right)\left[\left(\frac{n_{21}}{N}\right)-\left(\frac{n_{12}}{N}\right)\right]$, as long as $b_1 > 0$ and $b_2 > 0$

6. Spatial topology: Endogenous environment, with random search

(a) Source reproduction of rules

$$\frac{d}{dt}E(n_{12}) = \left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[\left(\frac{b_1}{B}\right)-\left(\frac{b_2}{B}\right)\right]\left[9-\left(\frac{b_1}{B}\right)\left(\frac{n_{21}}{N}\right)-\left(\frac{b_2}{B}\right)\left(\frac{n_{12}}{N}\right)\right]$$

(b) Target reproduction of rules

$$\frac{d}{dt}E(n_{12}) = \left(\frac{1}{64}\right)\left[\left(\frac{b_2}{B}\right)\left(\frac{n_{21}}{N}\right)^2-\left(\frac{b_1}{B}\right)\left(\frac{n_{12}}{N}\right)^2\right]\left[8+\left(\frac{n_{12}}{N}\right)-\left(\frac{n_{12}}{N}\right)^2\right]$$

7. Spatial topology: Endogenous environment, with either selective or random search

(a) Either source or target reproduction of <u>balls</u>

$$\frac{d}{dt}E(b_1) = \left\{\left(\frac{7}{8}\right)\Big/\left[1-\left(\frac{1}{8}\right)^2\right]\right\}\left[\left(\frac{n_{21}}{N}\right)-\left(\frac{n_{12}}{N}\right)\right]$$

drifts in random-walk fashion until it eventually crashes into the absorbing states of either $n_{12} = 0$ or $n_{12} = 1$.

Here we compare analytic results with simulation outcomes to verify that our model is behaving as expected. For hypercycles more complex than the two-element dyad, we can no longer derive solutions analytically. But the simulations show this basic dyadic finding to be true more generally. Target reproduction generates higher rates of hypercycle survival than does source reproduction for all corresponding spatial settings. To repeat, the mechanism that generates this sizable superiority is the direct "altruistic" repair of complementary rules by each other. Target reproduction repairs hypercycles without intending to do so, once given the precondition of spatial "symmetry breaking," which induces the distinction between altruistic and selfish in the first place.

This conclusion is consistent with anthropological emphases on gift giving in primitive economies (Mauss 1967, Sahlins 1972). It is also consistent with sociological observations about the "strange" persistence of generous behavior in modern economies (Macauley 1963; Granovetter 1985; Uzzi 1996, 1997; Padgett and McLean 2002). Our explanation for generosity may not be the only explanation. However, repair is one evolutionary reason for the natural selection of this behavior in competitive economies of all sorts. Altruistic learning stabilizes the reproduction of distributed technological skills, on which all depend.[19]

3. When altruistic learning is not present for whatever reason, then stigmergy (the endogenous construction of resource environments) is second best. Entomologists (e.g., Bonabeau et al. 1999; Camazine et al. 2001) have shown that stigmergy can flexibly coordinate sophisticated collective behavior among myopic social insects. We have shown that stigmergy also can regulate the cancerous growth of selfish learners, keeping even long chains of distributed skills alive.

   Adams (1966a,b) has long argued that cities are crucial to the history of technology. His exemplar case is Mesopotamia, where spatial feedbacks between settlements and rivers guided the joint emergence of urban concentrations, irrigation technologies, and the shapes of the rivers themselves.[20] Even though our model is far too minimalist for real history, it may illustrate one reason why the spatial reorganization of land into cities and the development of complex technologies proceeded hand in hand. Technology causes cities, as we all know; less obviously, the spatial products

---

[19] This may be news to some rational choice theorists, but it will not come as a surprise to parents.

[20] One would not expect the mechanisms of our model to explain the invention of writing. But it is worth noting that writing, too, was implicated in these co-evolutionary developments.

of technology channel and regulate the social forces that produce them. To put it simply: cities stabilize selfishness.[21]

CONCLUSIONS

We imported a few simple tools from chemistry and developed them for systematically investigating the co-evolution of distributed technology and social organization. Extreme assumptions about the absence of consciousness are implied by our specification. The payoff of such extreme simplification is the discovery of three social-organizational principles that enable technological evolution. How robust such principles are when applied to alternate specifications is important to explore in the future. Regardless of the answer to that question, however, we have demonstrated, at a minimum, that complex cognition is not necessary for the emergence and functioning of complex economies, just as March and Simon (1958) argued long ago.[22]

REFERENCES

Adams, R.M., 1966a, *The Evolution of Urban Society: Early Mesopotamia and Prehispanic Mexico*, Chicago, IL: Aldine.

Adams, R.M., 1996b, *Paths of Fire: An Anthropologist's Inquiry into Western Technology*, Princeton, NJ: Princeton University Press.

Antonelli, C., and R. Marchionatti, 1998, "Technological and Organizational Change in the Process of Industrial Rejuvenation: The Case of the Italian Cotton Textile Industry," *Cambridge Journal of Economics* **22**:1–18.

Boerlijst, M.C., and P. Hogeweg, 1991, "Spiral Wave Structures in Pre-biotic Evolution: Hypercycles Stable against Parasites," *Physica B* **48**:17–28.

Bonabeau, E., M. Dorigo, and G. Theraulaz, 1999, *Swarm Intelligence: From Natural to Artificial Systems*, Oxford, U.K.: Oxford University Press.

Camazine, S., J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau, 2001, *Self-organization in Biological Systems*, Princeton, NJ: Princeton University Press.

Cohen, M.D., R.L. Riolo, and R. Axelrod, 2001, "The Role of Social Structure in the Maintenance of Cooperative Regimes," *Rationality and Society* **13**: 5–32.

---

[21] Given the absence of cities in the prehistory of mankind, perhaps one is as justified in speculating about the evolution of selfishness as one is in speculating about the evolution of altruism. More reasonably, we should support the efforts of the evolutionary economists cited at the beginning of this article and other more historically oriented economists (e.g., Herrigel 1996; Sabel and Zeitlin 1997) who try to specify alternate social-organizational paths to technological complexity.

[22] March and Simon (1958) went beyond our point to also argue that social structures enable human cognition, by mapping the world down to levels we can comprehend.

Coriat, B., and G. Dosi, 1998, "Learning How to Govern and Learning How to Solve Problems: on the Co-evolution of Competencies, Conflicts and Organizational Routines," in A. Chandler, P. Hagstrom, and O. Solvell, eds., *The Dynamic Firm: The Role of Technology, Strategy, Organization and Regime,* Oxford, U.K.: Oxford University Press.

Davidow, W., 2000, *Social Tornadoes: Interconnection, Instability, and Social Policy*, unpublished manuscript.

Dosi, G., R. Giannetti, and P.A. Toninelli, 1992, *Technology and Enterprise in a Historical Perspective,* Oxford, U.K.: Clarendon Press.

Dosi, G., R. Nelson, and S. Winter, 2000, *The Nature and Dynamics of Organizational Capabilities,* Oxford, U.K.: Oxford University Press.

Eigen, M., 1971, "Self-organization of Matter and the Evolution of Biological Macromolecules," *Naturwissenschaften* **58**:465–523.

Eigen, M., and P. Schuster, 1979, "The Hypercycle: A Principle of Natural Self-organization," New York, NY: Springer-Verlag.

Fontana, W.F., and L.W. Buss, 1994, "The Arrival of the Fittest: Toward a Theory of Biological Organization," *Bulletin of Mathematical Biology* **56**:1–64.

Granovetter, M., 1985, "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology* **91**:481–510.

Hayek, F., 1948, *Individualism and Economic Order*, Chicago, IL: The University of Chicago Press.

Herrigel, G., 1996, *Industrial Constructions: The Sources of German Industrial Power*, Cambridge, U.K.: Cambridge University Press.

Hofbauer, J., and K. Sigmund, 1988, *Dynamical Systems and the Theory of Evolution*, Cambridge, U.K.: Cambridge University Press.

Hughes, T.P., 1983, *Networks of Power: Electrification in Western Society, 1880–1930*, Baltimore, MD: The Johns Hopkins University Press.

Hughes, T.P., 1987, "The Evolution of Large Technological Systems," in W. Bijker, T. Hughes, and T. Pinch, eds., *The Social Construction of Technological Systems*, Cambridge, MA: MIT Press.

Hutchins, E., 1995, *Cognition in the Wild*, Cambridge, MA: MIT Press.

Jain, S., and S. Krishna, 1998, "Autocatalytic Sets and the Growth of Complexity in an Evolutionary Model," *Physical Review Letters* **81**:5684–5687.

Kauffman, S.A., 1986, "Autocatalytic Sets of Proteins," *Journal of Theoretical Biology* **119**:1–24.

Kauffman, S.A., 1993, *The Origins of Order: Self-organization and Selection in Evolution*, New York, NY: Oxford University Press.

Landesmann, M., and R. Scazzieri, 1996, *Production and Economic Dynamics*, Cambridge, U.K.: Cambridge University Press.

Macauley, S., 1963, "Non-contractual Relations in Business: A Preliminary Study," *American Sociological Review* **28**:55–67.

Malerba, F., and L. Orsinigo, 1993, "Technological Regimes and Firm Behavior," *Industrial and Corporate Change* **2**:45–71.

March, J.G., and H.A. Simon, 1958, *Organizations*, New York, NY: Wiley.

Maturana, H., and F. Varela, 1980, "Autopoesis: The Organization of the Living," in *Autopoesis and Cognition*, Dordrecht, Netherlands: D. Reidel.

Mauss, M., 1967, *The Gift: Forms and Functions of Exchange in Archaic Societies*, New York, NY: Norton (translation of original work published in French in 1899).

May, R.M., 1974, *Stability and Complexity in Model Ecosystems*, Princeton, NJ: Princeton University Press.

McKelvey, M., 1997, "Co-evolution in Commercial Genetic Engineering," *Industrial and Corporate Change* **6**:503–532.

Nelson, R.R., 1994, "The Co-evolution of Technology, Industrial Structure, and Supporting Institutions," *Industrial and Corporate Change* **3**:47–64.

Nelson, R.R., 1995, "Recent Evolutionary Theorizing about Economic Change," *Journal of Economic Literature* **33**:48–90.

Nelson, R.R., and S.C. Winter, 1982, *An Evolutionary Theory of Economic Change*, Cambridge, MA: Belnap Press of Harvard University Press.

Padgett, J.E., 1992, "Learning from (and about) March," *Contemporary Sociology* **21**:744–749.

Padgett, J.E., 1997, "The Emergence of Simple Ecologies of Skill: A Hypercycle Approach to Economic Organization," in W.B. Arthur, S.N. Durlauf, and D.A. Lane, eds., *The Economy as an Evolving Complex System II*, Reading, MA: Addison-Wesley.

Padgett, J.E., and P. McLean, 2002, "Economic and Social Exchange in Renaissance Florence" (unpublished manuscript).

Powell, W.W., 1996, "Inter-organizational Collaboration in the Biotechnology Industry," *Journal of Institutional and Theoretical Economics*, **152**:197–215.

Prigogine, I., and P. Glansdorff, 1971, *Thermodynamic Theory of Structure: Stability and Fluctuations*, New York, NY: Wiley.

Rosenkopf, L., and M.L. Tushman, 1998, "The Co-evolution of Community Networks and Technology: Lessons from the Flight Simulation Industry," *Industrial and Corporate Change* **7**:311–346.

Sabel, C.F., 1994, "Learning by Monitoring: The Institutions of Economic Development," in N.J. Smelser and R. Swedberg, eds., *The Handbook of Economic Sociology*, Princeton, NJ: Princeton University Press.

Sabel, C.F., and J. Zeitlin, eds., 1997, *World of Possibilities: Flexibility and Mass Production in Western Industrialization*, Cambridge, U.K.: Cambridge University Press.

Sahlins, M., 1972, *Stone Age Economics*. New York, NY: Aldine.

Stadler, B.M.R., and P.F. Stadiler, 2002, "Molecular Replicator Dynamics," Working paper 02-09-049, Santa Fe Institute, Santa Fe, NM.

Uzzi, B., 1996, "The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect," *American Sociological Review* **61**:674–698.

Uzzi, B., 1997, "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness," *Administrative Science Quarterly* **42**:35–67.

Warlien, M., 1995, "Hierarchical Selection and Organizational Adaptation," *Industrial and Corporate Change* **4**:161–186.

DISCUSSION:

NETWORK COMPLEXITIES

(October 4, 2004, 9:45 to 11:45 a.m., Session 2)

Chair and Discussant: *Fabio Rojas, Indiana University*

Empirical Estimation and Multi-agent-based Simulation of a Discrete Choice Model with Network Interaction Effects: An Exploration of the Role of Global versus Local and Social versus Spatial Networks in Transportation Mode Choice Behavior in the Netherlands

Fabio Rojas: Originally Shah Jamal Alam of Saarland University was scheduled to speak today. He is not able to come, so today we'll be having David Sallach in his place. And our new schedule for today will start with Elena Dugundji and Laszlo Gulyas, who will be talking about empirical estimation and multi-based simulation of a discrete model, choice model. Then we'll have John Padgett, *Doowan* Lee, and Nick Collier talk about economic production as chemistry, and then we'll finish up with David Sallach, who'll be talking about indexicality in agent-based modeling.

Elena Dugundji: This is work together with myself and Laszlo Gulyas from the Hungarian Academy of Sciences. I'm at the University of Amsterdam.

[Presentation]

Unidentified Speaker: Would you just repeat what you mean by endogenous specification?

Dugundji: No. This is within transportation mode choice. For the first example I showed you, I used one thick set of perimeters. And I applied that set of perimeters to all the same networks. For the others, I let, for each network, the perimeters be re-estimated. So how you could think about that intuitively, for econometricians in the group, would typically be to do what I'm showing you right here, that you would do your estimation based on the data, in this case, this created network for each individual network.

The other case would be something like you had estimated your data and then there was an exogenous shock to the system. The system didn't have a chance to recalibrate itself, and therefore it operated on the basis of the coefficients that you had estimated in the first pass.

So this is what is done in the current transportation work. You do some either cross-sectional or dynamic estimation of your data, and then for the data when you make a prediction, you change the X variable slightly and you use the betas that you had estimated as input. That's what you just saw.

Now I'm allowing them to be re-estimated every time, and here you get very different behavior. And I will come back to that in the conclusions.

[Presentation Concludes]

Rojas: Our next presenter is Doowan Lee. While he is setting up his laptop**,** does anybody have a quick question …

Unidentified Speaker: So basically your findings, the conventional modeling approach for the discrete choice, I mean, the conventional logic property model, what kind of econometrical problem were [they] having in their models? For example, would the coefficient, the responsible coefficient, be unstable?

Dugundji: The issue with the current models is that they're not taking this interaction effect into account. And in particular, adding this coefficient automatically adds a dynamic to the model, because as soon as your choice is dependent on other people's choices, once you've estimated your model, then you have a feedback immediately, because people might have changed their choice. And then that changes the explanatory variable, which is why you get this feedback automatically. In these simple examples without the negative feedback, there are only two options: You either get a bell-shaped curve or you get this two-peak behavior, where one option is chosen or the other one. Adding a nonlinear negative feedback would be very important, because clearly this type of behavior is maybe not so realistic. What you would expect to see is maybe some more limit cycle type of behavior. And we expect that adding this nonlinear negative feedback might be able to give us this limit cycle behavior, which is what you see more in real life.

Now, the issue with the typical current models that they don't include this interaction could have very important effects on your implementation of a new policy. If this is an important variable, which we think it is, and you're leaving it out of policy decisions, you're basically not taking into account the fact that society is influencing your choice. So your choice is not just an individual decision that doesn't affect what your neighbors do.

Clearly, for car-pooling policy to have any kind of effect at all, you really do have to be able to take into account other peoples' decisions. I mean, this is a very extreme example, because if you want to car-pool, you can't just say, "Oh, I'm going to car-pool" if there isn't anyone in your environment or workplace that also wants to car-pool. So your choice is critically dependent on other people.

You can expect that car-pooling models will certainly be incorrect if you're leaving this out of the model. So that's the rationale for wanting to include these types of variables.

Economic Production as Chemistry

Rojas: Doowan Lee is from The University of Chicago.

Doowan Lee: This work is a sequel to John Padgett's 1997 article on hypercycles. In that paper, he was trying to see what kind of conditions are most conducive to the emergence of hypercycles, or the emergence of skill-based organizations. There he had skills and agents trying to find some compatible neighbors and then producing some kind of hypercycles or a chain of organizations.

[Presentation]

Unidentified Speaker: A skill is like one to two, two to three, so how long is the chain?

Lee: Actually, we have only binary rules at this point. So each rule will have only two integers, input and output. In fact, increasing the chain of the rules is an extension we are thinking about for the next project.

Unidentified Speaker: What's the complexity?

Lee: Two, tonight. For example, if you set your simulation to two, you'll have only two types of skills, right? One to two, two to one. If you set the integer at three, you'll have one to two, two to one, three to one, and so on. So at this point, our complexity is limited to binary rules, as opposed to increasing the chain of those rules.

Unidentified Speaker: So the complexity is just the number of rules.

Lee: So this setting, in terms of the complexity of rules, is called "Solo H," because we're going to allow only one-way permutation, as opposed to back two, three to one, two to one, etc. So that's what we're using here. In terms of width size, it's set at 10 by 10, so there will be 100 agents. The interaction mode, as I said before, is more neighborhood. So each agent will have eight neighborhoods in which to look for a compatible form. The number of goals is set at 200 throughout all our experimental settings. There isn't a set of 200. Instead we have to see whether the variation of the volume of each type of rule has an impact on the probability of hypercycles.

Probably this has to be more dynamic in the future. For example, when you have more types of skills, perhaps it might make sense to increase the total volume of goals or rules allowed in the system.

So, for example, when our maximum goal type is at full, we're going to set the number or rules at each type at 50, so we'll have 200 total volumes. Target production is that the recipient will reproduce its rules and in terms of the resource mode, once. This has sort of been a very harsh environment, because I started with only ones, or since there'll be a lot of different types of rules. Even though you start with only ones, the products will be very diverse as it goes on. So this is the basic setting.

Unidentified Speaker: What do the numbers in each circle mean again?

Lee: Those are rules, actual rules. So when you see two/three, that means it's transforming two or three. And whenever there's a kind of successful transformation, the green lights indicate they are networked.

Rojas: This is headed for death.

Lee: Hmm, no. This is just a first pass. Of course, we'll show you more, some of the statistics of our experiments that will better represent the findings we got out of the simulation.

What I'm trying to show you here is that the simulation is actually doing what we intend it to do, and actually there's no pattern with our analytic results, where we use only two types of rules and then try to show that our findings are analytically attainable as well. And that's I think what John mentioned previous in the morning — you start with a simple setting and increase the complexity of your setting and see whether you have a strong analytic foundation to increase the complexity, because if you go the other way around, you might have very interesting findings, but it's really difficult to show whether those findings are analytically attainable as well.

So this is a basic test run of our model. I'm going to show you our statistics later, but these are two snapshots. The first one shows you the complexity of skills is at full, is endogenous, a rich environment that's set for target reproduction. The bottom one has the same settings, except it's source reproduction. As you can see, there's a drastic difference between these two snapshots.

[Presentation Concludes]

So any questions, comments?

Rojas: Does anybody have a question?

Shu-Win Chen: Could you define the emergence of form again?

Lee: The forms are not the basic unit of our analysis at this point. They are given. For example, we start with a 10×10 grid, and then you'll have 100 spots, and every spot is populated with bins. And those bins will have randomly distributed skills, depending on the integer we defined for the complexity of skills in the system. So in that sense, forms are completely given in this setting. What we're really interested in is, whether these conditions form any types of industrial districts, if you will, or some kind of production chains, like a fundamental dependant variable. I don't think we actually study the concept of the form here with this project.

Stephen Guerin: Right now, the production rules are one input/one output. Have you considered more chemistry of multiple inputs/multiple outputs?

Lee: Actually, we haven't. There are several reasons we stayed with this very simple permutation at this point. First we were aiming to finish this paper in our given amount of time, so we decided to start with the most simple setting, which is just one-way permutation, although we have, as you see in the parameter window, possible permutations option. This is a little bit more complex, but the problem is it'll just take a lot of computer power to process those kinds of settings. So that is something that we have in mind, but in terms of whether we have implemented multiple inputs and multiple outputs, no. But that's something probably we can consider in the future.

Rojas: How long did it take both of you to do your simulations?

Lee: I really appreciate the question, because that that's when we were sort of pulling our hair a lot. In terms of the necessary computer power, it depends on the complexity of the skills, a two or three. It'll take just one or two days to go through all permutations of our experiment of variables. But when we went up to nine, I think one setting took about three days on a

Pentium IV, in a 2.4 gigahertz computer with 1 gigahertz of memory. So I think for our old settings, like from two to nine integers, I think it took about two months on this one computer.

Rojas: Right. Elena, how long did it take you to do your simulations?

Dugundji: It's very comparable. I don't know if you remember, you saw lots of and lots of little histograms. An individual line is quick, but when you want to scan over 35 densities, times four different modules, times 50 networks per histogram, times 20 different networks, then you have thousands and thousands of runs, and I think it took about three months just to run through all of them.

Rojas: And you were using similar hardware as Doowan, like a good desktop?

Dugundji: Yes, 1 gigabyte of memory and 2.4 gigahertz. And it was just continuously running batch mode for about three months. An individual run is very quick, but to do all the different permutations takes a long time.

Rojas: So both of these projects are similar in that way, that the individual simulations are pretty short, but going through all the parameters of the model actually …

Dugundji: Particularly because, one, just doing one network isn't sufficient. You really want to do at least 20 networks, and then you want to use 50 different random seeds per network. Then you need 2,000 iterations until you get to the steady-state behavior.

Lee: And that's not counting all the erroneous [runs].

Rojas: We're going to hear from David Sallach of Argonne National Laboratory and The University of Chicago.

INDEXICALITY IN AGENT-BASED MODELING[1]

David Sallach: Because members of society must actually have shared methods for achieving social order, and every situation has a different pattern of order that is required for a coherence of action within that situation, there are characteristics that are reflexive and relational in the production of meaningful social practices. One of the ways that you recognize the coherence that's inherent in the situation is in terms of taken-for-granted expectations.

So my point is that this is a pretty high standard for agent simulation, to begin to do that. And I don't mean to suggest that we come anywhere close to approximating that. Infinitely deep is a long way! But I do mean to say that it's useful to review this as a standard and then think "to what extent can these capabilities be approximated?"

Situation theory can be viewed as a variety of information theory. Information theory, as you know, is quantitative. It has to do with the size of information flow, the size of information

---

1 David Sallach's presentation was a late addition to the program to replace the planned presentation by S. Alam, who did not attend the conference. Because there is no corresponding paper for Dr. Sallach's presentation during this session, his comments are provided in full.

store, the error rate, and things like that. The original Shannon and Weaver documents recognize the importance of meaning, but they also recognize that it was a tough problem and that it wasn't something that they were prepared to address at that point in time. But as we've seen, all social processes are inherently meaning-oriented; that is, why people do what they do at a given point in time is dependent upon the meaning that they attribute to the situation. There are continuous and countless examples of this: was that remark, joke, or insult that, in the right situation, could lead to war or peace?

There's a famous Cossack wink that I cite, which is that at the beginning of the 1917 Revolution in Russia, the women were out demonstrating for bread, and one of the Cossack troopers winked at them. The women they interpreted this to mean that they weren't going to ride them down this time as they had in the past, and this event was followed by massive demonstrations and, ultimately, the February Revolution.

Now, if he was really just trying to flirt, that's a meaning attribution that got out of control. Also, just think of all the attributions that go on about why the market was up or down today: "Well, there was nervousness about crop failure in Australia." "Well, the market is down today, but it was just profit-taking." I'm not saying such interpretations have no substantial basis, but I am saying that there are thousands of actors that make their own interpretations. All these interpretations do not necessarily manifest a consensus. On the contrary, in a sense, each is using his or her own member methods. And so these kinds of processes are really what we talk about when we consider social processes.

Situation theory gives us the potential to address the *meaning* of information. It gives us a formalism that's strongly contextual, supports indexical processes, and maybe can be used to talk about how social outcomes are negotiated and how locally effective meaning attribution ends up creating a basis of attraction for action, if you want to think of it in system dynamics terms. And, with this formalism, we can relate action to social theories, but also to information theory, to complexity theory and to systems theories.

So let us just quickly go over what situation theory is. Situation theory is a formalism for meaning, or sometimes it's said it's a formalism for semantic content, or for information flow. In situation theory, situations are first-class objects that are also "rich." One of the things that distinguishes situation theory is that there is, within situation theory, a formal object called a situation. Therefore, it can be reasoned about, it can be parsed, it can be communicated, inferences can be drawn about it, and so forth. The fact that it's logical allows it to be used in formulaic expressions. The fact that it's "rich" means that it is built into situation theory so that a situation cannot be fully defined or described. And the nice thing about that is that it's a formal object that's also open-ended. It's a formal object that is reasonable to use it as part of a social process. Then there are resources that can be invoked, and you can learn something about a situation that you didn't know before. And if you learn something about a situation that you didn't know before, that may change your motivation, your impulse to action, and so on. It is only a formalism, so in that sense there's no particular theory associated with it, and so this kind of discussion concerns what kind of substantive theories would be most appropriate to associate with situation theory. But it does provide a formalism for reasoning about context and situations within theory, where it is presumed that situations are individuated or intuited by the agent, and where "individuated" means "brought to consciousness" (in this case, perhaps there's an internal model of the situation), and where "intuited" means that the agent may respond in the context of a situation without fully cognizing it. Thus, for example, a dog may love to run in the woods without having a concept of the woods.

There was originally an ecological aspect of this, saying that information exists prior to language. In the unit of study as the organism situated in its environment, meaning is located in the interaction of living things and their environment. This kind of slides over into biosymbiotics: that organisms anticipate events and, thereby, thrive. Organisms can survive, thrive and adapt better by categorizing situations, which is one of the reasons we have to have this skill — at least that's the argument — and that situation recognition thereby facilitates adaptation. So, you see, it arises from a kind of ecological realism.

However, that's a philosophical point, and there are competing philosophical points that can be made. Barwise and Perry in the ecological interpretation say that meaning is discovered by animals, not invented or created by them. And that's a nice simple basis to start from, and maybe a good focus for certain kinds of simulations. But if we want to invoke Peirce, we say truth is always "relative to a context of interpretation," and I think that situation theory can be applied within that context as well. That's more or less what Terry Winograd did as part of a situation-theoretic discourse, saying that ultimate grounding is in "the potential of continued coherent discourse," and, thus, we take the attribution of reality as provisional. The nice thing there is, whatever your philosophical orientation, if you're a constructionist, a realist, either way, you're not precluded from utilizing situation theory.

One interpretation of meaning within situation theory is that it's a relationship between two situations. You can particularly see this for causal relationships. If you have a situation in which a seismograph shows a certain pattern, maybe your next situation is that you're going to have an earthquake. Thus, the meaning of the first situation is determined by what it predicts about the second situation. But this is by no means an exhaustive definition of situational meaning. I think that from a social standpoint we're better off at looking at the way meaning is interactively constructed, but state transition does suggest one way that situated meaning can be interpreted.

As a formalism, situation theory is actually much like a relational model and relational database where you can have a series of arguments (technically infons) that define a situation. But, there is an exception: situation theory adds polarity, where polarity indicates truth. The interesting aspect of this relative to formal logic is that in the latter truth is viewed as syntactic and universal, whereas in situation theory, it's subject to circumstance. Something is true under one circumstance when an event or a change occurs it's no longer true. It's also something that is attributed by an actor, and so you're able to express situated actor disagreement as to what is true rather than to (only) indicate universal truth.

Beyond semantic indexicality, I want to emphasize indexicality of action. It is quite striking how action is subject to the same types of interpretation that semantics are. The meaning of an action depends on who the actor is, when and where the action takes place, actor connections, resource situations, what resources are available, and what can be brought into play in particular circumstances. And it's efficient in the sense that the same actions mean different things because they're taken by diverse actors in distinct ways, in different times and places. That's part of what one would like to capture.

I should also talk about the social aspect of indexicality. There's an interpretive indexicality, in which the meaning of a given act or, of a given communicative act, is subject to the particular circumstances in which it's produced. And so we hear people say that something is true, but sometimes we take that as information about what's going on: they're in a room that you're not in, they're talking to you on the telephone, they tell you what's going on, and from that you infer what's happening in that room. But other times, we take what they say as information about how accurately they *can* see, whether they're physically able to see, but also whether they bring some biases to the situation, so you can then take their report as an indication about their biases. Which of those interpretations you use is very much an indexical property.

Information is indexical in some fundamental ways. One of the points that I wanted to make relative to the discussion we had yesterday with Keith Sawyer and David. There was an undertone in that conversation that indexicality was something esoteric, that every now and then it pops up, especially in exotic cultures, but we don't regularly use the concept to handle situations. But actually indexicality is used in discourse about every five seconds: every time we use a pronoun, he, she, we or they, it's instantiated by context. The same thing happens every time we say "here" or "there," "then" or "now." Our entire communicative apparatus is suffused with indexical interpretation.

Randall Collins makes a point that social structure is also indexical. Property is indexical. And this is a direct quote, "It mostly takes the form of here and now, inexpressible in general, except by pointing to a concrete situation; one points implicitly with one's body when one occupies, with a sense of appropriating it, one's own home, one's place of work and so forth. This is usually all that it takes, not a lot of high-level processing to keep the macro-structure intact."

There are millions of situations that happen every day. From childhood on, we remember and categorize dozens or hundreds of situations a day, and we might think about what methods or skills we have available to us (methods in a programming sense, skills in a human sense), to facilitate that categorization. Then we can see the situational interpretation as a basic skill.

So now let us turn to action selection. I've noted that situational interpretation is a basic skill that we use all the time, that we might like to provide our agents with deeper capabilities, and that situation theory provides one mechanism for doing that. Joanna [Bryson] talked yesterday about action selection mechanisms, which *are* being used in agents, and whereby there is the possibility of interruption or a change of focus that requires a reprioritization. So it fits very naturally into attempting to implement a situational focus. We can say in general that an agent has a field of orientation. You could have a cognitive field, but we should also add an emotional field as well. Intentionality is multi-focal and attention fluctuates among focal points and purposes.

This is just a quick view of Joanna's multi-tier reactive hierarchy, where you've got drives that are constantly competing for predominance, and you may in fact find ways of combining multiple needs simultaneously. You have competences that are context-relevant. When a given context occurs, that competence arises as relevant. And then, from a programming point of view, you have action sequences that are not subdivided, like strong habit or unconscious skills. They just kick in and operate, giving the system better performance capabilities.

So reactive plans or structures that specify subsequent contextual acts, mean, in particular, that a mechanism is set up so that it handles interruptions, combinations of alarms, requests, and opportunities. And when that happens, it may require a current sequence to be reordered; it may require shifting to a different plan.

Relative to this basic reactive plan, I want to call your attention to a priority that's used so that information can be compared to other priorities, a releaser. This is the case where you have to be in a particular context to be able to accomplish something. I mean, if the action sequence is to eat, then you must be at a place where there's food. If there's not food, then that action sequence would not come into play. That's the role that the releaser plays. (And then you have a number of retries and other implementation details.) And I'm just pointing out how that information is stored, and for now just take a note that this is very similar to the structure of an infon, so that there's a potential of this working closely in a situation theoretic context.

And then there is the outer drive, which is also very similar. They have a priority and a releaser and a current action. They keep track of what action is going on, but it's not their action. The basic plan, that action is part of the plan. This is just saying, "Well, no, where are we at? What are we doing here?" But unlike a basic reactive plan, you can have multiple drives that are effectively active simultaneously.

From this, we can make a general point about the representation of situated action, and a strategy for implementing situated action. The situation-theoretic infons, which we've seen before, and the action selection tuples, which we just considered, bear some interesting character similarities to the relational data model. The relational model also deals with relations where you have a named relation that carries a number of arguments, each containing information. However, in some ways, the relational model is much richer in terms of what you can do with it, because it gives you operators, so you can SELECT, PROJECT and JOIN, which gives you not only ways of combining information, but ways of creating sets of richly documented information, generating them on the fly. It gives you constraints so that you can make sure that it's an entity of a given type before a particular operator would be effective in relation to it.

I think that truth table may be too high [referring to a slide], but maybe it should be down below RM/T, but Codd has then extended it to do semantic modeling. And what is happening here is that it supports graph-theoretic relations. I'm talking about the ability to build networks here, so there actually is an arguable basis for being here. But so that you can build type/subtype relations, aggregation relations, event precedence relations (which I think are very interesting in this context), universal predecessors, universal successors, alternate predecessors, alternate successors, and so forth. And therefore you can have dynamic data structures that enable situated responses. And so one way of approaching this would begin to try to say, if you have the richness of the relational model and all that's defined and immediately available, then how can it be best utilized? If you began to articulate situation and use what I would call orientation theoretic operators using that as a model, if you begin to have some rich operators indeed. It might then be possible to build a fairly high-level declarative environment that would support situated action.

You could use that environment, as I've already illustrated. That's the definition of situation theory. It would be very straightforward to simply implement that using the relational model, especially using RM/T with its graph theoretic operators, and then have them completely integrated.

I'd like to give you a concrete example, so here is an attempt to modify or represent social structure, in particular social slavery in the U.S. in 1850 and 1870. It actually represents the abolition of slavery. You have Situation 1 and Situation 2. It is based on a concept of social structure where you have functionality and hierarchy, and the overlap in the first case and not in the second case. And it also has the polarity that says it's true in one case, as well as an emotional valence. So this might be an internal mapping within an agent in which you'd want the lower-level one (and not omniscient mapping of external objects); different agents might have different non-omniscient mapping. Since they're not omniscient, they don't know all things equally, they may get some things wrong, etc. So, you have an internal structure in which the situation would be represented differently — broadly like social structure and narrowly like the next context — by different agents.

I don't have time to dwell on this … what I'm really doing is wending my way back to why this might be interesting to do in the first place and to say there's been a whole century of interaction theory that's very rich that I think is very relevant to agent simulation. And I believe that there are enough mechanisms here, and that we can specify it sufficiently that we can do some interesting kinds of experiments along this line.

Panel Discussion

Rojas: I have a couple comments, for all three of the presenters. I will go through them quickly and have ample time for questions and discussion.

My first comment overall about the entire panel was that these three projects together represent two different dimensions along which simulation studies could vary. One is what I like to think of as "how complex is each agent?" So, for example, you have a small number of variables describing whether people use certain kinds of transportation. Or, for example, you have certain production rules describing what happens in each room. That's pretty straightforward. On the other hand, David's presenting an entire research program for how you would describe how one person processes information. And so that's one way in which you can vary the simulation: by how complex the agents are.

Last year, Kathleen Carley talked about simplicity versus complexity. She called it veriticality versus transparency, in the sense that computer simulations can have so many parts. They can be such large projects. It's very hard for people to understand how they're put together, to understand how the inputs lead to the outputs. But, on the other hand, such complex models are very useful for policy-making; because they're so complex, they can actually capture a lot of the nuances of what happens in actual data and actual things that you observe.

And so I think if people really carry out David's project, we might end up with an extremely complex, veritical kind of model. And that's okay, because how people construct meaning is a very complex process. However, on the other hand, it's very easy for Doowan or for Elena to explain to us very simply how their models work, and you can get a lot of punch out of it, although in some ways they are not realistic models.

A very nice thing about the first two projects is that they're thinking of things in terms of basic social science, but also they have direct policy implications. And that's a very nice aspect of these papers; they're both hitting on basic social science topics: the nature of individual choice, instability of patterns of interaction. These are things that have a rich history in social science. And this is not just interesting for sociologists such as myself, but it could be interesting for engineers and policy-makers and leaders.

So let me go through a couple specific comments that I had for each of the panelists.

On the Dugundji-Gulyas paper, one thing that I liked in the write-up, which didn't come out quite in the talk today, was that we talk a lot about the small world's property and how sometimes choosing one property of the network to measure is sometimes easier than getting a measurement of the entire network itself. So when you present that to other audiences, that might be nice to emphasize, saying that, "Well, you don't need to know what every single connection is. You need to know these kind of more aggregate measures of the network," and that really gives you a lot of punch in your model.

Dugundji: It's a very good point. The reason why I didn't emphasize that so much is that the preliminary version of the paper that you read, and on which we drew these conclusions, was the first half of the talk, where I was describing the small-world network and the random network before we added the additional explanatory variables. So the result that we found was that this small-world property or the small-world threshold. This applies for the small-world network, and then the other critical parameter was the 1 over N, which was the moment when the giant component emerges in the random network. These two parameters were sufficient, dependent on whatever of the 20 different networks we did and the 50 different random seeds that we did per each of the 20 networks. So 100 total per density or per rewiring. These two parameters, the small-world threshold in the small-world network case and the giant component emerging in the random network case, were sufficient when we didn't add the additional explanatory variables.

So the conclusion that we put in the preliminary paper applies only when you have agents that are all in some way homogeneous. And they're heterogeneous in that they are connected to different people, so their X is varying across, but there are no additional explanatory variables such as travel cost, travel time, gender, business trip or not, social recreational trip or not. When we added those variables into the model, then this result didn't hold anymore, and that's when you saw that this is very similar to actually the approach where you saw those envelopes of those curves of how the beta parameter varied with density. And with some of those networks, the beta was high enough that we got that two-peak behavior. And for the others, the beta wasn't high enough and we got the single-peak behavior.

So the real punch line is, when you add the additional explanatory variables, this original result does not hold anymore. So that's the real punch line. And then a really a lot more work is necessary to understand the effect of all of these additional variables.

We have some hypothesis that when you add these additional explanatory variables, then the clumpiness of the network becomes important, and you have a random network, but there still may be some slight variations in the clumpiness of one random network to another. And that would be a useful parameter that we're actually in the process of trying to map this and see if this hypothesis is indeed true. So these slight variations actually matter.

Rojas: Well, let me just run through my other comments very quickly, so I can get everybody's comments in. One really nice aspect of this paper is that you vary the network structure, the network interaction structure. One of my hobby horses is telling people not to live in a grid world, because very few things actually look like grids; even geographical surfaces that we see aren't really grids.

Oh, I guess this is another criticism of what I read, but was not presented. There's a promise of a conceptualization of in-household interactions, and I really wanted to see that, because that would be interesting. So, for example, if you have a child, maybe some sort of negotiation at some times, you have to drive the kid to soccer or whatever, and that changes whether you'll use car-pooling or not, because it's a lot harder to car-pool if you have other household responsibilities.

And then also about the whole mean field framework. When I was rereading some of the basic literature on diffusion this summer, [I was reminded of] this idea that there are trend-setters, high-status individuals who may set trends for the adoption of innovations. Treating everybody the same is a nice assumption imported from physics, but maybe it's worth thinking about how you might drop in trend-setting in a parsimonious and succinct way.

We talked a little bit yesterday evening about the bifurcations, going from the unimodal to the bimodal distributions and the histograms. It'd be nice to have some data showing that this actually happened, because that'll be a nice way of saying, "Wow, we're getting somewhere on this."

