From Factors to Actors:
Computational Sociology and Agent-Based Modeling[1]

Michael W. Macy

Robert Willer

Cornell University

August, 2001

Contents:

73 citations (2.5 published pages), 1 figure (1 page), 8522 main-text words (19.5 published pages)

Total length: 23 published pages

[1] The first author expresses gratitude to the National Science Foundation (*SES-0079381*) for support during the period in which this review was written. We also thank James Kitts, Andreas Flache, and Noah Mark for helpful comments and suggestions.

From Factors to Actors:

Computational Sociology and Agent-Based Modeling

Introduction: Agent-Based Models and Self-Organizing Group Processes

Consider a flock of geese flying in tight formation. Collectively, they form the image of a giant delta-shaped bird that moves as purposively as if it were a single organism. Yet the flock has no "group mind" nor is there a "leader bird" choreographing the formation (Resnick 1994). Rather, each bird reacts to the movement of its immediate neighbors who in turn react to it. The result is the graceful dance-like movement of the flock whose hypnotic rhythm is clearly patterned yet also highly non-linear.

If we tried to model the global elegance of the flock, the task would be immensely difficult because of the extreme complexity in its movement. Yet the task turns out to be remarkably easy if instead we model the dynamics of local interaction. This was demonstrated by Craig Reynolds (1987) when he modeled the movement of a population of artificial "boids" based on three simple rules:

- Separation: Don't get too close to any object, including other boids.
- Alignment: Try to match the speed and direction of nearby boids.
- Cohesion: Head for the perceived center of mass of the boids in your immediate neighborhood.

Reynold's computational method is called "agent-based modeling." Had Reynolds chosen instead to write a "top down" program for the global behavior of the flock, he might still be working on it. By choosing instead to model the flock from the bottom up, based on agent-level interaction, he was able to produce highly realistic flight formations using very simple rules that imposed relatively small computational demands.[2] Note that Reynolds did not model the flock, nor did he model isolated birds. He modeled their *interaction*, at the relational level.

Agent-based models (hereafter ABMs) of human social interaction are based on this same theory-building strategy. Like flocks of birds, human group processes are highly complex, non-linear, path dependent, and self-organizing. We may be able to understand these dynamics much better not by trying to model them at the global level but instead as emergent properties of local interaction among adaptive agents who influence one another in response to the influence they receive.

Despite growing interest in relational modeling and computational methods, sociologists have not fully appreciated the potential for ABMs as tools for theoretical research. This review of recent developments is intended to demonstrate how this technique can provide sociologists with a more rigorous method for specifying the microfoundations of global patterns at the relational level. We begin with a brief historical sketch of the shift from factors to actors in computational sociology that shows how agent-based modeling differs fundamentally from earlier sociological uses of computer simulation. We then review recent contributions focused on two problems, the emergence of social structure and social order out of local interaction. We conclude with a set of recommendations, including a criticism of the microsociological bias in agent-based modeling and identification of a research strategy that can bring "factors" back in to agent-based computational sociology.

---

[2] Reynold's "boids" were so realistic that they provided the starting point for bat swarms in the movies *Batman Returns* and *Cliffhanger*. You can see the "boids" in action at www.discovery.com/area/life/life1.3.html.

Historical Development of Agent Based Models

Computer simulation is more tractable (but less generalizable) than mathematical modeling and more rigorous (but less nuanced) than natural language. Gilbert and Troitzsch (1999) identify three periods in the development of social simulation over the past half-century: dynamical systems, microsimulation, and adaptive agent models. In the 1960's, the first wave of innovation used computers to simulate control and feedback processes in organizations, industries, cities, and even global populations. The models typically consisted of sets of differential equations that described changes in system attributes as a holistic function of other systemic changes. Applications included the flow of raw materials in a factory, inventory control in a warehouse, urban traffic, military supply lines, demographic changes in a world system, and ecological limits to growth (Forrester 1971; Meadows 1974).

Beginning in the 1970's, computer modelers introduced the use of individuals as the units of analysis but retained the earlier emphasis on empirically based macro-level forecasting. In striking contrast to the holistic approach in models of dynamical systems, "microsimulation is a 'bottom-up' strategy for modeling the interacting behavior of decision makers (such as individuals, families and firms) within a larger system. This modeling strategy utilizes data on representative samples of decision makers, along with equations and algorithms representing behavioral processes, to simulate the evolution through time of each decision maker, and hence of the entire population of decision makers" (Caldwell 1997). However, the models do not permit individuals to directly interact or to adapt. Nor are the models designed or used for basic theoretical research; the primary orientation is toward applied research, mainly forecasting macro effects of public policies that alter individual behavior. In that sense, these models remain equation-based, like the earlier dynamical systems models.

The third wave in social simulation, agent-based modeling, coincided with the advent of the personal computer in the 1980's. Like microsimulation, these "bottom-up" models explored the micro-foundations of global patterns. The difference is that, unlike the socially isolated actors in microanalytical simulation, the agents now interact. More precisely, ABMs impose four key assumptions:

1. *Agents interact with little or no central authority or direction.* Global patterns emerge from the bottom up, determined not by a centralized authority but by local interactions among autonomous decision-makers. This process is known as "self-organization" (Kaufman 1996).
2. *Agents are interdependent.* All ABMs assume that agents influence others (directly or indirectly) in response to influence that they receive. Some models go further to assume that agents are also strategically interdependent. This means that the consequences of each agent's decisions depend in part on the choices of others.
3. *Agents follow simple rules.* Global complexity does not necessarily reflect the cognitive complexity of individuals. "Human beings," Simon contends (1998, p. 53), "viewed as behaving systems, are quite simple." We follow rules, in the form of norms, conventions, protocols, moral and social habits, and heuristics. Although the rules may be quite simple, they can produce global patterns that may not be at all obvious and are very difficult to understand. Hence, Simon continues, "the apparent complexity of our behavior is largely a reflection of the complexity of the environment." ABMs are designed to explore the minimal conditions, the simplest set of assumptions about human behavior, required for a given social phenomenon to emerge at a higher level of organization.

4. *Agents are adaptive and backward-looking.* When interdependent agents are also adaptive, their interaction can generate a "complex adaptive system" (Holland 1995, p. 10). Agents adapt by moving, imitating, replicating, or learning, but not by calculating the most efficient action (Holland 1995, p. 43). They adapt at two levels, the individual and the population. Individual learning alters the probability distribution of rules competing for attention, through processes like reinforcement, Bayesian updating, or the back-propagation of error in artificial neural networks. Population learning alters the frequency distribution of agents competing for reproduction through processes of selection, imitation, and social influence.