Okay. Here are my comments for the Padgett-Lee-Collier paper. I found it to be a very powerful example of how to translate a physical model into a social science model. And there are a lot of papers I've read in recent years where that transition is less than perfect. The authors have a very clear idea, a very good grasp of what the original model's about, and a very best way of translating it into at least some simplified version of an economy. And, yes, they do admit it's a very minimalist approach to how they're doing it, but I found it a very plausible and persuasive approach, and so I think a lot of people who are starting out might have a lot to learn by reading this paper, if they want to translate a physical science model into a social science model. So that's kind of an esthetic comment.

One criticism I had was the assumption about forgetting of skills, how when one is adopted, some of the others are killed off. I could see how, for example, a firm might die. If a firm gets customers, then the other firms go under, but this, even as a simplifying assumption, seemed very counter-intuitive to me, especially since we think of firms and organizations having all sorts of mechanisms for memory, that something may go into a library somewhere or some sort of storehouse. Maybe I'm mistaken and this is actually very intuitive assumption. [I would add] a footnote or something in the presentation, a little bit more discussion of why that's a plausible assumption.

I also think it's really nice that, in future work, your group is going to start looking at different kinds of interaction topologies and get away from the grid. I would like to see many different researchers start from like a grid or other very baseline model, and then change the interaction topology and see consensus formation models. How do those change? That's something I'm working on right now. When do you see changes in the topology, the establishment of hypercycles in patterns of behavior? How does that change when you relax the grid assumption? We could go through all sorts of classical models and see how robust findings are with respect to the basic topology where people are interacting.

And then I had one complete out-of-left-field comment for this working group. And this is something that I just thought about last night. When you think about what this paper is really about, you have a bunch of products and you have a permutation of a number of products. So for the people who take abstract algebra out there, what you have is product number one, product number two, product number three and product number four, and then you have these chains: one goes to two, two goes to seven, seven goes to three, and so forth. That's called a permutation. And the set of permutations has an extremely rich algebraic structure.

And then you have a decision rule for what kinds of permutations survive over time. I could be completely out of my field, and this is a very vaguely worded question or challenge: what is the connection between the process of survival, survival of cycles, and the underlying algebraic structure? If you were able to actually come up with a concrete link to analytically show, or even to numerically show that, for example some subgroup or subset of permutations on N elements has a very specific profile of survival in this model, then that would be really a gigantic result, which would be something like John Nash showing that noncooperative games could be analyzed by solving an underlying fixed-point problem. So there's some sort of underlying algebraic structure which is under there somewhere, and then you have the selection mechanism going through this algebraic structure.

What is the link between these two processes? I have no idea. That could take you a hundred years to solve or it could take you one year. But if you do that, that would really be on the level of showing a noncooperative game as really a fixed-point problem. Just show a very profound link between one kind of mathematical process and another, abstract algebra and this algorithmic thing happening over here. I'd be interested to see if anybody could solve that, and it might be easier than it sounds.

And then I have two comments for David Sallach.

First, just a comment I've made before, that I really like how you've introduced me and a number of other people to situation theory, because in the social sciences we talk all the time about meaning. So when I was teaching my Economic Sociology course this week, we talked about how actors, how the meanings that people have in markets, really affect what they do. So I talked about Liars' Poker, the book about bond traders and how the whole idea was that nothing was stopping people from trading certain kinds of securities for a long time, so it's just specifically mortgage bundling. But it was specifically someone who said, "I'm introducing this new concept called bundling of mortgages, which means that you can now create this new market."

But, on the other hand, social scientists haven't been very good at formalizing meaning, and introducing this idea to me and to other people, I find extremely valuable, so I wish you all the success, and I recommend that other people who are interested in these topics start thinking about it.

And then I kind of had a general social theory comment question for David that, when we think about meaning in a lot of contemporary organizational theories, especially neo-institutional theory it's like a brick wall in sense, right? And if you really read the hard-core early '80s institutionalists, these enforcing mechanisms, as well as Dick Scott's institutional carriers, are all this diffusion process, trying to push us towards one solid block of meaning, that we just have to deal with it, or if we don't do it we're just out of the game. It's as simple as that.

But in the start of your talk, you talk about almost like the infinite pliability and playfulness of meaning and creation of meanings and situations. So it seems that, in modern sociology, we have two contradictory approaches to meaning. One is the kind of macro-institutionalist approach of, for example, the institution of marriage or the institution of property, and if you don't like it, tough. It's this way. And then you have this very kind of ethnomethodological approach, which has basic elements, but they're so pliant. They're not completely plastic that there's no structure at all, but it's so playful and there's so much variation and subtlety to it. And those are two very different approaches to this whole issue of situations and meanings. I think that would be a really interesting issue to talk about or to comment upon.

And those are my comments. And, as usual, I will type up my comments and e-mail them to the presenters later, which is my normal habit.

So I'll open it up to whoever wants to respond or to people from the floor who have their own comments that they would like to introduce.

Unidentified Speaker: I'd like to respond just to the comments that you gave to us, which I find really sort of, very interesting. Your three comments were forgetting, network topology, and algebra.

Let me start with the last, the sort of the left-field one, because actually that's not so left field. There's a guy at Santa Fe Institute, Walter Fontana, who thinks exactly the way that you're talking about, and he has a much richer conception of rules, lambda calculus, which generates rules and so forth. So I understand how what you say would actually be implemented, because I talk to someone who does that.

This gets back to an issue that I see come up over and over again, which is this KISS versus verisimilitude issue. Walter, when he did the rich algebra approach, generated incredibly deep and profound ideas, but they were all of an existence-proof nature, just as this example. So he would have a rich topology. They would interact, and very algebraic closed systems would result, and then you would show the existence of the ability of very primitive chemistries to generate very complex organizational systems.

We actually went much more obviously in the KISS direction, and the reason is that it was much clearer to generate actual results, that is to say, "Target reproduction has this effect. Spatial things have this effect. Endogenous environments have this effect." I mean, this exercise we went through was possible in our stripped-down model where it was not possible in Walter Fontana's set-up because things were so difficult and interactive. And so it goes back to the verisimilitude versus transparency issue. Everything keeps coming back to that.

Our approach is transparent. You can see, but at the cost of working in an incredibly stripped-down sort of type setting. So I know that going in the direction of your last comment is very, very important and very, very difficult, but it gets us back to this issue of how do you actually generate results in this rich environment.

So that's not a defense of my position. I'm just saying there are these esthetic tradeoffs. It's hard to resolve them.

On the other points, reasoning backward to your second point, network topology, I'm totally in favor of everything you said. Let's not get stuck on the grid and so forth and so on. So I'm in favor of moving on, the way [Elena] has done and others. But the problem, the issue that you run into, is the following, that Michael Cohen and a lot of other people in evolutionary game theory have done this exercise in a prisoner's dilemma, population sort of world — nonspace grid, small worlds, clusters. After going through the usual array of candidates, they found in their set-up, which is not our set-up, that there was this huge effect between nonspatial random interaction and *any* form of network. Beyond that are other effects but marginal compared to the big step from nonspace to space. So from structure to nonstructure, huge effects and from there on, effects of second-order character. That was the background for me putting off this exploration of network topology to the future rather than now, because I think that the same thing's going to happen in our case, which is not to say that we shouldn't do it, but as primary effects go, we think we did what we did.

And back to the first point, about forgetting, there's one defense and one complete agreement with your criticism. The defense is that forgetting is interesting. I mean, randomly killing off rules based on no performance measure is interesting, because that generates powerful selection effects in spite of the fact that there's no actual selection for any optimal sort of thing. So it goes back to the idea of, what is fitness? Fitness in an economist sense is a performance measure. Fitness in a biologist sense is just relative birth rate. That's it. And we have fitness in a relative birth rate and we sort of push away fitness — there's just complete random death, and we show that even in that sense there's powerful selection going on. So in a sense it's a rhetorical reason why we do what we do, not a realism reason.

We show that you don't need selection in the economist sense to generate powerful self-emergent organizations. But other than that defense, which is a rhetorical defense not a realism defense, I think we need more robust exploration of different mechanisms.

So I agree with your criticism, but my rhetorical point was to do this minimalist exercise to show that you can have a lot of powerful evolution, even without systematic selection.

Rojas: Yes, it's almost like a vector. You don't have to actually be rational to do rational behavior. You don't need to be selected by this one specific mechanism to still have that kind of interaction.

Were there any other responses from the audience or speakers?

Sallach: I think that where meaning resides and how stable it resides is really a question of what your perspective is. I mean, when you're flying at 20,000 feet, you don't see each individual field, you don't see each individual farmer and so forth. You see a big area where they're growing wheat. And I think if you look more closely, you get down to the very specific kind of detailed action.

But note that any of the people, the most ethnomethodological you could ask, say, take Garfinkel — it's out of his perspective that the idea of embodied practices arises in which a lot of things aren't constantly up for grabs, in which it's done at a relatively nonconscious level. And yet he still finds the patterns to be nonrepeatably unique. For example, he encourages his students to study people standing in line, because every line witness, and the actions taken to maintain that line, are different.

I guess what I would say is that meaning construction processes are always potentially open, down to infinite detail.

Political Processes

AGENT-BASED RESIDENTIAL SEGREGATION: A HIERARCHICALLY STRUCTURED SPATIAL MODEL

D. O'SULLIVAN,* J. MACGILL, and C. YU
The Pennsylvania State University, University Park

ABSTRACT

This paper presents a variation on Schelling's model of residential location dynamics that combines two concepts of neighborhood: continuous and bounded. A brief description of Schelling's model is included, and the structure of the current model is described in detail. The effect on model behavior of varying the size of bounded neighborhoods, while also varying the balance between local- and regional-level effects on agent behavior, is explored; preliminary results are reported. The scale of bounded neighborhoods considered by agents in making residential location decisions has important impacts on overall model outcomes. The range of possibilities for further work is discussed.

Keywords: Residential location dynamics, bounded neighborhood, ABM

INTRODUCTION: SCHELLING'S MODELS OF RESIDENTIAL SEGREGATION

Schelling's simple model of residential segregation dynamics (Schelling 1969, 1971, 1978) is rightly regarded as a seminal example of multi-agent simulation in social science (Macy and Willer, 2002). In the fifteen years from 1988 to 2002, Schelling's "dynamic models of segregation" (1971) has been cited 125 times; 70 citations occurred from 1999 to 2002 (ISI, 2003). The model's persistent popularity derives from its simplicity and its compelling demonstration of the emergence of stable, aggregate, socio-spatial patterns from local interactions between household agents. In Schelling's model, households of two types make decisions to remain at or leave their current residential location depending on dissatisfaction with that location. Dissatisfaction arises when a household has either too many neighbors of the opposite type or too few neighbors of its own type. By using this framework, Schelling shows that strongly segregated large-scale residential patterns can arise even when two groups are relatively tolerant of one another's presence.

Neighborhoods are conceived in two different ways in Schelling's work: continuous and bounded. In the *continuous neighborhood* case, households occupy locations on a lattice or 'checkerboard' (Sakoda, 1971), and decisions are made with regard to the types of households in adjacent locations on this lattice. Agents demand that some fraction of their immediate neighbors on the lattice is of the same type as themselves. Agents that are unhappy under these criteria move to a nearby location where their residential preference requirements are satisfied. Unsurprisingly, when households demand many neighbors similar to themselves, the result is dramatic segregation of the lattice into large regions occupied exclusively by households of only one type. Schelling assesses the extent of segregation by counting the average number of like neighbors in the final stable pattern that results (Schelling, 1971, pp. 157–158). He finds that

* *Corresponding author address:* David O'Sullivan, The Pennsylvania State University, Department of Geography, 302 Walker Building, University Park, PA 16802; e-mail: dbo2@psu.edu.

"the resulting segregation [is] a rapidly rising function of demands [for like neighbors] in the range from about 35% to 50%" (Schelling 1971, p. 159). This pattern occurs because increasing demand for like neighbors leads to additional discontent in random initial patterns, a greater likelihood that moving households will displace previously content households, and a higher concentration of households in regions of the checkerboard already densely populated by other households.

Schelling also investigates the effect of different combinations of household tolerance profiles in two populations, with respect to *bounded neighborhoods*. A bounded neighborhood is a 'container' populated by a number of households. Schelling uses the bounded neighborhood concept to analyze the relationship between residential preferences in the population and the stability properties of different combinations of numbers of households of each type in a *single* bounded neighborhood in isolation. By analyzing plausible — albeit hypothetical — tolerable ratios between two groups, he demonstrates that the only stable states in many cases are exclusive neighborhoods where all the residents are from one group or the other.

A number of aspects regarding Schelling's models deserve comment in the current context. First, in the continuous neighborhood case, household behavior is governed by a demand for like neighbors, whereas in the bounded neighborhood case, antipathy toward different neighbors is the driving force. Although these mechanisms can be combined, it is difficult to do so without introducing numerous arbitrary parameters (demand for like, tolerance of different, and so on). In the present model, we adopted the tolerance/antipathy approach for household behavior at both local and regional scales. Thus, it is the presence of too many households of a different type whether locally or in a larger bounded neighborhood that causes household decisions to relocate.

Schelling's experiments in the continuous neighborhood case seem to have been conducted by hand, although this is not clear. The most important effect of this on the operation of the model is vagueness about the order in which households are considered for relocation. Thus, in Schelling's description of the rules of movement, he says, "Identify the discontents […] and, *in some order* move them to where they are content" (Schelling 1971, p. 156). This tenet is followed, almost immediately, by an acknowledgment that this makes a difference in detail, but not in general: "The *particular* outcome will depend very much on the order in which discontented [households] are moved, the *character* of the outcome not very much" (Schelling 1971, p. 156).

Similar comments apply to vagueness in the rules of movement such that a household relocates to the "nearest" vacant spot within a neighborhood that is acceptable. Vagueness regarding these points makes it impossible to replicate Schelling's simulations in a computational simulation. We therefore use random ordering both of household relocation decisions and of consideration of equidistant vacant locations to minimize effects that seem likely to arise from any more structured sequencing of relocation events.

The continuous neighborhood formulation of Schelling's work has been widely acknowledged in the multi-agent social simulation community. This acceptance is perhaps because this approach is suggestive of common devices in contemporary multi-agent work, particularly the grid-based space in which agents interact (see, for example, the Sugarscape model [Epstein and Axtell, 1996]). The continuous neighborhood approach is also consistent with notions from complexity science about the efficacy of purely local interactions in producing larger global structures.

Schelling's work is extremely insightful and thought provoking. The important finding in the bounded neighborhood case (i.e., that stable racially integrated neighborhoods are unlikely for many combinations of tolerance profiles) has been confirmed on the basis of empirical data (Clark, 1991).

COMBINING CONTINUOUS AND BOUNDED NEIGHBORHOODS IN A HIERARCHICAL MODEL

Description of the Hierarchically Structured Model

We present a hierarchical version of the Schelling model that combines his two neighborhood types. Household agents not only consider the type of immediately neighboring households in a lattice of residential locations (the continuous or local neighborhood), but also the aggregate nature of the bounded neighborhood (or district) that contains their residential location. The model consists of a number of bounded neighborhoods, each containing a number of residential locations at points on a lattice. This structure is illustrated in Figure 1.

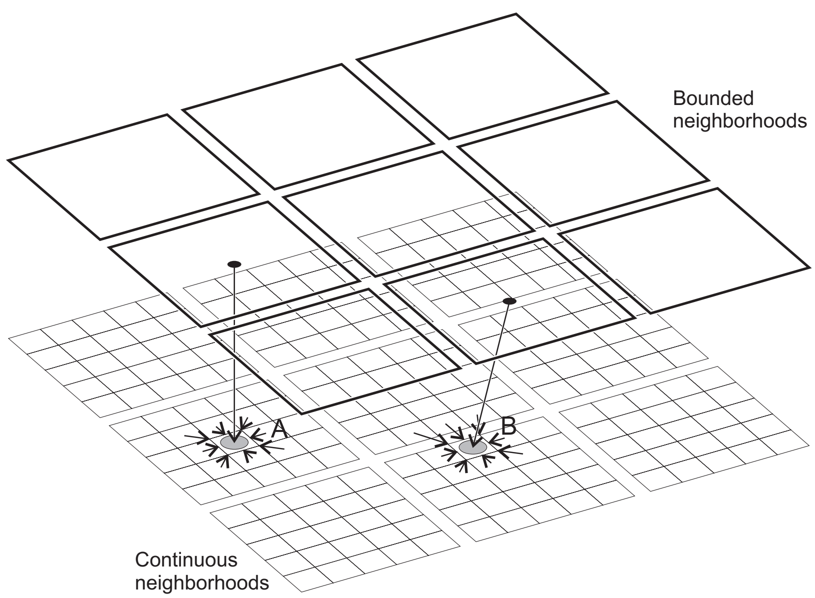

FIGURE 1 Structure of the hierarchical Schelling model (The small grid cells represent residential locations with a *continuous neighborhood* structure. Residential locations are contained in *bounded neighborhoods*, whose aggregate state is also considered by agents in residential decision-making behavior. Arrows show relations of influence on decision making.)

Decisions of agent A, at the center of a bounded neighborhood are affected by immediately neighboring agents, and by the aggregate state of the containing bounded neighborhood. An agent B, at the edge of its containing bounded neighborhood, is affected by neighbors both inside and outside the bounded neighborhood, but only by the containing bounded neighborhood at the aggregate level.

As far as possible, we retain the simplicity of Schelling's formulation of households' rules of movement. In assessing the level of 'happiness' with a current or potential location, household agents determine local happiness h_L and regional happiness, h_R. This assessment is made with respect to a single agent parameter called tolerance T, which varies from agent to agent. Tolerance is a real number between 0 and 1, indicating the fraction of occupied neighboring locations whose agents may be of a different type from the agent without negatively affecting its happiness and prompting it to seek an alternative location. Local happiness is the difference between an agent's tolerance and the fraction of occupied locations in the agent's continuous neighborhood occupied by agents of a different type. Similarly, regional happiness is the difference between agent tolerance and the fraction of occupied locations in the agent's bounded neighborhood occupied by agents of a different type. Formally, an agent A_i has tolerance $T(A_i)$ and type $t(A_i)$, where

$$0 \leq T(A_i) \leq 1 , \tag{1}$$

and

$$t(A_i) \in \{\mathrm{RED}, \mathrm{BLUE}\} . \tag{2}$$

Color is a convenient visual marker for agent type, but any discrete valued variable will suffice.

If we denote the set of agents in the local (continuous) neighborhood by N_L, and the set of agents in the district (or region or bounded neighborhood) by N_R, we can determine local and regional happiness for the agent from the following equations:

$$h_L(A_i) = T(A_i) - \frac{\left\| \{A_j : A_j \in N_L \wedge t(A_j) \neq t(A_i)\} \right\|}{\left\| N_L \right\|} , \tag{3}$$

and

$$h_R(A_i) = T(A_i) - \frac{\left\| \{A_j : A_j \in N_R \wedge t(A_j) \neq t(A_i)\} \right\|}{\left\| N_R \right\|} . \tag{4}$$

It is simple to combine happiness values by using a single model-wide parameter local-regional balance b_{LR} to determine an overall happiness h for the agent, according to

$$h(A_i) = (1 - b_{LR})\, h_L(A_i) + b_{LR} h_R(A_i) , \tag{5}$$

where b_{LR} = 0 results in agent happiness depending only on agents in the continuous neighborhood, while b_{LR} = 1 means that happiness depends only on agents in the same bounded neighborhood.

Overall, an agent's rule of movement is to determine overall happiness, based on neighboring agent types both locally and in its bounded neighborhood. If the agent has an overall negative happiness score, it is unsettled and tries to relocate. This move involves examining available vacant locations at successively greater distances in the lattice until one is found where the agent's happiness score would be positive. Potential locations at the same distance from the current location are considered in random order to ensure no directional bias in agent relocation. As soon as a suitable location is found, the agent moves to that location. It is possible that no suitable location is available, in which case the agent does not resettle.

Agents are considered for relocation one at a time in random order in one sweep through the agent population. Each sweep of the population occurs in a different random order.

Implementation Details

The model described above was implemented in the Repast agent modeling toolkit (Collier, no date). Repast provides a simple bridge to the GeoTools open source package for displaying and analyzing geographically referenced datasets, and given our interest in exploring the impact of geographic perceptions on models of segregation behavior, this selection was a natural choice.

To facilitate future investigation of more complex spatial patterns, the model's geographic structure is initialized by reading two geographic information system (GIS) files, one that represents residential locations in the continuous neighborhood layer and one that represents bounded neighborhoods. Geographic processing is applied to determine both the continuous neighborhood relations among residential locations (i.e., the lattice structure), and the nesting of residential locations in the continuous neighborhood layer within containing bounded neighborhoods. The resulting neighborhood relationships are stored in a graph data structure that records *adjacency* relations between locations and *containment* relations between bounded neighborhoods and locations. This design allows agents to retrieve information concerning the numbers of agents of their own or opposite type in their continuous neighborhood and in the bounded neighborhood. This approach allows for future investigation of model dynamics with irregularly shaped locations and districts, although such examples are not considered in this paper.

RESULTS

Model Input Parameters

The values of the various model parameters used in the reported experiments are summarized in Table 1.

TABLE 1 Summary of model parameter settings

Parameter Name	Settings Used	Comment
Local-regional balance	0.0, 0.1, 0.2, …, 0.9, 1.0	Varied through full range to study effects of variation in local versus regional behavior
Bounded neighborhood sizes	4×144 locations 9×64 locations 16×36 locations 36×16 locations	Varied to study impacts of different bounded neighborhood sizes on behavior
Tolerance	0.2 to 0.4	Agents initialized from a random uniform distribution
Occupancy rate	0.75	Fixed
Fraction blue	0.5	Fixed

Model Output or Measurement Parameters

A number of summary statistics are used to track model progress. The dissimilarity index D is reported with respect to the set of bounded neighborhoods. The value of D is a measure of residential segregation for population count data reported for zones, which indicates the extent to which two population groups are not similarly distributed among the zones (Duncan and Duncan, 1955; Taeuber and Taeuber, 1965, 1976). Given two population groups with total populations R and B, the counts of the groups living in each zone i can be denoted r_i and b_i. These values are combined across all n zones, to give

$$D = 0.5\sum_{i=1}^{n}\left|\frac{r_i}{R} - \frac{b_i}{B}\right| , \tag{6}$$

where D has a value of 0 if two populations are distributed identically across a set of zones. It has a value of 1, if they are completely segregated (i.e., if all blues are located in zones that contain no reds, and vice versa).

Because D is calculated with respect to a set of bounded neighborhoods, it is possible even when D indicates little segregation, for agents to be locally segregated such that agents have neighbors in their continuous neighborhood predominantly of the same type as themselves. Such local segregation is measured using an average fraction of like neighbors statistic S_L (for locally similar). For each agent, the fraction of occupied neighboring locations whose occupying agents are of the same type is averaged across all agents.

Both D and S_L are pattern measures calculated at any point during a model run. The remaining two output parameters are cumulative measures of model dynamics over each sweep through all the agents. The fraction of agents unsettled p_U and fraction of agents resettled

p_R record, respectively, the fraction of all agents in the model that were unsettled and tried to relocate, and the fraction of all agents in the model that successfully resettled, during a sweep of the whole agent population. Note that p_R is less than p_U by definition, since only unsettled agents attempt to relocate.

Final Stable Patterns

Initially, we observe final stable patterns in the model to see the differences in outcomes relative to the patterns of Schelling's continuous neighborhood case. As expected, with the local-regional balance parameter set to 0, outcomes are identical to Schelling's examples (see Figure 2), and the bounded neighborhoods make no difference.

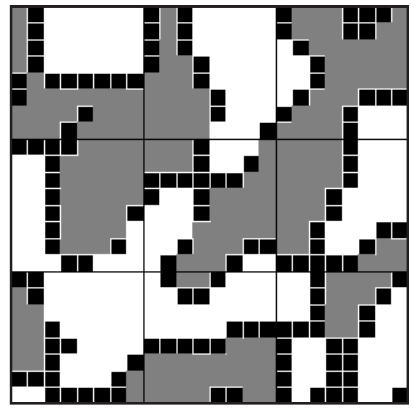

FIGURE 2 Typical outcome with the local-regional balance parameter set to 0 (fully local) (Agent states are shown in gray and white. Black locations are vacant.)

When the local-regional balance is increased to 0.5, final configuration is reached, as illustrated in Figure 3, for two different sets of bounded neighborhoods. Agent responses to the composition of bounded neighborhoods result in a bimodal distribution of bounded neighborhoods — either predominantly red or predominantly blue. The change in agent priorities also means that dissimilar agents may be tolerated as immediate neighbors in the continuous neighborhood, along boundaries between bounded neighborhoods with different majorities. Examples of the opposite effect are also apparent: single isolated agents of the "wrong" type are found in some bounded neighborhoods because the tolerable (empty) configuration of their continuous neighborhood allows them to ignore the majority of unlike agents in the bounded neighborhood.

(a) (b)

FIGURE 3 Two outcomes with the local-regional balance parameter set to 0.5: (a) a 3 × 3 grid of bounded neighborhoods, each with 64 locations, and (b) a 4 × 4 grid of bounded neighborhoods, each with 36 locations

Other settings of the local-regional balance parameter result in different balances in the outcome patterns between the tendency to local segregation on the one hand, and to bounded neighborhood segregation accompanied by tolerance for unlike neighbors across district boundaries on the other, as illustrated in Figure 4.

When the local-regional balance is set to 1 (completely regional), greater variation in outcomes is observed. This effect is discussed in greater detail in the following section.

Varying Bounded Neighborhood Size and Local-regional Balance

In this section, we report preliminary findings from experiments where the sizes of bounded neighborhoods and the local-regional balance parameter were varied. Results are shown in Figures 5 through 8, as the local-regional parameter is varied from 0 to 1 in increments of 0.2.

The first point to make about these figures is that there is considerable continuity in the model behavior through all the results shown. The dominant behavior is for the model to segregate, and to do so rapidly. When the model stabilizes, agents are (usually) no longer unsettled and are content to stay where they are. Segregation behavior occurs in almost all cases.

Two differences are observed as the local-regional balance parameter is increased. First, the final stable state exhibits patterns that are increasingly segregated as measured by the dissimilarity index and decreasingly segregated as measured by the average fraction of like neighbors. This phenomenon was already noted above, whereby increasing emphasis on the bounded neighborhood allows agents to have dissimilar immediate neighbors across district boundaries. This fact is a direct result of the presence of bounded neighborhoods "steering" local segregation to fit inside the boundaries, so that a higher dissimilarity index is observed.

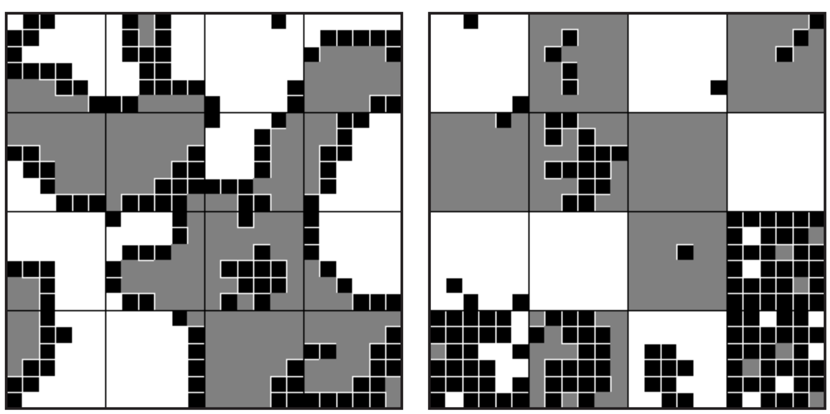

(a) (b)

FIGURE 4 Typical outcomes with a 4 × 4 grid of districts and the local-global balance parameter set to (a) 0.25 and (b) 0.9 (Note that all the images in Figures 2 through 4 are based on the same initial random number generator seed setting of 1061992058618.)

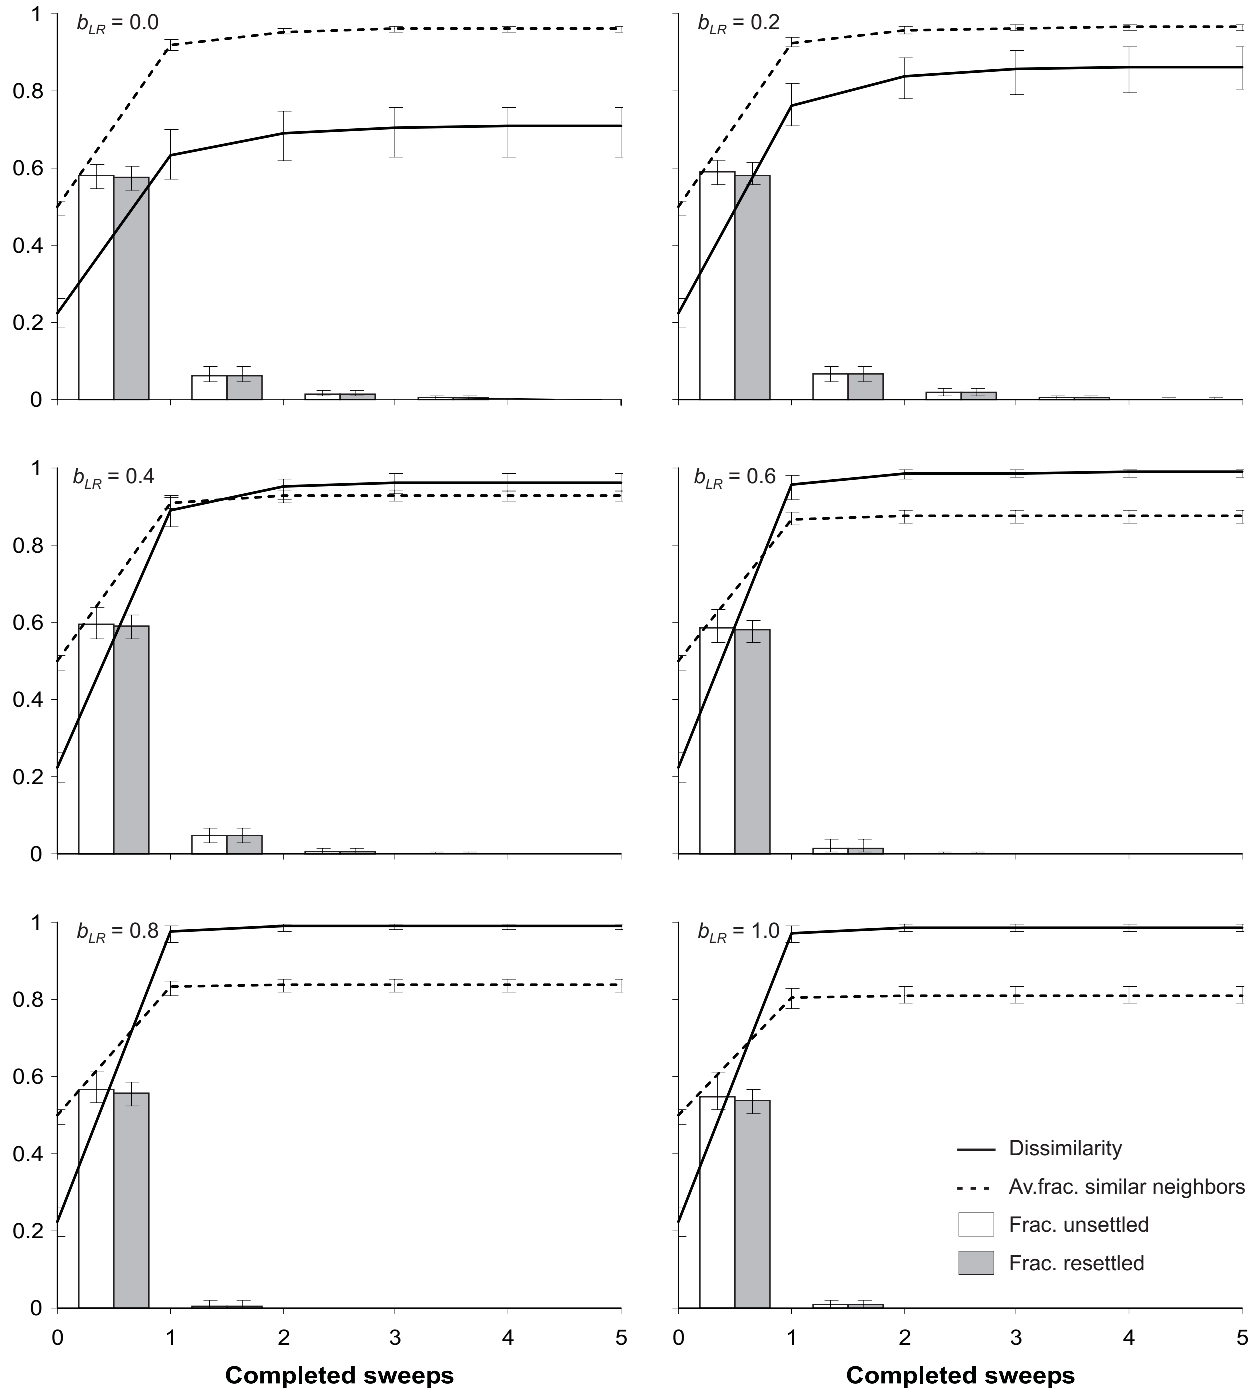

FIGURE 5 Summary results for 16 location bounded neighborhoods, as local-regional balance is varied from 0 to 1 in steps of 0.2 (Pattern measures [dissimilarity and average fraction of similar neighbors] are shown as line graphs with values recorded at the end of each sweep through all agents. Dynamic summary measures [fraction of agents unsettled and fraction resettled] are shown by bars, and record these values summed over a sweep through all agents. All four statistics include an error bar indication of the range of values between the 10th and 90th percentile over 100 randomly generated runs.)

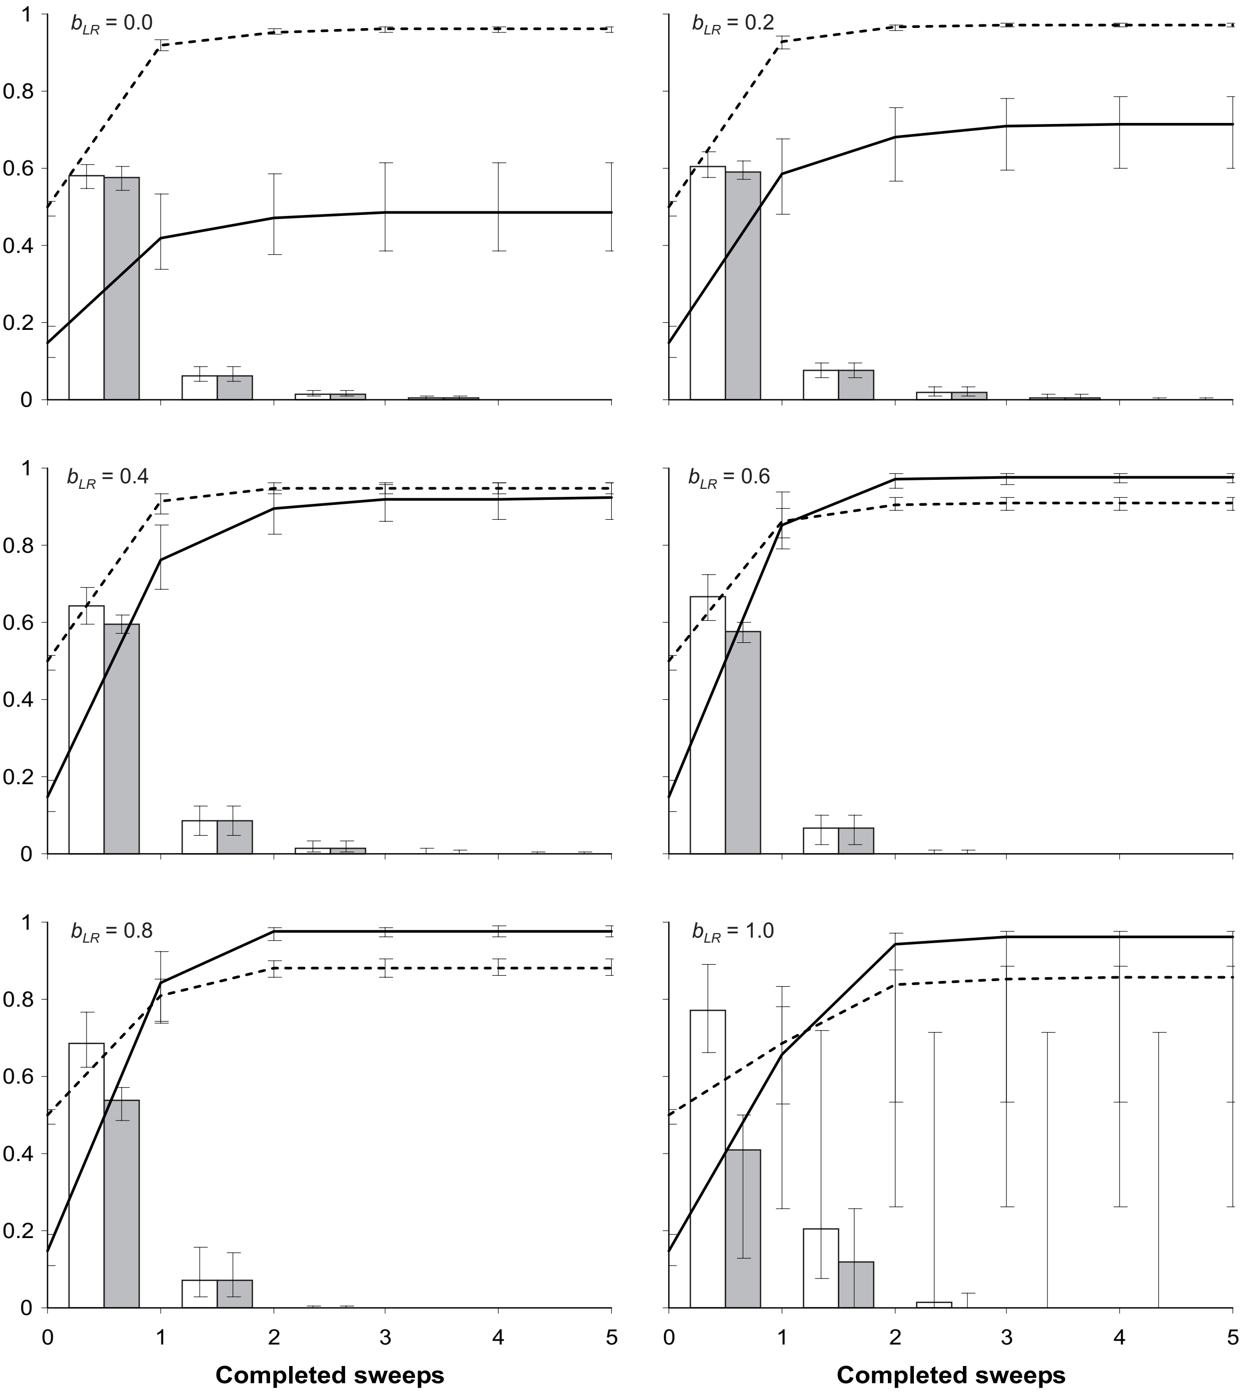

FIGURE 6 Summary results for 36 location bounded neighborhoods as local-regional balance is varied from 0 to 1 in steps of 0.2 (See Figure 5 caption for further explanation.)

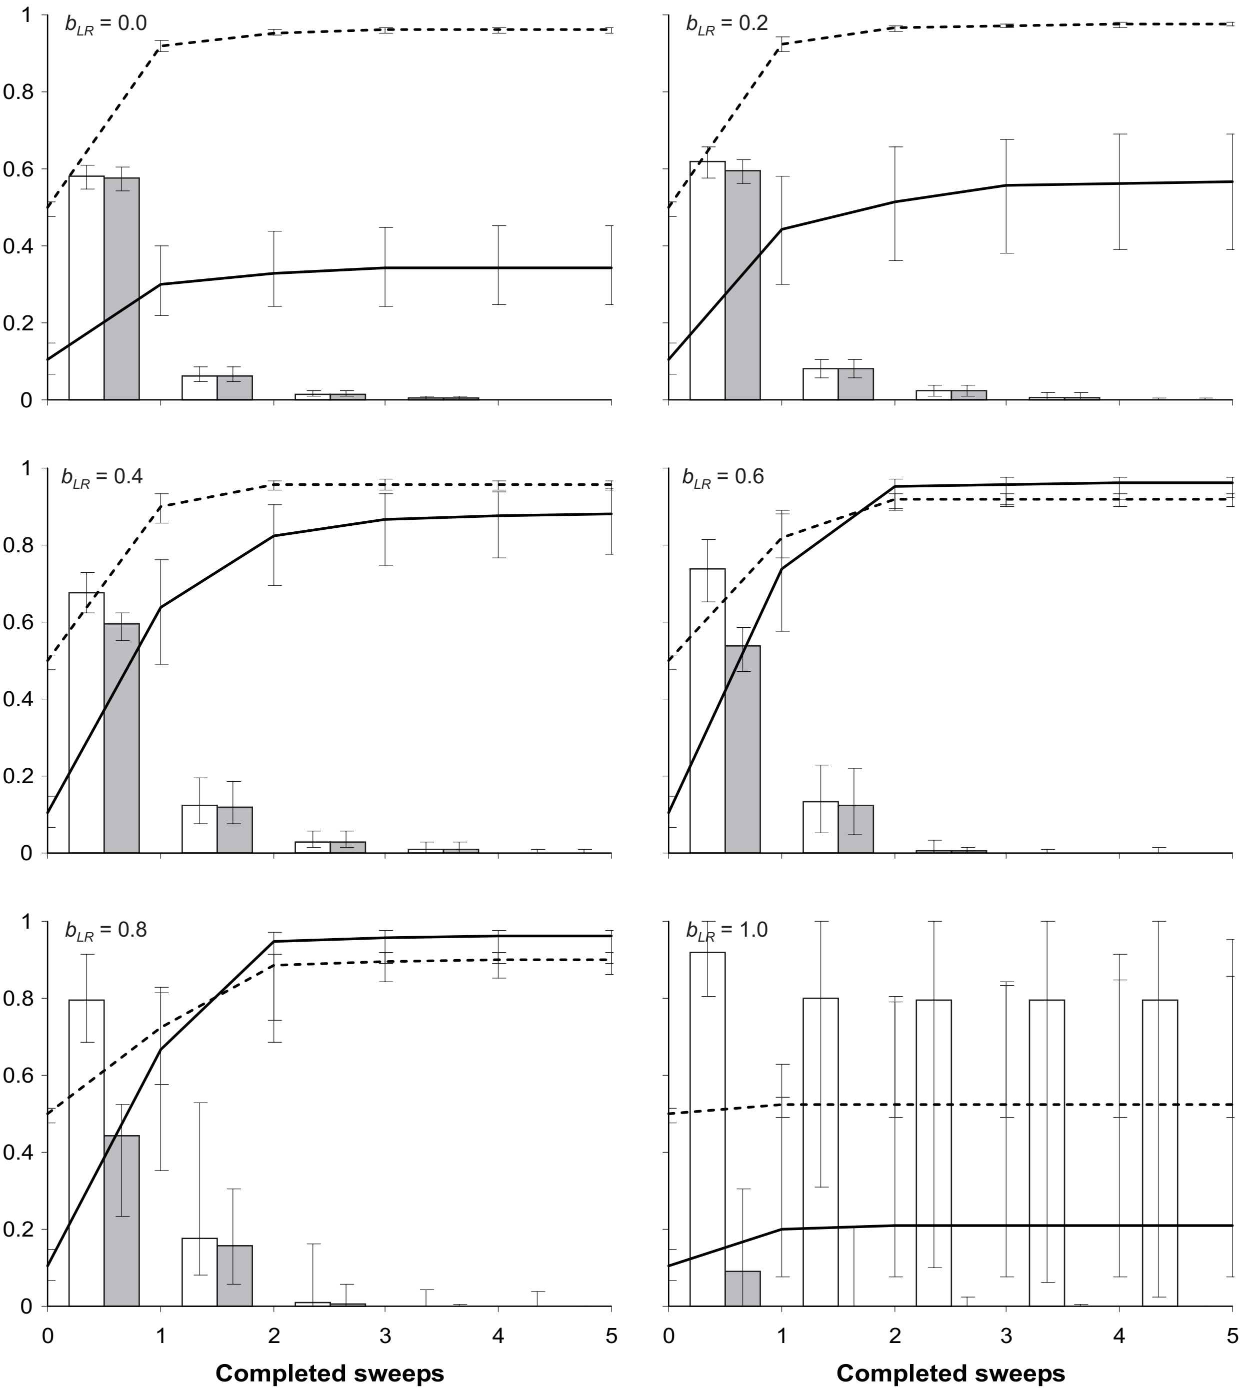

FIGURE 7 Summary results for 64 location bounded neighborhoods as local-regional balance is varied from 0 to 1 in steps of 0.2 (See Figure 5 caption for further information.)

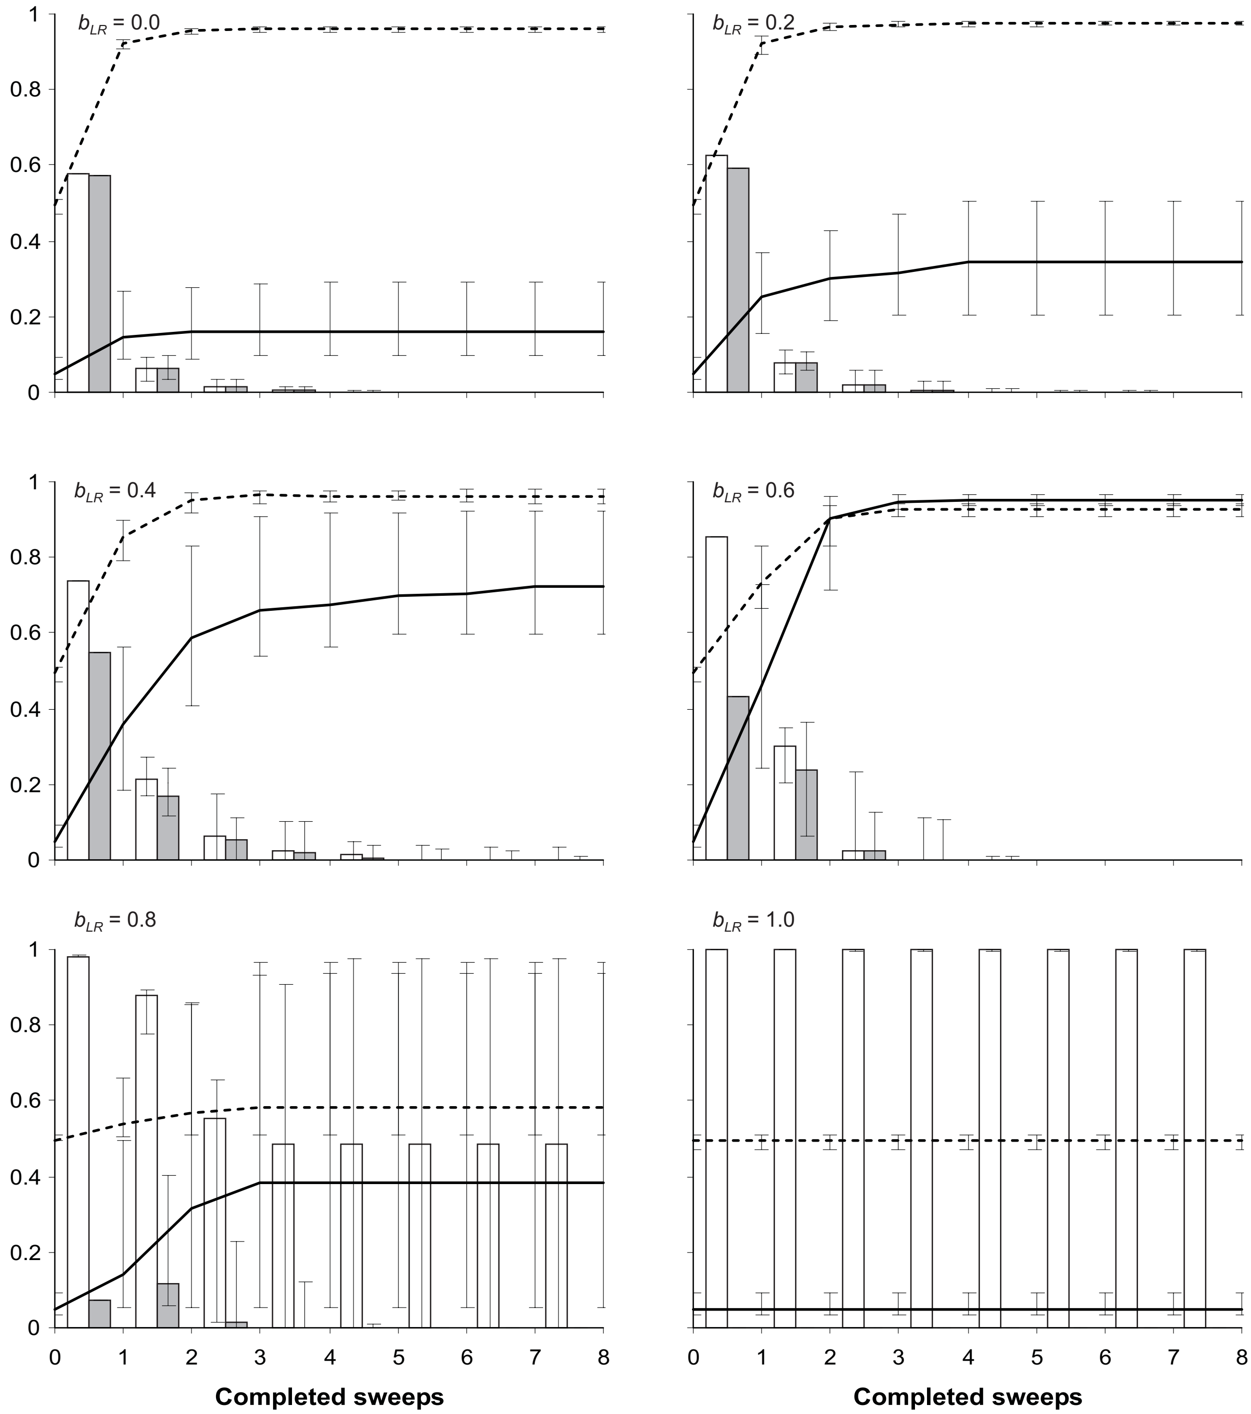

FIGURE 8 Summary results for 144 location bounded neighborhoods as local-regional balance is varied from 0 to 1 in steps of 0.2 (These cases have longer time sequences, and the results are based on running only 50 random sequences.)

The second point is that as the local-regional parameter increases the time taken for the model to stabilize decreases. Thus, in the first row of all diagrams with b_{LR} set to 0 or 0.2, there is relocation activity over about four sweeps of the agent population; in the second row (b_{LR} equal to 0.4 or 0.6), relocation activity occurs for only around three sweeps; and in the third row (b_{LR} equal to 0.8 or 1.0), stabilization happens after only two sweeps of the population in most cases. This is a result of the combination of bounded neighborhoods and regional-level behavior causing agents to relocate to districts that are tipping into a state where all agents are of the same type. Once settled in such locations, agents do not move again. When more local considerations are dominant, it is possible for an agent to initially resettle in a location that is locally congenial, but which subsequently becomes less desirable as the bounded neighborhood starts tipping into the opposite type of agent.

Clearly, this relatively neat picture of the model's behavior breaks down in the last plot in Figure 6 (with a 36-location bounded neighborhood, and $b_{LR} = 1.0$) and is similarly inadequate for high values of b_{LR} in both the 64- and 144-location bounded neighborhood cases (Figures 7 and 8, respectively). With these combinations of settings, a wide variation in outcomes across the sets of random runs is evident. In a significant fraction of cases, the model becomes stuck in a configuration where bounded neighbors are incompletely sorted so that the dissimilarity index D is not near 1. In these situations, many household agents remain unsettled but are unable to find preferable locations and so do not resettle. At present, it is unclear if any statistic anticipates this outcome, although large numbers of unsettled households failing to resettle during the first agent population sweep is a promising candidate predictor.

The last plot in Figure 8, which has a large bounded neighborhood (144 locations) and the local-regional parameter set to 1, exhibits the most extreme form of this behavior — no relocation is seen at all. This result occurs because agents care only about bounded neighborhood states; with only four neighborhoods to choose from, there is a strong probability that randomly initialized bounded neighborhoods will be judged the same. Large bounded neighborhoods make all model locations effectively the same, so that little or no relocation is observed even though virtually all agents are unsettled.

CONCLUSIONS AND SUGGESTIONS FOR FURTHER WORK

Both bounded neighborhoods and variations in agent local-regional behavior have significant effects on Schelling-type model dynamics and on the resulting stable spatial patterns observed.

Small bounded neighborhoods have a relatively limited impact on the model, except to alter the details of final stable patterns if agents are attentive to bounded neighborhood effects. As agent attention to bounded neighborhoods increases, however, the speed with which the model stabilizes increases because of preferential movement into neighborhoods that are tipping into exclusive occupation by one or the other type of agent. This observation could be partially confirmed by measuring the average distance moved by relocating agents to see if agents move farther as the local-regional balance parameter is increased.

For larger bounded neighborhoods, it is possible for the model to get stuck in a configuration where agents are unsettled but unable to relocate because no alternative location is judged preferable. This dilemma appears to be a result of initial preferential relocation into tipping neighborhoods, leaving unsettled agents with a choice of locations in the remaining neighborhoods, which are all judged similar to one another. These effects occur only when the local-regional balance is tilted toward regional effects because no local preferences enable habitable niches to be established by a series of local relocations.

Clearly, this model is extremely abstract, so that interpretation of these findings is tricky. Perhaps the most useful way to think about the results is in terms of communication processes among agents. When local behavior is dominant, segregation is slower (but surer) because agents only attend to nearby locations and local niches can be established gradually that enable eventual complete segregation. When regional-scale behavior is dominant, segregation is more rapid (but less sure) because preferred new locations are rapidly identified. However, depending on the scale — the bounded neighborhood size — overattention to only larger-scale neighborhoods can prevent segregation from occurring completely.

In this light, the model seems to be a useful vehicle for exploration of the important role of information in residential location decision making. In the current version of the model, bounded neighborhoods are fixed, but we intend to remove this limitation in future developments so that the dynamics of emerging local property markets can be explored. Modeling of the behavior of other agents operating at different spatial scales in the residential location context (realtors and banks, in particular) is also planned. In fact, in the current implementation agent relocation is handled by a global realtor class to facilitate exploration of these aspects in revised versions of the model.

In addition, abstraction from the current model to a more general class of geographic agent models using the Repast architecture is planned. The difficulty of understanding the behavior of even this relatively simple model may be greatly reduced by closer integration with dynamic visualization environments. We intend to pursue this direction in the medium term.

REFERENCES

Clark, W.A.V., 1991, "Residential Preferences and Neighborhood Racial Segregation — A Test of the Schelling Segregation Model," *Demography* **28**(1):1–19.

Collier, N., no date, *RePast: An Extensible Framework for Agent Simulation*, Social Science Research Computing, University of Chicago. Available at http://repast.sourceforge.net/docs/repast_intro_final.doc (accessed June 24, 2002).

Duncan, O.D., and B. Duncan, 1955, "A Methodological Analysis of Segregation Indexes," *American Sociological Review* **20**:210–217.

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom Up*, Washington, DC: Brookings Press & MIT Press.

ISI, 2003, *Web of Science*. Available at http://www.isinet.com/products/citation/wos/ (accessed May 24, 2003).

Macy, M.W., and R. Willer, 2002, "From Factors to Actors: Computational Sociology and Agent-based Modeling," *Annual Review of Sociology* **28**:143–166.

Sakoda, J.M., 1971, "The Checkerboard Model of Social Interaction," *Journal of Mathematical Sociology* **1**:119–132.

Schelling, T.C., 1969, "Models of Segregation," *American Economic Association Papers and Proceedings* **59**(2):488–493.

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

Schelling, T.C., 1978, *Micromotives and Macrobehavior*, New York City, NY: Norton.

Taeuber, K.E., and A.F. Taeuber, 1965, *Negroes in Cities: Residential Segregation and Neighborhood Change*, Chicago, IL: Aldine Publishing.

Taeuber, K.E., and A.F. Taeuber, 1976, "A Practitioner's Perspective on the Index of Dissimilarity," *American Sociological Review* **41**:884–889.

IMPLICIT COOPERATION IN CONFLICT RESOLUTION FOR SIMPLE AGENTS

P. SCHERMERHORN* and M. SCHEUTZ, University of Notre Dame

ABSTRACT

Conflicts over resources can be resolved in many ways, from fighting to sharing. We introduce a very simple mechanism for implicitly taking turns, called the 2-*turn-taking rule.* Agents adjust their tendencies to fight over a resource on the basis of the outcomes of previous encounters. Agents that possess this mechanism are shown to be effective in competition with agents that lack the mechanism, indicating that there is some benefit to fairness, particularly when it comes at such a low computational cost.

Keywords: Conflict resolution, turn-taking, cooperation

INTRODUCTION

In real life, agents (humans and animals) have different needs and desires of varying urgency, which they attempt to satisfy. These needs and desires can range from very basic ones in all animals, such as the need to eat, survive, or procreate, to more complex ones, such as the desire to be respected or the need to have social relationships. Needs and desires typically involve resources (the objects of the need or desire). When fewer resources are available than there are agents who need or desire them, agents would likely be in conflict over these resources. In its most general form, a conflict would end in one of three ways: (1) some agents win, and others retreat; (2) nobody wins the resource (everybody loses); or (3) the resource is shared (everybody gets some, but not all, of the resource). In this paper, we examine encounters of the first kind and study a mechanism that allows agents that implement item 1 to reach item 3 (over multiple encounters) while avoiding item 2; skipping item 2 would be of benefit to the entire population.

Previous work with agents that display their action tendencies (whether to continue an encounter or to abort it) has shown that taking other agents' displayed action tendencies into account leads to better group outcomes (Scheutz and Schermerhorn, 2004). For example, if it is obvious that an opponent is very likely to continue to fight over the resource (i.e., it has a high action tendency to fight) and ultimately win the encounter, it is not in an agent's best interest to enter the fight when it is less likely to continue to fight (i.e., it has a lower action tendency to fight) and win the encounter, thus wasting resources fighting while gaining no benefit. Retreating immediately may also be costly, but compared with the cost of prolonged fighting, it is in the agent's best interest to retreat. Furthermore, it is in the more aggressive agent's best interest for its opponent to leave early because prolonged fights reduce the net benefit of the resource being contested.

* *Corresponding author address:* P. Schermerhorn, Artificial Intelligence and Robotics Laboratory, Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556 USA; e-mail: Paul.W.Schermerhorn.6@nd.edu.

So-called “rational” agents that can determine immediately who is likely to win an encounter (by comparing their own action tendencies with their opponents’ displayed tendencies) tend to have higher survival rates than agents whose ability to predict is worse, given a random distribution of action tendencies (Scheutz and Schermerhorn, 2004). This advantage is especially pronounced in competition with “asocial” agents (i.e., agents who do not attempt to take their opponents’ action tendencies into account). In an evolutionary environment, however, perfect prediction can lead to problems. Perfect predictors with high action tendencies win more encounters than agents with lower action tendencies, and, therefore, have access to more resources than others. They are then more likely to reproduce and pass on the higher action tendencies to their offspring, causing the average tendency of the population to skew upward. As time passes, successive generations continue this arms race until every agent has an action tendency of 1, at least in the absence of external pressures. Once every agent has a 100% probability of fighting, however, fights continue until one or more agents is unable to go on. Extended conflicts are very costly, and the population either declines or becomes extinct as a consequence.

One strategy that could overcome this problem is turn-taking (Neill, 2003). In a two-agent environment, for example, if each player would gain the resource benefit every other turn while paying only the retreat cost every other turn, both would benefit while avoiding costly fights. Thus, players obtain the benefit of perfect prediction (i.e., fewer resources spent on fighting) *and* the added benefit of fairness, so that the most aggressive agents would not hoard resources, and the arms-race scenario described above would be avoided. Generalizing the turn-taking strategy to environments with more than two agents is made difficult by the fact that agents do not enter conflicts with the same opponents every time. It is possible that both agents in a conflict took turns at losing in their respective previous encounter, so both would expect to win in the current encounter. An effective turn-taking strategy must resolve this issue. This paper describes a strategy for implicit turn-taking based on modifying action tendencies in response to wins and losses for multi-agent environments.

THE TURN-TAKING RULE

A fair turn-taking rule can be implemented in number of ways (see, for example, Iizuka and Ikegami, 2002). An agent can dedicate memory and processing resources to remembering with whom it has interacted and whose turn it is to win next. Over its lifetime, however, an agent can interact with dozens or hundreds of other agents. Requiring an agent to set aside resources for all of these agents in such an explicit turn-taking mechanism can be quite burdensome, especially for very simple agents. A mechanism that ensures fairness without consuming a substantial proportion of the available resources can flourish in “selfish” populations, given its potential benefits and low cost.

We introduce the 2-turn-taking rule (2TTR), which allows agents to keep track of their wins and losses by using the simple computational procedure described in the following.

Definition: Let r be the rest value of agent A, and let m be the action tendency. Then $2TTR(m)^+$ is defined (for losses) as:

If $m \geq r$, then $2TTR^+(m) = m + (1 - m)/2$.

If $m \leq r/2$, then $2TTR^+(m) = 2*m$; else $2TTR^+(m) = r + (2m - r)\ (1 - r)/2r$ (this maps values in the interval $(r/2, r)$ into $[r\ (1 - r)/2])$. Similarly, $2TTR(m)^-$ is defined (for wins) as follows:

If $m \geq r + (1 - r)/2$, then $2TTR^-(m) = m - (1 - m)$.

If $m \leq r$, then $2TTR^-(m) = m/2$; else $2TTR^-(m) = r/2 + r\ (m - r)/(1 - r)$. This maps $[r, (1 - r)/2]$ into $(r/2, r)$.

Corollary 1: Let $2TTR^{+,n}(m)$ $[2TTR^{-,n}(m)]$ denote the n-fold (recursive) application of $2TTR^+$ $(2TTR^-)$ to m. Then $2TTR^{-,n}\ [2TTR^{+n}(m)] = m$ and $2TTR^{+,n}\ [2TTR^{-,n}(m)] = m$, where $2TTR^{+,1}: = 2TTR^+$ and $2TTR^{-,1}: = 2TTR^-$.

Proof: We show only the first part by induction on n, the second being analogous. For $n = 1$, we consider all three cases:

(1) If $m \geq r$, then $2TTR^{+,1}(m) = m + (1 - m)/2 \geq r + (1 - r)/2$;
hence, $2TTR^{-,1}\ [m + (1 - m)/2] = m$.

(2) If $m \leq r/2$, then $2TTR^{+,1}(m) = 2 * m < r$; hence, $2TTR^{-,1}(2 * m) = m$.

(3) If $2TTR^{+,1}(m) = r + (2m - r)(1 - r)/2r$; hence, $2TTR^{-,1}\ [2TTR^{+,1}(m)] =$
$r/2 + \{r[r^+\ (2m - r)(1 - r)] - r\}/(l - r) = m$.

Now suppose the statement is true for $k = n - 1$. Then:

$$\begin{aligned} 2TTR^{-,n}[2TTR^{+,n}(m)] &= 2TTR^{-,1}(2TTR^{-,k}\{2TTR^{+,k}\ [2TTR^{+,1}(m)]\}) \\ &= 2TTR^{-,1}[2TTR^{+,1}(m)] = m, \end{aligned}$$

by definition, induction hypothesis, and base case, respectively.

For the following, let P denote a population of rational agents with the 2-turn-taking rule, and let $|P| = n$ be its size. Furthermore, let $@A$ denote the action of an agent A, and let $@(t)$ denote the set of action tendencies of P at time t, called "configuration." Finally, we assume that the action tendencies of the initial population of agents are at their rest values, all of which are between $2TTR^+(\min)$ and $2TTR^-(\max)$, where max is the largest and min the lowest action tendency/rest value in P, and we define region 0 to be the interval given by $[2TTR^-(\max), 2TTR^+(\min)]$. Positive regions $k > 0$ are then defined inductively by $[2TTR^{+,k}(\min), 2TTR^{+,k+1}(\min)]$, and similarly, negative regions are defined by $[2TTR^{-,k+1}(\max), 2TTR^{-,k}(\max)]$. The regions are defined such that when an agent competing against another agent in the same region k loses, the 2-turn-taking rule updates the agent's action tendency such that the losing agent's action tendency is in region $k + l$, and the winner's action tendency is in region $k - 1$, as shown by Corollary 2.

Corollary 2: Let a be an action tendency in region k. Then $2TTR^+(a)$ is in region $k + 1$, and $2TTR^-(a)$ is in region $k - 1$.

Proof: Let a be an action tendency in k. We distinguish three cases. Suppose $k > 0$, then $\min^k = 2TTR^{+,k}(\min) \leq a < \min^{k+1} = 2TTR^{+,k+1}(\min)$. Then applying $2TTR^+$ to all parts of the inequality given that $2TTR^+(a)$ is strictly monotone,

$$\min^k = 2TTR^{+,k+1}(\min) \leq 2TTR^+(a) < \min^{k+2} = 2TTR^{+,k+2}(\min);$$

that is, $2TTR^+(a)$ is in region $k + 1$. The other cases are shown analogously.

We can now show that the 2-turn-taking rule in combination with the rational agents is fair in a clearly specified sense: the difference between wins and losses is bound by $int(n/2) + 1$, where $|P| = n$ is the size of the population of competing agents. First, observe that all action tendencies are less than $2TTR^{+,int(n/2)+1}(\min)$ and greater than $2TTR^{-,int(n/2)+1}(\max)$ for rational agents. (This lemma essentially uses the rational agent's decision rule and is not true of other agents, e.g., probabilistic agents.)

Lemma 1: For every agent A in population P size $|P| = n$ of rational agents, all action tendencies $@A$ are less than $2TTR^{+\ int(n/2)+1}(\min)$ and greater than $2TTR^{-,\ int(n/2)+1}(\max)$.

Proof: The following lemma shows that the spread is at least $n/2$ in each direction.

Lemma 2: For each n there is exactly one configuration, in which each positive and negative region $n/2$ inhabits exactly one agent (with the 0-region also inhabiting one agent for odd n and empty for even n) and the configuration can be reached from the initial configuration (with all agents inhabiting region 0).

To see that the spread is at most $n/2$ in each direction, suppose that there is an agent A whose action tendency $@A=a$ is greater than $2TTR^{+,\ int(n/2)+1}(\min)$, i.e., $a \in$ region $int(n/2) + 1$. Then A must have lost a fight against another agent with a higher action tendency than $2TTR^{+,int(n/2)}(a)$ in region $int(n/2)$ given the rational decision rule. However, that means that two agents were in region $int(n/2)$, which is not possible. Suppose it were possible, then by backwards induction using the above argument, that at least two agents must have inhabited region $int(n/2) - 1$ at some point, and so forth. Eventually, after $int(n/2)$ steps, we reach region 1, which also must have two agents in it, but that is impossible, since it means that there is a total of $2 \cdot (int(n/2) + 1) > n$ agents (given that positive and negative regions are symmetric for the maximal spread).

Now we can prove that the difference between wins and losses is bounded by a fixed parameter d for all agents for any number of interactions.

Lemma 3: For an agent population P of size n, a d exists such that $|wins - losses| < d$ for all agents in P for any configuration.

Proof: First, observe that $losses - wins$ denotes the region the agent is in at any time, given that the agent started out in region 0 (as we assume for all agents). Whenever an agent wins/loses, the agent is put in a lower/higher region and its number of wins/losses is increased/decreased. At the best/worst, the agent can be in the highest/lowest region $int(n/2)/{-int(n/2)}$, corresponding to $int(n/2)$ losses/wins. Hence, the largest possible difference between wins and losses, $|wins - losses|$, is bounded by $int(n/2)$ for every agent.

This fixed bound *d* implies that in the long run the difference between the agent's accumulated utility is at most the constant given by $int(n/2) * (w - l)$, where *w* is the utility of winning, and *l* the utility of losing. This leads to a definition of what it means for an agent's action tendency update rule to be *balanced* for an agent group *P*.

Definition: Given an agent group *P* with a uniform update rule, $R = @A$ of *A*'s action tendency for every agent $A \in P$. *A* is balanced if a fixed bound *d* exists such that for any number of random interactions of agents in the agent group $|wins - losses| < d$, for all agents in *P*.

The above definition captures the best possible balance in an agent population that is continuously engaged in competitive interactions. Although the best distribution would be achieved if all agents inhabited region 0, which would imply that $|wins - losses| = 0$ for all agents, this is not always going to be the case, given that competitions occur. The worst possible distribution would be one in which $|wins - losses| = d$ for all agents, but this is fortunately only possible in the two-agent population (where there are only three classes). In any other population, only two agents can have *d* at any given time (the others must lower their differences). If interactions are chosen at random, in principle, nothing prevents an agent population from assuming any of the possible configurations determined by the initial distribution of agents and their population sizes.

We now can state the main theorem of this section, which follows from the above lemmas and corollaries:

Theorem: The 2-turn-taking rule is balanced for rational agents.

The theorem is valid for *rational agents* as described above who make a determined decision based on the action tendencies of both agents in a conflict. These agents basically treat the action tendency as a counter that keeps track of how many wins or losses an agent has. The *asocial agents*, on the other hand, use the action tendency as a probability that they will decide to fight. For this reason, it is still possible for both agents to flee, for both agents to fight, or even for the agent with the lower action tendency to stay while the other retreats. The 2-turn-taking rule serves as a nondeterministic place-keeper whose behavior could, in the short run, appear unfair, but, in the long run (as the number of encounters approaches infinity), should be fair. In fact, experimental results described below indicate that even over the relative short run, in which agents average only a handful of conflicts over their lifetimes, the probabilistic version of the turn-taking mechanism provides its holders with an advantage over "selfish" agents. The details of the formal argument are currently under investigation; however, at this time, the results outlined below support this hypothesis.

EXPERIMENTAL RESULTS

A series of experiments were run to test the effectiveness of the turn-taking mechanism. The experiments took place in the SimWorld artificial life environment (Scheutz, 2001). SimWorld provides a world in which agents forage for food, procreate asexually when they are sufficiently old and have enough energy stores, and interact with one another during the pursuit of these goals. Agents in close proximity are considered to be in conflict over whatever resources are present (if any). In a given cycle during a conflict, an agent may decide to retreat or fight. Each action carries with it a substantial cost. The benefit of retreating is that the cost is paid only once, whereas the cost of fighting is paid as often as both agents decide to fight. Agents who fight until their opponents retreat (or die!) obtain the benefit of whatever resources are in the immediate vicinity.

Agents decide how to react in an encounter on the basis of their action tendencies for conflicts. Each action tendency is mapped onto the range from 0 to 1, and can be thought of as the likelihood that an agent will decide to stay and fight in an encounter. Some agent types (so-called “asocial” agents) consider only their own action tendencies when making decisions about fighting, so their action tendencies map directly onto their probability of fighting. Others, however, mediate their probability of fighting based on a comparison of their own tendencies with those of their opponents. For these agent types, the probability of fighting increases when the agent’s action tendency is higher than that of its opponent and decreases otherwise. The “rational” agent is an extreme case of this type that will fight if and only if its action tendency is higher than that of its opponent.

In addition to these rational and asocial agent types, turn-taking versions of each type based on the 2-turn-taking rule are tested. The first set of experiments tests the agent types in homogeneous environments. Each agent type is placed in an environment containing only agents of its own kind. The environment is unbounded; however, food is randomly generated in only a $1{,}440 \times 1{,}440$ area; agents may wander out of this area, but must return to replenish their energy stores. New food sources are generated at random locations within the food region with a probability of 0.5 per cycle. Agents reproduce in roughly 350-cycle generations; the simulations were run for 10,000 cycles. Twenty agents are generated and placed at random locations within the food region at the beginning of each experimental run; performance is measured by counting the number of survivors at the end of an experimental run. The numbers given here are averaged over 80 experimental runs with differing random initial conditions.

Each agent type was tested using two methods to determine initial action tendencies or action tendency rest values (for non-turn-taking agents and turn-taking agents, respectively). The first method assigns a value randomly chosen in a Gaussian distribution centered at 0.5 to ensure diversity in the action tendencies and rest values throughout the experimental run. The results from experiments using this method are given in Figure 1. The first four columns depict the average number of survivors of each of the four types (asocial, asocial turn-taker, rational, and rational turn-taker) in homogeneous environments where they compete against only agents of their own kind. For both asocial and rational agent types, the normal agents outperformed their turn-taker counterparts. In the absence of turn-takers, it would appear that the unfair approach is more efficient.

Placing normal and turn-taking agents in the same environment yields mixed results. The six columns on the right in Figure 1 are paired to indicate that these results are from heterogeneous environments in which two agent types competed. The first pair (asocial vs. asocial turn-taker) compares asocial agents with asocial turn-takers; the turn-taking agents enjoy a pronounced advantage over the normal asocial agents. The second pair (rational vs. rational turn-taker) shows that no significant difference is evident between rational agents and rational turn-takers. The third pair (rational vs. asocial turn-taker) is of interest in light of previous results (Scheutz and Schermerhorn, 2004) in which normal asocial agents failed to average even one survivor against normal rational agents. The improvement to averaging more than 15 survivors with the addition of the turn-taking mechanism is testimony to the benefit of the mechanism.

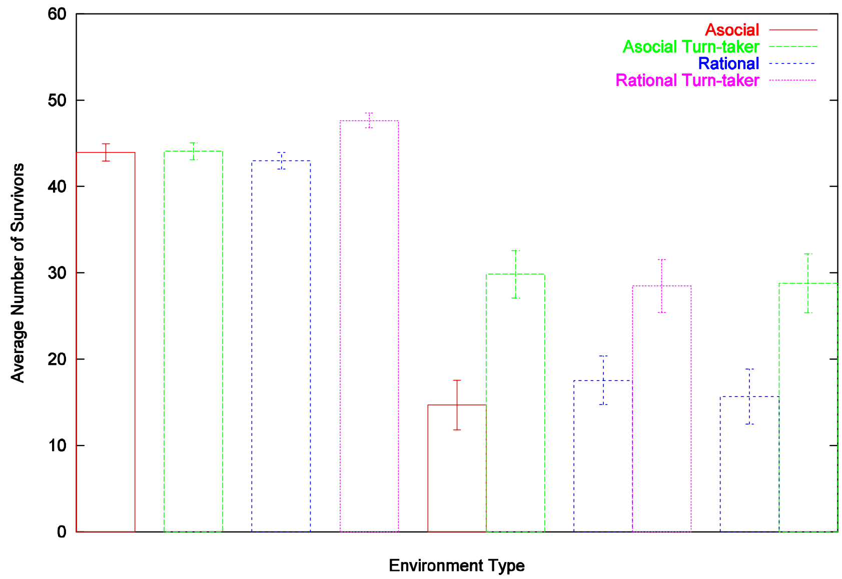

FIGURE 1 Survival rates for Gaussian-distributed action tendencies for asocial, rational, asocial turn-taker, and rational turn-taker (the four bars on the left indicate results from homogeneous environments, whereas the six on the right are results from heterogeneous environments)

Initial experiments used only the Gaussian method for determining action tendencies and rest values. We predicted, however, that in an evolutionary context, agents would engage in a self-destructive arms race, driving their action tendencies up to the point where the population could not sustain itself. To explore this possibility, further experiments were conducted in which agents of both kinds inherited their tendencies and rest values. This experiment gives the normal agents a chance to raise their own action tendencies, albeit over a longer time scale than the turn-takers, and in all likelihood in only one direction. The agents in the initial group of a simulation run were given values using the same Gaussian distribution, but thereafter the values were inherited by their offspring (without modification). The idea is to test whether this turn-taking mechanism is really an effective method of conflict resolution that could avoid the trap of a destructive arms race.

Figure 2 presents the results of this set of experiments, which we find to be encouraging. The advantage of the normal agents has vanished in the homogeneous environments. In fact, rational turn-takers outperform normal rational agents, which is in line with our prediction that an unrestrained increase in the average action tendency would eventually prove costly. In the mixed environments, the asocial turn-takers retain their advantage over normal asocial agents (although the gap closes somewhat). The other two pairings are more dramatic. Rational turn-takers now perform significantly better than their non-turn-taking counterparts. Asocial turn-takers also now perform significantly better than normal rational agents (recall that normal asocial agents perform very poorly against rational agents). These results support the hypothesis that fair turn-taking can be an effective strategy for avoiding inflation of action tendencies to destructive levels.