From Forecasts To Thought Experiments

Unlike dynamical systems and microsimulation models, whose value depends largely on predictive accuracy, adaptive agent models are "much more concerned with theoretical development and explanation than with prediction" (Gilbert 1997:2.1). They are used to perform highly abstract thought experiments that explore possible (if not always plausible) mechanisms that may underlie observed patterns. As such, these models do not necessarily "aim to provide an accurate representation of a particular empirical application. Instead, the goal of agent-based modeling is to enrich our understanding of fundamental processes that may appear in a variety of applications" (Axelrod 1997:25). When simulation is used to make predictions or for training personnel (e.g., flight simulators), the assumptions need to be highly realistic, which usually means they will also be highly complicated. "But if the goal is to deepen our understanding of some fundamental process," Axelrod continues, "then simplicity of the assumptions is important and realistic representation of all the details of a particular setting is not." On the contrary, making these models more “realistic” might add complexity that could undermine their usefulness as tools for theoretical research.

Nevertheless, many sociologists remain highly skeptical about the validity of simulation results when computational models are used for theoretical exploration rather than empirical prediction. As noted (and lamented) by Sawyer (2001), recent survey articles on sociological simulation neglect agent-based modeling and focus primarily on the earlier equation-based method of macrosimulation and social forecasting (Halpin 1999; Hanneman, Collins, and Mordt 1995; Meeker and Leik 1997).

It is ironic that sociological interest in ABMs has lagged behind that of the other social sciences, for sociology may be the discipline best equipped to develop a methodology that bridges Schumpeter’s (1909) methodological individualism and Durkheim’s rules of a non-reductionist method. Durkheim anticipated the concept of emergence: “The hardness of bronze lies neither in the copper, nor the tin, nor in the lead which have been used to form it, which are all soft or malleable bodies. The hardness arises from the mixing of them” (Durkheim 1901/1982, pp. 39-40). The principle applies as well to sociology, he continued. “(Social) facts reside in the society itself that produces them and not in its parts – namely its members.”

Here Durkheim oversteps. While the principles of emergence and self-organization imply that properties of the larger system are not properties of the components – and may not resemble nor be intended by any of the constituent actors – these principles also incorporate an essential insight of methodological individualism, the idea that societal patterns emerge from motivated choices and not from “social facts” external to individuals. Global properties are *sui generic* but they also emerge from the bottom up, through local interactions. Without a model of the microfoundations of emergent properties, path-dependent self-organizing processes (such as informal social control) are likely to be mistaken for institutions that are globally coordinated (such as bureaucratic controls in formal organizations). In short, ABMs defy classification as either micro or macro but instead provide a theoretical bridge between them (Saam 1999).

Clearly, not all problems can be usefully viewed from the bottom up. Agent based models are most appropriate for studying processes that lack central coordination, including the emergence of institutions that, once established, impose order from the top down. The models focus on how simple and predictable local interactions generate familiar but highly intricate and enigmatic global patterns, such as the diffusion of information, emergence of norms, coordination of conventions, or participation in collective action. Emergent social patterns can also appear unexpectedly and then just as dramatically transform or disappear, as happens in revolutions, market crashes, fads, and feeding frenzies. ABMs provide theoretical leverage where the global patterns of interest are more than the aggregation of individual attributes, but at the same time, the emergent pattern cannot be understood without a "bottom up" dynamical model.

In surveying recent applications, we found that most were congregated around two problems, the self-organization of social structure and the emergence of social order. The two problems are highly complementary. In one case, the clustering of social ties is the explanandum and in the other it is the explanans.

1. *Emergent structure*. In these models, agents change location or behavior in response to social influences or selection pressures. Agents may start out undifferentiated and then change location or behavior so as to avoid becoming different or isolated (or in some cases, overcrowded). Rather than producing homogeneity, however, these conformist decisions aggregate to produce global patterns of cultural differentiation, stratification, and homophilic clustering in local networks. Other studies reverse the process, starting

with a heterogenous population and ending in convergence: the coordination, diffusion, and sudden collapse of norms, conventions, innovations, and technological standards.

2. *Emergent social order*. These studies show how egoistic adaptation can lead to successful collective action without either altruism or global (top down) imposition of control. A key finding across numerous studies is that the viability of trust, cooperation, and collective action depends decisively on the embeddedness of interaction.

Despite a common focus on two central problems, there has been little effort to provide a meta-analysis of how results differ depending on the model designs. To that end, we have grouped studies by substantive application in order to highlight methodological differences that may explain conflicting results. These differences emerge through a series of interrogations about model design:

1. Is interaction global or local, that is, is the population fully connected or is interaction constrained by the structure of social ties?
2. If interaction is local, is the network spatial or social?
3. Is tie formation elective (through movement, exit, or assortative mating) or is interaction forced?
4. Is adaptation based on *learning* (which modifes the probability distribution of strategies in each agent's repertoire) or *evolution* (which modifies the frequency distribution of strategies across the population of agents)?
5. If evolution, does reproduction involve competition for survival or social influence?
6. If influence, is this limited to external states of the agent (e.g., behavior) or do agents copy other agents' internal programming, even though this cannot be directly observed?

7. Is influence based on attainment (success, fitness, payoffs, status) or familiarity (proximity, frequency)?
8. Is the model used as an *experiment* (parameters are manipulated to test for predicted differences) or a *demonstration* (parameters are manipulated only to test for robustness)?
9. If used experimentally, are the manipulations mainly of agent-level parameters (to test a micro theory) or population-level parameters (to test a macro theory)?

Figure 1 classifies the papers we review in a typology based on answers to these nine questions. The articles we included are not intended to be exhaustive. The field of social simulation is now too large to survey in a single article. We have therefore narrowed the focus to agent-based models of emergent structure (differentiation and diffusion) and emergent order (cooperation and collective action), written by sociologists or published in sociological journals[3] in the past five years.[4]

[Figure 1 about here]

Although sociology has lagged behind other social sciences in appreciating this new methodology, a distinctive sociological contribution is evident in the papers we review. They show how ABMs can be used to perform "virtual experiments" that test macrosociological theories by manipulating structural factors like network topology, social stratification, or spatial mobility. Used in this way, ABMs can close the circle by moving back again, from actors to factors.

---

[3] Reviewed articles were published in *American Sociological Review* (6), *Computational and Mathematical Organization Theory* (5), *Journal of Artificial Societies and Social Simulation* (5), *American Journal of Sociology* (3), *Rationality and Society* (3), and *Journal of Mathematical Sociology* (2). In addition, referenced articles on simulation were published in *American Sociological Review* (3) and *Computational and Mathematical Organization Theory* (2).

Emergent Structure: Models of Convergence and Differentiation

In models of structural differentiation, interest centers on the global self-organization of the population into local networks or clusters based on simple rules of local interaction. Applications include residential segregation, density-dependent organizational survival, group formation, and cultural differentiation.

These models often study clustering within spatial networks, using “cellular automata” (CA), a technique first proposed by Stanislaw Ulam (Coveney and Highfield 1995:94-96). Hegselmann and Flache (1998) provide a lucid introduction to and history of CA in the social sciences. The agents usually live on a checkerboard (either flat, or a donut-like torus), and the state of each agent depends on the states of its neighbors. Simple rules of local influence or spatial movement sometimes generate surprising results and lead to unexpected insights. They illustrate a key advantage of the CA approach: two-dimensional visual representation of diffusion and clustering across a spatial network.

Schelling’s (1971) model of neighborhood segregation is one of the earliest and best known ABMs based on movement in a spatial network. Red and green agents are randomly distributed on a lattice and move to empty locations if the number of in-group neighbors falls below a certain threshold. The model shows how extreme segregation tends to arise even in a population that prefers diversity, as agents relocate to avoid being in the minority.

Another classic model is Conway’s “Game of Life” in which the survival of each agent depends on the density of its neighborhood.[5] Although Conway was not a sociologist, his design

---

[4] For an *Annual Review* of related papers published more than five years ago, see Bainbridge et al. 1994.

[5] Like Schelling’s, Conway’s model was originally created without a computer, using a game board. The model is now implemented in Java and interested readers can experiment on-line at www.math.com/students/wonders/life/life.html.

has immediate application to problems in organizational ecology, in which the agents are supra-individual. For example, Lomi and Larsen (1998) study the interaction between network structure and the lagged effects of population density on organizational survival. They use a cellular network in which the survival of each cell depends on the number of occupied cells in its "Moore neighborhood" (the eight adjacent cells). The model is very simple; agents do not make decisions or interact strategically or differ in success, they merely live, replicate, and die, based on local density. Lomi and Larsen then explore the implications for organizational survival of alternative hypotheses about the effects of density delay based on simple rules that regulate the appearance, survival and demise of individual organizations. Using event history analysis, they identify structural features that can generate organizational life histories that are qualitatively consistent with those observed in empirical organizational populations.

Another prominent contemporary example is Epstein and Axtell's (1996) "Sugarscape" in which a spatially distributed population of simple rule-based agents develops a culture, an economy, and a class structure. Agents move around on a grid and exchange with others to gain access to valued resources on which their survival and reproduction depend.

One criticism of spatial networks is that they preclude both structural equivalence (no two nodes can have identical sets of interactants) and relational heterogeneity (every node has an isomorphic relational pattern). Flache and Hegselmann (2001) relaxed the latter constraint by using "irregular grids" that allow the number and strength of social ties to vary randomly over the population. However, the models still permits the population to self-organize into clusters through spatial movement.

Another criticism of ecological models[6] is the assumption that adaptation occurs through the death (by starvation in Sugarscape) and reproduction of agents, based on relative fitness. This assumption is appropriate if the agents are organizations competing for resources or members. If the agents are individuals in a modern welfare state, however, a more broadly applicable assumption is that replication occurs through "imitation of the fittest." Agents are not replaced by the top performers, they simply copy their observed behavior.

Social Influence and the Paradox of Mimetic Divergence

Several recent studies depart from the ecological metaphor of death and reproduction and instead assume that adaptation operates via social influence. This in turn relaxes the assumption that selection pressures are performance driven. Although some influence models continue to posit selection of role models based on relative success, others assume that influence is density dependent, based on familiarity, popularity, or spatial proximity. For example, Latané's (1996) "social impact model" uses a rule to mimic one's neighbors in a two-dimensional lattice. From a random start, a population of mimics might be expected to converge inexorably on a single profile, leading to the conclusion that cultural diversity is imposed by factors that counteract the effects of conformist tendencies. However, the surprising result was that "the system achieved stable diversity. The minority was able to survive, contrary to the belief that social influence inexorably leads to uniformity" (Latané 1996, p. 294).

Following an earlier study by Carley (1991), Axelrod (1997; see also Axtell et al. 1996) makes the paradox of mimetic divergence even more compelling. Carley's and Axelrod's models

---

[6] Ecological models are those in which the spatial or frequency distribution of agents depends on rules that govern survival and reproduction. These models are often characterized as "evolutionary," but strictly speaking, the latter requires the possibility

couple local influence (the tendency for people who interact frequently to become more similar over time) and homophily (the tendency to interact more frequently with similar agents). This closes the loop: the more agents interact, the more similar they become, and the more similar they become, the more likely they are to interact. More precisely, neighboring agents on a two-dimensional lattice interact with a likelihood determined by the similarity of their cultural traits (given by a simple, randomly-assigned string of numbers). Interaction, in turn, reduces remaining differences. Axelrod expected this self-reinforcing dynamic would lead inexorably to cultural convergence and homogeneity. Again the result was surprising. He found that "local convergence can lead to global polarization" and that unique subcultures can survive in the face of a seemingly relentless march towards cultural conformity. Stable minority subcultures persist because of the protection of structural holes created by cultural differences that preclude interaction, thereby insulating agents from homogenizing tendencies.

Axelrod's model also reveals a surprising effect of population size. Intuitively, one might expect larger numbers of stable subcultures to emerge in larger populations. However, Axelrod found a nonlinear effect, in which the number of minority cultures first increases with population size but then decreases. This counter-intuitive result illustrates the principle of "gambler's ruin." Large populations allow for larger cultural movements that can survive random fluctuations in membership better than smaller competitors. As the big get bigger, the number of minority subcultures diminishes.

Axelrod begins with a heterogeneous population and shows that heterogeneity persists. But how does the initial heterogeneity arise? Axelrod also assumes spatial networks that restrict

---

for entirely new types of agents to appear that were not present at the outset.

interaction to nearby neighbors. Will differentiation persist if the spatial restriction is removed and interaction is based only on similarity?