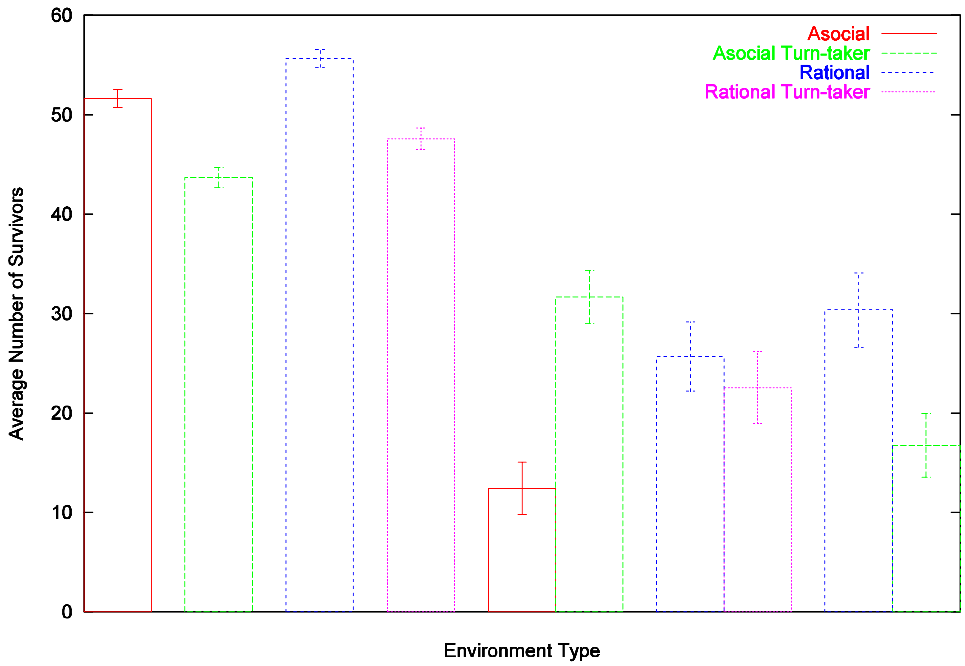

FIGURE 2 Survival rates for inherited action tendencies for asocial, rational, asocial turn-taker, and rational turntaker (the four bars on the left indicate results from homogeneous environments, while the six on the right are results from heterogeneous environments)

CONCLUSIONS

We describe a simple method for implicit cooperation via turn-taking. The 2-turn-taking rule specifies that an agent should modify its behavioral tendencies in response to winning or losing an encounter with another agent. The 2-turn-taking rule is guaranteed to be fair for encounters between two agents that employ a deterministic decision procedure (i.e., who use their action tendency values as counters to keep track of wins and losses).

When implemented with simple artificial agents, the mechanism provides an advantage to its possessors, particularly when agents are allowed to take part in an "arms race" of inheriting the action tendencies of their parents. Such an arms race leads to progressively higher action tendencies, on average, which is destructive to non-turn-takers who have no mechanism to back down from time to time, and so become engaged in costly extended conflicts.

The results of our simulation indicate that fairness can be a winning strategy against selfish agents, and that, moreover, the proposed mechanisms—the 2-turn-taking rule—can even be implemented by agents who do not or cannot take other agents' action tendencies into account, while retaining its benefits for the entire agent population.

It is important to note that the turn-taking mechanism is a very simple one whose cost is small enough that, for all practical purposes, it can be ignored. In the evolutionary environment, this is key because it indicates that the substantial benefits outlined above will not be offset by equally substantial costs as would be the case for a more sophisticated turn-taking mechanism, one that, for example, relied on memory to recall the history of an agent's interactions in order make good decisions about whose turn it is to win. On the basis of Figures 1 and 2, this extremely low-cost method of ensuring fairness would almost surely invade any unfair population in which it arose.

We are currently exploring further the theoretical properties of the probabilistic version of the 2-turn-taking rule and expect that long-term evolutionary investigations (including the ability to mutate) will reveal its advantage in a great variety of environmental settings and conflict types.

REFERENCES

Iizuka, H., and T. Ikegami, 2002, "Simulating Turn-taking Behaviors with Coupled Dynamical Recognizers," in M.A. Bedau, R.K. Standish, and H.A. Abbass, eds., *Artificial Life VIII, Proceedings of the Eighth International Conference on Artificial Life*, Cambridge, MA:MIT Press.

Neill, D., 2003, "Cooperation and Coordination in the Turn-taking Dilemma," in *Proceedings of the 14th International Conference on Game Theory at Stony Brook* [www.sunysb.edu/gametheory/ Conf03/ttdconf.pdf].

Scheutz, M., and P. Schermerhorn, 2004, "The Role of Signaling Action Tendencies in Conflict Resolution," *Journal of Artificial Societies and Social Simulation*, **7**(1), Jan. [http://jasss.soc.surrey.ac.uk/7/1/4.html].

Scheutz, M., 2001, "The Evolution of Simple Affective States in Multi-agent Environments," in D. Cañamero, ed., in *Proceedings of AAAI Fall Symposium*, Falmouth, MA: AAAI Press, pp. 123-128.

SECTIONALISM INDEX FOR ORGANIZATIONS: ANALYZING SECTIONED RANDOM NETWORK MULTI-AGENT MODEL

K. YUTA* and K. SHIMOHARA, Advanced Telecommunications Research Institute International, Kyoto, Japan

ABSTRACT

A new model is presented for indexing the sectionalism, or "clanism," in a unit of a section. This model involves the metaphorical walls between sections, which should be regarded as pathology of an organization. It is beneficial to understand organizational conditions among sections to assist in conflict management and to enhance cooperation between them. A sectional conflict in an organization is thought to result from cultural gaps between sections. These gaps are based on differences of specialty and communication density inside and outside those sections, respectively. The model is designed by using multi-agents, which consist of sections and have sectional specialty, and an alternative network that is sectioned in a random network. Empirical results show that the proposed index is superior to conventional indices with regard to capturing the sectioned organizational network conditions. Furthermore, the model clearly illustrates the effect of cross-sectional links on sectionalism by following the so-called "power law."

Keywords: Sectionalism index, conflict management, sectioned random network, network multi-agent model, power law of cross-sectional link

INTRODUCTION

Use of the term "sectionalism" in the context of human organizations implies the phenomenon of a sectional "wall." When sectionalism is a problem, the members of the section think and act only for their own section's benefit. This problem raises issues about optimization based on biases of members, especially those biases reinforced by sectional division. It is like excess adaptation. If there is no need to cooperate with other sections, a section performs well because of local optimization; this is far from global optimization, which should occur across the entire company. Although it is important that sections specialize to increase efficiency, a human network divided by sections with apportioned work results in overspecialization.

Bureaucratic entities and large traditional companies in Japan are often cited as prime examples of this sectionalism problem. Although their importance is recognized, little is known about organizational models of sectionalism. To study sectionalism, it is necessary to analyze the sections that constitute a company, not the individual nor the entire company.

For example, suppose a company has three sections: Sales, Manufacturing, and Operations. The signs of typical sectionalism are discussed below. Members of the Sales Department say that they actually earn the organization's money. In their view, the

* *Corresponding author's address:* Kikuo Yuta, Human Information Science Laboratories, ATR International, 2-2-2 Hikaridai, Keihanna Science City, Kyoto, Japan 619-0288; e-mail: kikuo@atr.co.jp.

Manufacturing Department merely produces the goods. Sales would be unlimited if the products were more customer-oriented. On the other hand, members of the Manufacturing Department say that they produce the products; thus, benefits are derived from their work. They think the Sales Department performs poorly, and, with sufficient sales power, their quality products would sell in unlimited quantities. Experience proves that indirect communications between Sales and Manufacturing relayed by Operations is apt to increase this type of bias.

As for organizational issues, work has been performed in various research fields, such as motivation, management, and economics (Milgrom and Roberts, 1992; Lazear, 1998). If related to conventional research, sectionalism would fall under the domain known as "conflict management." Little research on conflict management, however, has been conducted using a quantitative network approach. Sociologists have performed most of such research on organizations as conventional human networks, and analysis of important people by such means as sociograms has produced plenty of useful knowledge (Yasuda, 1997). Whereas individual and entire network analysis is very important, as reported by Torenvlied and Velner (1998), even that study is lacking in its investigation of sectionalism.

In recent years, models such as the Small World Network (SWN) and the Scale-free Network (SFN) have helped to expand knowledge of networks (Albert, et al., 2000; Barabási, 2002). These can be qualitatively called homogeneous networks. In our research, a conservative random network is designed to improve sectioned structures, each node has the property of belonging to its own section, and nodes are set as in a simple multi-agent model (reaction only, no learning nor evolution).

By using this network model and applying a multi-agent model as reported by Weiβ (1999), this study focuses on sections and networks in organizations. We also design an index to grasp the emerging output of this structure, called "sectionalism," in excessive case. Computer simulation is issued to complete three tasks:

1. Verification of whether a simulation expresses the behavior of sectionalism,

2. Comparison of conventional network indices, and

3. Investigation of the influences of cross-sectional links for sectionalism.

NETWORK AND INDICES

This section briefly and informally describes the type of network and its index related to this research. In this study, the network is characterized by three conventional indices for comparing a presented index. One of these is the average degree D. The other two are the average path length L and the cluster coefficient C, as reported by Watts and Strogatz (1998).

The average degree D is the arithmetic average of the edge, which the nodes of all networks have. This index indicates how many nodes and links are linked per node.

The average path length L is the average of the course length $d(ij)$ for the group of nodes of all networks. It is the number of edges that the path length $d(ij)$ needs to connect node i and node j ($i \neq j$) along the shortest path. By using the average, it is possible to search for the average path length $L(i)$ of node i. Therefore, $L(i)$ is calculated for all nodes, and the average path length L of the entire network is acquired by averaging them. These formulas are as follows:

$$L \equiv \frac{1}{N}\sum_{i=1}^{N} L(i)\,,$$

and

$$L_i \equiv \frac{1}{N-1}\sum_{j \neq i}^{N} d(ij)\,.$$

The cluster coefficient C quantifies the gathering condition of the node that constitutes the network. A cluster consists of three nodes whose edges are linked to form a triangular structure. If node i has $k(i)$ edges, a realizable cluster is formed from a number of combinations that choose two nodes from $k(i)$ nodes. Then, the cluster coefficient of node i is defined by dividing the number of clusters $E(i)$ that actually exist in a certain node by the number of possible clusters. The cluster coefficient C of a network is obtained by performing it for all nodes and calculating an average. These formulas are as follows:

$$C \equiv \frac{1}{N}\sum_{i=1}^{N} C_i\,,$$

and

$$C_i \equiv \frac{2E(i)}{k(i)[k(i)-1]}\,.$$

The most fundamental structure in network research is a random network as proposed by Erdös and Rényi (1959). This model has links by probability p_r to all nodes. The average degree of all nodes is calculated when the total number of nodes is N and the average degree is set to $N \times p$. The degree of each node becomes a Gaussian distribution centering on the average degree. As a character, the cluster coefficient is small and its average path length is short. The trait of this random network is that there are very few cliques because members become acquainted uniformly and broadly; moreover, one can become acquainted with others with just a few relays.

DESIGN OF ORGANIZATION MODEL OF SECTIONALISM

Architecture

We first describe the sectioned network design and architecture of agents and simulation. To initialize a simulation, a sectioned network structure is constructed by extending the random network design method. A multi-agent simulator then obtains information on link, section, and initial data, which depend on sectional traits where the agent belongs. Interaction is repeated until the sectional change converges. Finally, deviations of agent data in each section are totaled to obtain the sectionalism index.

Design of a Simple Sectioned Random Network (Simple SRN)

It is better to design a model that is as simple as possible to avoid confusion. Since the network considers interaction among those sections, at least three sections are needed so that influence can extend indirectly. (Recall the example company above, which includes three sections: sales, operations, and manufacturing.) A simple SRN needs at least 100 nodes to enable calculation of the probability of links with an error of 1%. The organization design is then set with three sections of 100 nodes each, giving a total of 300 nodes for the simplest model.

Sectioned random networks need two types of probability. A self-section link probability is fixed first, followed by two other probabilities of links to cross-section nodes set to half of the remaining non-self-section link probabilities from 100%, as shown in Figure 1. At this time, when the self-section link probability is set at 33.3%, it is equal to two remaining cross-sectional link probabilities. This uniformity is important for a sectioned network. Consequently, a uniformity coefficient *UC* is introduced. The formula below exists for the general case wherein a symmetric sectional random network has *N* sections. A *UC* of 100% means a conventional nonsectioned random network. A sample image of SNR's adjacency matrix is shown in Figure 1 for uniformity coefficient *UC*:

$$UC \equiv (1 - Pself\,)\frac{N}{1-N}\,.$$

The property of this simple SRN is shown in Figure 2. A decrement in *UC* causes the cluster coefficient to increase, but the average path length does not change. Degrees are designed to maintain the average. Each degree is then totaled and checked, with results showing a Gaussian distribution.

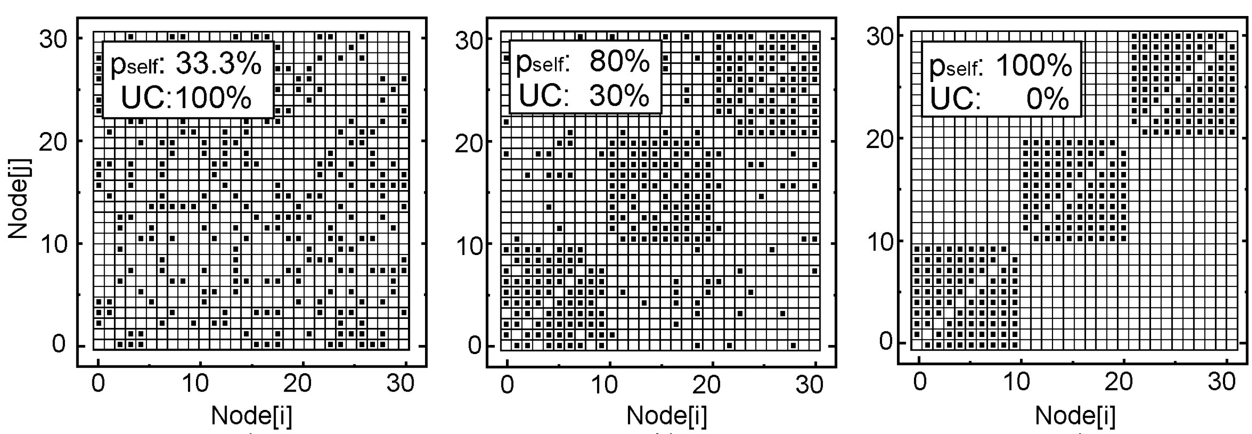

a) b) c)

FIGURE 1 Simple SRN adjacency matrices, with setups of (a) 100%, (b) 30%, and (c) 0% of *UC*, respectively (The black dots were plotted when nodes *i* and *j* are linked. *UC* is the uniformity coefficient, and *Pself* is the probability of a self-sectional link.
Shown is a 30-node matrix for easy understanding.)

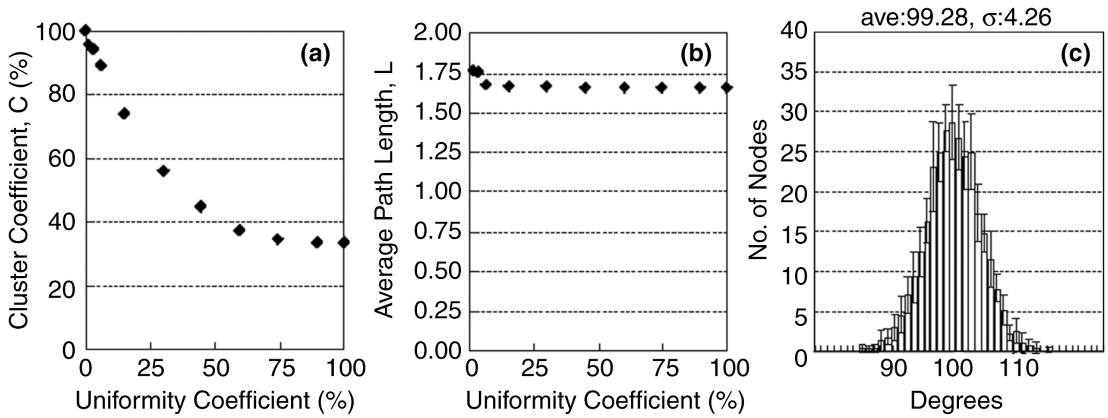

FIGURE 2 (a) The relation between a cluster coefficient *c* and uniformity coefficient *UC*, (b) the relation between average path length and *UC*, and (c) the distribution of degrees (The standard deviation of the average degrees from 10 trials was 4.26 from a degree average of 99.28.) Shown is a 100-node simulation.

Design of Agents and Interaction

This section describes the agent and its interaction mechanism. An agent's initialization acquires link information first. An agent then has two additional pieces of information: a section code and a sectional culture set. For example, a company that consists of three departments — sales, operations, and manufacturing — has three typical sectional culture factors, and these three factors make one sectional culture set. All agents have their own sectional culture set initializing [100, 100, 100], meaning each section culture [sales, operating, manufacturing]. The section code indicates the agent's job and sets the tendency of his own sectional culture factor to increase. After initialization, the simulation starts to act on management and continues until 1,000 turns have been completed, resulting in an experientially obtained convergence turn number.

At each turn, all agents have two opportunities to change their sectional culture set. The first is under the influence of the section to which the agent belongs. The agent can be infected with its own sectional culture factor, which enlarges the factor value and causes a 10% probability per one turn. The second opportunity is interaction with a linked agent. An agent interacts with another agent at each turn. Both agents simultaneously recognize the difference in their sectional culture set. Thus, they share a 25% decrease in the difference between each other's set. A total of 300 sectional cultures are maintained. The impact of the self-sectional tendency is 30%. An example of infection by a self-sectional culture is as follows. Sections A, B, and C have sectional culture factors *fa*, *fb*, and *fc*, respectively. Each agent has a culture set of [*fa*, *fb*, *fc*]. Agent 'Agt' belonging to Section A is initialized Agt[100, 100, 100]. If this agent obtains the self-sectional tendency at the first turn, it would be Agt[130, 100, 100]. The sum total is 330, which is the overflow capacity of one agent. This number is then normalized and reset to 300 Agt[118, 91, 91]. Finally, *fa* increases by ~18.

An example of interaction between two agents is described next. When agent *A*[150, 100, 50] and agent *B*[120, 80, 100] interact with each other, they recognize the difference *d*[–30, –20, +50] between them. If *A* puts *d* in its self-sectional culture set, *A* joins the same set as *B*. This means that *A* has been completely influenced by *B*. If *d* is 50%, it is a completely shared model. In this research, we set 50% as a completely shared model. This percentage represents a 25% change for each agent. In this process, the agent does not care about its target section or its own section. This mechanism has a strong effect when the difference is large and a weak effect when it is small. Against an agent in same the section, the effect of the self-sectional culture tends to spread quickly. Meanwhile, for interaction with another section's agent, it takes a role of rectifying the deviation. Figure 3 shows the simulation result of agent interaction in a Simple Sectioned Random Network Multi-agent Model (simple SRNMAM).

Since the simple SRNMAM has a symmetric structure, charts for Figure 3 show Section A only. For cases with no cross-sectional link, other sectional cultural factors (scf's) are washed out by the 300th turn. The tendency for uniformity also increases as cross-sectional links grow in number. Considering the indication of the deviation of the scf in a section's unit, it is a perfect sectionalism index. This chart also shows that 1,000 turns is enough for convergence to occur at this setting.

Sectionalism Index (SecSD)

In this section, we describe the design and calculation of a sectionalism index. For a symmetrical network structure, such as a simple SRNMAM, it is simple to understand sectionalism by using the standard deviation. As a result, the average of each scf is the same as for the total of all sections: [100, 100, 100], as shown in Figure 4a.

This approach, however, is not available to measure in asymmetric network structures because it is strongly influenced by a constituent number and its links. Furthermore, deviations arise in *fa*, *fb*, and *fc* for the entire company average; this is the so-called corporate culture (see Figure 4b).

To grasp the section level accurately, several procedures are followed. We focus on the variances of each factor with respect to the average for the entire company. First, variances of all factors are totaled for each unit in a section. Second, the total variance is divided by the number of agents in a section. Third, that number is divided by three, which is the number of sectional cultures. Finally, the square root of the number is taken and considered as the sectionalism index of this research, and this number is the extended standard deviation. We now refer to this index as "Sectionalism Standard Deviation" (SecSD), the standard deviation by which a section is

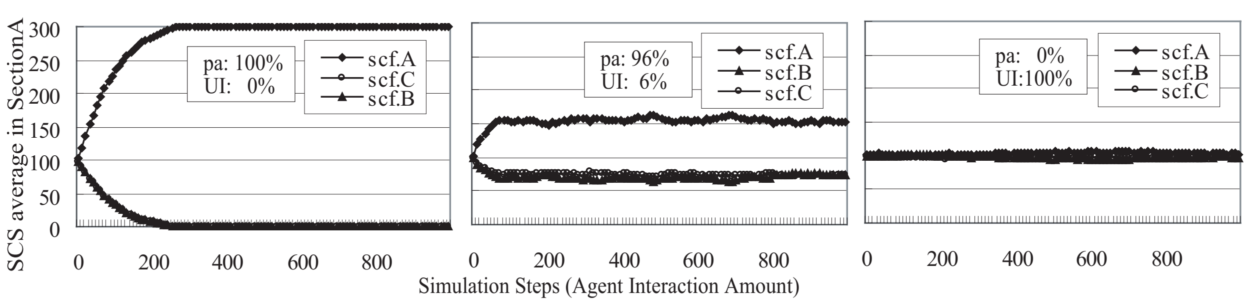

FIGURE 3 Simulation result of Simple SRNMAM (SCS = sectional culture set; scf.A = sectional culture factor of Section A)

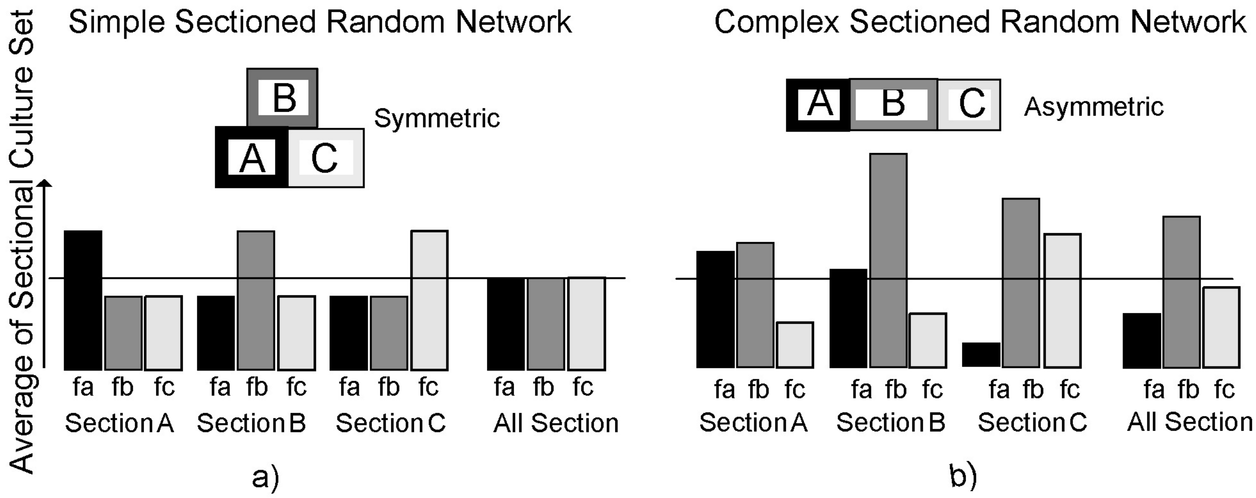

FIGURE 4 (a) Image of symmetric network simulation result, such as a simple SRNMAM, and (b) image of asymmetric network simulation result, such as a complex SRNMAM

characterized in this study. The organization: $O \equiv (A_1, A_2, A_3, \ldots A_N)$ has N agents. An agent $A(i) \equiv (a_{i1}, a_{i2}, \ldots, a_{is})$ holds s sectional culture factors. The average of all agents in an organization is set to

$$\bar{a} \equiv \left(\frac{1}{N}\sum_{i=1}^{N} a_{i1}, \frac{1}{N}\sum_{i=1}^{N} a_{i2}, \frac{1}{N}\sum_{i=1}^{N} a_{i3}, \ldots, \frac{1}{N}\sum_{i=1}^{N} a_{is} \right) \equiv (\bar{a}_1, \bar{a}_2, \bar{a}_3, \ldots, \bar{a}_s).$$

Then, a group G that includes O has g agents. The sectionalism index of group G is calculated as follows:

$$G \equiv (A_1, A_2, A_3, \ldots A_g), V_G \equiv \sum_{i=1}^{g}\sum_{j=1}^{s}(a_{ij} - \bar{a}_j)^2 ,$$

$$V_{\bar{G}} \equiv \left(\frac{V_G}{g}\right), V_{\bar{G}}: \text{ variance per agent in group,}$$

and

$$SecSD \equiv \sqrt{\frac{V_{\bar{G}}}{s}} \qquad SecSD: \text{ sectionalism standard deviation as “sectionalism index.”}$$

Design of Complex Sectioned Random Network (Complex SRN)

In this section, the network structure that can verify the performance of an index of grasping sectionalism in a section’s unit is designed manually as shown in Figure 5.

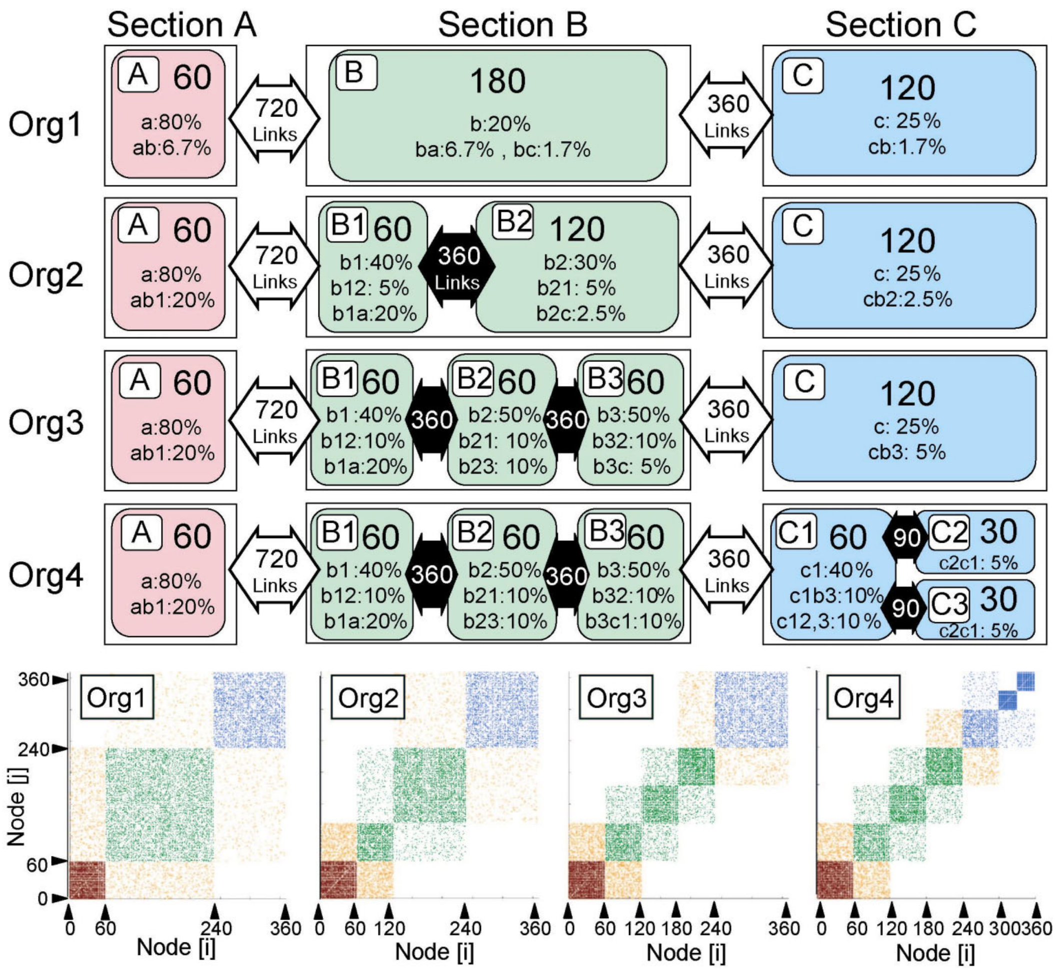

FIGURE 5 The complex SRN design from Org1 to Org4 with each adjacency matrix (manual setting of section, links, cross- or self-sectional link probability)

The purpose of this research is to optimize the capability of this index. It consists of three sections: Section A has 60 nodes, Section B has 180 nodes, and Section C has 120 nodes, for a total of 360 nodes. There are 720 interacting links between Sections A and B, 360 between B and C, and no links between A and C (Org1). The general structures between these sections are invariant. Furthermore, each section has internal links set between subdivisions. For example, Org2 divides Section B into two smaller classes — 60 + 120. Org3 divides Section B into three even smaller classes — 60 + 60 + 60. Org3 has separated subsections of B from both A and B. Org4 contains additional divisions to Org3, where Section C is divided into three subsections. In Figure 5, A, B, C, etc., show the connection probability of a link, respectively, and b, c, b1, b2, and b3, and c1, c2, and c3 show the self-section link probability. The symbol ab shows the connection probability from a to b.

EXPERIMENTAL RESULT

Validation of Sectionalism Index (SecSD) via Complex SRNMAM Simulation

The simulation result of complex SRNMAM is shown in Figure 6. The degree of sectionalism in each section or the entire organization can be verified. When the composition of the set in each section is investigated, it is found that there is no difference between the compositions of the entire organization and the organizations Org1-4. The composition of sections, however, reflects differences in internal structures.

Although the entire organization acquired the same composition with a sectional culture set in each of Org1-4 (see Figure 6c; each Org1-4<ALL> is the same), SecSD proves there are differences. In this model, each agent interacts without losing all its sectional culture set. Therefore, the final total average depends on the number of constituents in each section. Otherwise, an internal structure is affected by the diffusion speed of the sectional culture and the volume of relay that changes the sectional deviation between sections (see Figure 6c, pentangle-attached bar of each 1-SecC ,…, 4-SecC).

This sectionalism index measured the different sizes of organizations simultaneously (see Figure 6 where each Org <ALL> SecA, B, C is compared in the same graph). When the entire organization is put into a larger organization, the former organization becomes a section. Therefore, it is meaningless if the index becomes impossible to use as a means of comparison between sections and entire organizations. The index's important features are its flexibility for size and consistency for measurement. The index proposed here avoids the problem by asking for the deviation of sectional culture set per single agent.

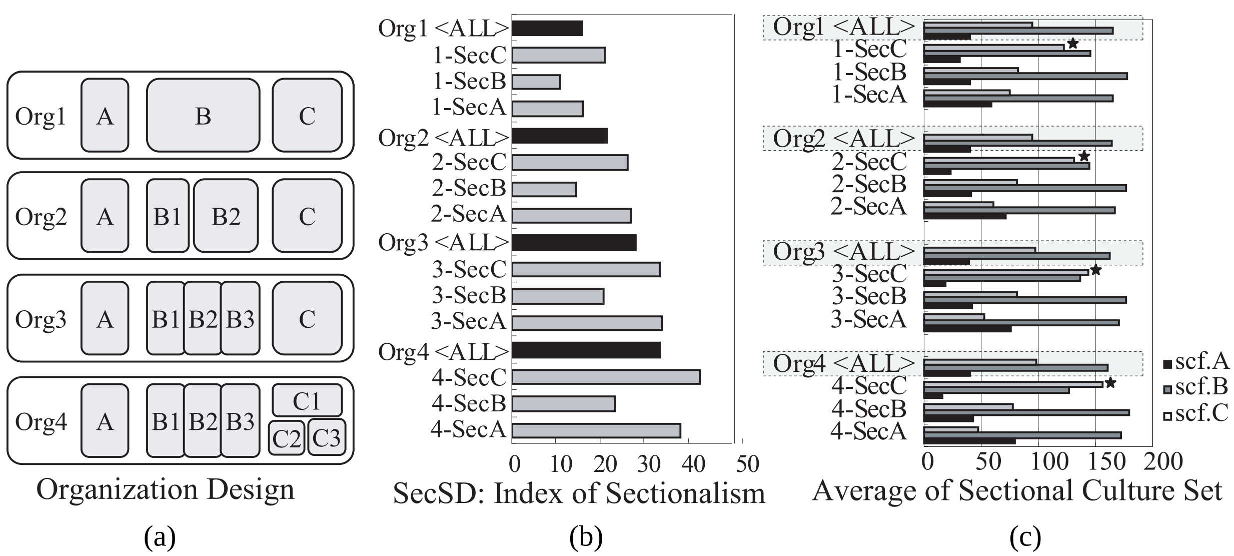

FIGURE 6 Complex SRNMAM simulation result for (a) the organization structure summarized for reference, (b) the value of SecSD for each section and for the entire company, and (c) the composition of each sectional culture factor average in each section

Table 1 and Figure 7 show the comparison between SecSD and the conventional index. They verify that the proposed index is far more sensitive than the usual network index for determining organization sectionalism behavior. SecSD has the additional influence of the cluster coefficient. The structure of a cluster causes deflection between sections, which has a greater effect on SecSD. Consequently, SecSD also reflects the differences in sectional cultures. In the case of structural change from Org1 to Org4, the separation between Sections A and C has a more considerable effect on SecSD than the cluster coefficient, especially in this organization design, which keeps the total links between sections for the entire Org1-4. This means that the division in Section B decreases the speed of sharing these sectional cultures. The internal share in Section B acts as a relay between A and C. The decrease in relay increases the deviation between Sections A and C in the sectional culture.

TABLE 1 Network index of four organizations of complex SRNMAM

Organization	Cluster (%)	toOrg1[a]	Length	toOrg1[a]	SecSD	toOrg1[a]	Degree
Org1	24	1.00	2.05	1.00	9.17	1.00	43.11
Org2	29	1.20	2.19	1.07	12.41	1.35	42.77
Org3	32	1.31	2.41	1.18	16.20	1.77	43.01
Org4	43	1.76	2.56	1.25	19.42	2.12	42.95

[a] In this table, toOrg1 indicates the rate of sensitivity to each index of Org1.

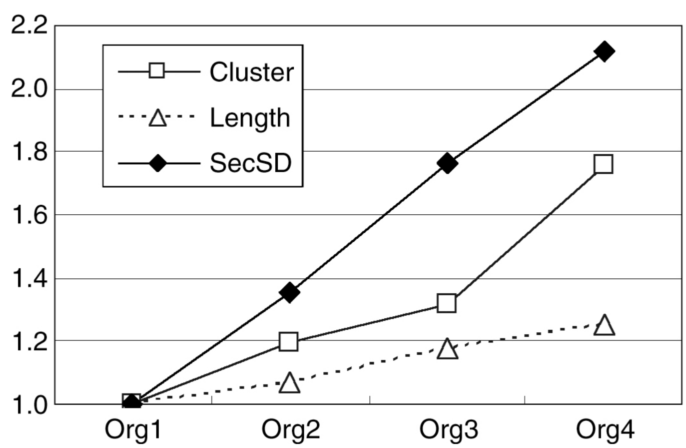

FIGURE 7 Comparison of indices' sensitivity performance to different organizational structures (Note that the structure becomes more complex from left to right.)

Power Law of Cross-sectional Link Effectiveness for Sectionalism

Incremental cross-sectional links are effective for reducing the sectionalism tendency in this model as well as in real companies. The cross-functional team focused on the field of management as reported by McDonough (2000) and is considered to be one example of sectionalism reduction. We present the simplest model to study this reduction. The relation between SecSD and *UC* is investigated in this simple SRNMAM. Results show that at a *UC* of 10% or less, the index value rises rapidly, but a *UC* greater than 20% is enough to reduce SecSD. Focusing on this matter, when both axes in the graph are changed to a logarithmic scale, a straight line — known as the power law — is observed, as shown in Figure 8.

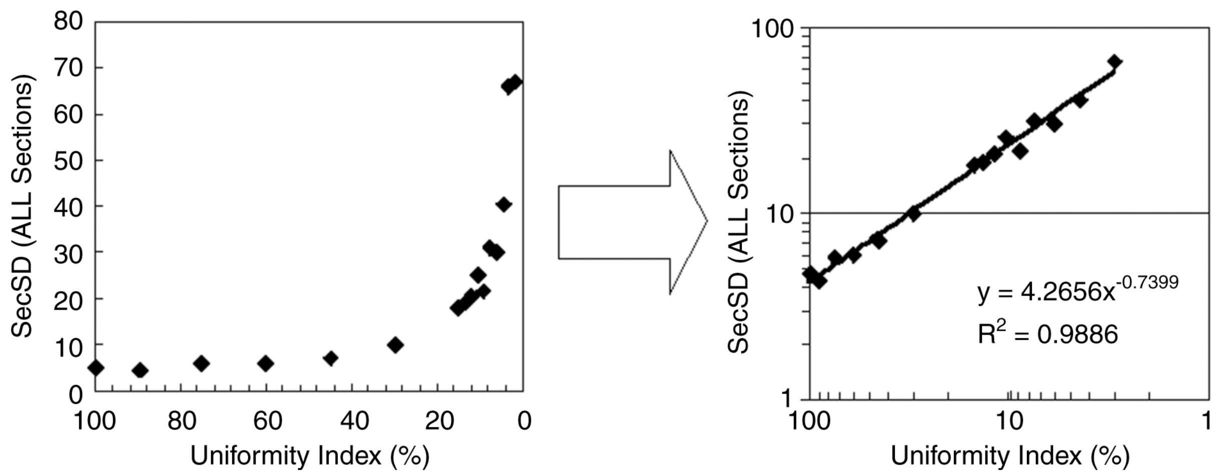

FIGURE 8 Power law of effect of cross-sectional links to sectionalism index, SecSD

Why did the uniformity coefficient carry out the power law distribution to SecSD? There has been no explicit setup in which an exponential effect is derived. The power law is thought to influence the effect of section structure on a network. Since the members inside a section greatly influence each other, the influence of a single agent is shared at high speed. Where little connection occurs between crossover sections, then one new link of a crossover section will influence all links of the target. The power of the second link seems to serve as an exponential influence.

DISCUSSION AND FUTURE DIRECTION

The power law is important in organizational design in an actual company. This model shows the effectiveness of power of the cross-sectional link, in which 15% of the links initially reduce the 85% risk of sectionalism like a Pareto law. This reduction effect is exponential. In the human resource management division of a real company, these types of cross-sectional communications and events imply costs in time, money, and workforce. We believe that this type of indicator would help cut such costs.

We plan to expand this model to include organizational hierarchy, such as manager-subordinate interaction; change from an undirected graph to a directed graph to represent the network; and addition of weights to links. When the above knowledge is synthesized with a real survey, this model also expresses "the robustness of headhunting" and "the degree of organizational openness."

CONCLUSIONS

In this research, we present a sectioned random network multi-agent model and sectionalism index for organization with four achievements:

1. An illustration of sectionalism behavior,
2. A sectionalism index in order of section level and size-free,
3. A higher performance of SecSD than usual, and
4. Finding of the power law for determining the effectiveness of a cross-sectional link in reducing sectionalism.

ACKNOWLEDGMENTS

We are grateful to Wataru Souma, Yoshi Fujiwara, Osamu Katai, and Keiki Takadama for provocative discussions about network and multi-agent modeling. Special thanks are also due to Masaki Yasui and Masashi Hada for practical discussions about management. This research was conducted as part of "Research on Human Communication" with funding from the Telecommunications Advancement Organization of Japan.