Mark (1998) addresses these questions in a paper that explains social differentiation "from first principles," starting from homogeneity and without spatial constraints on interaction. Agents initially interact randomly, with anyone in the population, and then later with a probability determined by cultural similarity. Mark finds that a self-reinforcing dynamic based on homophily is sufficient to create an emergent network with local patterns of interaction among distinctive subcultures. Contrary to Axelrod, Mark also finds that population size decreases cultural homogeneity, due to the absence of spatial restrictions on interaction in the model.

One limitation in most social influence models is the assumption that influence is only positive. However, social relations can also have negative valence, such that the state of an agent tends toward maximal distinctiveness rather than similarity. Contrary to theories of homophily, dissimilarity does not always weaken the social tie; rather, it may sometimes *strengthen* the negative relation (or enmity). Structural differentiation based on positive and negative influence has been studied using "attractor neural networks," a cognitive modeling technique developed by Hopfield (1982) and applied to social influence by Nowak and Vallacher (1998; see also Kitts, Macy, and Flache 1999).

Artificial neural networks are a simple type of self-programmable learning device based on parallel distributed processing (Rummelhart and McClelland 1988) and modeled after the nerve systems of living organisms. In elementary form, the device consists of a web of neuron-like units (or "neurodes") that fire when triggered by impulses of sufficient strength, and in turn stimulate or inhibit other units when fired. The effect of an impulse (as stimulus or inhibitor) depends on the sign and strength of the synaptic connection between the two neurodes. The network learns by modifying these path coefficients. In “feed-forward” networks, learning is based on environmental feedback that propagates backward through the network. In “attractor” networks, the path coefficients are updated based on the similarity between the states of connected nodes.

While feed-forward networks can be used to model agent cognition, attractor networks provide an intriguing alternative to standard network models of social interaction. Attractor networks not only allow for both positive and negative influence, they also incorporate basic principles of structural balance (Cartwright and Harary 1956) and network transitivity. These models go “beyond the usual depiction of the similarity-attraction relationship” by modeling “dynamics of any given dyad in the context of other dyads in a larger social structure” (Nowak and Vallacher 1998:21).

Feed-forward devices have been used to model cognitive social differentiation, based on self-affirming stereotypes. Vakas-Duong and Reilley (1995; see also Bainbridge 1995) study the emergence of irrational racial hiring preferences that are less profitable than purely meritocratic selection. In their model, employers learn to make intuitive hiring decisions based on what connotations come to be associated with the traits exhibited by job applicants, while applicants associate traits with relative success. The simulation results showed how irrational beliefs can easily become self-sustaining following an early accident of association that sows the seeds of racial preference in an employer’s mind. This in turn makes it difficult for talented members of the same race to gain employment and this diminishes that race’s access to emerging status symbols.

Other models of cognitive social differentiation focus on the self-reinforcing dynamics created by stereotypical beliefs that change the behaviors on which the beliefs are based. Orbell et al. (1996) model self-organizing stereotypes in a population of 1000 adaptive agents playing Prisoner's Dilemma games with an option to exit. Prisoner's Dilemma is a two-person game in which the best move is to "defect" (e.g., cheat), no matter what the partner is expected to do, but when each defects, the outcome is deficient for both. However, the authors' interest is focused not on the problem of cooperation but the formation of groups and group stereotypes. Agents are assigned a "tag" that indicates their membership in one of two groups. That is the only difference among the agents. Agents update their tag-specific expectations of cooperative behavior based on the outcomes of interactions. Agents also become more likely to cooperate with members of groups they expect to cooperate. In one experiment, each agent cooperates sixty percent of the time with members of both groups**.** The two groups are equal in size but one group enjoys a slightly better (but undeserved) initial reputation. They find that agents from both groups converge on an increasingly strong preference for interaction with members of the initially preferred group. In a second experiment, the two groups have identical initial reputations but differ slightly in size. Again, agents from both groups converge on a preference for interaction with members of the larger group. In both experiments, the mechanism is straightforward. If you interact with one group more than the other (either because you expect cooperation from this group or there are more of them in the population), you update your expectations for this group more than for the other. Since both groups are relatively cooperative, the updating always causes expectations of cooperation to increase. Therefore, the more you interact with one group, the more you expect its members to cooperate, and in turn, the more you cooperate with them. The more others cooperate with them, the more they cooperate as well, thereby affirming the expectation on which the preferential interaction is based.

Diffusion of Innovation

The models considered so far all explore emergent networks based on structural differentiation. However, social influence models can also be used to study self-reinforcing dynamics that lead to diffusion of innovations, coordination of conventions, emergent norms, and cultural convergence. Diffusion models start with some distribution of practices and a rule by which agents decide whether to abandon current practice in favor of one used by another agent. In contrast to models of differentiation that typically assume prior interaction as the basis of influence, diffusion models often assume that influence is based on popularity (either directly or indirectly) or success or some combination of the two.

Rosenkopf and Abrahamson (1999) studied diffusion where influence derives from popularity, without regard to prior interaction history or relative success. This creates a "positive feedback loop where adoptions by some actors increase the pressure to adopt for other actors" Rosenkopf and Abrahamson (1999: 361). However, influence was weighted by reputations (which were exogenous to the model) and combined with information about the unprofitability of innovations. The network was fully connected, that is, each agent had access to the decisions and reputations of all other agents in the population. They found that "bandwagons occur even when potential adopters receive information about others' unprofitable experiences with the innovation" (Rosenkopf and Abrahamson (1999: 361). Their model shows how agents can converge on inefficient practices but not how conformity collapses.

Bullnheimer and Zeller (1998) studied the diffusion of innovation among firms based on the assumption that firms evaluate the relative performance of alternative technologies based on their own experience with a technology as well as the performance of others. They assume that firms have knowledge of the adoptions and performance of all other firms. They found that firms who combine imitation with their own experience outperform “both pure imitators and nonimitators in production efficiency as well as in profits.” This model also explains stable homogeneity but not how conformity might then collapse.

Strang and Macy (2001) bridge these studies by showing how a decision rule similar to those Bullnheimer and Zeller found to be highly efficient can nevertheless trap firms in a fad-like bandwagon of adoption and abandonment of innovations that are worthless or nearly so. They assume firms evaluate current practice based on their balance sheets, and if dissatisfied, turn to “best practices” for new ideas. In a series of computational experiments, they manipulate the intrinsic value of innovations, the stratification of the market, and the skepticism of managers to see how these affect the fad-like pattern. Results show that fads are most likely in stratified markets where innovations have a modest effect on performance and managers are not so skeptical that they cannot see the performance differences.

Emergent Order: Models of Collective Action, Trust, and Cooperation

In models of structural differentiation, agents influence others in response to the influence they receive, leading to spatial or social clustering, such as Reynold’s flocks of “boids.” Interest centers on the self-organization of dynamic structural configurations, and not on their consequences. Models of emergent order, in contrast, focus attention on the ways in which network structures affect the viability of prosocial behavior. Four network properties have been shown to promote cooperation and participation in collective action:

- *Relational stability*: on-going relationships lengthen the “shadow of the future” (Axelrod 1984).

- *Network degree*: the coordination complexity of local norms increases with the number of social ties (Macy and Skvoretz 1998).
- *Homophily*: agents tend to interact with partners who use similar strategies (Cohen, Riolo, and Axelrod 2001).
- *Transitivity*: an agent's partners tend to interact with each other. This in turn affects:
    - Diffusion of reputations (Takahashi 2000).
    - Bandwagons caused by threshold effects (Chwe 1999).
    - Monitoring and enforcement of conformity to prosocial norms (Kim and Bearman 1997).

Relational Stability

The classic study of emergent order is Axelrod's (1984) *Evolution of Cooperation.* Although defection is the dominant strategy in a single play of Prisoner's Dilemma, that is not true when the game is played repeatedly in an on-going relationship. However, this does not guarantee cooperation. In fact, there is no dominant strategy in repeated play, and game theory cannot predict which of countless possible conditionally cooperative strategies will emerge as an equilibrium. To find out, Axelrod organized a computer tournament in which agents played a round robin iterated Prisoner's Dilemma. He invited leading game theorists to submit strategies, and each submission was assigned to one of the agents. The winner was the simplest contestant, Anatol Rapaport's "Tit for Tat," a strategy that never defects unless provoked and quickly forgives. However, the success of Tit for Tat depends on the prospect of an on-going interaction, what Axelrod calls "the shadow of the future."

Tit for Tat always cooperates unless provoked, and then always retaliates. Variations on Tit for Tat use less-strict accounting. De Vos, Smaniotto, and Elsas (2001) compare the evolutionary viability of two types of reciprocity, based on strict vs. loose accounting. Payoffs determine whether an agent survives, and strategies reproduce in proportion to the survival rate. Agents ask for help when it is needed and decide to whom to give it (if anyone), based on past exchanges. Some agents insist on keeping the books strictly balanced, while others favor commitment to old partners, even if they are in arrears. In a population that includes non-givers, simulations demonstrate the importance of committing oneself to an on-going relationship. Ironically, the results suggest that a looser "commitment strategy," based on long-term balancing of the books, is superior to a strategy of strict reciprocity that is less vulnerable to being cheated, a result similar to that reported by Kollock (1993) based on a similar ecological competition. However, Kollock's agents were paired randomly, without the option to select their partners. He also found that loose accounting is superior, but only if the environment is noisy (with occasional mistakes and misinformation). Strict reciprocity is then prone to needless recrimination that can be avoided by looser accounting systems.

One problem with this modeling strategy is that the outcome of an evolutionary tournament may be an artifact of a theoretically arbitrary set of initial contestants. This led Axelrod to use a genetic algorithm (GA) to see if Tit for Tat would evolve in an open-ended population in which strategies could evolve from a random start (1997:14-29). Working with John Holland, Axelrod found several strategies similar to Tit for Tat that proved to be highly robust.

Genetic algorithms are strings of computer code that can mate with other strings to produce entirely new and superior programs by building on partial solutions. Each strategy in a population consists of a string of symbols that code behavioral instructions. These symbols are often binary digits (or “bits”) with values of 0 or 1. A string of symbols is analogous to a chromosome containing multiple genes. A set of one or more bits that contains a specific instruction is analogous to a single gene. The values of the bits and bit-combinations are analogous to the alleles of the gene. A gene’s instructions, when followed, produce an outcome (or payoff) that affects the agent’s reproductive fitness relative to other players in the computational ecology. Relative fitness determines the probability that each strategy will propagate. Propagation occurs when two mated strategies recombine. If two different rules are both effective, but in different ways, recombination allows them to create an entirely new strategy that may integrate the best abilities of each “parent,” making the new strategy superior to either contributor. If so, then the new rule may go on to eventually displace both parent rules in the population of strategies. In addition, the new strings may contain random copying errors. These mutations restore the heterogeneity of the population, counteracting selection pressures that tend to reduce it.

The GA can be used to discover both optimal and likely solutions. Where the aim is to discover what agents should do to optimize performance, the models typically assume global search. This means every agent has complete knowledge of the strategies and fitness of every member of the population and plays against every member of the population with equal probability. Where the aim is to find what agents are likely to do, models often assume local rather than global interaction and knowledge (Klos 1999). Local search can be implemented by embedding the GA in a spatial or social network.

Network Degree

Macy and Skvoretz (1998) embed the GA in a social network to test Weber's theory that Protestant sects in colonial America provided cultural markers needed for trusting strangers in physically dispersed markets. The problem in Weber's argument is that the need for economic growth does not guarantee the evolution of the means for its realization. Macy and Skvoretz's simulations show that a system of "telltale signs" is highly fragile, even with unrealistically generous assumptions about cultural diffusion. However, the robustness can be greatly improved when exchanges are embedded in social structures comprised of a large number of small communities, precisely the conditions that Weber identified in colonial America.

Homophily

Smith and Stevens (1999) model the formation of psychological support networks in which agents seek out relationships with others that will help them manage anxiety. In their model, agents decide with whom to form relationships through a process of assortative mating. They find that agents form relationships with partners who are similar to themselves in their ability to manage stress, creating homophilous clusters. They also discover an exception to Granovetter's (1973) theory of the strength of weak ties. In needy populations, support networks form with stronger attachments but lower transitivity than in populations with less need for social support.

Several other recent studies also suggest that the viability of cooperation is greatly improved when populations can self-organize into locally homogenous clusters (Lomborg 1996). For example, Pedone and Parisi (1997) use socially embedded artificial neural networks to show how altruistic behavior can arise among similar agents and conclude that similarity conveyed by culture may be what allows altruism to evolve in natural settings.