REFERENCES

Albert, R., H. Jeong, and A.-L. Barabási, 2000, "Error and Attack Tolerance of Complex Networks," *Nature* **406**:378–382

Barabási, A.-L., 2002, "Linked — The New Science of Networks," Cambridge, MA: Perseus Publishing.

Erdös, P., and A. Rényi, 1959, "On Random Graphs," *Publicationes Mathematic (Debrecen)* **6**:290–297.

Lazear, E.P., 1998, *Personnel Economics for Managers*, New York: John Wiley and Sons, Inc.

McDonough, III, E.F., 2000, "Investigation of Factors Contributing to the Success of Cross-functional Teams," *Journal of Product Innovation Management* **17**(3):221–235.

Milgrom, P., and J. Roberts, 1992, *Economics, Organization and Management*, Englewood Cliffs, NJ: Prentice Hall Inc.

Torenvlied, R., and G. Velner, 1998, “Informal Networks and Resistance to Organizational Change: The Introduction of Quality Standards in a Transport Company,” *Computational & Mathematical Organization Theory* **4**(2):165–188.

Watts, D.J., and S.H. Strogatz, 1998, “Collective Dynamics of ‘Small-world’ Networks,” *Nature* **393**:440–442.

Weiβ, G., 1999, “Multiagent Systems — Modern Approach to Distributed Artificial Intelligence,” Cambridge, MA: The MIT Press.

Yasuda, Y., 1997, *Network Analysis*, Shin-yo-sha Co., Ltd. (in (Japanese).

DISCUSSION:

POLITICAL PROCESSES

(Saturday, October 4, 2003, 1:00 to 3:00 p.m., Session 1)

Chair and Discussant: *Richard Cirillo, Argonne National Laboratory*

Agent-based Residential Segregation: A Hierarchically Structured Spatial Model

Richard Cirillo: Good afternoon, everyone. I think we have some interesting papers this afternoon to discuss the application of agent-based modeling to political processes, which I think is something that we're all particularly interested in, particularly as we get into an election year.

The first speaker will be David O'Sullivan, from Pennsylvania State University.

David O'Sullivan: Before I get started properly, I should probably just say that the real motivation for the model that I'm talking about today was sort of methodological rather than applied, in the sense of, we're interested in the processes that this model looks at in terms of neighborhood formation in urban settings and residential segregations of populations and all of those sort of longstanding applied substantive issues. But we're also interested in developing, from a relatively well-known starting point, a model that we can build on over time to explore some methodological issues having to do with getting more substantial geography into agent-based models. So in some senses this paper could have sat fairly comfortably in the "Methods, Models and Toolkits" sessions on Thursday.

[Presentation]

Cirillo: Thank you, David. Do we have any questions? Comments?

Greg Madey: That was very interesting. A couple questions or suggestions. Have you or anybody, I'm just curious, looked at asymmetric tolerance?

O'Sullivan: You may have noticed in the demo that we had tolerance down there as something we could play around with. And we initially when we were getting the thing up and running we were playing around with that, and we sort of steadily realized that we needed to keep some things fixed. And so we haven't really explored any of that. Schelling's work, that bonded neighborhood stuff with the intersecting parabolas, the whites and the blacks, he explores different assumptions about the two different populations' tolerances for one another, how that affects the state space and its dynamics and sort of shows that there are stable combinations, given sufficient tolerance of one population or the other. For our work, we've just kept the tolerance stuff fairly fixed.

Madey: Then one more quick question. I remember a few years back in some communities I lived in, the local city council passed laws that said you couldn't put "For Sale" signs on your property. And so the other thing to look at is the velocity of the communication. How fast did the information transfer or become aware [that people are leaving], so even before someone left, people know someone's planning on leaving. Or maybe a fuzziness where maybe you're not even sure. And then I guess multiple populations, more than two colors.

O'Sullivan: Oh, certainly. There's dozens of people messing around with models of this kind. There's a guy at Texas A&M, Mark Fosset, who's actually got commercial software emerging out of stuff that he hasn't really published. And he has multiple different groups and this whole kind of configuration of stuff. We wanted something that we could build ourselves and work up from, rather than borrowing from that.

Cirillo: One more question?

Unidentified Speaker: You made a comment about lattices. Are you implying that you're considering three-dimensional structures so that you can account for social hierarchy as well? Is that what your suggestion is?

O'Sullivan: Well, I guess, in much the same way that Tom Howe talked about and relations being potentially, you know, arbitrary. What we're focused on is spatial rules for building the lattice, things like contiguity, neighboring within some certain distance, and relations of containment between spatial units in a hierarchy, but there's no reason in general why those need to be kind of planar or two-dimensional. So by implication, yes, you can move on to that three or higher dimensional types of lattice.

Cirillo: One more?

Unidentified Speaker: I'd be interested in seeing something where you combine some of what you're doing with some of the stuff that Ed MacKerrow talked about earlier, the dry grass model and the grievance model. There are some areas in Chicago where pretty much everybody knows that if you're of the wrong race you're in deep yogurt if you wind up in the wrong area. "Wrong" meaning black in white, or maybe white in black, for all I know. It has the obvious connection with the level of dissatisfaction and stuff like that.

David O'Sullivan: Right.

Unidentified Speaker: And you could measure, you know, propensity to nasty incidents, whether racial or similar to the Palestinian/Israeli conflict or any other number of groups.

O'Sullivan: There's actually a group in Israel at Tel Aviv who don't seem to be known much outside the geography community, Vival Portugali and Itzhak Beninson, and they've done really quite detailed modeling. They started out 10 years ago with reds and blues and greens, and then they kind of came out of the closet with what they were really doing, when they made it geographical. They have this very detailed plot level data for Yafo, which is a suburb of Tel Aviv, and they have Jews and Arabs in that setting. And they have some very similar dynamics going on there, but it's very much their own particular tool, and they have access to very, very detailed individual-level data for Tel Aviv which they've been working with. So, yes, you're exactly right.

They also explore something which is very interesting, which I didn't really get to talking about, the sense in which as a neighborhood changes, how people's tolerances change: they learn different tolerance from living in a certain kind of neighborhood. So the individuals are altered by the nature of the space, which is a really interesting sort of feedback.

We're interested also in how what are generally perceived as coherent neighborhoods might change over time. As you get persistence over time of one group dominating a particular area, and maybe around the edges of it, well, then, perhaps it starts to be perceived as a bigger space or a smaller space. So that sort of ghetto formation, or gentrification would be the opposite sort of a fact.

Implicit Cooperation in Conflict Resolution for Simple Agents

Richard Cirillo: Our next speaker is Paul Schermerhorn, from University of Notre Dame.

Paul Schermerhorn: Thanks. I guess this is a little bit of a change of pace, because what we are talking about here are very simple agents. And what we want to talk about specifically is a mechanism that we've used to impose an implicit cooperation upon them.

[Presentation]

Cirillo: Do we have any questions?

Dan Kunkle: If the turn-taking might be robust against collusion and group-formation in other agents, could that group of colluding agents take away the benefit that turn-taking agents seem to bring to the system?

Schermerhorn: I haven't thought about that. I'm not sure. You probably have an idea about how that would work?

Kunkel: Yes, a suggestion for future research. It would be based on a coalition of selfish agents who agree to, say, not hurt each other and share with each other, but try to take out the turn-taking agents. Some such agreement could overcome the system.

Schermerhorn: That could be. I guess the only thing that I would say about that, is that these are really simple agents. They're basically reactive agents that just map their perceptions onto actions. So if they see food, they go to food. So I would worry about the cost of doing that, as far as whether that would be plausible or not. But that is a really interesting idea. Thanks.

Aaron Frank: Have you considered something along the lines of pack behavior, as agents form groups and hunt in packs but also perhaps change their character over time; that is, so younger agents can be more aggressive than older agents. And so over time you can see what part of a life cycle agents are most aggressive, or specifying and seeing how changes over time within the agent affect its behavior.

Schermerhorn: No, we haven't thought about that. That's also a very interesting idea.

All that we've looked at so far is the general idea of whether an agent becomes more sated or something like that if wins build up and it becomes less likely to be interested in fighting; whereas, if an agent has lost for awhile, maybe it's hungry or something like that and becomes more desperate. So it's all in a much shorter time scale than what you're talking about. But your idea would be very interesting. Thank you.

Sectionalism Index for Organizations: Analyzing Sectioned Random Network Multi-agent Model

Richard Cirillo: Our next presentation is by Kikuo Yuta from the Human Information Sciences Laboratory, Advanced Telecommunications Research Institute.

Kikuo Yuta: This research is on a governmental supported research project named Communication Mechanism, and what I would like to talk about is sectionalism index for organizations.

This sectionalism does not mean provincialism or localism in national politics. Here this sectionalism means sectional war in the companies, like a company and organizations.

First, I will talk about the mechanism on the behavior of sectionalism. And after I'll show the simple model using a random network agent model. Next, I'll show the more practical sectionalism indexes and the variations.

[Presentation]

Panel Discussion

Cirillo: I'd like to ask the authors to join me up front for a minute. This is the fourth conference that we've had on agent-based modeling in social simulation. Four years ago when we started this, one of the questions asked at the conference was, is agent-based modeling and simulation ready for prime time? That is, is the technique developed enough to be able to provide useful information and insight to key decision-makers who would use that type of result? Is the agent-based modeling approach good enough to give to a policy-maker or decision-maker?

At that time, there was some skepticism as to whether the methodology was far enough along to be able to do that. In the presentations we've seen this afternoon, there is obviously a movement toward taking some of this technique and methodology and putting it to work in very practical applications. In one case, the issue of housing segregation, in another case the issue of conflict resolution, and in another case the issue of organizational structure and communication between organizational units.

And so I'd like to ask each of the authors if they would comment on how far along you believe the approach and the methodology that you're using is, in order to be able to provide new insight and new information to decision-makers in the particular area that you're working in.

David, in your case, for example, can you provide new insights into the city of Philadelphia in developing their housing or desegregation policy? Paul, in your case, can you provide information to the State Department and the Defense Department on how to approach conflict resolution? And Mr. Yuta, in your case, can you provide information to Mitsubishi Corporation on how to use these results in structuring its organizational system?

Each of you has, I think, made some comments in the course of your presentation that indicate that there is some indication of the applicability of this technique to these very real-world problems. David, when you started off, one of the comments you made was that the results in the early Schelling work seemed to indicate that people did not have to be super-intolerant in order to develop a segregated housing pattern. Paul, in your presentation, you made the comment that it seemed the more the agents took the others' displays into account, the better they were able to perform in terms of their survival. And Mr. Yuta, you made the comment that companies spend a lot of money on enhancing the communication between different sections and departments in their organizations.

So these are all very real-world problems, and we as modelers would hope to be able to shed some new light and some new insight into that. And so I'd like to ask each of the speakers if you would comment on how your particular approach or model would be able to provide a very real type of decision-making information.

O'Sullivan: Okay. I guess I'm not sure of the answer. I think going back to Shelling's original work, there's certainly genuine insight there. This whole issue of super-segregated residential areas isn't necessarily indicative of any micro-level motivation toward that. And that's a possibly useful understanding, although, of course, the model doesn't prove that people aren't intolerant, because very intolerant people will also segregate in that way. So in terms of whether it provides useful policy guidance, I'm not really sure. I think certainly the model that we're developing needs to have complexity added in terms of the mechanism — I mean, the most obvious missing aspect of the model is that there's no actual housing market. Everybody who wants to move can look around; if there's somewhere they would be happy they can move there. There's no question of how much money they have. Can they get a loan? What's the value of their property? You know, there's a whole market dynamic that we'd need to be on top of , and issues of segregation by income. So, as I say, I'm not really sure.

I think beginning to understand that there are points in the neighborhood dynamics at which things are kind of on an edge, where relatively small injections of money or relatively simple policy interventions could make big differences — this whole kind of tipping phenomenon. Possibly, if you evolve these kinds of approaches a lot further, it might enable identification of neighborhoods and places that are approaching those edges. So I guess there's possibly something there in a very general sense. I'm not sure if the model that we're working on is the platform on which that would be achieved. And certainly there are people building much more complicated, fully realized models of housing markets. I'm agnostic at the moment on that.

Cirillo: Would it be fair to say that one of the things that could be perceived from just the results of the simple model that you're using, that the awareness of local effects versus regional information has an impact on the decision? Would that imply to a decision-maker that information to people is an important factor in helping them choose their housing location?

O'Sullivan: I think it's certainly important, but it's something that doesn't necessarily lead to outcomes that would be agreed to be better, because more information can allow you to behave in ways that are perceived as being less helpful, as well as in ways that are perceived as being more helpful.

For me, one of the points of departure for starting to look at this model was sort of a casual observation that there's an important role for the concept of neighborhood, and for perception of neighborhood — shared perception of the nature of different neighborhoods — in the way that housing markets develop. You see it in gentrifying neighborhoods: realtors intervene and invent new neighborhoods that nobody knew existed. Something like Soho in New York. It [started as] a marketing fiction by realtors that now has a real existence because people responded to the marketing. So I guess there's a point of leverage there, that you can create new phenomena by saying that they exist, publicizing their existence, and drawing certain populations into them.

So I think the whole issue of information about the nature of places has a big role to play. And I could certainly imagine a city government waking up to that fact. I saw somewhere in the literature speculation that Washington, D.C., should be marketing itself as a place that's ripe for gentrification, as a way of co-opting people who have money into regenerating the place. That's an idea I'm interested in exploring. I'm not sure if an agent model is the only useful tool for doing that, though. I think you need much more detailed local empirical study to really get inside what's going on there. I think that the modeling tools are a part of things, but I don't think that's the whole story.

Cirillo: Anyone care to comment or add to that?

Unidentified Speaker: I think that you might be right to discount agent models at this point. But I think when you can introduce some of the things that you said were missing, like costs, expectations and market forces, I don't see any intrinsic reason why that can't become part of the model.

O'Sullivan: I'm quite skeptical about the classical notion of prediction from models like these. I see the role as being much more to do with policy people developing a sense of the dynamics of the systems that they're engaged in trying to manage through interacting with models of these types. So they build up experience — without screwing up the real world — of how systems might respond to certain kinds of intervention. I think that's one of the important roles for these kinds of models.

SimCity™ is a "jokey" example of the kind of thing I'm talking about. But people develop a sense of the role of being a city manager through messing about with a game like that, how developing different kinds of land use affects how people feel about the city, how taxation policies and spending policies affect the system.

Now, the rules in SimCity are hidden and slightly Darwinian and so on, so we might want to build models that are more grounded in observation and allow policymakers to play games with those to develop the kind of understanding that then enables them to see opportunities that they might otherwise miss.

But I think in getting too obsessed with the models, we lose sight of [the fact that] people engaged in policy understand cities pretty well. They have experience of working with people in communities and so forth, and there is a risk in getting so obsessed with this kind of technical fix that we lose sight of the knowledge and the different ways of knowing that are available. I mean, statistics and math and quantification aren't the only way of understanding how things work. And those are important types of knowledge to make use of as well.

Desmond Saunders-Newton: I think as I listen to the responses to your question about whether agent-based models or any models actually reflect agency, I think you're selling yourself short in a variety of ways. Number one, probably one of the world's biggest users of models and simulations is the U.S. Government, period. Right? If there's any community which is willing to embrace and to actually make that a part of their work process, or their contemplation process, it would be the U.S. Government, at the federal level.

Having worked also at the state and local level, there has been an increasing move within various governing structures to actually be much more consistent by using these more rational decision-making processes. They're much more defensible within the context of how decisions are made. I worked with this with the Commonwealth of Virginia when I was a research methodologist for the state, and also when I was over the activities for the LACD Board of Education, which had a budget of $8.4 billion. So my job was to make sure that seven locally elected officials didn't look particularly bad in the context of how it is that they couldn't spend money in alternative ways. And how did we do it? We did it with models.

So I think in some ways the real issue here is how you actually create a better bridge between this particular community of individuals who are involved in the modeling process and the individuals who actually really have a desire and need to have this type of rationalization process associated with them.

Cirillo: I think the whole process of taking any sort of modeling activity and modeling results and putting that in a form that a decision-maker or policymaker can use is a whole field unto itself as to how you use that.

Okay, Paul, [assume] you're now advising Secretary Colin Powell and Secretary Donald Rumsfeld. What new insights can you offer?

Schermerhorn: I'm not really sure that the stuff that we talked about is particularly applicable. And I'm not sure that it's so much because it's not ready, although you might argue that, too, but more because the agents that we're looking at are such simple agents. I made the comment that you could throw a lot of resources at the problem, but that this wasn't an appropriate response for very simple agents, because that cost is going to be too high relative to whatever resources the agent has.

Well, we don't have that problem, right? We typically *can* keep track of who it is that we've lost to, who owes us what and who we owe, right? And, furthermore, the way that we resolve conflicts is a lot more complex than what we're modeling. So our stuff, you know, I'm not sure that I have much to say to Secretary Powell.

As far as the applicability in general of agent-based models to political problems, I don't see any problem in principle with that. It's just this specific work uses much simpler agents. And in order to extrapolate from where we are, I think you have to make some really hard arguments that I think are actually going to be impossible to make.

Cirillo: Any comments? Okay, moving to the corporate world. Kikuo, you're advising Mitsubishi or the Nippon Telephone Company. What results would you be able to give them, or what insights?

Yuta: About Sony and Mitsubishi and Toyota, there are a lot of sections, 500 and more, 200 or more. So it is very difficult to understand or comprehend what's going on about sections' status, especially for the closed or open. So it's very difficult to comprehend. This method I presented, it is possible to measure the all sections in only the one graph. It is possible to explore which party is most dangerous and risky and which party is okay. Of course, I have to develop the model more, but there is a possibility to do such kind of things.

Ecological Interactions

RECIPROCAL VERSUS GROUP ALTRUISM AMONG VAMPIRE BATS

M. PAOLUCCI,* LABSS, ISTC-CNR, Rome, Italy
G. DI TOSTO, CESSBA, ISTC-CNR, Rome, Italy
R. CONTE, LABSS, ISTC-CNR, Rome, Italy

ABSTRACT

This paper explores two interpretations of altruism — group selection and reciprocal altruism — in populations where exchanging help is necessary to face an infrequent, but lethal, scarcity of food. This study, which was inspired by the food-sharing habit of the vampire bat, examines the role of groups as units of selection and reproduction. Findings indicate that when different groups compete for reproduction, altruistic, rather than nonaltruistic, groups are fitter. Grouping contributes to the fitness of the altruistic population because cheaters, although they perform better than altruists, favor the extinction of all groups that contain them.

Keywords: Social simulation, agents, altruism, reciprocity, groups

INTRODUCTION: THE THEORETICAL CONTEXT

Following Group Selection Theory (GST), it must be possible for biological evolution to operate on groups, not only individual organisms. Aggregates of individuals are said to work as units of selection and reproduction. Although popular during the 1970s (Williams, 1971), GST was subjected to severe critiques from sociobiologists. Inspired by the principle of inclusive fitness, in which individuals are seen as vehicles for genetic reproduction (Dawkins, 1976), sociobiologists explained altruism among non-kin in terms of reciprocal altruism (Trivers, 1972) (i.e., the probability of donors being reciprocated when needy).

Recently, GST has been proposed again by Sober and Wilson (1999). The main argument for the revival of GST is that if individuals are the recipients of genes (e.g., machines for their reproduction), groups (or other high-level entities) are the recipients of the recipients of genes. Much like individuals, groups can be characterized in terms of a genetic pool to which all individuals contribute to a different degree. In addition, groups can compete on the same evolutionary stage and act as units of selection. A given habit or trait that characterizes one group can increase its fitness and therefore its preservation, but as the following discussion shows, the presence or absence of groups can influence or even dictate the performance of a population. As the group-selection argument goes, adaptive groups flourish and eventually produce a constellation of new groups that maintain the same characteristics and share the same genetic pool. Conversely, groups that do not have the adaptive trait or habit decline until extinction or disaggregation (its members migrate, dispersing the genetic heritage).

* *Corresponding author address:* Mario Paolucci, Laboratory of Agent Based Social Simulation, Institute of Cognitive Science and Technology, CNR, Viale Marx 15, 00137, Roma, Italy; e-mail: mario.paolucci@thinkingolem.com.

The controversy around GST remains harsh (see, for example, Palmer's [2002] review of Field's book [2001]), sometimes because of an equivocal and collectivist interpretation of the theory. Rather than sharing this interpretation, we prefer to use the definition provided above; that is, groups are high-level units of biological selection.

ALTRUISTIC BEHAVIOR IN NATURE

Examples of altruism abound in nature (see, for example, Brembs, 1996). Interspecific mutualism has been documented among lycaenid butterfly larvae and ants (Leimar and Axén, 1993). A predator inspector in shoaling fish is another well-known case in which different individuals leave the shoal together and swim toward the predator, gathering information about its precise location and current motivational state (Milinski, et al., 1990). Through cooperation, they share the cost of actions that cannot be performed alone.

Among mammals, the most famous example of prosocial behavior is blood-sharing in vampire bats, a behavior that favors starving, unlucky hunters (Wilkinson, 1984). In addition, many controversial examples exist among primates and humans. Ethological observations and even simulation experiments, such as those by Wilkinson (1990), are often interpreted as supporting the inclusive fitness explanation, based on the evidence that animals (in this example, vampire bats) will help individuals they recognize (possibly, once donors).

The evidence in favor of reciprocal altruism, however, is not fully satisfactory. Is it possible to use natural experiments to differentiate between individual versus in-group recognition if a group is small enough so that individuals are allowed to meet all others at least once?

A brief explanation helps to capture the rationale of reciprocal altruism. This phenomenon, and its related theory, is a twofold concept. First, reciprocity can occur in the direct form (i.e., when the current donor receives help from its current recipient). Second, reciprocity can occur in the indirect form (i.e., when the current donor receives help from someone that received help from a current recipient). In its indirect variant, reciprocity circulates in the group, increasing the fitness of donors. This second variant of reciprocal altruism poses two questions:

- Circularity makes reciprocity inherently fragile: at any step, the chain can be interrupted either by cheaters or by accident (noise). Direct reciprocity, however, is more robust because the number of steps is reduced, and cheaters (non-reciprocators) are immediately found out after one defection.

- Consequently, indirect reciprocity appears to be more irrational from the individual point of view:

    - Why should self-interested agents give away a share of their current probability of reproduction if the chance of getting it back is remote and uncertain?

    - What kind of guarantee is needed for agents to participate in chains of reciprocity?

To our knowledge, no fully convincing answer to the latter question is available. Essentially, the sociobiological answer is as follows: agents act altruistically because they are programmed to do so. If altruists reproduce to a higher degree than nonaltruists, their genes spread over the population. Consequently, the altruistic behavior not only survives, but also finds increasingly less hostility because cheaters tend to be replaced by the more prolific altruists. Under this perspective, reciprocity (either direct or indirect) is an emergent effect of altruism, rather than a condition for its execution, that reinforces its occurrence. Altruists do not aim to obtain reciprocity nor calculate its probability. If altruism spreads, altruists survive and reproduce. Emergent reciprocity should work so that donors and/or their future generations are refunded. In this case, the altruistic act increases the donors' individual fitness.

Under what conditions does reciprocity emerge? Does reciprocal altruism account for this emergence? If not, which theory provides a complementary or alternative explanation? Can group selection play a role in this sense?

Simulation — proposed by Axelrod (1997) as a new way of doing science — offers an experimental instrument for testing both theories, refining existing hypotheses or, possibly, formulating new ones.

THE SIMULATION EXPERIMENT

Our simulation experiment is based on the vampire bat example. The species studied by Wilkinson lives in Central America in small groups (a few dozen individuals) that share the cavity of trees. We call this basic unit group a "roost."[1] The group's daily diet consists of ingesting fresh blood, which they suck from herbivores. Each night, however, about 8% of the adults do not find prey to parasite. On those occasions, the bats can survive because of luckier fellows who regurgitate for them a portion of the food ingested. Wilkinson (1984) actually stated that such behavior "depends equally and independently on [the] degree of relatedness and an index of opportunity for reciprocation." Now, which thesis receives stronger support from this evidence — inclusive fitness or group selection?

By using the Repast simulation platform,[2] we constructed a pilot experiment that mimics the vampire bats' behaviors in what we perceive as their essential traits. We introduce two different algorithms: altruistic (food sharing) and selfish or cheating (no food sharing). The former mimics the behavior of lucky hunters that give away an extra amount of the blood ingested, if any, to the benefit of starving fellows asking for help. The latter reproduces the behavior of selfish animals that refuse to help their unlucky fellows, which then starve to death.

In the simulations, only starving animals are allowed to ask for help; they receive help from their addressees if these are both altruists and satiated. No bluff is allowed. Agents have no memory of past interaction and cannot calculate the probability of reciprocation. No explicit mechanism for punishment of cheaters is implemented. In such conditions, how can reciprocity

---

1 To be more precise, real vampire bats move in subgroups around several cavities, creating a fluid and territorial group system. Usually, roosts contain only one alpha male, plus several other males and females in a rigid hierarchy, but we do not model this level of detail in our simulation.

2 For information on Repast, visit http://repast.sourceforge.net.

emerge as a mere “objective” effect, implying neither computation nor deliberation on the side of altruists?

To answer this question and explore the effect of groups on the evolution of altruism, we ran simulations with mixed populations (both algorithms in variable combination) initially distributed over a given number of roosts. During the simulation, roosts can grow or collapse, depending on the survival and reproduction rates of their members, which in turn depend exclusively on social attitudes (whether altruistic or not). Ecological conditions are equal for all roosts. At given times, roosts give rise to new roosts if the number of young individuals reaches a given threshold. This option was meant to be an operational simplification of the notion of group selection and reproduction.

THE SIMULATION MODEL

The agents — bats — are modeled as objects. In nature, bats live in roosts, which are physical sites (usually tree cavities). They return to their roosts in the day after hunting during the night. Bats reproduce and perform social activities (nursing, grooming, and sharing food) inside the roosts. In our simulations, the roost is a social space that contains any number of bats. “In-roosts” are allowed to share food and groom. No other social activity has been modeled.

Each simulation turn (in Repast language, each tick) corresponds to a 24-hour period and includes one daily and one nightly stage. In the day, the simulated animals perform social activities (grooming and food sharing). At night, they hunt. In our model, “hunt” is defined as an ecological parameter. In accordance with real-world data (Wilkinson, 1990), its default value is set to 93%. In substance, each night 93% of the population finds food, which permits them to survive until the next hunt. The remaining 7% begins (or continues) to starve unless they receive help from some fellow (under the form of regurgitation). Vampire bats do not accumulate resources: they hunt only for short-term food consumption. In addition, although the average lifetime of these animals is about 14 years, starvation and death are a constant threat because each good hunt gives them no more than 60 hours of autonomy. As a consequence, for a bat in isolation, two failures in a row are fatal. These harsh conditions characterize the life of bats, which face infrequent (in the simulation, about 1.65 episodes of double unsuccessful hunts per animal per year), but lethal, food scarcity. The only way to prevent starvation and death is to receive help from fellows, which is what these animals appear to do in nature.

As for daily activities, the rationale of grooming is at least twofold: animals familiarize themselves with other bats thanks to and during grooming and assess their respective physical shape. Because satiation increases body volume, a lucky hunter may grow to even twice its normal size, which is easily detected by its grooming partners. A starving bat is also likely to be recognized. Bluff would immediately be found out in such extreme life conditions.

Each day, animals choose one partner from the roost population. In our model, as in the real world, grooming has the effect of increasing the probability for food sharing among in-roosts: a starving bat will turn to grooming partners for help and will avoid death if a partner is found to be full (having had a good hunt). Because of the bat’s metabolism, the donor loses much less than the recipient gains. In the simulation, the agent loses an amount of energy worth six hours; this amounts to losing a chance to ask for help during the last day, given two failures in a row following donation.[3] In this set of experiments, we limit the number of partners per day to one.

The numbers obtained are in accord both with what is known by ethological observations in the presence of mutual help (the annual mortality for adults is about 24%) and with the results of a simulation carried out by Wilkinson (1990) in the absence of help. As noted, help is rare, but critical: in roosts in which all individuals deny help, population is reduced 80% per year. Figure 1 shows typical simulation results with and without help. In one-year experiments, both with and without food sharing, no roosts are lost, meaning that each roost still has at least one vampire; the situation with food sharing gives figures similar to the mortality rate among adult

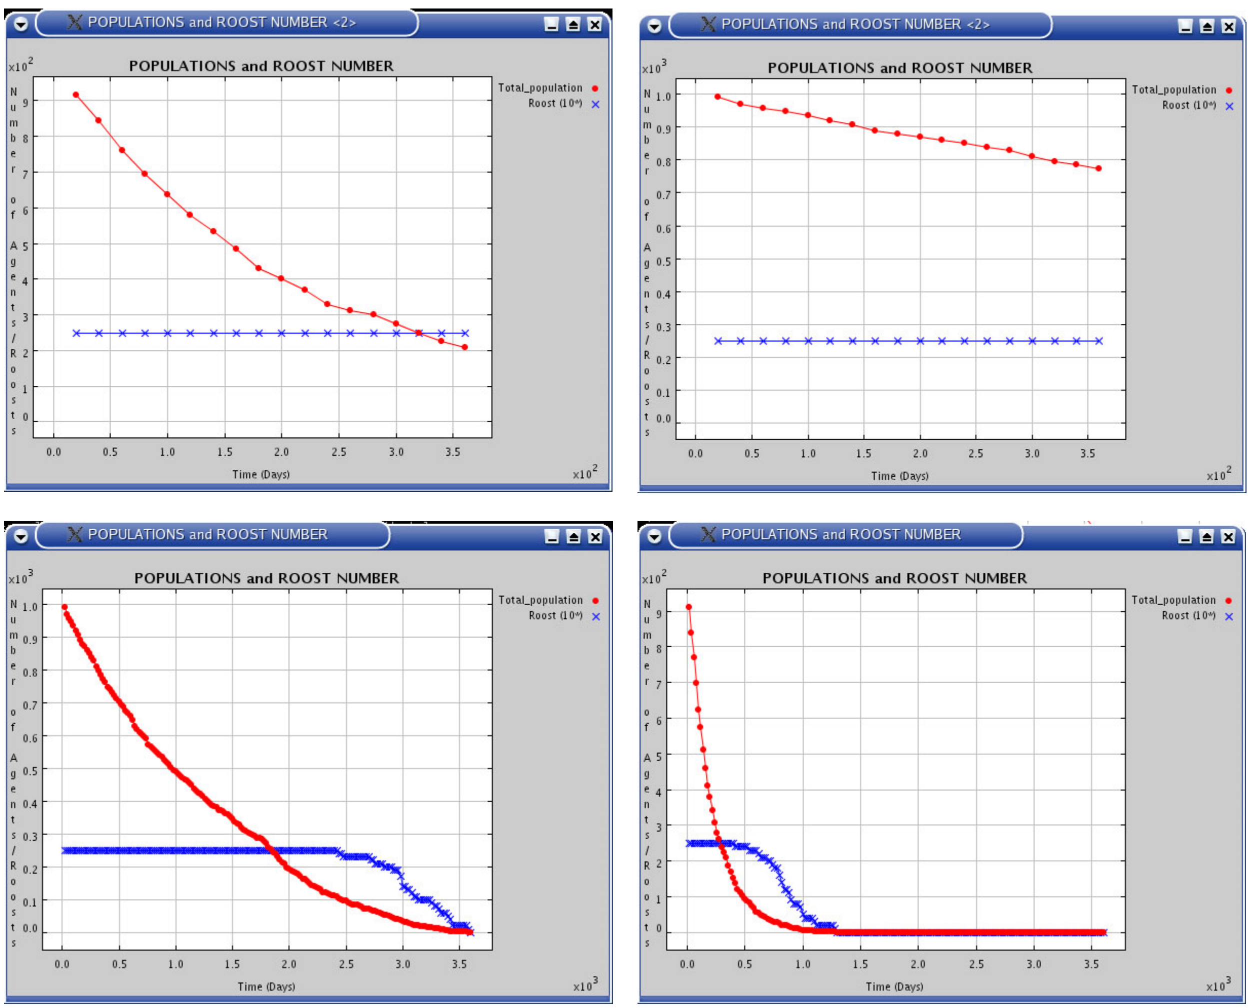

FIGURE 1 Four typical simulations. Left: The reciprocal help system is deactivated. Above: 360 ticks; below: 3,600 ticks. Total population and roost numbers are shown. Although not frequent (agents look for help 1.65 times per year), help is critical for the population.

3 The brief amount of time lost can be misleading. In our simulations, approximately 50% of total mortality comes from donations, followed by a couple of failures to find food.

bats, while the other one reproduces Wikinson's simulation results. After 10 years, both populations are substantially extinguished, but the one with help has been around much longer.

The case under study presents two key characteristics:

- It is exceedingly difficult to accumulate any kind of wealth. Energy coming from a meal is dissolved after two nights, so there is no such thing as a wealthy individual. The lucky hunter of today has the same chances as everybody else of starving tomorrow.

- Direct retaliation is simply impossible in the present setting. The victim of cheating dies on the spot; asking for help is the last resort, and given our restriction of one helping partner per night, a cheater (i.e., an agent that refuses to help) is a dangerous killer that is very difficult to find out.

These points make the experiment exceptionally significant for investigating the conditions under which altruism spreads. Indeed, the impossibility of direct retaliation seems to make the system extremely fragile to the introduction of cheaters. The only way out would consist of third-party enforcement, which is usually obtained by means of cognitive artifacts, such as image or reputation (see Conte and Paolucci, 2002). Is added cognitive complexity the only way out? To answer this question, we describe in greater detail the relevance of groups and roosts.

Reproductive Algorithm and Roost Formation

In nature, female vampire bats may give birth to one child at a time; they reproduce about every 10 months. Newborns leave the roost as soon as they are able to care for themselves, but they are never found in isolation; the chances of survival for a lone individual are extremely low.

In the simulation, individuals are identical at birth and sexless. They reproduce by cloning every 10 months, starting with the 20th month, in order to model the juvenile phase. To obtain a reasonable rate of reproduction, at each occurrence each agent has a 50% probability of cloning. Although poorly realistic, this minimal condition allows for roost formation, which is the focus of our study.

As for roost reproduction (scission), we avoided the additional complexity of entrance criteria, which should be met for a newcomer to be accepted in an established roost. In our model, the formation of new roosts is allowed only when a critical mass of new individuals is reached. We have fixed this threshold at 20 individuals. The rationale underlying roost formation is reproductive success: the more the in-roosts, the higher the number of new roosts formed.

HYPOTHESES

The objective of this study was to test different interpretations of the evolution of altruism by using simulation. The reference example in the real world is food sharing by vampire bats. As previously noted, this species offers clear evidence of the advantages of altruism on life expectancies. Wilkinson's simulation findings, also reproduced in our simulation, show that the probability of survival each year is around 78% of the population when food sharing is activated, as opposed to a mere 20% with no food sharing. How does one interpret these findings?

In line with the sociobiological theory of reciprocal altruism (Trivers, 1972), one could say that bats survive to such a greater extent when sharing food because of reciprocity (Dawkins, 1976), which adds to the individual fitness of donors much as altruism adds to the fitness of recipients. In other words, giving help acts as an "investment," although neither deliberate nor acknowledged, on the part of the altruist, which accumulates credits to be refunded by reciprocity. Since bats do not accumulate food, donors that are reciprocated later on in their lives will survive longer than if they had not performed an altruistic act. Whereas the initial donation caused a mere reduction of the time interval before starvation, the following reciprocation prevents immediate death!

However, it is unclear whether and to what extent bats take measures against cheaters. Wilkinson's findings refer to the comparison between a condition in which all bats cooperate versus a condition in which all bats are defectors. What happens in intermediate conditions? What is the minimal share of altruists for obtaining an increase in the survival rate with regard to the all-defector condition? Moreover, does the increase in the rate of survival effectively correspond to an increase of donors' fitness, or is it redistributed over the entire population? And if so, are individual donors always refunded or do they sustain a share of the costs of redistribution?

The latter question is crucial because if donors are not always reciprocated in person or along their future generations, there is reason to question the reciprocal altruism interpretation and to look for concurrent explanation. One good candidate is the group-selection theory. In this theory, aggregates of non-kin individuals are considered as units of biological selection and evolution. Under this perspective, a given trait, such as altruism, is accounted for in terms of its contribution to the fitness of the group rather than to the fitness of individual members. Consequently, food sharing vampire bats can be seen as a habit that evolved as a result of the positive effects on the fitness of roosts taken as a whole, rather than on the individual fitness of donors.

In short, the reciprocal altruism theory proves adequate if donors are almost always reciprocated in their lives or in the lives of their offspring. In such a case, the altruistic gene spreads because the genes of donors survive and replicate through generations.

Instead, the concurrent group-selection theory proves right if (1) the survival rate increases in altruistic roosts although donors are poorly reciprocated both in their lives and in their offspring, provided (2) the altruistic roosts as wholes are fitter than their nonaltruistic competitors. How can one measure a roost's fitness? Roost formation is a possible solution. In this sense, the higher the number of roosts that are formed from an original roost, the fitter the latter roost, provided the rationale for roost formation is reproduction. The higher the number of reproductions of a parent roost, the higher the number of offspring roosts.

EXPERIMENTAL CONDITIONS

In the experiment, the same conditions (stated above) apply, except that food sharing is always allowed in all roosts, and reproduction is possible. Moreover, a percentage of cheaters is introduced to check the robustness of the altruistic strategy and to obtain insights about its evolutionary stability. Cheaters never give help when asked, even if they are full; unlike altruists, they sustain no costs. Because a retaliation mechanism is not modeled, a first expectation might be that cheaters prosper, thus reducing the efficiency of the system as a whole. Indeed, this happens in the short run, but when longer simulations are considered, the scenario changes dramatically.

In addition to the effects of cheating, we were also interested in a measure that would allow discrimination between group selection and inclusive fitness. To this purpose, we tracked the lineage of the agents from the beginning of the simulation. Reproduction by cloning allows for clear tracks; if the mortality rate of one lineage is equal to or lower than the average in the same roost (but this produced a significantly higher number of offspring roosts than under the control condition), then group selection seems an adequate interpretation of vampire bat altruism. Instead, if donors' lineages show a mortality rate that is significantly lower than the average in the same roost, and the number of offspring roosts is not significantly higher than in the control condition, then reciprocal altruism provides a more adequate interpretation.

FINDINGS

Simulations were run for a number of cycles corresponding to around 40 years, which includes about four generations of vampire bats. Figure 2 shows some typical examples of what happens during the run for different initial shares (from 5% to 40%) of cheaters.

Clearly, the food-sharing condition has the reproductive advantage. Selfish bats are sure to die in a few generations, leading also to collapse of their roosts. They play a destructive role by gradually reducing the reproductive capacity of their roosts until global extinction occurs. When either the distribution of the two behavioral modalities (cheating and altruism) is such that altruists far exceed cheaters, or some demographic catastrophe (triggered by cheaters themselves) leads to earlier extinction of cheaters, however, the reproduction of altruists begins again and the number of roosts grows in proportion. This happens after a critical period during which cheaters extinguish, and the global fitness of the entire population is nearly on the verge of collapse. After cheaters have been totally extinguished, the population starts to grow rapidly and indefinitely.