Other studies explore the effects of homophilic clustering in spatial networks. Eshel et. al. (2000) use spatial clustering on a one-dimensional array where agents play Prisoner's Dilemma. Their agents have only two possible strategies, cooperate or defect. Agents interact strategically with nearest neighbors and imitate those (in a somewhat larger neighborhood) who are most successful. Because they cannot reciprocate, there is no advantage to cooperating even in an on-going relationship. Nevertheless, when the game is spatially embedded, they find that cooperation is "a stable strategy that cannot easily be eliminated from the population."

Flache and Hegselmann (1999) explore the macro implications of alternative assumptions about agent cognition in a social support game played on a torus. Rational agents make the choices prescribed by analytical game theory, while adaptive agents respond to experience through reinforcement learning. In each case, agents migrate on the grid, selecting neighbors from whom to request help and deciding whether to service the requests of others in a sequential-play asymmetric Prisoner's Dilemma, where payoff asymmetry reflects difference among agents in the need for and ability to help and sequential play represents the problem of generalized exchange. The forward-looking model makes strong assumptions about information: Each agent knows all players' locations, their payoffs, and their level of need. Flache and Hegselmann find that both forward- and backward-looking agents self-organize into solidary clusters based on class position. However the stratification structure differs. Rational egoists tend to form an "onion-shaped structure of solidarity" with the wealthiest at the center of a large heterogenous cluster, surrounded by rings of increasingly needy agents. In contrast, backward-looking agents migrate into distinct homogenous clusters with much greater class segregation than found among rational egoists.

In sum, numerous studies converge on the conclusion that prosocial strategies thrive on both spatial and social embeddedness, due to homophilic interaction (the tendency to interact with similar strategies while avoiding contact with predators). However, Cohen, Riolo, and Axelrod (2001) point out that what appears to be the effect of homophily may actually be due to the effects of relational stability and transitivity. Transitivity (or “clustering”) means that “paired agents have neighbors who are themselves paired” (Cohen et al. 2001: 11). For example, in a Moore neighborhood, each of an agent’s eight neighbors interacts with two of the other seven. Relational stability (which they call “context preservation”) means that agents continue to interact with the same partners across many periods, creating a “shadow of the adaptive future.” The smaller the neighborhood, the greater the chance of interacting with a previous partner. In short, the effect of embeddedness may not be due to the tendency for local interactions to be with partners who are *similar*, as most studies have assumed, but the tendency for partners to be *correlated* (due to network transitivity) or *familiar* (due to local pairing).

The authors tease apart these effects by manipulating network structure in a population of 256 agents who play a four-iteration Prisoner’s Dilemma game with each of four different partners in each period, for 2500 periods. Agents are programmed with three (initially random) probabilities for cooperating under three conditions: the first move of the game, after the partner cooperates, and after the partner defects. Between periods, agents adopt the strategy used by the most successful of their four partners, based on the payoffs accumulated during that period. In this model, many different strategies can produce identical behavior under given conditions, and the authors do not explain how agents know which of these was actually responsible for the behavior they observe, but somehow the agents guess correctly about ninety percent of the time.

Using controlled computational experiments, the authors observe the independent effects of homophily, on-going relations, and transitivity (“clustering”). They find that on-going relations greatly improve the viability of cooperation while clustering alone has little effect. However, the effect of relational continuity is not due to the prudence of being nice to those one expects to meet again. Because agents imitate their partners, ongoing relations increase the chances that an agent will interact with a partner using a similar (if not identical) strategy. They conclude that “friendly” strategies do well so long as they can generally avoid those that are not.

Diffusion of Reputations

Takahashi (2000) uses an evolutionary model to study the emergence of generalized exchange, in which agents give and receive help but not to one another directly (as in the Kula ring). Takahashi challenges previous studies that assumed that these exchange systems require either altruism or centralized enforcement of the rules of exchange. He then uses an evolutionary model to show that exchange systems can self-organize based on norms of generalized reciprocity (giving selectively to those who give to third parties). He programs agents with two genes that control compliance with norms of generalized exchange and enforcement of compliance by others. The first gene controls the amount the agent gives to others and the second gene controls reciprocity, based on the recipient’s reputation for giving to others. Giving and receiving determines each agent’s relative fitness or chances for reproduction. Reproduction copies the agent’s genes with a small probability of mutation. However, with only two genes, there is no need for recombination, so Takahashi does not use a genetic algorithm. Simulations show that a system of generalized exchange can evolve in a population that is initially non-generous, assuming agents have perfect information about the past behavior of other agents.

Takahashi then relaxes this assumption by positioning agents on a two-dimensional grid, restricting their knowledge, interaction, and reproductive competition to their Moore neighborhood. Thus, agents continue to have perfect information about all their potential exchange partners, of whom there are now only eight (instead of 19). Generalized exchange emerges within each of the overlapping neighborhoods, but Takahashi did not test to see if generalized exchange could evolve between members of different neighborhoods when reputational knowledge remains local.

Castlefranchi, Conte, and Paolucci (1998; see also Conte and Castelfranchi 1995) examine the effect of reputations on deterrence of aggressive behavior on a two-dimensional grid where agents compete locally for scarce resources and adaptation operates through evolutionary selection. They find that a pro-social strategy can thrive in a homogenous population but suffers as contact with aggressors is increased. However, the aggressor's advantage is diminished if agents can exchange information on the reputations of others. Saam and Harrer (1999) used the same model to explore the interaction between normative control and power. They find that systems of informal social control can tip toward either greater equality or inequality, depending on the extent of inequality at the outset.

Bandwagons

Network transitivity (clustering) becomes much more important when outcomes depend on the flow of information through the network. Chwe (1999) proposes a threshold model of collective action in which agents choose to participate depending on the number of neighbors expected to participate. Expectations of neighbors' behavior depends in turn on expectations of neighbors' neighbors' behavior, and so on. In the base condition, 30 agents are randomly assigned two partners with whom they remain attached for the duration of the simulation. Chwe then manipulates transitivity by increasing the number of partners and the bias toward selecting the partners of one's partners. High transitivity avoids an endless regress because an agent's neighbors and the neighbors' neighbors are likely to be the same people. Transitivity is especially important in populations with low thresholds that can be triggered by local knowledge about the behavior of members of densely tied but relatively small local clusters. This may explain the importance of overlapping social ties for Freedom Summer participation reported by McAdam (1988). Conversely, Chwe demonstrates the strength of weak ties in populations with high thresholds. Low transitivity facilitates the diffusion of information about participation of distant agents. Watts' (1999) "small worlds" model suggests that the optimal configuration may be a highly clustered network with a few ties between each small and densely tied locality.

Social Pressure

Chwe's thresholds correspond to agents' concerns about the efficacy of participation in collective action (see also Marwell and Oliver 1993; Macy 1991). Thresholds can also represent agents' responsiveness to social pressures to conform to an emergent norm, as in "bandwagon" models of self-reinforcing popularity. Kim and Bearman (1997) model collective action among agents whose interest in the public good is heavily influenced by social pressure from other participants in their local network. This causes interest in the collective action to spread like a contagion through network channels. The authors find that participation spreads most effectively within densely clustered subnetworks comprising a critical mass of highly-interested primary contributors.

Kim and Bearman's study reflects the conventional wisdom that social pressure to participate is needed to overcome the temptation to "free ride." It follows that dependence on the group for social direction promotes compliance with group obligations, as argued by Homans (1974). Yet a number of ethnographic studies of "deviant cliques" have shown that conformist pressures can also undermine normative compliance, leading to badly suboptimal outcomes for all group members, including the deviants (Shibutami 1978; MacLeod 1995; Willis 1977). This led Kitts, Macy, and Flache (1999) to explore the possibility that dependence on peer approval can backfire, leading to collective action failure rather than success (see also Flache and Macy 1996). They modeled self-organizing social relations using an attractor neural network similar to Nowak and Vallacher's (1998). In these models, pressure to imitate the behavior of other agents increases with the strength of the connecting tie. They added the innovation that agents respond not only to normative pressure but also to the lessons of direct experience (similar to the back-propagation in feed-forward neural nets). Agents in a group-rewarded task group decided whether to work or shirk and whether to approve of other group members. The authors then manipulated agents' tendency to conform to others' behavior. Computer simulations revealed a surprising result – a curvilinear effect of conformity on compliance with prosocial norms. Moderate doses of conformity reduce the coordination complexity of self-organized collective action and help the network achieve satisfactory levels of cooperation. High doses, however, undermine the agent-based learning required to find cooperative solutions. Increasing group size also diminished compliance due to increased complexity, with larger groups requiring more conformity to overcome the coordination problem.

Conclusion

Agent-based modeling is a new tool for theoretical research at the relational level, with particular relevance for sociologists as a bridge between the micro and macro levels. Nevertheless, sociology has lagged behind the other social sciences in recognizing and exploiting this methodology. Computational sociology has traditionally used models that are highly realistic and macrosociological, but agent-based models are highly abstract and tend to be microsociological. This has led to confusion about the appropriate standards for constructing and evaluating agent-based computational models. We conclude our review with a series of recommendations – to referees as well as authors – for realizing the rich sociological potential of relational modeling.

1. *Keep it simple.* Pressure to make agents more cognitively sophisticated is misguided if models become so complex that they are as difficult to interpret as natural phenomena. When researchers must resort to higher order statistical models to tease apart the underlying causal processes, the value of simulation is largely undermined.
2. *Avoid reliance on biological metaphors.* Chattoe (1998) has raised probing questions about modeling cultural evolution as a genetic analog. What is the mechanism that eliminates poor performers from the population and allows others to propagate? "Imitation of the fittest" may be more applicable than starvation and reproduction, but this mechanism replicates only observed behavior and does copy the underlying rules. The distinction between behaviors and the underlying rules has not been given sufficient attention. For example, in repeated Prisoner's Dilemma games, it is easy to imagine that an agent observes and then copies the cooperative behavior of a successful neighbor, but

how does the agent know that this behavior is based on a rule (or supergame strategy) like "Tit for Tat" and not "Win-Stay, Lose-Shift" or "Always Cooperate"?

3. *Experiment, don't just "explore."* Agent-based modeling is an experimental tool for theoretical research. While important discoveries can be made by open-ended exploration of theoretical possibilities, researchers need to resist the temptation to become freewheeling adventurers in artificial worlds. Careful, systematic exploration of a parameter space may be less engaging but it makes for better science. This requires theoretically motivated manipulation of parameters, based on careful review of current theoretical and empirical knowledge, and a clear statement of the hypotheses that guided the experimental design.
4. *Test robustness*. Although simulation designs should use experimental rather than post-hoc statistical controls to identify underlying causal processes, that does not mean researchers should avoid statistical analysis of the results. On the contrary, ABMs, especially those that include stochastic algorithms, require replications that demonstrate the stability of the results. Where possible, replications should include variation in parameters that are theoretically arbitrary or of secondary interest. Authors then need to be careful to distinguish between experimental manipulations (where results are expected to change with the parameters) and robustness tests (where they are not).
5. *Test external validity*. Virtual experiments test the internal validity of a theory, without which there is no need to test the external validity. However, this does not mean there is *never* such a need. ABMs are often used to "grow" familiar macro level patterns, as a way to identify possible causal mechanisms (Epstein and Axtell 1996). When this

succeeds, researchers need to think about ways these mechanisms can be operationalized and tested in laboratory or natural conditions.

6. *Test domain validity*. Axtell et al. (1996) argue that “If computational modeling is to become a widely used tool in social science research, it is our belief that a process we will call ‘alignment of computational models’ will be an essential activity. Without such a process of close comparison, computational modeling will never provide the clear sense of "domain of validity" that typically can be obtained for mathematized theories. 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 results” (1996: 123).
7. *Avoid micro blinders.* The “bottom-up” approach in ABMs might seem to imply that these models can only be used to test microsociological theories. That is a tragic misunderstanding because it precludes what is most exciting to sociologists about this methodology. An artificial world populated by computational agents is a laboratory in which researchers can manipulate structural conditions to test macrosociological theories. However, unlike an earlier generation of holistic equation-based models of dynamical systems, ABMs avoid the reification of causal factors at the macro level. Changes in population density or network structure can not directly affect the emergence of norms or the collapse of conventions. However, macro experiments in virtual worlds provide a rigorous methodology for studying the microfoundations of macro dynamics, in the way advocated by Coleman (1990: 8). The shift from factors to actors is clearly one of the most exciting developments in computational social science. Bringing factors back in will realize the full potential of agent-based modeling, especially in sociology.

References

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

Axelrod, R. 1997. *The Complexity of Cooperation*. Princeton, NJ: Princeton University Press.