This observation leads to an appraisal of the role of roosts. In fact, if the entire population were sharing one roost (see Figure 3 for an example), the presence of cheaters would lead them to certain extinction. With single roosts, most simulations converge to zero after some time with or without later resurgence. In any case, the presence of cheaters increases until they cause a catastrophic lowering of the population, after which they start to increase again until extinction. No reciprocity could emerge in a world in which cheaters are allowed to repeatedly exploit others, without incurring retaliation or isolation.

Under these extreme conditions, in which help exchange is vital, after having exploited their altruistic in-roosts to death, cheaters find no way to face adversity and are soon bound to share the same fate. Those few altruists that might survive the extinction of cheaters soon take off, repopulate the roost, and produce new ones. If no one survives cheaters, which is the most likely event since they survive longer than their good fellows, the roost extinguishes.

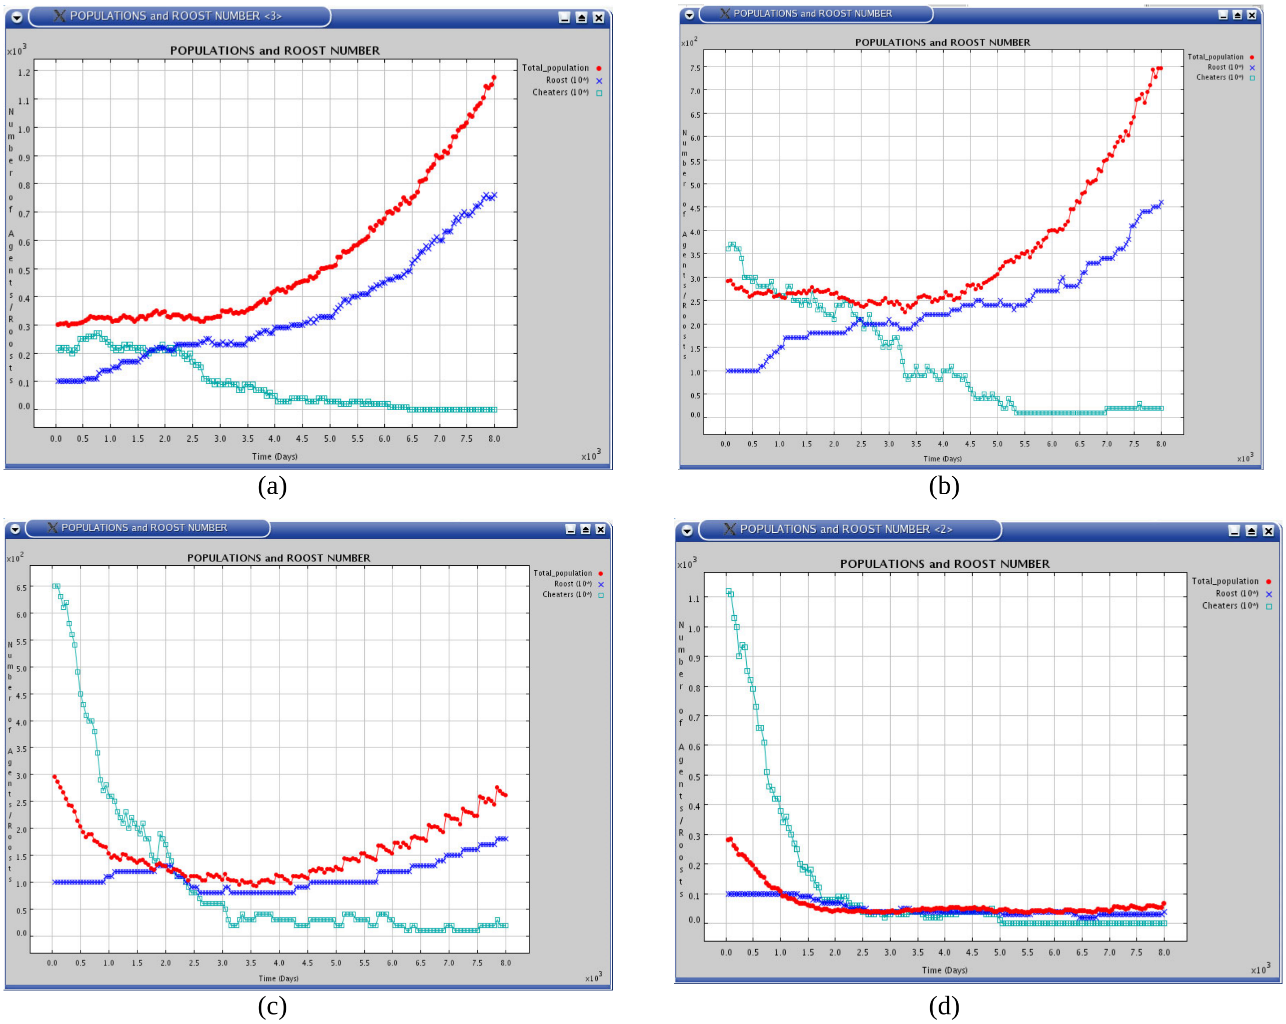

FIGURE 2 Four typical experiments showing total population, number of roosts, and number of cheaters (multiplied by 10) for (a) 5%, (b) 10%, (c) 20%, and (d) 40% cheaters at startup. Each simulation starts with 300 agents in 10 roosts for 6,000 ticks.

On the contrary, the phenomenon of roosting radically modifies the situation. Due to the presence of cheaters, most of the roosts disappear. If at a certain point, however, any roost without cheaters appears, it grows and repopulates the world. This phenomenon is shown in Figure 2 after a demographic catastrophe.

To determine whether the evolutionary advantage of food sharing is better explained by inclusive fitness or group selection, several methods of discrimination were tried. A clear signal in favor of group selection would be given if we were to find that helpers (i.e., agents that often helped others during their lifetime) had less offspring than other altruists. The problem here is subtle because agents that have had more chances to reproduce (i.e., have lived longer) automatically have had more chances to help others. So the correlation between help and fitness — defined as both the number of descendants and the number of living descendants — is

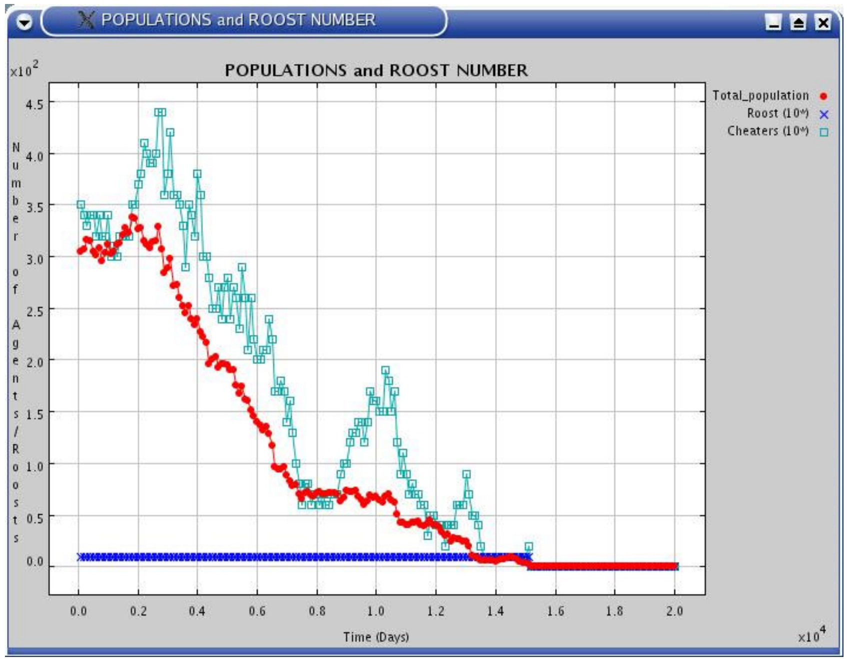

FIGURE 3 Single roost with 300 agents, no new roost formation, 10% cheaters, and 20,000 ticks.

obviously positive and therefore of no significance in distinguishing the two explanations. In further trying to pinpoint the effect, we selected a subset of the initial agent population, characterized by death, old age, and not being cheaters. In the simulation, 10-year-old agents are automatically removed from consideration. By tracking their descendants, we can look for a correlation between the number of helps and the size of the offspring.

We ran a set of 100 simulations, containing 10 roosts with 30 agents each. The number of cheaters is 10%; simulations stop at 5,000 ticks, and we extract only values for agents that were present in the beginning and that lived to the maximum age possible in the simulation. For them, we track the number of times help was given and the number of agents in the entire lineage up to the end of the simulation. Only a very low factor (0.07) of positive correlation was found between these measures. In substance, there is not a linear relation between the number of times help is given in life and the total number of offspring in the simulation for the subset we analyzed.

DISCUSSION

The findings presented above point in two distinct directions. The first direction concerns the fact that altruistic roosts survive longer and reproduce by far more than mixed roosts, which include cheaters. Indeed, cheaters tend to cause the extinction of their roosts, while extinguishing themselves. This result is not particularly original, since it essentially confirms what Wilkinson (1990) found through his simulation study of food sharing in vampire bats. Although distribution of cheaters over the roost population makes a difference in terms of the reproductive success of the entire roost, even one cheater may be enough to lead the roost to extinction. This fact occurs because of two reasons: first, cheaters usually survive longer than their fellow altruists; second, in a competitive environment, any roost that by chance finds itself cheater-free will generate more offspring, both in terms of roosts and in terms of individuals. In this sense, the roost is a critical unit of selection, whose efficiency in finding out and eliminating cheaters is amazing.

The relevance of roosts could be lessened by a different kind of cheater, not included in this study: a roaming cheater that distributes the cost of its presence over the roosts it has access to. In nature, however, individuals requesting entry in a new roost find it very hard to be accepted by the new roost inhabitants; an entrance fee could then have co-evolved to protect the altruistic mechanism.

The second direction in our findings concerns the role of the roost in the evolution of altruism. In this respect, grouping, or better roosting, seems to matter. If altruistic groups survive longer than nonaltruistic ones, the reverse is also true: altruists have better chances of survival if they roost together. In a one-roost world, inhabited by cheaters and altruists, the chances of survival of altruists — and of the entire population — would be proximal to zero. Conversely, in a multi-roost world, where altruists happen to co-exist in variable distribution with cheaters, an interroost competition for reproduction occurs. Since cheaters lead to the extinction of their roosts, only altruistic ones will survive to reproduce, and these will soon populate the world. This perspective seems to support the group-selection argument. Groups act as units of selection and reproduction, much like individual organisms.

However, this finding per se does not say much about the internal rationale of altruism. If it supports group selection, it does not disclaim the concurrent sociobiological theory of reciprocal altruism. Indeed, precisely because no rule for reciprocity is explicitly represented in our simulation model, the only way for altruists to survive is to roost together, waiting, so to speak, for cheaters to extinguish. In this sense, and rather tautologically, reciprocity emerges only when cheating disappears.

Our findings at this point also indicate that actual donors do not reach a significant higher rate of survival and reproduction than the rest of the population; indeed, the correlation factor is too low to make conclusions in this matter. No negative correlation has been found, giving no definitive support to the thesis of group selection. Apparently, then, the global increase of fitness of the roost population is not obtained at the expense of one share of it (the actual donors). Conversely, the final generations do not necessarily include the lineage of the actual donors. This finding might be due to the simplicity of the algorithm, which does not allow for a specific rule of reciprocity. On the other hand, it corresponds to the simple rationale of reciprocal altruism, for which agents neither aim nor calculate the probability of reciprocity, which should be an emergent effect of the altruism fitness. If actual donors are not reciprocated, however, their fitness decreases to the benefit of the global fitness. If this is the case, as it appears to be in our findings, inclusive fitness does not account for the spread of altruism. Indeed, group selection theory gains the ground that is lost by the reciprocal altruism theory. Grouping matters and helps altruists to survive and reproduce even in the presence of cheaters. Under the shelter of their roosts, animals helping each other increase their chances of reproducing, although some, finding themselves in roosts with high numbers of cheaters, will pay dearly for such good behavior.

ACKNOWLEDGMENTS

The authors are grateful for contributions from the Fifth Framework European Projects ALFEBIITE and FIRMA.

REFERENCES

Axelrod, R., 1997, *The Complexity of Cooperation*, Princeton, NJ: Princeton University Press.

Brembs, B., 1996, "Chaos Cheating and Cooperation: Potential Solutions to the Prisoner's Dilemma," *Oikos* **76**:14-24.

Conte, R., and M. Paolucci, 2002, *Reputation in Artificial Societies. Social Beliefs for Social Order*, Boston, MA: Kluwer.

Dawkins, R., 1976, *The Selfish Gene*, Oxford, U.K.: Oxford University Press.

Field, A.J., 2001, *Altruistically Inclined? The Behavioral Sciences, Evolutionary Theory, and the Origins of Reciprocity*, Ann Arbor, MI: The University of Michigan Press.

Fog, A., 2000, "Simulation of Evolution in Structured Populations: The Open Source Software Package Altruist," *Biotech Software & Internet Report* **1**(5) 226–229.

Leimar, O., and A.H. Axén, 1993. "Strategic Behavior in an Interspecific Mutualism: Interactions between Lycaenid Larvae and Ants," *Animal Behavior* **46**:1177–1182.

Milinski, M., D. Külling, and R. Kettler, 1990, "Tit for Tat: sticklebacks (*Gasterosteus aculeatus*) 'Trusting' a Cooperating Partner," *Behavioral Ecology* **1**:7–11.

Nowak, M.A., and K. Sigmund, 1998, "Evolution of Indirect Reciprocity by Image Scoring. The Dynamics of Indirect Reciprocity," *Nature* **393**, June 11.

Palmer, C.T., 2002, "Review of Field, 2001," *The Human Nature Review* **2**:92–94, March 11.

Sober, E., and D.S. Wilson, 1999, D.F. *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Cambridge, MA: Harvard University Press.

Trivers, R., 1972, "The Evolution of Reciprocal Altruism," *Quarterly Review of Biology* **46**:35–57.

Wilkinson, G.S., 1984, "Reciprocal Food Sharing in the Vampire Bat," *Nature* **308**:181–184.

Wilkinson, G.S., 1990, "Food Sharing in Vampire Bats," *Scientific American* **2**:64–70.

Williams, G.C., 1971, *Group Selection*, Chicago, IL: Aldine.

Wilson, D.S., and E. Sober, 1994, "Reintroducing Group Selection to the Human Behavioral Sciences," *Behavioral and Brain Sciences* **17**(4): 585–654.

AGENT-BASED MODELS OF LAND USE AND COVER CHANGE IN THE ATLANTIC COAST REGION OF NICARAGUA: EXAMINING THE AGENTS OF TROPICAL DEFORESTATION*

L.E. FERNANDEZ, University of Michigan

Previous studies of deforestation dynamics primarily focused on modeling approaches that used remotely sensed data at regional scales. Few efforts integrated household-level social data with regional-scale spatial models. This paper presents results from a study that used an integrative approach, combining logistic regression analysis, social science surveys, and agent-based land-use modeling. It focused on the lowland rain forest region of Pearl Lagoon, Nicaragua. Its objectives were to (1) identify eco-physical and aggregate socioeconomic variables and their effects on the pattern of land-use change, using logistic regression analysis; (2) determine key factors used by farmers to select land, through social surveys; (3) model key decision factors in an agent-based land-use-change model; and (4) determine the strengths and shortcomings of these approaches and present linkages between them in an integrative framework offering the advantages of each approach. The paper presents results from a land-cover-change detection analysis of the study landscape over 37 years and logistic regression analyses of several independent influencing factors associated with forest loss over this period. A main hypothesis is that an integrated approach combining the high frequency and cost-effectiveness of remotely sensed data with more culturally accurate community-level data provides more accurate results in multicultural developing regions undergoing rapid changes in ethnic composition. The development of agent-based models (ABMs) to simulate land-use decision-making processes of peasant farmers allows exploration of this hypothesis in an innovative, effective manner. The ABMs, developed in Swarm and NETLOGO, effectively model the complex economic, social, and cultural networks in which farmers operate. They are used to study how land-cover patterns change in several possible policy and migration-pattern scenarios.

---

* At the time of publication, the full paper for this presentation had not been received.

AGENT-BASED MODELING FOR VEGETATION SUCCESSION AFTER OPEN-CUT MINING: STAGE 1, A STUDY IN PROGRESS

X.F. SU* and J.A. DUGGIN, University of New England, Armidale, Australia

ABSTRACT

Agent-based modeling and geographic information systems are used to describe plant colonization and community development for species diversity and changes in community structure and composition over both spatial and temporal scales. The model is used to develop a predictive pattern for ecosystem succession following open-cut mining. A trial box cut on Leard State Forest on the northwestern slopes of New South Wales, Australia, is used as the case study. The model was built to represent dynamic interactions between individual plants with different capacities for habitat occupation and competition on two different substrates — one covered with topsoil and the other bare overburden. Both environmental influences and biotic factors were incorporated into the model. The results of the simulation show that two contrasting plant communities develop, one similar to the surrounding forest, the other a shrubland likely to persist for considerable time if no further disturbances occur. Thus, the use of topsoil in rehabilitation produces an effective vegetative cover, but it also creates a problem in that shrublands, rather than forests, develop. Further developments in the model are outlined to make it more interactive for mine rehabilitation management purposes.

Keywords: Vegetation succession model, forests and shrublands, design process and parameters, predictive patterns, mine rehabilitation, agent-based modeling

INTRODUCTION

Open-cut mining causes significant environmental changes. The removal of soils and vegetation and the disposal of mining waste have long-lasting effects on ecosystem development (Fox, 1990). Early mining activities had little regard for the environment (Farrell and Kratzing, 1996), and this practice resulted in many abandoned mine sites through Australia that need rehabilitation (DEST, 1996). The goal of mine rehabilitation is to return the disturbed area to a vegetated and productive condition that is ecologically sustainable over the long term. Recovery is a complex process and takes considerable time to demonstrate ecological sustainability. During this process, the community, scientists, and the industry at large strongly emphasize that rehabilitation should be monitored to make sure it satisfies the aims and purpose of use for which it is intended (Brooks, 1981; Unwin, 1985, 1988). Therefore, an ecological model is needed that has strong simulation and predictive abilities that can address different management strategies and ecological states within the successional process. Results from such a model will provide basic information to support decision making during site rehabilitation and management. Classical ecological models are unable to fit this requirement.

* *Corresponding author address:* Xianfeng Su, Ecosystem Management, University of New England, Armidale, NSW, 2351, Australia; e-mail: xsu@metz.une.edu.au.

The aim of this study is to use agent-based modeling and the technology associated with geographic information systems (GISs) to describe species colonization and community development over both spatial and temporal scales by means of simulating the dynamics of invasion, establishment, competition, and adaptation of individual species to a changing environment. In this way, the model can be used to predict patterns of ecosystem succession. The study uses information gained from periodic measurements of vegetation development over the past 20 years on different soil materials, preliminary studies on soils and overburden for plant growth, together with results from germination, survival, and early growth of different species to develop the parameters and rules used in the model. In addition, results from the scientific and technical literature are used to provide additional basic information relevant to parameter estimations.

METHODS

Model Framework

Repast, which was created and developed by Social Science Research Computing at The University of Chicago, provides a library of Java code for creating, running, displaying, and collecting data for agent-based simulations (Collier, 2003). It has been used predominantly for social science simulation but has the potential to be used for ecological studies. Repast has been selected as the framework for simulating vegetation succession after mining and has been coupled with GIS technology to import space objects into the model and to store and display the results.

Study Area Location, Vegetation, Soil, and Climate

Leard State Forest, on the northwestern slopes of New South Wales, Australia, was used as the case study. Two experimental spoil heaps were left after a trial box cut into coal seams was made during mine exploration and the assessment of the potential for commercial development. The spoil heaps are approximately 3.1 ha in area, and were mostly covered with stockpiled topsoil to depths varying from 0 to 20 cm, leaving some patches of bare overburden. A fence was constructed around the research site to exclude kangaroos and wallabies from the area during the early years of rehabilitation.

The vegetation in close proximity to the spoil heaps is a dry sclerophyll, open forest and woodland, dominated by mixed communities of eucalypt and cypress pine (Croft, 1979). The soils are alkaline or sodic duplex soils (Wiram, 1979). The long-term average annual rainfall is 607 mm with a summer dominance. The summers are hot while the winters are mild with an average of 35 frosts over winter.

Experimental Design and Data Collection

Six transects (10×30 m) were established in 1981 to monitor the recolonization by native shrubs and trees (Figure 1). Germination, height, and mortality of all individuals were recorded periodically from 1981 to 2002 to develop records of life histories for each species. The distribution of each individual was mapped in the earlier stages and later recorded by global positioning system (GPS) and then transferred into the GIS for mapping. In addition, all individual trees across both overburden heaps were recorded by species and their growth measured periodically. The study concentrated on the main species developing on the overburden heaps. Included were the tree species *Eucalyptus pilligaensis*, *E. crebra*, *E. albens*, *E. populnea*, *Callitris glaucophylla*, and *Casuarina cristata*, and the shrub species *Acacia deanei*, *Cassia nemophila*, and *Dodonea viscose*.

An additional seven transects (20×50 m) were established in the natural forest surrounding the spoil heaps in 2002, and data were collected to describe its structure and composition. Height, diameter, and stratum of individual trees and shrubs were recorded by species.

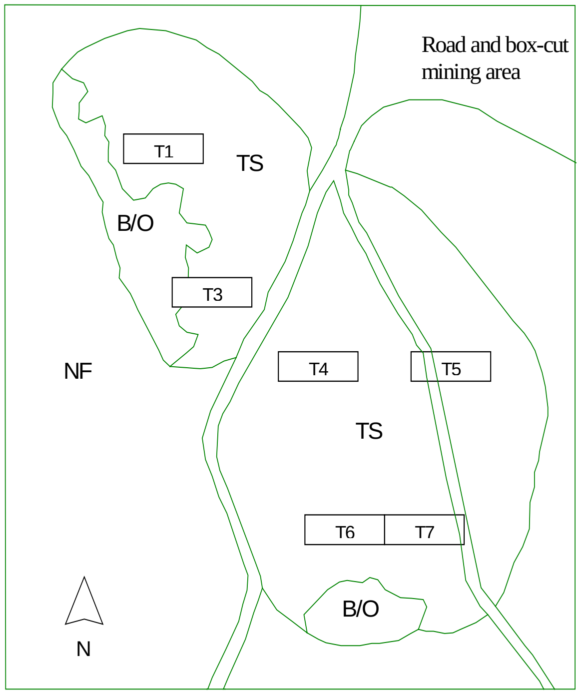

TS – Topsoiled area, B/O – Bare overburden area,
T – Transect, NF – Natural forest

FIGURE 1 Overburden site and the location of transects used for vegetation monitoring

MODEL DESIGN AND OUTPUT

Model Design and Establishment

A succession model was developed based on individual life histories for each species. The model was based on simulating invasion, establishment, persistence, and competition (both intra- and interspecific competition) for each species. Since Repast is an agent-based model and is object oriented, the basic elements are the objects. Consequently, for this application, the objects include the seed, plant, species, forest, soil, and space (Figure 2). Ecosystem development and spatial and temporal dynamics are then described by combining the objects for the study area.

The individual plant (tree or shrub) is the basic agent used to describe and develop the dynamics of vegetation succession. Each individual plant in this model changes over time as a result of germination, growth, adjustment to the environment, competition between individuals of the same and different species, and mortality. This is based on the knowledge of the life history of each species. Because each plant changes its status over time, so too do the condition of the vegetation and the environment that surround it.

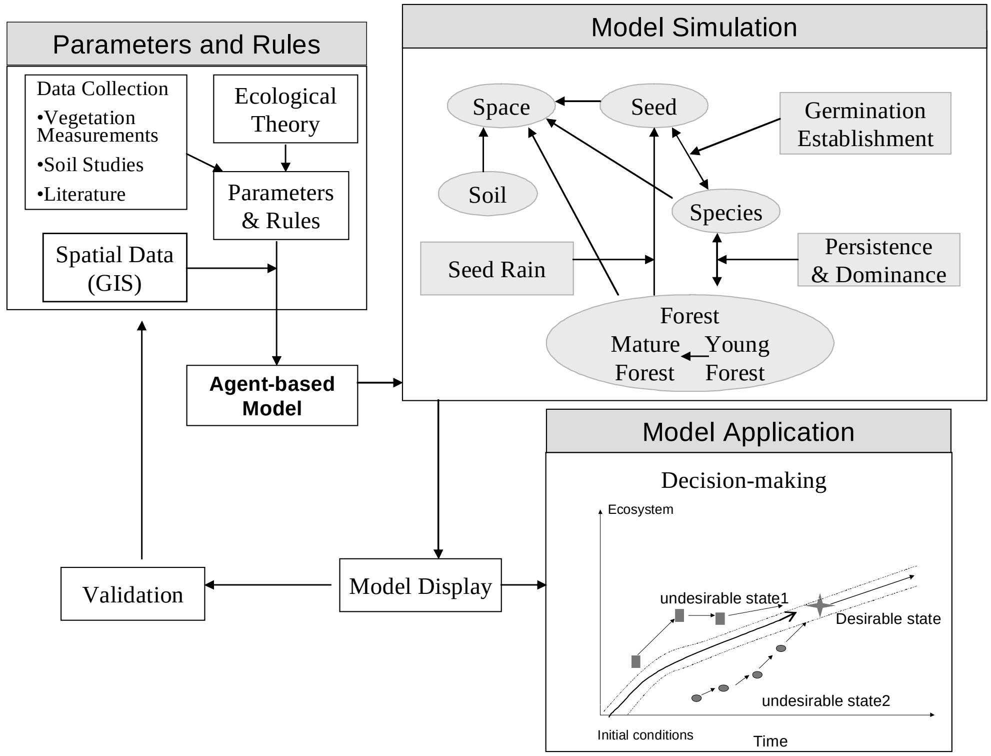

FIGURE 2 The model design showing the development of parameters and rules; model structure with the simulation component, emphasizing the relationships between objects (ovals) and functions (rectangles), and outputs; and the validation and use of the output (The model application example shows how the simulation can be used to predict ecosystem development over time using several starting points, the first moving from the initial condition and other two starting from undesirable states.)

The seed object is used to keep track of the species name and has the parameters of seed production, viability, dispersal, and germination for each species. The species object is a description of the life history for each species under study and includes information on seed rain, seed germination, seedling establishment, and growth rates within the space object.

The space object, where individual agents live, is treated as a two-dimensional lattice of grid cells of 2×2 m. The soil object and its characteristics (namely, whether the area was covered with stockpiled topsoil or was bare overburden) is described in the GIS and can be imported into the model to overlie the space object. This was then used to describe the environment in which the plant develops. Results from the fieldwork and literature were used to develop different parameter estimates for each soil type for such components as seed germination rates, seedling establishment rates, and seedling mortality.

The forest object is divided into a young forest and a mature forest. The young forest transfers to a mature forest over time as the dominant plants reach maturity. Old trees die as they reach their life expectancy, and they leave space for disseminated seeds to germinate and develop into saplings. Intra- and interspecific competition is involved in the forest object and can result in premature deaths through competition and the creation of additional space for recruitment of new individuals.

Parameters used for each species object in the model are described in Table 1. Values for the parameters were derived from research results involving periodic measurements over the past 20 years on vegetation development on different soil materials, from soils-overburden studies, as well as germination, survival, and early growth studies of each species (Gouvernet, 1980; Duggin, et al., 1982; Grigg, 1987). These data were then supported where necessary with results presented in the scientific and technical literature.

A schematic diagram of the model is shown in Figure 2. The first component of the diagram identifies the parameters and rules used in the model for each object. Spatial data are stored in the GIS and can be retrieved and imported into the agent-based model as an object. The second component is the model simulation and shows each object and its relationship as well as the functions that interlink those objects. The third component addresses model application, initially by providing graphical and map outputs followed by its evaluation and validation and then its use in management for mine rehabilitation purposes. The model application diagram shows ecosystem development over time starting from its initial state and moving through a range of conditions to the desired state along the successional trajectory (Hobbs and Norton, 1996; Grant, et al., 2001). It also shows two examples of undesired states and how their successional development is simulated over time to come back onto the desired trajectory.

TABLE 1 Parameters for each species used in the model

Parameters		Explanation
Seed rain	Seed production	The amount of seed produced by a species in a good seed year (seeds per unit area)
	Viability	The proportion of viable seeds (% viability)
	Viable seed yield	The number of viable seeds produced (production $\times$ viability)
	Seeding frequency	The time from one good seed production event to the next (years)
	Dispersal distance	The dispersal distance from parent tree (m)
Seed germination and establishment	Germination rate	Seed germination rate for each soil type (% germination); the soil type is read from soil object.
	Seedling survival	Number of seedlings surviving over the first three years; soil type affects seedling survival.
	Competition mortality	Competitive ability of two plants in a cell; affected by species, height, and age
Species persistence and dominance	Growth	Annual increment in height (m)
	Height	Cumulative height growth; height influences competition
	Age	Increases with each time step (yr)
	Maximum height	Growth terminates when the height of a plant reaches maximum height (m).
	Maximum age	The plant reaches its life span and is then removed from the model (yr).
	Shade tolerance	A criteria for growth and competition (a ranking of relative tolerance)
	Mortality rate	Random mortality applied to mature plants (5%); species dependent

The model is designed to incorporate the ecological processes of seeding, germination, growth, mortality, and competition (both within and between cells of the space object) and simulates how a disturbed site regenerates and develops over both spatial and temporal scales. Measurements and observations on vegetation development during the first two decades showed that the topsoil provided a seed bank dominated by shrubs, while the nearby natural forest provided the seed source for tree species. Consequently, tree density across the overburden heaps decreases with distance from the seed source. Different buffer zones were used in the GIS (Figure 3) in calculating seed rain onto each space object and then transferred into the model. When individual plants reach maturity for that species, they provided an additional seed source that could be dispersed according to the parameter estimates used in the model for that species.

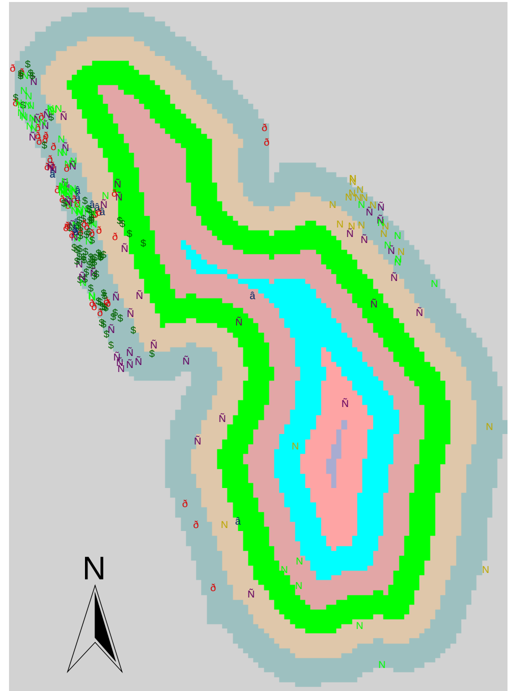

FIGURE 3 Seeds dissemination from the surrounding natural forest across the overburden heaps (The 10-m buffer zones used in calculating seed rain are shown in different shades of grey, while individual trees are shown as different colored dots.)

Simulation Results

The objective of this model is to simulate the natural regeneration process and to predict spatial and temporal patterns of different successional states that can then be used in decision-making strategies for mine rehabilitation options. One-year time steps were used in the simulation. Each space cell can contain up to one mature species, one sapling and many seeds, or alternatively two saplings with many seeds. Dynamic changes for every individual plant in each time step follow a set of rules based on the ecological processes for each species.

Results of simulations after 100 and 200 years highlight the spatial and temporal variations in community structure (Figure 4) and floristic composition (Figure 5). In general, trees and forests developed around the margins of the site, while shrublands dominated the interior and persisted for considerable periods of time. The abundance of species across the study site suggests that particular species of shrubs are driven by pulses of regeneration, while others show an initial increase and then become relatively stable (Figure 6). Consequently, community structure and composition appear to stabilize over long time periods, particularly in the absence of exogenous disturbances such as fire, extremes of climatic variations (droughts or abundant moisture), wind damage, or heavy grazing. These findings imply that the initial floristic composition is a strong determinant of community structure and composition over long time periods (Deutschman, et al., 1997).

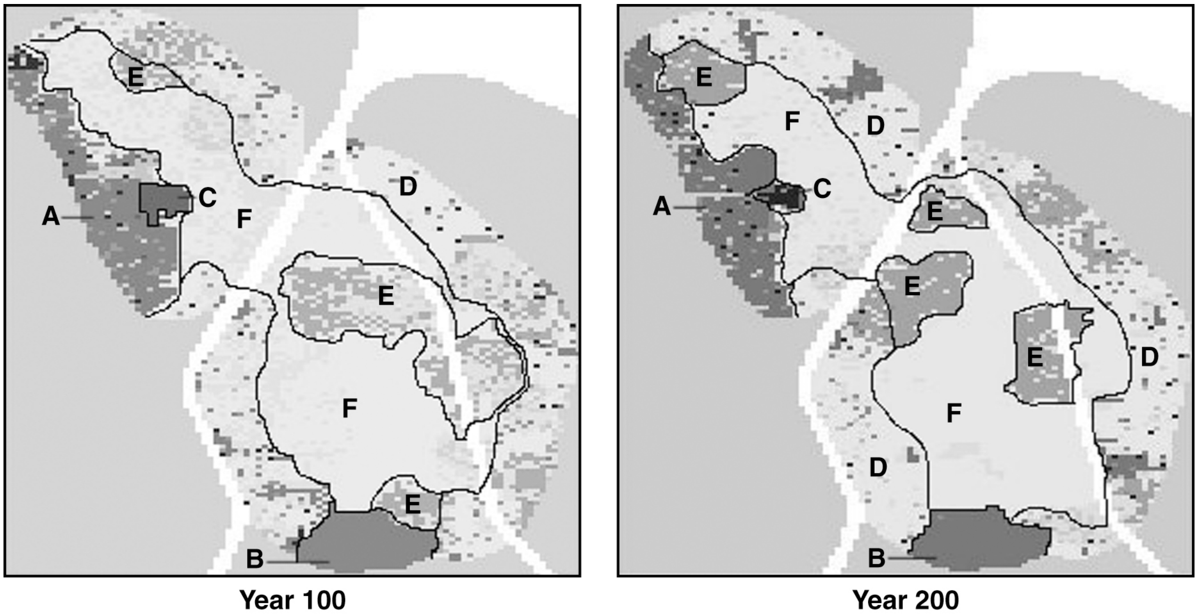

A	Forest	Mixed forest of *Callitris glaucophylla*, eucalypts, and casuarinas
B	Forest	*Callitris glaucophylla* forest
C	Forest	Eucalypt forest (note the species dominance changes from year 100 to 200)
D	Woodland	Mixed tree species with a shrubby understory
E	Shrubland	*Cassia nemophila* shrubland
F	Shrubland	*Acacia deanei* shrubland

FIGURE 4 A map of community structure after simulating vegetation development for 100 (left) and 200 years (right)

Sensitivity Analysis

A variety of information sources are used to estimate model parameters, such as seed dispersal distance, seedling mortality, seeding time and frequency, height and growth rate. Seed dispersal and time of seeding appear to be important parameters that influenced the initial floristic composition in this model. Most of the study site was covered with stockpiled topsoil, which contained an important seed bank for shrubs. When environmental conditions are suitable, the seeds of shrubs germinate and seedlings establish over the topsoiled area. However, tree seeds are dispersed from the natural forest that surrounds the site, and their germination and growth depends on the time of seeding and dispersal distance. After only three to five years, shrubs begin to produce seeds and disperse them in the immediate vicinity of parent plants. Over several generations, they begin to occupy larger areas. Trees tend to grow on the bare overburden site where shrubs are absent and on some topsoiled areas where shrub development has been sparse. Competition between trees and shrubs restricts tree growth on topsoiled areas, particularly if there is no further disturbance. Sensitivity analyses need to be completed in Stage 2 of this project to evaluate and understand which interactions make significant contributions to the development of the structure and composition of the plant communities.

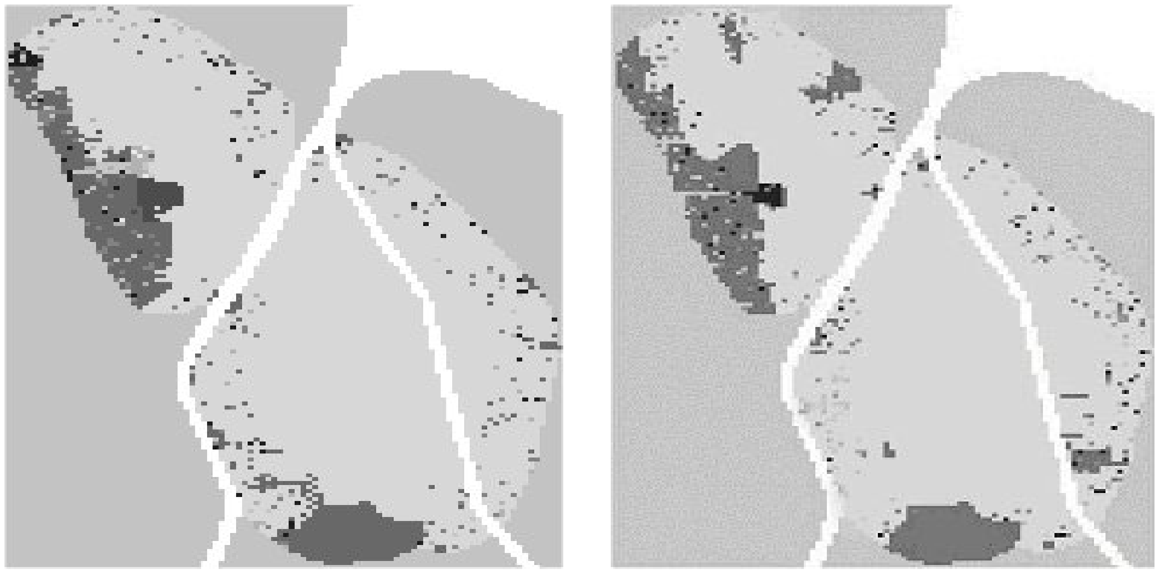

Year 100 Year 200

FIGURE 5 Overstory species composition in the forests and woodlands after simulating vegetation development over 100 (left) and 200 years (right) (Different shades of grey represent different tree species emphasizing the restricted distribution around the margins of the site.)

STAGE 2 – FUTURE DEVELOPMENTS

Stage 1 results demonstrate that the concepts and principles of ecological succession can be modeled in Repast and provide meaningful results. However, additional work is required to refine the model and experiment by modifying conditions to evaluate their impact on community structure, composition, and spatial distribution.

Further developments to be considered in Stage 2 include:

- Develop rules and functions for competition between individuals in adjacent cells, so that large individuals can occupy more than one cell.

- Refine parameter estimations through additional quantitative analysis of ecological data.

- Evaluate the potential to introduce stochasticity for important parameters.

- Continue the sensitivity analyses for parameters, rules, and functions to evaluate those interactions that most contribute to the spatial and temporal dynamics for each of the recognized vegetation communities.

- Attempt to isolate sources of uncertainty emanating from errors associated with parameter estimations and stochasticity.

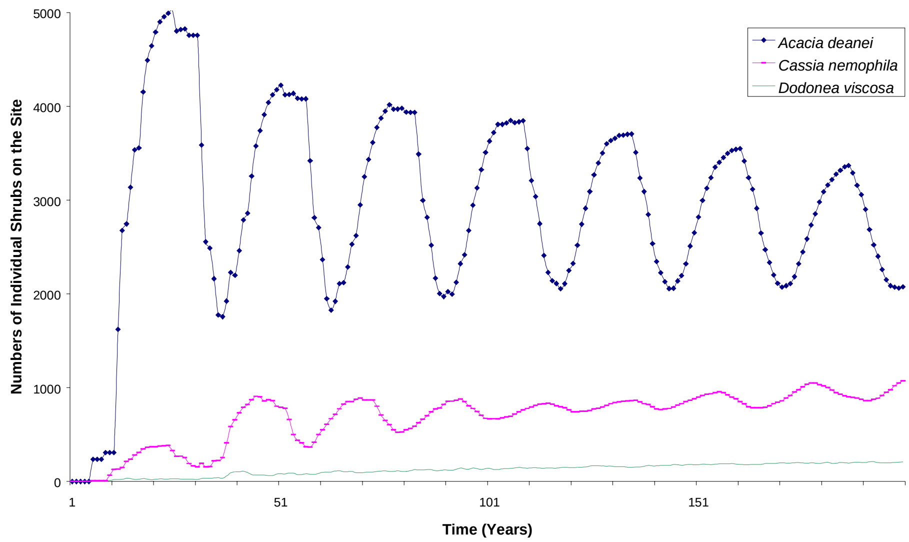

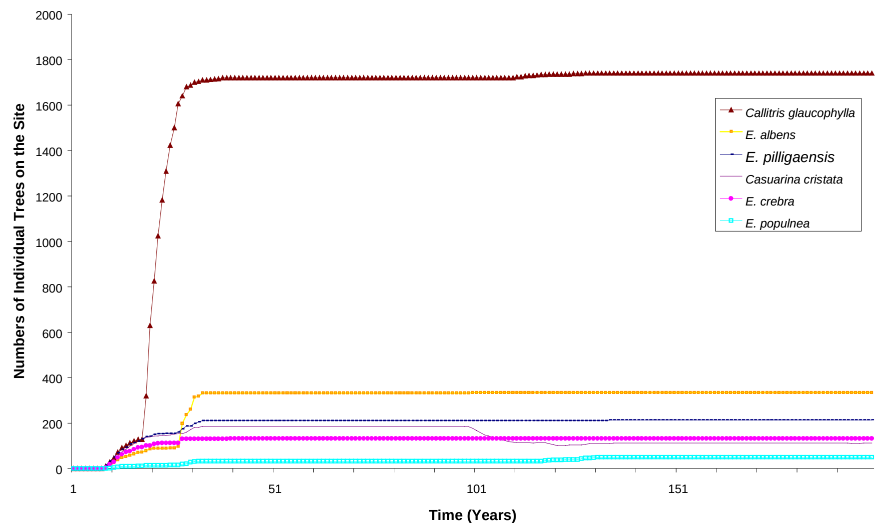

FIGURE 6 The abundance of shrubs (above) and trees (below) across the study area after simulation for 200 years

- Validate output against the conditions observed and measured in the surrounding forests and ensure that the simulation predictions are ecologically sound and consistent with contemporary knowledge of forest and shrubland dynamics.

- Develop a protocol to transfer the output from the model into the GIS and map community structure and floristic composition.

- Vary the initial conditions at the commencement of the simulation to represent different landform designs (e.g., slope, aspect, overburden, types) and management strategies (e.g., site and soil preparation, introduction of seed mixes of different composition, introduction of vegetation manipulation strategies).

- Introduce disturbance regimes into the model by using different kinds of disturbances over a range of frequencies and intensities.

- Make the interface more interactive and user-friendly (menu-driven) so users, such as rehabilitation officers without programming skills, can use the model to evaluate the likely end result from adopting different management strategies to suit local conditions (such as site variations, different forest endpoints to match the pre-mined forests, and antecedent climatic conditions).

DISCUSSION AND CONCLUSIONS

Agent-based modeling, coupled with GIS, can be used to simulate vegetation succession across a range of sites with different environmental conditions. In this study, only one environmental character was used (soil condition) but others can be incorporated, particularly if they are influential in driving ecological processes and vegetation development (e.g., slope, position on slope, aspect elevation, soil moisture regimes, soil fertility, and climatic variables). The procedures developed here highlight the way in which spatial information in a GIS can be incorporated into the agent-based model. Likewise, the buffer zone function in the GIS can be used to determine spatially dependent parameters, such as seed dispersal from natural forest boundaries into the study site. Parameters and rules must be developed to explain how each environmental variable will influence and control each species object and affect competition within and between species.

Establishment of individuals depends first on seed rain onto each space object across the study site, then the ability of each species to germinate and establish, and finally on the ability of each species to persist through competition and become dominant in that space object. Further development will be needed to expand competition of individuals from within the cells to between cells in the space objects so that community structure and composition can then be identified.

Species composition on topsoiled area differs from that on the bare overburden as a result of initial propagule composition differences and the competitive ability between trees and shrubs. The topsoil introduces a seed bank dominated by shrubs, whereas the tree seed is poorly represented and mostly comes via dispersal from parent trees. Once the shrubs are established, they generally outcompete the less tolerant tree species. Trees can occupy gaps created by the death of a shrub, but again, they have to compete with shrubs for the site. Thus, the use of topsoil in rehabilitation is an appropriate procedure to gain vegetative cover and stabilize the site, but the shrublands will persist for long periods of time and restrict forest development unless management intervention strategies are adopted.

The simulation results indicate that the developing vegetation has several prominent forms — forests, woodlands, and shrublands. Once established, these forms tend to be relatively stable. The floristic composition within each community also stabilizes and shows little change over the period of simulation. These results are similar to other models for forest succession. For example, Deutschman, et al. (1997) used the SORTIE model in southern New England forests, where the spatial distribution of species demonstrated relatively slow dynamics over the 1,000 years of simulation. However, when disturbance was introduced into the model, the floristic composition varied and favored those species able to take advantage of the gaps created. The introduction of disturbance to our model may well change the development of different communities in terms of both structure and floristic composition, and a reduction in shrublands, with a subsequent increase in mixed forests.

For rehabilitation after mining, different strategies can be adopted to develop communities suitable for the desired endpoint and subsequent land use. The development of this model can then be used to predict the likely communities that may develop with each management strategy. For example, the model may be used to predict community development on areas sown with different seed mixes (varying in species composition and abundance). Alternatively, the site may be prepared by using different techniques and substrate materials so establishment conditions can influence the success of different species. Also, fire could be introduced at various stages in ecosystem development. In this way, management decisions can be made as to the strategy that would best achieve the desired ecologically sustainable endpoint, and at the same time avoid development of undesirable states.

ACKNOWLEDGMENTS

We wish to acknowledge Dr. Lalit Kumar for assistance in developing some of the rules used in this model and Cate MacGregor for assistance with the GIS. Dr. Kumar and Ms. MacGregor are from Ecosystem Management, University of New England.

REFERENCES

Brooks, D.R., 1981, “The Planning and Logistics of Mineral Sands Rehabilitation,” in *AMIC, Environmental Workshop*, Canberra, pp. 80–92.

Collier, N., 2003, “RePast: An Extensible Framework for Agent Simulation;” available at http://repast.sourceforge.net/#download.

Croft, J.B., 1979, “Report on the Botany, Wildlife and Ecology of Leard State Forest,” Newcastle (unpublished).

DEST, 1996, *Australia: State of the Environment 1996*, State of the Environment Advisory Council, Department of the Environment, Sport and Territories, Canberra.

Deutschman, D.H., S.A. Levin, C. Devine, and L.A. Buttel, 1997, “Scaling from Trees to Forests: Analysis of a Complex Simulation Model,” *Science Online*. Available at http://www.sciencemag.org/feature/data/deutschman/index.htm.

Duggin, J.A., J.M. Gouvernet, M. Fosdick, and T.R.G. Watts, 1982, *Rehabilitation Studies for the Proposed Boggabri Open-cut Coal Mine Amax: BHP Boggabri Coal Project*, Natural Resources Project Report P.R. No. 19, Department of Ecosystem Management, University of New England, Armidale, NSW, Australia.

Farrell, T.J., and D.C. Kratzing, 1996, “Environmental Effects,” in D.R. Mulligan, ed., *Environmental Management in the Australian Minerals and Energy Industries: Principles and Practices*, pp. 14–45, Sydney, Australia: UNSW Press.

Fox, B.J., 1990, “Two Hundred Years of Disturbance: How Has It Aided our Understanding of Succession in Australia?” *Proceedings of the Ecological Society of Australia* **16**:521–529.

Gouvernet, J.M., 1980, *Rehabilitation Research for the Proposed Boggabri Open-cut Coal Mine*, BNatRes Honours Thesis, University of New England, Armidale, Australia.

Grant, C.D., J.A. Duggin, I.K. Meek, and M. Lord, 2001, “Endpoint Criteria and Successional Pathways for Manganese Mining Rehabilitation on Groote Eylandt, Northern Territory,” in *26th Annual Minerals Council of Australia, Environmental Workshop — Back to the Future,* Adelaide, SA, October 14–17, 2001, pp. 129–142.

Grigg, A.H., 1987, *Tree and Shrub Recolonization of Coal Mine Overburden — Boggabri, NSW*, BNatRes Honours Thesis, University of New England, Armidale, Australia.

Hobbs, R.J., and D.A. Norton, 1996, “Towards a Conceptual Framework for Restoration Ecology,” *Restoration Ecology* **4**:93–110.

Unwin, G.L., 1985, “Mineral Sand Mining in NSW, Australia to 1986,” in *AMIC Environmental Workshop*, Canberra.

Unwin, G.L., 1988, “Rehabilitation of Coastal Ecosystems after Mining within Crowdy Bay National Park, NSW, Australia,” in *AMIC Environmental Workshop*, Darwin.

Wiram, V.P., 1979, *Soil-overburden Characterisation of Boggabri Coal Prospect Area, Boggabri, NSW, Australia*, Report for AMAX B.H.P. Boggabri Coal Project, AMAX Coal Company, Indianapolis, IN.

DISCUSSION:

ECOLOGICAL INTERACTIONS

(Saturday, October 4, 2004, 1:00 to 3:00 p.m., Session 2)

Chair and Discussant: *Pamela Sydelko, Argonne National Laboratory*

Simulating Land-use Entelechy Using the Multi-agent-based Environmental Landscape (MABEL) Model[1]

Pamela Sydelko: I think we have a really great session ahead of us. As a set-up, I thought one thing I would mention is that it seems to be that there's a mix of using agent-based modeling in the digital, virtual laboratory sense, to actually uncover what I think is to basic understanding of ecology, and then there are others who are using it more as policy-making or decision-analysis types of tools. And I think that's interesting, because that tends to be in this field two different philosophies that tend to sometimes clash.

And so I'm interested in seeing the way that we'll kind of grapple with that, because I think ultimately modeling a simulation has always been thought of as prediction.

To start us off, Kostis [Alexandridis] is going to present his group's paper. I will assume that if there are any co-authors in the audience, you can go ahead and introduce them yourself. And we're talking about simulating land-use entelechy using the multi-agent-based behavioral economics landscape, MABEL, model.

Konstantinos Alexandridis: I would like to thank the organizers of this conference for this great opportunity to present some of our work. I would like to introduce the co-authors of this paper: Dr. Brian Pijanowski, who's the greater organizer of this effort, because it requires a lot of collaborative effort in that, and Zen Lei, who is our developer of the group we have. We used to be at Michigan State University, but now some of us are going to Purdue.

[Presentation]

Sydelko: We have time for questions. Does anybody have some questions? Yes.

Mike North: You're saying that you're using the Microsoft belief network to do Bayesian analysis behind the scenes. One thing I was wondering is whether the graphs were all the same for similar type of agents, so the farming agent would have the same belief network?

Alexandridis: The graphs are for agent classes. So they correspond to land-viewed classes. We chose to group it that way because that corresponded with a Swarm construction of MABEL. So it's easier for each class of agent to call a specific belief network. And I think it makes sense in reality. We can construct, if you have the data and the whole process, belief networks for even smaller classifications of agents.

---

1 This topic was also presented during the workshop on Thursday, October 2; see page 65 for the full paper.

North: Right. But under the circumstances you'd really have no information to make those networks more specific, I would think. So it sounds like a reasonable thing to do.

Alexandridis: Yes, and sometimes you have to put a line and say this is the obstruction that I have, the hypothesis that I have, or the assumption I have to make to create a realistic model.

Zien Lei: When you search for the optimal policy, what's your performance index, or otherwise, what's your objective function?

Alexandridis: Well, to solve the mathematics, after you go through the belief network assessment and the fitter, you define a list of actions that are optimized actions. To choose the action and maximize utility, we have a set of static ways, since we ran a dynamic programming Markov decision process that converges into the action that maximizes the utility and has the relationship with the transition metrics and the reward faction.

Claudio Cioffi-Revilla: I am familiar with only one other model of land use and cover change, which is Fearless of the McCauley Institute of Gary Bohill and Nick Gotz, and it strikes me that your model is far more complicated in the internal structure than theirs is. What additional phenomenology is this ABM producing that the far more simple model of Fearless can't do?

Alexandridis: I think it has to do with the modeling approach you take. I think how we build the model is by first sitting down and assessing what's important and how we get there. Then we put equal emphasis on the modeling component as well as the realistic group representation of reality. So I think it's a very good question. I think it emerged at some point.

Brian Pijanowski : That's an excellent question. In our lab, we've got our MABEL model and then we have our Land Transformation Model, which is this kind of black box, neural net model. We're not simulating any process — we're interested in just fitting a curve. So in the next year we're going to be spending a lot of time trying to figure that question out, because it intrigues us, all of us.

Unidentified Speaker: So the evolution of this is not a projection of the future?

Alexandridis: Yes, it is. We embedded the agents with some properties, but they drove their properties from the data. That's what the Bayesian classifier is doing. It takes the data, classifies them and gives some meaning to the data according to some other variables that, as I said, are associated with belief. So the agents are adoptive in that when you update, you do an inference back to the network. The agents define their own variables, a variable that's 10 time sets were significant to identify the agent behavior might end up not being significant after 10 more sets.

Cioffi-Revilla: Okay, so then the land use change that we observe on that region is a future evolution of the system, not the replication of the past.

Alexandridis: No. The historical data are used on the model initialization state.

Xu-Hen Su: Have you also generated transaction data, and if so, can the transaction data be actually comparable to the reality you observed?

Alexandridis: Well, that's the sensitivity analysis that needs to be done. We start talking about running the model a thousand times and having all the range of results and compare what kind of properties you see there and how you can group them together for comparison. That's the next stage—unfortunately, building a model like that takes a lot of time and effort and data mining processes. So that's one of the things that we are working with.

Unidentified Speaker: I saw lots of computational intelligence stuff being part of your building of the sophisticated agent. For example, you have a belief network that can maybe replaced by some other CI techniques. Have you considered alternatives?

Alexandridis: Well, yes, a lot of those techniques were used traditionally on robotics and artificial intelligence. The earlier version of the model didn't have the data fusion. We had the problem of addressing long-term goals while projecting step by step. When faced with problems like that, you search and find the best solution, so it's a trial and error process, and I think a lot of modeling processes are like that. But I'm sure there might be other things that we haven't tried, and we are very open to exploring those horizons.

Pijanowski: I think you made a very good point. One of the things that we're very interested in doing is modularizing the different tools. And so we've created an environment that will allow us to do that and to do it fairly efficiently, because it's in a nice clean client-server mode and everything's managed on the server side. So what we can do is just plug in a new tool, just to see what that tool could do for us in terms of understanding the system, increasing predictability, and studying the patterns that emerge from that type of simulation environment.

Reciprocal versus Group Altruism among Vampire Bats

Gennaro Di Tosto: My name is Gennaro Di Tosto, and co-authors of this work are Mario Paolucci and Rosaria Conte, of the Institute of Cognitive Science and Technology of Rome. The main topic of this work is altruistic behavior, and we designed a simulation using a multi-agent system and tried to point out some of the important issues concerning altruistic behavior.

[Presentation]

Sydelko: Do we have any questions for Gennaro? Mike?

Mike North: That was a very interesting paper, particularly the vampire blood-sucking angle. One question I have, though, as you mentioned at the beginning and then alluded to at the end, and that is to say, how do groups of cheaters form in the first place? In particular, in the studies that Wilkinson and others did, was it typical in a roost to have a mixture of cheaters and altruists?

Di Tosto: In nature, things are not so simple. Let's say that between groups, it's possible to have some kind of migration, and it is also a fact that we can simulate, but inside a group, we have that altruistic behavior is common and is the rule for defining an agent. The question is how is it possible for a simple animal like vampire bats to avoid exploitation by shooters?

In nature, Wilkinson had studied that there is particular mechanism to avoid the newcomer, which can enter freely in a roost. So there is a cost to sustain, to become a member of a group. But without this, the question is still there. How a group composed only by altruists can pop up in a population. So do we have to account for a role played by groups in this mechanism? Or is it possible that, well, the simple interaction between agents can account for this?

North: But it seems that the answer you're providing is that it's actually a random process, in the sense that groups will die out until you happen by chance to get a group that's completely altruist. And then that group will then survive. And so it's kind of a random watch.

Di Tosto: It's not a random process, because we have a mechanism for selection.

North: Well, right. But what I was trying to say, though, is that given your mechanism for selection, you start out with a given population, and essentially anyone, any group that had — considering groups now as individuals — a cheater in it is doomed basically. Then, whichever one happens to sort of win the lottery and be all altruistic is going to succeed, and in that sense the mechanism is basically random.

Di Tosto: In groups, I always establish a cost for the presence of a selfish agent. So the problem is how to avoid that kind of agent. And, yes, without any mechanism based on the complexity of the agent, like memory and individual recognition, we have to find another way to give account for that mechanism. And maybe group formation is one of the possible answers.

North: That's very good, thanks.

Sydelko: We might have room for one more question, a short one.

Cioffi-Revilla: As you know, the empirical social science research on within-group and intergroup conflict and cooperation is enormous. It's really huge. Think just for example in terms of in anthropology, all the data that exists in the human relations area files and things of that nature.

How is this research project planning to incorporate what we know empirically about the way in which groups, roosts work together. Is there a plan to build on that empirical base for validation purposes and so on?

Di Tosto: Well, this work was born after another simulation in which we tried to put a different kind of agent. We tried to implement in that agent an accounting for the emergence of offers based on the reciprocal altruist mechanism. We presented that work at the conference for the European Social Simulation Association, and it's possible, according to you, to obtain the same result without that mechanism, but only by a simple underlying structure. Well, we try to give an answer to that question with this other experiment.

The path forward will be to integrate this process and provide more sophistication for the model of the behavior of these simple animals, including also parental caring that we do not consider in this simulation and also migration between groups.

Agent-based Models of Land Use and Cover Change in the Atlantic Coast Region of Nicaragua: Examining the Agents of Tropical Deforestation

Luis E. Fernandez: I would just like to thank the organizers of this group for a wonderful meeting. I've met a lot of interesting people, and I'm especially pleased about the range of topics that are covered under this conference here. Also, thanks to Kostis for letting me borrow his high-tech device so I don't have to be tethered to my laptop.

This is my dissertation work from the University of Michigan. I'm just going to very briefly go over the overall project — the agent-based model is just one component of the project — and then go into the description of the social system I studied, specifically deforestation in Central America. Then I'll talk about the agent-based modeling, some "philosophical constructs" that I used in designing the model (I think I wanted something a little more empirical), the environment, the agents and the interactions between the agents, and finally some patterns in the model runs and how they relate to other results and other parts of the project.

[Presentation]

I'm out of time, so I'm just going to throw out some acknowledgments. Thank you, University of Michigan, and the people I've worked with there, and some Fulbright recipients. I got some money from UCA. These people who are our workers are Nicaraguans who don't really get the kudos that they should; they are in very under-funded agencies. I'll be happy to take any questions.

Sydelko: Who wants to start?

Unidentified Speaker: I have two questions. What kind of a spatial regression method do you use? What's the window or the size of data you use?

Fernandez: The data that I used, and actually I have a listing of them, are part of a larger presentation. The data were from various sources, so I had landscape, land site data, soil, gradient, roads, rivers; there's a greater list actually. And they were all digitized, rasterized, and brought to 30- × 30-meter resolution. I ran logistic regression when I took a look at whether it was forested or deforested. I also did multivariate to take a look at what the direction or the trajectory of land use change was. And I did it for three periods. That was 1959 to 1986, 1986 to 1989. In the interim, there was a large hurricane that basically went right over the area. It was extremely catastrophic, deforesting the area to a certain extent, or defoliating it, I should say. And then from 1989 to 1996. And these also correspond to political periods. One is the Samosa period, which is a dictatorship; the second is the Sandinista period, which was the very famous socialist-communist regime; and then the last period was the rise of the neo-liberal model and the switch from a communist-socialist central system to the neo-liberal model that's in place now. And I have another presentation that goes into how that affects the forest tremendously. The neo-liberal model actually just seems to correspond with a tremendous increase in deforestation.

I tried to "segment down" spatial autocorrelation as well. I tried to "detrend" the data a little bit. Those were the results that I showed: those four—there were actually six, but since I wanted to be a little conservative, I wanted them connected with the information that I got from the survey.

Unidentified Speaker: The second question is really, have you had a chance to look at the deforestation process from the percolation?

Fernandez: I haven't actually, no.

Unidentified Speaker: So in your study area, is there any forest law that imposed …

Fernandez: There isn't, actually. The only thing they have is forest concessions, very large blocks of land that are drawn in Managua and sold off or given to curry political favor, for development. The Mezzo-American Biological Corridor Project is something that's being started to try to maintain these areas. So I didn't show the borders of them, but most of this area actually falls within an ecological preserve. But it's one of those famous paper parks or preserves, where there really isn't any protection at all. Theoretically, no one should be living there, no one should be using it, but nobody that lives there actually knows that it's a preserve.

Unidentified Speaker: So are you going to involve some kind of policy intervention in your model?

Fernandez: I go back and forth with CITCA, which is my counterpart there, but they don't have an awful lot of power, and the folks in Managua don't, and I've talked to people at the Agricultural Ministry in Marena, which is sort of the natural resource agency, and I guess they're a little jaded because there are lots of agencies that are doing studies. I haven't seen that much [interest in policy studies].

Agent-based Modeling for Vegetation Succession after Open-cut Mining: Stage 1, A Study in Progress

Sydelko: Our last presenter this afternoon is Xung-Fung Su, who is going to be talking about agent-based simulation for vegetation succession in open-cut mining.

Xung-Fung Su: Thanks. My name is Xung-Fung Su, from University of New England, Australia. And this topic is a part of my Ph.D. Here I should say that I have a Chinese *and* Australian accent, so if you can't hear clearly, please feel free to ask.

Okay, the topic is about agent-based simulation of vegetation succession after open-cut mining. Just now we heard about a very attractive and big model, but here I'd like to give you a more simple model in a smaller area.

[Presentation]

Sydelko: Okay. Are there any questions?

Alexandridis: Something wasn't very clear about your intentions on the agent-based modeling. Are you studying just the vegetative cover or are you studying the ecosystem as a whole, because I noticed you said there are study areas that have been isolated from animate life, like animals and kangaroos and things like that. So the question is what are you really studying and is it realistic to say there will be succession without having any animals moving in and out.

Su: Good question. Most mining sites were just left there, without any management or other disturbance, so the question is "how can these sites be recovered?" from natural forest to mining production and then back to the natural forest. And so the fence is used to keep the kangaroos and the wildebeest out of the site, to prevent another disturbance, because the kangaroos can eat the seedlings, which would involve another problem for the site. Also it should be noted that there are farmers who just go inside and cut down trees for the wood; if there is no fence, the recovering vegetation is too easy to destroy.

Sydelko: Along the same lines as that, is there any thought on trying to look at feedback loops. I don't know if you use soil fertility as initial data to look at establishment, but it seems to me if you're doing a 200-year simulation, there certainly are some feedback loops between establishing the soil as well as the seed bank. It also provides nutrients and organic matter, and over time, you'd think the soil itself would then actually evolve and you wouldn't be using the initial soil conditions anymore. You'd actually have a different soil over time, and if there's any thoughts on how you'd introduce those kinds of changes, those dynamics into the model.

Su: Companies are required to first remove the soil and keep it in a particular site other than the mining site. After logging and mining is complete, the waste, we call it burdened waste, is disposed of on the site and the top soil is then laid on top. So, yes, it's true that I only used the soil as any other factor. If they haven't got enough soils, some burdened waste is left [exposed] and then that site has properties more like the waste and chemicals.

Sydelko: Right. I would be interested in the physical properties, too. I'm just specifically interested in this. I did my thesis on strip mine reclamation, and when you put the soil back out and mostly the compaction's a very big deal, and over time that becomes less of a deal. And it could be that part of your selection for shrubs and grasses is that they are more able to establish themselves because of compaction, but over time the tree species might actually have a better competitive advantage in a swell that is not as compacted, because over time, you're seeing some feedback loop in the sense that there's some establishment of a certain kind of vegetation that will actually change the soils, and that's part of succession that actually opens it up for another stage of succession, being the forest, because the conditions are changing also.

So something that I think would be interesting to add, some of those dynamics just to see what would happen to your model.

Su: Yes, in the future this change of conditions will be introduced more and more … There are no computations from the shrub to the tree species, because there is no top soil. There is a forest and a few shrubs growing there and dropping seeds in there, but because there was no top soil in there, seeds couldn't even germinate to become established in that site.

Sydelko: Oh, that's interesting.

Su: And, yes, but inside the topsoil sites there are many kinds of many trees and shrubs in as little as two years — they germinated just when the environment was suitable for them to quickly germinate and occupy that area.

Sydelko: Obviously, the different kinds and methods of putting the topsoil back on causes a great deal of compaction below the topsoil layer and then on top again, so there's layers of compaction, then there's relatively little, and then there's compaction again. And it's very difficult for establishing roots in those kinds of soils. And it also seems that over time we noticed that as organic matter is built up again and the texture was changing, that then it became a natural thing, that the soil itself, even without ripping or anything, would establish a better environment.

Closing Panel

CLOSING PANEL

(Saturday, October 4, 2003, 3:45 to 4:45 p.m.)

David Sallach, The University of Chicago
Desmond Saunders-Newton, University of Southern California
Claudio Cioffi-Revilla, George Mason University

Moderator: Welcome everyone to our closing session of Agent 2003. We have a distinguished panel to address the challenges in social simulation, which, of course, is the theme of our conference.

We have on the panel David Sallach from the University of Chicago and Argonne. David is also the Vice President of the North American Association for Computational Social Organizational Sciences, NAACSOS, as well, in case that wasn't mentioned earlier in the conference.

We have Desmond Saunders-Newton, whom many of you know of … as he said earlier in the day, he wears several hats: DARPA, DIA and USC, University of Southern California. He is also the Associate Editor for the *Social Science Computing Review*.

And then, thirdly, we have Claudio Cioffi-Revilla of George Mason University. Claudio is the Director of the Center for Social Complexity at GM.

So with that, I'm turning things over to the panel.

David Sallach: All of us are going to be very brief, because what we really want to do is open it up to a general discussion, because what we're thinking about is that the theme has been the challenges in social stimulation, and we should all discuss what those challenges are, what our greatest needs are, and where we'd like to see the most progress. You know, when we meet again next year and over the next several years, what would we consider to be progress, success, and so forth. So we're all just going to say a few things to focus the discussion initially and then open it up.

So I just have a couple of slides that have to do with, first, substantive and theoretical progress; second, technical innovation; and, third, institution building.

In the area of substantive progress, I would like to see us begin to have our papers focus more on ontological exploration. That's not too surprising that I would think that, since I think our ontologies are not consistent with how progress has historically been made in areas like quanta and genes and tectonic plates, and so forth, but I also think that there are some ontological advances that we might [achieve] by getting to the right type of explanation at the right level of abstraction.

I'd like to see us have an increasing emphasis on endogenous social dynamics, a greater emphasis on socially coordinated and self-organizing processes, dynamically blended entities as opposed to discreet entities. And that might include agent capabilities for conceptualization and classification.

I'd like to see us intensify the discourse among social and computational scientists. I think that the qualitative richness of substantive social science should be incorporated into complex models. Increasingly, models need to be driven by and provide feedback to theories, and those theories might include social theories, complexity theory, information theory, and system theory, but also the dialogues among those theoretical frameworks.

In order to do some of these things, we need technical innovation as well. And I think that one thing that might be helpful is for us to identify the kinds of areas where we need it for the substantive social science, but it will actually be an achievement in computer science as well. And I think there are areas like that, and those are areas that are naturally fertile for dialogue.

This means that we need to raise our abstraction level for modeling and program design and things like that, including ensemble representation and high-dimensional visualization. In my view, high-dimensional visualization would be very, very nice.

I think that we have underutilized constraints to this point, because constraints give us constraint programming, gives us a fairly high declarative form of specification, and there are needed areas in constraint programming that we have not developed at a computer science level. And I'm thinking of things like context-sensitive constraints, which could go a lot further than we've seen to this point, and endogenous constraints. I've seen nothing on endogenous constraints, and yet think how nice that would be to have that available as a tool.

And, of course, since I said something about it once or twice, you know that I'd like to see development of semantic and pragmatic processes. They pervade perception, communication, action, and interaction. And the technical area there might involve meta-reflection and auto-instantiation along multiple or complex dimensions.

Regarding ontological experimentation via generative dynamics … Roger has, I know, been active in this area. And I think that when we talk about generating families of models, that's one way of pitching everything at a higher level of abstraction, whether that's at the architectural level, where it would be nice to have support for extensive prototyping of simulation types; at the design level, where perhaps agent identities and preferences can be generated via frequency distributions and, of course, parameter sweeps; and at the agent level, where abstract action types can be indexically invoked. My basic argument here is that the realization of some social science goals will require computational advances.

We also need to think about institutional support. Frankly, in academic and other institutions, the support for computational social sciences is fairly fragile and peripheral. So we need to think about how that can be addressed. I think that some help can come from professional societies and professional associations like NAACSOS, which has an annual meeting in June 2004 in Pittsburgh. The call for papers has not been announced yet, but I encourage you to keep an eye on that. But this conference was, as you know, partially organized by the SIGs from NAACSOS: the Methods, Techniques and Toolkits SIG, the Simulation Application SIG, Computational Social Theory SIG and Computational Organization Theory SIG. And within that context we could do things like organize model curriculums.

In addition to that, to the extent that it's feasible, we should look for opportunities to do program-building in the area of social computation. Social science students need cogent and integrated training in emerging methodologies, which would probably include data warehousing, data mining, spatial statistics agent simulation, and software development. Computer science and engineering students need and want effective exposure to substantive social science. So there's a need for blended or transdisciplinary programs and departments. They crop up under different names. Claudio heads one, so he'll probably be more concrete than I am, but I'm referring to social informatic, social complexity, and computational social science programs. I think it would be great to have a really solid Master's program that can either go in a technical direction or in a more substantive social science direction. There's all kinds of room, I think, for innovative doctoral programs as well. But they won't be created automatically or randomly. They need to be coordinated initiatives.

And one thing I might mention here is that computer science departments are morphing into schools and colleges. I'm working on a study of the creation of the new schools and colleges — this will be in the Communications of the ACM next year — of computing and information technology, where computer science is merging with a number of other schools, information science and so forth. It comes up under a lot of names. What you can see on the left [referring to a viewgraph] is the number of new schools and colleges that have been created in the last few years; on the right is some of the names that those programs have had and some of their academic coverage. But I think that as they expand, the exemplar here is bio-informatics. I mean, bio-informatics is one area that has been extremely active, and I think that social informatics or computational social science could do that as well.

I don't know if this is true, but I am struck by the fact that for a lot of the technology institutes — MIT, California Institute of Technology, Illinois Institute of Technology — their social science programs are kind of service; they provide service courses at the undergraduate level and so forth. But they exist; they're there. And I wonder given that they're fairly technical in their orientation, if that might not be a good place to build up a dialogue between the social sciences and the computer and information sciences and strengthen those social science programs, but strengthen them so that they're computational social science programs. Just a thought … not really my area.

So those are just a few ideas. And, as I said, if we could have progress in those areas, that would be great.

Now I'll turn it over to my colleagues.

Desmond Saunders-Newton: Good afternoon. As mentioned before, I wear multiple hats. And I'm actually going to focus on one of my other hats, the hats which actually are more of my practitioner activities in this particular presentation.

At DARPA, I am the senior science advisor for one of the technical offices. It was formerly known as the Information Awareness Office, but as of the end of September 30, it no longer existed, so it has been wiped from my memory and yours as well. However, I do still hold that position and am waiting for whatever that next office shall be.

I am also the Senior Program Manager for Advanced R&D, focused on modeling and simulation, at the Defense Intelligence Agency. So I'm kind of on loan from the University of Southern California.

What I'd like to talk with you is about some of the thoughts that I've actually picked up in reflections on today. A number of the presentations followed general themes: utility, agency, issues of environment, processes, interactions and dynamics. On my more academic days, it's a satisfying itch I get scratched, so to speak. But when I actually go back into the office where people start thinking about what will agent-based modeling do, than I have to kind of answer some interesting questions. And I'd like to kind of cast these into challenges and some things for you to think about in terms of people, tools, uses and the speculative.

In terms of individuals, some notion of finding individuals who we will call social simulationists, as opposed to simulationists in a general sense, is that they are in short supply, both in terms of how people are actually become parts of the academy, but also in terms of the practice.

Now, I use the particular phrase "simulationists," because I actually came from the military OR community before I started practicing this whole notion of agent-based models. In that community are actually a group of people who are called "simulationists." Some are system scientists, some are operations researchers, but they actually have this kind of interesting title; people who are transdisciplinary in the sense that they don't actually claim any one particular field, but they have a tool set that they like to use. And simulation models have actually been a really big part of defense industries in general.

And so part of the challenge is for agencies who are attempting to bring in people who do agent-based models … and when I say "agent-based models," it's not just about the tools that most of us have been talking about today, but also social network analysis and the use of physical science and analog models, for example, reaction diffusion, IC models, looking into migration, and issues like that. But there's been a real challenge in being able to find those types of individuals who can both do those techniques well and also can tightly couple, and I emphasize *clearly* couple, this with social science theory.

So the grounding in both the areas, whether it be CS, mathematical physics, or computational physics, with social theory is actually something of a challenge. And as it turns out, a lot of people who reside in these agencies don't have a true appreciation about the culture of a fragmented academic institution. They believe that all university people are laid back, that they get together and all talk to each other. But as it turns out, people *don't* talk to each other, and it's really hard for these agencies to find the right places to connect and bring these individuals in to work on these types of problems. So the mere fact that a physicist wouldn't be spending quality time with a sociologist doesn't strike them as being something that's really understandable, particularly given they're all on the same campus.

And, again, that leads to my next notion of what is an agent-based model? … because we all have different definitions. And that's generally methodologically driven, much less in terms of this notion.

Now, for *my* definition, and I just offer this because this is Desmond's definition as opposed to anything that lies on everyone else's radar, which is that the big distinction, in terms of the tool sets that have been typically associated with the practice of the craft known as social science research, is that we are moving from doing research on the level of the aggregate, where we aggregate things on large numbers of individuals, down to the notions of agency. So whatever methodologies you choose, this whole notion of agency-based modeling, as opposed to agent-based modeling, may be the way to think about this. And also in terms of how it is that we think about training and individuals that we connect with, and how do we actually imply notions of individual-level activities or individuals acting in groups of various types; so whatever groupings that are of interest to you, whatever types of institutions that they form.

And then this whole notion of how do you bridge between these various practice communities? If you have individuals who view themselves as practitioners, they really want to practice. They have things they wish to accomplish. For those who actually are much more concerned with living in a life of mine, where that's actually a much more leisurely timeframe, there are some real interesting challenges. And we discover quite often when we try to have contractual relationships with folks in the academy is that those timelines are actually very different than individuals who may be in the private sector or the public sector. And so this whole notion of bridging these two communities is actually also a challenge as well.

The tools. Well, I'm not just going to focus on individual, specific models. However, I do want to think about this whole notion of development environments. And that actually turns out to be a really interesting challenge. As of late, we've been focusing on at least a number of agencies with whom I've had involvement, are focusing on what I call the LCD approach, Lesser Competent Developer. Some of the models that we've actually been seeing here actually require people who are actually quite adept, either in coding or in terms of inferencing. They can see things in their model that the average mortal can't see. That is not the average analyst who resides in a lot of these institutions. We have to have individuals who are going to create these models, but essentially the tool itself is going to have to provide flags to support them.

One of the models, at least one of the structures that people were thinking about using, is much akin to the tertiary curve model. As it turns out for those of you who have the unfortunate need to work with a physical therapist or occupational therapist, you quickly discover the person who does most of your work is either a PT assistant or a certified occupational therapy assistant, because the medical model thus far has moved individuals away from using the people who are the most adept. Their job is now to manage.

Well, this is also happening in the analytic community, within the intelligence and the operational research community in DOD, for example. Basically, you have one master or well-adept analyst looking over the shoulder of some erstwhile bachelor's degree-holding person who basically says, "Well, I don't understand this piece." Or asking the right questions, like, "Why is your data corrupted?" and those types of issues. So that's what we're actually looking at, which argues for different types of development environments for the models that we're talking about here.

What's also important to remember, when you have communications between those who view themselves as being experts or sage-level with respect to modeling, is that the second language of the analyst, while it may be Spanish or Arabic or Pashtun, is not models and it is not computer code? Bridging that gap is actually quite interesting. And this whole notion about learning curve in terms about what it is to craft models and how do models fit with the real world is actually something of an interesting challenge. And right now there are people thinking that tools to develop an environment specifically is one way of actually kind of bridging this gap. Is it the right way? I'm not certain. But we are talking about a credible investment activity here.

Issues of use. Not to rework too much of what Steve did yesterday, because he and I both resonate quite greatly on this, but it should be emphasized on the following. Number one, one model or one discipline does not solve all the problems. More specifically, as it turns out — and I'll pick on economists for a moment … economists actually have a lot of leeway in terms of the models that they bring to the table. So one of the areas of particular interest to me is conflict prevention. You start looking for root causes why conflicts happen; if you talk to people who are at the World Bank, their issue with that is all about the economics, right? If you had more resources, this would work out.

Now, it turns out that if you left for the World Bank and you just go directly east for about seven miles, you'd be at the University of Maryland with Ted Gurr, who actually believed that all this is related to ethnic conflict. Whether you believe it or not is another issue. But if you left from Ted Gurr's shop and went someplace else, another individual would argue that other structural components are responsible for conflict.

As it turns out, all these people were right and they're all wrong. Right? Because basically these disciplines actually only cover a small part of the picture. So the question is, how do we actually begin thinking about appropriate ways to use models from various disciplines concurrently? How do we actually infer from this meta-model structure, things that arise from running multiple models. I know people are thinking about issues of docking. But that doesn't resolve the entire issue about, how do you do the right mix of the different types of model structures and how do you do the right collection of computational experiments? Do you actually combine results from various models of similar problems?

Another issue for us in terms of use is this notion of flipping the collection analysis presentation cycle. Back in the late '70s or early '80s there's a study with the intelligence agency, and it applies for other types of analysis as well, about how people spend their time. And as it turns out, it's something of an interesting U-shape, and we call it the tub. Collection takes a large chunk of time, analysis we don't give much time, presentation takes a large chunk of time.

There's something wrong with this picture. A large portion of our work in terms of investment, and particularly in collaborative analysis, is actually flipping this over. So basically most of the time of the analyst is spent on analysis. Data collection becomes a small portion of this, and of course presentation becomes smaller as well. Given a preponderant number of PowerPoint charts, particularly in the government, the templates are pretty much the same. It shouldn't take long to actually figure out the right format that people want to see. But there are some interesting things around presentation and thinking about notions of multidimensional data and multidimensional outcomes.

Now, all that being said about some of the challenges, I would like to at least note for you that there are really some interesting opportunities and events which are happening. One is the Preconflict Management Technology Program, for which I'm the deputy director. We're actually looking at how to couple together fairly advanced models of various types of preconflict, as well as models of elites, and coupling that with a collaborative environment to suggest how it is that we can actually prevent conflict, and if we can't prevent it, at least reduce the amount of time spent in conflict and move it toward more successful post-war outcomes.

This is a real use of this, and we've actually been in this particular program for at least eight months. And this is a test of principle within Central Asia. And it's actually one of the most interesting things, which is one of the reasons I felt somewhat obligated to state, we shouldn't sell ourselves short, because we're actually invested in this. And at this point we're not trying to get the perfect model, or the best model; we just want some models. And we just get better at the process along the road. This is a spiral development cycle, and this is how we're going to approach this.

On the activities in terms of collaborative analysis, we're actually beginning to think about how to help people on a variety of ends with this particular U shape that I mentioned to you before, when in terms of what's real-time data collection, and what does that mean in context of performing analysis? You think about the typical data sets that we work with in social science. They're generally about five, six, seven, eight, 19 years old. They've been vetted, they've been cleaned. But what if you had a real-time data stream that isn't necessarily clean, but it could produce some really interesting results, and that were based on indicators, which are actually taking advantage of the real-time data stream.

In terms of the analysis and collaborative analysis, how do I actually bring together other nodes in this whole collection of individuals, of people who are actually serving almost like processors? Analysts are actually working with data. They have their own inference patterns, they have their own cognitive frames for taking in information. And it's actually quite interesting, because now what if I have competing analysts? I can actually think about how to use the analyst in competition amongst individuals in terms of what's a better analysis or how people are asking their questions — what's their underlying assumptions? So we're thinking about that as well.

On the presentation piece … what we're thinking about, we jokingly call it the "What if Harry Seldon was real?" For those of you who are not science fiction fans, that's an allusion to *The Foundation* by Asimov. But some of the more recent work is thinking about these huge multidimensional spaces where, in my vision, someone can basically sit in the midst of it and touch certain parts of time or space. They can actually begin pulling down more data, because they have a much more graphically connected notion of what's possible. And our option space for analysts or for decision-makers is no longer three bullets on a decision memo. It now becomes maybe 250 options, in some graphical form, which is actually consumable. So that's some of the things that we're thinking about on that front.

And last of my notions is speculative … One of the challenges I think for us is to rethink the linkage between theory and methods. It is important to keep in mind that some of the methodologies that exist are truly a function of how we could … with respect to our theories. So if you think about why it is that linear aggression models used in the social sciences tend to be quite parsimonious, with very few explanatory variables, because they follow the KISS principle for the one reason. But it's also because we tend to not always have the data to actually flesh those out. And it gets more and more difficult to start thinking about 15 variables as opposed to three variables.

But one of the things it would be neat to consider is that, since we're moving away from these, we're actually thinking about other sets to choose besides just closed-form solutions and the like — are statistical twos strictly statistical twos? — is what would be an algorithmic social science? How would you recast social science theory in a context of using algorithms?

Rob Axel and I had a conversation about this a little while back, and I remember both of us saying that is there any equivalent social science problem that would suck up all available processing power the way that you can in the physical sciences? If you think about the grand challenges in terms of fluid mechanics and weather problems, is there some equivalent where you can actually just suck up all the processing power in the social sciences? We couldn't come up with one. It would be kind of neat to think about what it would be, particularly given the availability of great computing constructs that exist and will continue to exist.

On my last bullet, this is my kind of like true speculative bullet, we'll call this "the other world inferencing." Again, I'm a big science fiction fan. This is a series of books by Tad Williams. And basically these old guys decide they would instantiate themselves in the form of a computer and create their own special little worlds, and they called it "Otherworld." And so they have all their spaces they had to live in — one guy had his reign in Egypt and another wanted to live in a bug's life, as a bug.

However, somewhat similar to that and not so tongue-in-cheek, one of the big efforts now by the Defense Department is to take advantage of massive multi-player online gaming environments; not to run our own, just to be able to figure out how people form teams, how people form strategies in the context of EMOS. It's actually kind of an interesting notion. It's a different type of inferencing.

As it turns out, social science has a lot to offer this, because this is the type of work that many people who are anthropologists who have been involved in naturalistic research have always thought about. So the question is watching for patterns and thinking about this. Can we discern this? And the nice thing about EMOS is that it is always leaves a track of all the interactions that occur. Can you infer from it? Maybe, maybe not.

And along the same lines, what if we could actually solve this problem about not having enough data to deal with in different types of cultures. Artificial culture is something that is of interest to a number of people in DOD. And we haven't quite figured out what that means yet in terms of representing artificial cultures; somewhat akin to Nick Gessler's work out from UCLA, I think may be one way of thinking about this. It is based on this whole notion of, if we could create a thousand alternative worlds and situate them in a context of virtual reality and let them run persistently, where do they go? What's their trajectory? If they're that rich, what will we actually see? That might be an interesting thing to do inferencing from. I'm not quite sure what it is, but it's one of the things that we're actually just kind of playing with and saying "What if?" And that's what I get paid to do at DARPA, which is kind of fun.

So anyway, those are my thoughts, in the context of challenges. Thanks.

Claudio Cioffi-Revilla: You should know that we had no coordination on this at all. Okay? So these are three independent samples of ideas from this conference. But I would encourage you to start drawing the intersection of these comments, because I think it's not empty for sure. And I think that's quite interesting.

I just have a few ideas in terms of research teaching and organizational issues, motivated in large part by meetings here at the agent conference, this year and also some of the previous ones, as well as the NAACSOS meetings.

I think that a good idea in terms of the research program of this community is every so often set some explicit research challenges out there, sort of benchmarks against which to have a sense of progress. And the reason I say that is because, absent that, one sort of meanders and drifts into research directions, which has its function as well, but I think it is useful to have some fairly specific things that would be shared as important things to discover and to make progress on.

Another has to do with making sure that we maintain a great deal of high-standard quality control as we gain altitude and lift in this large project. That paid off handsomely in the earlier two ways of doing science. In statistical and econometric research programs and also in classical mathematical formal modeling approaches, where today, compared to what we had 50 years ago, there is really marked scientific progress, but it was won at a very hard and tough price. And I think that that should always be foremost in mind.

I think we should also test the formal modeling epistemologies that were so successful in achieving a great deal of progress in mathematical social science. We should test this in the ABM environment. I'm not saying that we should stick to it alone, and in fact new epistemological ideas are yet to be developed out of the computational social science research program taken as a whole. But some very classical and important ideas probably translate very well.

Just to mention one example, the ideas taken from Lakotas in terms of research programs. What constitutes positive and negative heuristics in a research program in computational social science? What constitutes a truly progressive problem shift, in terms of Lakotas and so forth? I think these are important ideas to consider that, again, I think are quite valuable.

I guess an early start on this was given by one of my former students, Chuck Tabor, who very clearly in that little green sage book on computational modeling leaned on Laban March, who in a different app book, writing only about mathematical models, had proposed a criteria of truth, beauty and justice. And he defends these criteria as also applicable to computational social science. So that's the sort of thing I mean.

By this I don't mean that epistemology is frozen, by no means. I think computational science will develop epistemology for a very long time, along directions that are not yet all that clear.

I think we also need to develop and strengthen research alliances with non-ABM computational, nonetheless computational social science. There are many examples of these. One that I like to point to often is work by folks like Phil Schrodt and Doug Bond, who are using certainly computational methods for carrying out event data analysis. Phil was one of the first users in the social science community of Holland classifiers. The credentials are pretty clear on that. Nonetheless, that's not commonly thought about as being computational social science, and I think that we should have that type of inclusive definition, notwithstanding the fact that agent-based modeling is obviously perhaps the major engine that is driving the discipline at this point.

And finally, along terms of research, we really need to support and focus on peer review publication. That's very important, and it has made, judging from the past in statistical and mathematical social sciences, a very, very big difference in having some truly fine journals, including the *Journal of NAACSOS* (the computational and mathematical organizational theory), the *Journal of Artificial Societies and Social Systems*, *Social Simulations*, and the *Social Science Computer Review*, which, as many of you know, began about 30 years ago as kind of an SPSS paper type of journal. But in recent years it has really followed some of the frontiers and begun to publish ABM and more modern computational research, which is great.

A few ideas on teaching … A lot could be said about this, and David and also Des mentioned many of these ideas already. We need to develop new core computational social science courses that include agent-based modeling, but as I mentioned earlier, also other varieties to expose our students to the landscape of computational social science. And part of this should also be historical, beginning with the roots of computational social science in systems dynamics days and even earlier than that. We need to exchange syllabi.

I have no idea how many syllabi already exist in computational social science; you know, probably less than 100, but I would be willing to bet more than 10. It's not too soon to begin some of this exchange and conduct more educational workshops. This meeting has done a wonderful job over the years in a dissemination pedagogical mode, even in some of the sessions, but certainly in the early days of the meetings, during the repast workshops and in the toolkit sessions on Thursday and so forth. I think more of that is necessary, even in an ad hoc manner; summer workshops and so forth. Some of this had been anticipated in an ITR grant that didn't work so well, but, nonetheless, we should pursue that. My center is certainly very supportive of this, and I'd be happy to collaborate with others in sponsoring these activities.

One thing that I think we should do is more collaboration in sponsoring events like this. This is cosponsored by the University of Chicago and by Argonne. Other similar cosponsored events ought to take place instead of putting the burden on a single institution to carry off these things.

I also think that we need to get to young kids at the high school level. I don't know about many of you, but, personally speaking, I made up my mind to become a scientist when I was in high school, not when I was at the university. I went to the university already having decided that I was not going to be a philosopher, a historian, or a musician. I wanted to be a scientist. But in those days if you wanted to be a scientist you have to be a physicist and that was it, and so on. And I didn't discover social science until when I was forced as a physics major to take some credits outside of physics and math or I would not be able to graduate. And then I did and I was captured and hijacked by social science. And my social science advisor recycled my physics and math into some social science modeling. But the point is that there are many kids at the high school level that we need to inform about the existence of this field. And I know that some of that is already taking place and should be pursued.

I would like to put a plug in for the wonderful book that Nigel Gilbert and Klaus Troitzsch have been preparing. It's coming out in a new edition very soon. And I have used it twice already, and it really works very well as a core graduate textbook around which you can build other readings. And Nigel is very helpful in providing additional ideas.

Finally, in terms of organization, how many of you are NAACSOS members already? Very good. That's a great start. So those of you who have not yet joined, tell your friends about it, have them come to meetings, visit the website. There are a lot of nice people in NAACSOS. It's a very supportive and friendly scientific society that has to be promoted. It creates a lot of public goods and positive externalities. Float that with your economist friends. NAACSOS also in organizational terms should look into the question of developing some kind of certification standards in the future. As the organization grows, I think that's a very useful thing it could provide. This may take place in a more or less formal way, but take place nonetheless.

We need to lobby the National Science Foundation much more vividly. They are obviously following very closely what's happening in computational social science. You all know that there is a new priority area in human and social dynamics. I ran into Aretha Caldwell a few weeks ago at the Italian Embassy at a party that the science attaché was giving. And I gave her an earful about NAACSOS and CSS and how wonderful it was that they had now funded the priority area in computational social science. She is personally very excited about CSS, and she took a great deal of pride in the role that she played in having the social and behavioral sciences initiate this new priority area. So she's not alien to this. She's not ignorant about it. She is, as you know, a biologist. And so that should be pursued and we should flood the NSF with grant requests. That's the major leverage that program directors and others have there. When they have a huge flood of proposals they use that to turn around and show their bosses the demand in this area. A lot of people are not aware of those kinds of behaviors, but they really are very effective.

We should do a better job of linking to colleagues in Europe, in Latin America, and in Asia. The European Social Simulation Association just had its inaugural meeting a few weeks ago in Groningen, and I hear reports it was very successful. And they are proceeding in parallel at a very, very fast pace. Eventually we'll need to create some kind of a world federation among these associations, and there is plenty of other social science experience in this area. Sociology, economics, and political science have major international federative structures. We don't need to reinvent the wheel from scratch, and we can simply collect some of the best ideas and be aware of some of the major pitfalls in terms of organizing the government of these institutions.

So those are my two cents on this, and I thank you all.

Moderator: … open it up? What are your thoughts …

Mike North: This is responding to one of Des's comments, and maybe a little more generally to the panel as well.

I really liked what you were talking about in terms of developing people, or developing toolkits, I should say, that are usable in a variety of different levels, and I think that's important. At the same time, and this also goes to the education comments, producing a range of different types of skills, people with different skills, is important. Not every analyst is going to need to build a model. Many are also going to be using either models or, even better, ensembles of models, as Steve had talked about. And so I think that's an important thing to keep in mind when we talk about levels of modeling. For most analysts, they're never going to have to write a model, but they need to understand how the models work and what they're doing with the models. And there'll be other people who are tasked with actually creating the models, but probably not using them as often.

Saunders-Newton: As a quick addition to that, one of the challenges is that as we attempt to get more of these tools on people's desks, this whole notion about making them quasi-user-friendly is one of the issues here. And I concur. As it turns out, there's a lot that you can get out of models because you've been trained in a right fashion for them. But, on the other hand, it doesn't do us a lot of good that you craft them, because no one's going to use them. So we have to make tradeoffs in always looking for the right balance on that.

Joanna Bryson: Mike's comment just made me realize that that particular goal ties into something I wanted to say about Claudio's idea, which is my personal vision, what I'd like to do before I retire or something. Right now we're trying to make these tools so that other researchers, nonprogrammers, can use them. But I would love to have tools in high schools. And I model modular theories of individual intelligence, so I was thinking about helping people to understand their own conflicting goals and things like that. And I started getting interested in this field and thinking about social interactions and things like this. But if we get that to the schools, it would have so much social benefit, just because of self-understanding, understanding friends, peers — all those sorts of things. But it could also help, I just realized, if we start much earlier with tools, basically, then we'll have more skills, and we'll have those analytic things.

Edward MacKerrow: We talked, David and I, a bit about this, but maybe I could just pose it to the panel. It seems to me that getting social simulation into the schools is a really important key step. And I'm wondering what you guys think it would take to get it into, say, an economics curriculum, maybe in grad school or maybe in undergrad, I don't know. But what do you think would have to happen for that to take place?

Saunders-Newton: I'm going to give you an example from public policy on that one, Ed.

Aaron Widavsky before he passed — he was at Berkeley at the time — wrote a book called *In Speaking Truth to Power*. And in that particular book what he spoke to was an interesting notion, how they would change the way we have been currently practicing government, particularly at the federal level. And he said basically that we'll send them master's students, chunks of Master's students who would ultimately take over government and so assure good government. In some ways, he achieved that, because in much of the analytic community, regardless of whether it's DOD or Department of State, you have lots of students who have Master of Public Policy or Master of Public Affairs degrees.

Now, as it turns out, these programs are always trying to figure out, what's the next wave? what's the tool set? I'm affiliated also with the Association of Public Policy Analysis and Management. Every year for our national conference we have an entire day dedicated to new additions to the curriculum. I don't think there's been a tool workshop, basically like agent-based models, in at least six or seven years. These people are practitioners; they will all actually go into the field to use these tools. I think it would be quite interesting to actually get students of this type, and your basically applied social science students, and it would be a part of either economics classes or the like. But it would be an interesting way … to take advantage of some of the ways of introducing this into the fields.

The other thing that we've done is to go door-to-door. We go to the University of Michigan, to what's now the Ford School, as well USC's program, where I teach in their school of Policy, Planning and Development. And basically, we just say, "You really should have these tools as a part of your class. And we'll make sure that you get people trained up." And that's one way of doing it, too. Now, that's actually fairly time-intensive. I don't have time to do that one as much. But I do think going to the major conferences may be a slightly less time-intensive, but it also may be quite helpful in doing that. So that's one suggestion, at least on the graduate school level.

Sallach: I was just going to say I wouldn't target a department type, economics, sociology, or that kind of thing, because the cross-disciplinary, transdisciplinary nature of computational social science is one of its great strengths, regardless of what subdiscipline or discipline you're in, GIS, data mining, data warehousing, agent simulation, and so forth. So I think that the goal is to build cross-disciplinary programs that to the existing social sciences look like a supportive methodology program; to the computer sciences it looks like a specialization in social algorithms of one type or another, and that within its own confines it begins to define a new set of goals that are neither totally defined by computer science or social science.

Cioffi-Revilla: You know, one of the fun things in high school was the lab, the physics, biology, or chemistry lab, with all the wizardry you could do there. We don't have labs in social science, and it's one of the things that makes social science studies boring in high school. I mean really, really boring.

I think an obvious application from a pedagogical point of view of computational social science and agent-based models is to provide a laboratory environment, to carry out social experiments *in machina*. And there also happens to be a fair amount of funding for science reaching down to the high school level, at least, that has been really under-utilized by the social sciences. So this is another tool that could be used, and I think to great advantage.

Saunders-Newton: On the undergraduate level, I've actually started a couple of tracks with the National Society of Black Physicists, as another one of those reformed physicists in the room, on sociophysics. Now the physics review letters are actually printing large chunks of mathematical physics around this, this is actually kind of an interesting time to talk to students about alternatives other than being in weapons labs. No knock on Los Alamos, mind you.

Bryson: May I have a quick follow-up question, about residence? … After I already said that, I realized that the answer to Ed's question is minus 10, right? that the residents have already been doing this for a decade. So has anything come out of that? Are those kids now super-modelers or something? Are they in university?

Cioffi-Revilla: I would extend this also to the systems dynamics community, by the way, which is alive and well. In July, they had their annual meeting. If you put in systems dynamics in Google and look at the annual meeting, there are thousands of people attending this world conference in New York. John Sternman just finished publishing a huge major textbook on this, on the side of business and management, etc., but nonetheless systems dynamics. I attribute a lot of that vitality to Stella and how it really facilitated the simulation of systems dynamics coming out of the dynamo tradition.

But there are people like Sternman, for example, who should be at these meetings. He came to the NAACSOS meeting; I thought that was very good. I've never seen Mitchell Resnick show up here. Perhaps a personal invitation from the NAACSOS president would be in order, or the vice president, in absence of the president. Things like that, I think, would be very valuable, because Resnick's experience in teaching pre-college modeling is enormous.

Sallach: Okay, one comment. Throw out minus 10; I'll throw out three.

If you look at all of the models, or many of the models that we present and talk about, everything is displayed in two dimensions. Our world, unfortunately, is not in two dimensions.

Lars-Erik has mentioned topography, and it makes a big difference in how things work. I think if you look at computer gaming, for example, real-time gaming systems, online gaming systems, three-dimension real-time computation is the thing of the now and the future. And three dimensions becomes important. I think computer gaming is the key to getting high school kids involved; there's no high school kid who is not into that.

One of the models we've been playing around with, for example, is we have these artificial Anasazis who run around in three-dimensional environments looking for things. And whenever high school kids come in, they just love to watch this guy, you know, running around. And it's just one guy. But they can do it in real time and they can sort of observe and they can participate. And participation is the key. You need to have a lab where kids can participate and learn, and three dimensions, visualizing beyond two, is really a necessity, which means our models — *games* they are to me, I guess — need to be real-time, they need to be fast, they need to be visually striking, and they need to have that third realistic dimension.

Moderator: Any other comments, or does anyone want to respond to that?

Saunders-Newton: As a quick addition to that, one of the areas that we're exploring across all the model use areas, is basically visualization of social model output. How do you actually deal with this? I mean, some of this relates to the ensemble of model issues raised by Steve, because basically it's a lot of information. How do you deal with this? And at this point, one of the challenges is that a lot of the tools that are actually available for it are actually spin-offs from past activities, like Starlight from Pacific Northwest Labs, which really wasn't designed for this type of work. It had a very different intent in terms of its creation.

Some of the work comes out of scientific visualization, where you're trying to look at spectral analysis and the like. The question is trying to craft those and think about how is it the average user, what resonates with them, and, importantly, how that differs across type of user. So as it turns out, the type of information that an analyst would desire is actually very, very different than what a one-star general wants.

We quickly discovered that. We call it the PowerPoint chart rule. When we used to go on the road, for any one project we would carry three PowerPoint briefings. One was basically a 50- or 60-page PowerPoint briefing for the techno dweebs. The second one was for the executive officer, or the general, which was probably about 10 to 15. And then for the general, it was just three. Title page, a page with a pretty picture on it that says something funny, and then the third one was basically, "What did you want?" I mean, that was it. So basically you get a 30-minute meeting, you get five minutes to do the three charts, and he chats for the rest of it. But this is the way it works out.

But this whole notion about what information you present to what type of user is actually incredibly important here. And this whole notion of 3-D is very important, because it's quite immersive.

Moderator: Other questions, or comments people have?

I'd add in one more quick comment or question, especially with returning to Joanna's minus 10. And I think that it's definitely true it has been out there, has been used. But I think it also depends on how deep the use has been, in the sense that if you look at any one program, was it a week-long activity? Was it a day? Was it an entire term, perhaps a series of years? And that makes a big difference. Also, I'd say availability's an issue. Can you come to a lab after school or something like that, or at other times?

Bryson: Most of that stuff, as I remember it, is relatively simple. It's not the model. You're not getting the charts and everything. But it was absolutely about trying to get kids to recognize things. I think it was all run in schools, and they still have the stuff ongoing.

In fact, they're particularly interested in disadvantaged schools in Boston and things like this. And I believe they run longitudinal classes at school. And there are some afterschool activities, a computer clubhouse thing. I think that's a separate project, but I assume it's the same software.

Unidentified Speaker: So I think that there was a wider range of use. And, of course, it depends on the number of states that are involved in the range.

Coiffi-Revilla: There's a precedent, a precursor of this that comes to mind, too, in the area — not in the computational area, but in the area of mathematical social science.

Years ago, the NSF funded a New England consortium called COMAP — I don't know how many of you remember that. It was a consortium for mathematics and applications. And they also had a journal called *UMAP*, *Undergraduate Mathematics and Its Applications.* I think COMAP and UMAP are still alive and well. They started out in Newton, Massachusetts. So these were modules that were used in high school programs. And they were field tested and very sophisticated and very effective in the dissemination of mathematical social science. There was one I remember on differential equations applied to arms races.

Bill Griffin: They also had a national contest that they would publish the results in the little books, the little subsequent publications. I subscribed to it, and I would, because it was simple enough, take it into a classroom of undergraduates and say, "Look at this." They did a very nice job.

Cioffi-Revilla: You know, not to spend all the time on pre-university teaching issues, but getting back to the visualization points that Bob was mentioning, on a purely research level this is something we need to think about, not just in terms of software, but also in terms, as the resources for this discipline become more substantial and as computational power increases and the size of the modeling that is feasible becomes more and more advanced, we need to give some serious thought to the actual scientific environments, the physical environments in which this research will be conducted 5, 15, and 20 years down the road.

We are not experiencing this kind of stuff in the social sciences. The only colleagues are perhaps the new breed of experimental economists that are designing laboratory facilities for carrying out this kind of work, but not really in social simulations.

I have a pressing need on this, and if anybody has any ideas along those lines, please share them with me. Our center will be moving to a brand new science computational facility in two years. I'm working with the architects and the engineers of that environment to make sure that we have a really efficient facility where you have not just the usual accoutrements of desks and tables and so on, but something which is really conducive to collaborative interdisciplinary computational research. And this requires a great deal of learning, certainly on my part, about things about data displays and hardware and control room engineering and things like this, that it's not too soon to start thinking along these lines. And there are many other people here that have a lot more experience about that than I do.

Sean has begun to give me some dos and don'ts about that sort of thing, and it's been very helpful.

James MacGill: I've been out of the agent side of things for about five years, so it's good to be coming back into it again. But in the meantime, what I've been doing, particularly at Penn State since I moved there in January, is looking at developing visualization toolkits and interaction environments and collaborative environments for dealing with real social data. So we're visualizing the census data, we're visualizing remote imaging. We're probing visualization environments to do those kinds of things, and the real world is a hell of a lot more complex than most of the models we're building. And yet, because we're simulating social models in most of those cases, the same toolkits that we're developing to do that can be mapped onto those tools. And, you know, there's a wealth of research in that direction.

Among the things we're doing, 3-D's been mentioned. We have environments for speech-gesture interaction with data sets. We have graphing and charting techniques for dealing with these kinds of data sets, and the chance to talk with the Repast team and look at how we can hook those things together means that we can take what we've done in terms of probing the real world and use the same things to probe the simulated worlds that we're doing.

Unidentified Speaker: That's great news. That is very encouraging.

Lars-Erik Cederman: I wanted to turn the attention away from the pedagogical dimension to the research puzzles. And since you all mentioned targets and challenges, it will be interesting to see whether people can identify these puzzles, because I think it may be very helpful for the whole field if we set up certain targets, if there were, as I say, disputes and debates about results and findings, organized around very clear topics, because in the past I would say, having surveyed the field on computational modeling, the most exciting work has probably been done in areas where you'd had a pretty clear problem definition.

If we take, for instance, Axelrod's early work on cooperation theory, that's something that's created a phenomenal cottage industry of studies going well beyond computational modeling. After all, the most interesting puzzles from the perspective of, as I say, diffusing computational modeling may actually be those that have an anchoring outside computational modeling. If you can show that with computational tools you can see things that were not obvious before, that, I think, is a much more powerful way of selling our tools here than almost any other, as I say, more supply-driven type of measures.

So I think actually it's possible to identify other candidates here. Certainly, within the network literature computational modeling has made a big difference, and people like Barabasi and Duncan Watts have helped, as I say, promote this kind of thinking, with analytical and/or computational tools. But there must be other areas.

Another type of format for inquiry would be, who is going to be the first one to create a model that exhibits the emergence of X, whatever X may be? For instance, the first actor, truly emergent actor, including now its boundaries, its rule sets, self-consciousness or whatever. It can be almost as ambitious as you want. This will be, as I say, the social science answer to A-life. But I don't think we have been, how to say, bold enough in setting those challenges. And although I find that the A-life literature can sometimes become a bit flaky and speculative, still there is something laudable about setting up completely utopian goals. And I think we have more work to do along those lines.

Cioffi-Revilla: I think of this in terms of two sources for that sort of challenge. One source of challenge is the areas where classical social science using the earlier two ways of doing science has stumbled. And those are opportunities.

For example, in actor interaction problems there are now known analytical solutions in closed form. This problem came up in international relations with end-country arms races that were so interactive that had no … so one thing is, those areas, those problems in social science where classical statistical and mathematical approaches have failed to make a breakthrough because they're simply incapable, there's no tractability. That's one generic source.

The other generic source, in terms of parsing all of these sources, is in the new puzzles and new questions that we can now address through computational methods. And here I'm thinking in terms of, for example, the same thing that happens in biology and in astronomy in pushing completely new frontiers that would have been unimaginable. These are not just problems in which astronomers were failing to produce new results, but they were completely new areas of astronomy that were inaccessible; for example, including rational mechanics of worlds that violate known laws of physics on purpose in order to understand those that actually do work the way they are.

Saunders-Newton: Lars-Erik, I have two thematic examples which would really get me excited if I actually saw people who were doing this.

The first one would be to revisit what used to be called the societal instability literature or modeling. Our new euphemism is social fragility, right? But actually revisiting these tool sets as opposed to, like, the dustbowl empiricism models from the '50s and '60s, but to actually revisit this whole notion of why are certain countries more frail, either as the result of the governing structure, a lack of infrastructure, whatever it is, but actually have an activity. And then the metrics that come out of that would allow us to do an interesting type of risk assessment, where we can actually make a strong argument for why it is better to act early than to act late, in terms of investing in countries so that they don't fall into civil war or that ilk. So that would actually be one.

In the second one, which falls into the more utopic category, is basically how do you grow utopic societies? Let's say that the desire is actually that as a globe that we would like to spend less than 1% of our global GNP on weapons. What would a structure like that look like? And it would be interesting to actually see whether there some type of modeling structure that we can address. Those would be interesting grand challenges, I think, to consider.

Mark Diggory: As a software developer, I wanted to possibly just bring out some basic concepts from my experience that relate to some of the issues related with these various toolkits. And something that I've learned here today from interactions with James and with geo tools and Sean, dealing with the Mason platform, is that we have to deal much better with interoperability, and primarily in the areas of not reinventing the wheel for every single framework. And with that in mind, I want to suggest that as separate projects working on separate platforms, we think a lot more about interacting with each other and the interfaces and capabilities that we want to provide to the community without the persona of "our project's here and it's the be-all/end-all and we're going to take over the world with it."

Bill Lawless: I think Lars-Erik raised a really good point. Some of the problems I see in that regard is that we've got — like in the case of Steve, he's got proprietary problems that are being solved, and it's very difficult for us to deal with the nuts and bolts of how that can actually be applied in general. Or in the case of Ed's model, which I think is really good, too, there's classified information there, so you can't really run out and do much with it.

But, nonetheless, I think you hit on a critical point that unless we can find a problem like launching the first aircraft at Kitty Hawk, and we can compete to solve that problem, and actually *solve* the problem, then I don't think we'll get that far, other than winding up with a wonderful mechanism for explanation, with great explanatory power, but very little demonstrative or predictive power.

And I think this is a problem that *will* be solved, so I'm not trying to cast a bleak outlook. I think it's a problem that will be solved; otherwise, we'll never have these social systems of agents going hot out in a battlefield or these systems of agents landing on Mars or other planets or going to other systems and actually taking on complicated, ill-defined problems and solving them by themselves.

So it will be a bleak future to me if we can't solve these, and I think you're right on. These are things that we have to solve. Mind you, we're not the only discipline with these kind of problems. In social psychology, my field, there are very few replicable group phenomena, social phenomena, that have stood the test of time. One of them is Bibb Latinay's diffusion of responsibility problem that is still today one of the very few accomplishments in social psychology. This has been around for over 30 years. But nonetheless, I think that we've got to find a problem that we can solve, and then other people have got to be able to go out and replicate that solution, and then somebody's got to win a Nobel Prize for that. And once that's done, hey, we're not going to have any problems attracting people into this field.

Moderator: Well, do we have other questions?

One thing I'd like to add to this, and it also goes to Desmond's question about what do you do with all the computational power that's available? Well, one thing I'd say is we do have existing models. They don't answer all of the grand challenge questions but things like Sugarscape, which really hasn't been fully analyzed. Now, you can do humongous parameter sweeps and there might be interesting things that are found. And if nothing interesting is found, that's a finding, too. So that would be one use that is available now. But I totally agree with you, in the long run we also have to find these other challenges and make them fairly explicit as to what we're trying to solve.

Other comments, questions before us? Chick?

Charles M. Macal: I'd like to thank very much our closing panel. And I'd like to thank the audience for the interesting and very thoughtful discussion that followed. Thank you very much.

Moderator: I have a few closing comments, just for the most part administrative and related items to cover.

On behalf of the Agent 2003 Executive Program Committee, I'd like to thank all the invited speakers that talked during the conference. The presenters and the authors of the papers, the discussants, the session chairs, and, again, the audience for your attendance. We appreciate the fact that you're here, and your comments that you've given throughout the conference I think were very substantive and allow us to come to some synthesis potentially of, and consensus even, perhaps, on some points that we can all agree on and hopefully be goals for us to move forward in the field. So, of course, one of the goals would be to not be here again next year talking about the same things in terms of what needs to be done in the future because we haven't made any progress in the past year, for example.

I also just have some other thanks to go around. I'd like to thank in particular Kathy Ruffatto for excellent administrative support in logistics.

I'd like to thank the Abstract Review and Selection Subcommittee of the Executive Committee. You know who you are.

I'd like to thank the AV Technical Support Committee, including especially Kay Sitarz, who is with us here. I'd like to thank the Conference Booklet Design and Preparation Committee, the Webpage Design Committee, and the Webpage Development and Update Committee, which was a separate committee.

I'd like to thank the Proceedings Editing and Publication Committee, especially Margaret Clemmons, who many of you have corresponded with, perhaps only by e-mail. And a note on that: we will be publishing proceedings as we have done in the past years. So if you have not provided your final copies of your papers to us, we will relentlessly pursue you. That is guaranteed. And we have 100% success rate, based on previous years' conferences, so you might as well just relent immediately and give up your paper now.

I'd like to thank the Speaker Interaction Subcommittee, again, of the Executive Committee. And that was led by David Sallach, of course, so we should give David a hand.

I'd like to thank Tom Wolsko, who is the Director of the DIS Division of Argonne for his executive support that made this whole effort possible. I'd like to thank especially Mike North and Tom Howe for teaching the Repast training course, which occurred earlier this week.

Again, I'd like to thank Mike North for putting together the NAACSOS-sponsored sessions on Toolkits and Methods Day, which occurred on Thursday afternoon, as well as the toolkit developers' meeting that occurred on Thursday morning. And there was huge amounts of after-midnight work that was going on all week, as you can imagine. And so, Mike, our hats are off to you for all your work and effort. Agent 2003 is officially adjourned.

List of Attendees

LIST OF ATTENDEES

Kostas Alexandridis
Michigan State University, Human-Environment Modeling Analysis

Li An
University of Michigan

Steven Bankes
Evolving Logic/RAND

Nicola Baxter
BT Exact

Sid Bhattacharyya
University of Illinois at Chicago

Inanc Birol
Illinois Institute of Technology

Jon Brewster
Wayne State University

Marc Broklawski
Naval Surface Warfare Center, Dahlgren Division

Marina Brukhovetsky
University of Wisconsin

Joanna Bryson
University of Bath

Roger Burkhart
Deere Company

Ana Carrie
Trinity College, University of Dublin

Lars-Erik Cederman
Swiss Federal Institute of Technology

Shao-Ping Chen
Los Alamos National Laboratory

Shu-Heng Chen
AI-ECON Research Center

Bin-Tzong Chie
AI-ECON Research Center

Ali Cinar
Illinois Institute of Technology

Claudio Cioffi-Revilla
George Mason University

Richard Cirillo
Argonne National Laboratory

David Collings
BT Exact

Noshir Contractor
University of Illinois at Urbana-Champaign

Guenter Conzelmann
Argonne National Laboratory

Ann Cutler

Jean Czerlinski
DRW

Peter Deadman
University of Waterloo

Kevin Desouza
University of Illinois at Chicago

Gennaro Di Tosto
CESSBA

Mark Diggory
Harvard Business School

Elenna Dugundji
University of Amsterdam

Roberta Evaristo
University of Illinois at Chicago

Luis Fernandez
University of Michigan

Aaron Frank
National Defense University

Terrill Frantz
Organization Simulations Asia Ltd

Dennis Friend
Adica Consulting LLC

Steven Gabel
University of Chicago

Andrew Garman
Rush University

Felimon Gayanilo
University of Miami, Rosentiel School of Marine Atmospheric Science

Sudip Ghosh
Grouper Labs, Purdue University

Laura Goadrich
University of Wisconsin-Madison

Mark Goadrich
University of Wisconsin-Madison

William Griffin
Arizona State University

Stephen Guerin
RedfishGroup

Bruce Hamilton
Adica Consulting LLC

Tom Howe
University of Chicago

Joel Jeffrey
Northern Illinois University

Xu Jin
University of Notre Dame

Ricardo Kanitz
Harvard Business School

Vladimir Koritarov
Argonne National Laboratory

Daniel Kunkle
RedfishGroup

William Lawless
Paine College

Zhen Lei
Michigan State University

Jorge Louca
ISCTE

Sean Luke
George Mason University

Charles Macal
Argonne National Laboratory

James MacGill
Pennsylvania State University

Edward MacKerrow
Los Alamos National Laboratory

Greg Madey
University of Notre Dame

Tom Montgomery
Ford Motor Company

Jacobo Myerston
University of Chicago

Michael North
Argonne National Laboratory

David O'Sullivan
Pennsylvania State University

William Parker
ThinkVine LLC

Debra Patek
ThinkVine LLC

Sinem Perk
Illinois Institute of Technology

Bryan Pijanowski
Purdue University

Snehal Pithadia
Department of Forestry

Panomsak Promburom
Multiple Cropping Center

Damon Ragusa
ThinkVine LLC

Robert Reynolds
Wayne State University

Seung-Kyu Rhee
Argonne National Laboratory; Korea Advanced Institute of Science and Technology

Derek Robinson
University of Waterloo

Fabio Rojas
Indiana University

Meredith Rolfe
University of Chicago

David Sallach
University of Chicago; Argonne National Laboratory

Kedar Samant
University of Illinois at Chicago

Desmond Saunders-Newton
DARPA

Keith Sawyer
Washington University

Paul Schermerhorn
University of Notre Dame

Craig Stephan
Ford Motor Company

Wesley Stevens
TeraTrends Inc.

Xianfeng Su
University of New England, Australia

John Sullivan
Ford Motor Company

Pam Sydelko
Argonne National Laboratory

David Sylvan
Graduate Institute of International Studies, Geneva

Eric Tatara
Illinois Institute of Technology

Mehmet Tetiker
Illinois Institute of Technology

Prakash Thimmapuram
Argonne National Laboratory

Joshua Thorp

Jesse Voss
University of Wisconsin-Madison

Jason Woodard
Harvard Business School

Kazuya Yamamoto
University of Tokyo

Levent Yilmaz
Auburn University

Gao Yongqin
University of Notre Dame

Kikuo Yuta
Advanced Telecommunications Research Institute International

Moira Zellner
University of Michigan

Lei Zhen
Michigan State University