Axtell, Robert, Robert Axelrod, Joshua M. Epstein, and Michael D. Cohen. 1996. “Aligning Simulation Models: A Case Study and Results.” *Computational and Mathematical Organization Theory*. 1(2):123-141.

Bainbridge, W.S. 1995. "Neural Network Models of Religious Belief." *Sociological Perspectives*, 38,4: 483-496.

Bainbridge, W.S., E. Brent, K. Carley, D. Heise, M. Macy, B. Markovsky and J. Skvoretz. 1994. “Artificial Social Intelligence.” *Annual Review of Sociology*. Vol. 20:407-36.

Bullenheimer, B., H. Dawid, and R. Zeller. 1998. “Learning from Own and Foreign Experience: Technological Adaptation by Imitating Firms.” *Computational and Mathematical Organization Theory*. 4(3): 267-282.

Caldwell, Steven B. 1997. “Dynamic Microsimulation and the Corsim 3.0 Model.” Ithaca, NY: Strategic Forecasting.

Carley, K. 1991. “A Theory of Group Stability.” *American Sociological Review*. 56: 331-354.

Cartwright, D., and F. Harary. 1956. “Structural Balance: A Generalization of Heider`s Theory.” *The Psychological Review*. 63: 277-293.

Castelfranchi, C., R. Conte, and M. Paolucci. "Normative Reputation and the Costs of Compliance." *Journal of Artificial Societies and Social Simulation*. 1(3).

Chattoe, E. 1998. "Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes?" *Journal of Artificial Societies and Social Simulation.* 1(3).

Chwe, Michael Suk-Young. 1999. "Structure and Strategy in Collective Action." *American Journal of Sociology*. 105(1): 128-56.

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(1):5-32.

Coleman, J.S. 1990. *Foundations of Social Theory*, Cambridge; MA: Harvard University Press.

Conte, R. and C. Castelfranchi. 1995. "Understanding the Effects of Norms in Social Groups through Simulation. In *Artificial Societies: the computer simulation of social life*. Eds. N. Gilbert and R. Conte. London: UCL Press.

Coveney, Peter and Roger Highfield. 1995. *Frontiers of Complexity*. NY: Fawcett Columbine.

de Vos, H., R. Smaniotto, and D. A. Elsas. 2001. "Reciprocal Altruism under conditions of Partner Selection. *Rationality and Society*. 13(2):139-183.

Durkheim, Emile. 1982 (1901). *The Rules of the Sociological Method.* London: Macmillan.

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

Eshel, I., D.K. Herreiner, L. Samuelson, E. Sansone, and A Shaked. 2000. “Cooperation, Mimesis, and Local Interaction.” *Sociological Methods and Research*. 28(3):341-364.

Flache, A. and M. Macy. 1996. “The Weakness of Strong Ties: Collective Action Failure in Highly Cohesive Groups.” *Journal of Mathematical Sociology*. 21:3-28.

Flache, A. and R. Hegselmann. 1999. “Rationality vs. Learning in the Evolution of Solidarity Networks: A Theoretical Comparison.” *Computational and Mathematical Organization Theory*. 5(2): 97-127.

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* (forthcoming).

Forrester, J.W. 1971. *World Dynamics*. Cambridge, MA: MIT Press.

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

Gilbert, N. 1997. “A Simulation of the Structure of Academic Science.” *Sociological Research Online*. 2(2): www.socresonline.org.uk/socresonline/2/2/3.

Goldberg, D.E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning.* Reading, MA: Addison-Wesley. pp 46-48.

Granovetter, Mark S. 1973. “The Strength of Weak Ties.” *American Journal of Sociology*. 78(6): 1360-80.

Halpin, Brendan. 1999. “Simulation in Sociology.” *American Behavioral Scientist*. 42(10): 1488-1508.

Hanneman, Robert A., Randall Collins, and Gabriele Mordt. 1995. "Discovering Theory Dynamics by Computer Simulation: Experiments on State Legitimacy and Imperialist Capitalism." *Sociological Methodology*, 25: 1-46.

Heckathorn, Douglas. 1989. “Collective Action and the Second-Order Free-Rider Problem.” *Rationality and Society*. 1:78-100.

Hegselmann, R. and A. Flache. 1998. “Understanding Complex Social Dynamics: A Plea For Cellular Automata Based Modelling.” *Journal of Artificial Societies and Social Simulation*. 1(3).

Holland, J. 1995. *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Perseus.

Homans, G.C. 1974. *Social Behavior: Its Elementary Forms*. New York: Harcourt, Brace, Jovanovich.

Hopfield, J. J. 1982. “Neural Networks and Physical Systems with Emergent Collective Computational Abilities,” Proceedings of the American Academy of Sciences, 79, 2554-2558.

Kaufman, S. 1996. *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford: Oxford University Press.

Kim, H. and P. Bearman. 1997. "The Structure and Dynamics of Movement Participation." *American Sociological Review*. 62:70-93.

Kitts, J., M. Macy, and A. Flache. 1999. "Structural Learning: Attraction and Conformity in Task-Oriented Groups." *Computational and Mathematical Organization Theory*. 5(2): 129-145.

Klos. 1999. "Decentralized Interaction and Co-Adaptation in the Repeated Prisoner's Dilemma." *Computational and Mathematical Organization Theory*. 5:147-165.

Kollock, P. 1993. "An Eye for an Eye Leaves Everyone Blind." *American Sociological Review*. 58:768-86.

Latane, B. (1996) "Dynamic Social Impact: Robust Predictions from Simple Theory." In *Modeling and Simulation in the Social Sciences from a Philosophy of Science Point of View*. Eds. Hegelsmann, Mueller, and Trovitzsch. Pp. 287-310. Boston: Kluwer Dorderecht.

Lomborg, B. 1996. "Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma." *American Sociological Review*. 61:278-307.

Lomi, Alessandro and Erik Reimer Larsen. 1998. "Density Delay and Organizational Survival: Computational Models and Empirical Comparisons." *Computational and Mathematical Organization Theory*. 3(4): 219-247.

Macy, M. 1991. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review*. 56:730-47.

Macy, M. and J. Skvoretz. 1998. "The Evolution of Trust and Cooperation Between Strangers: A Computational Model." *American Sociological Review* 63:638-660.

Mark, N. 1998. "Beyond Individual Differences: Social Differentiation from First Principles." *American Sociological Review*. 63:309-330.

Marwell, G. and P. Oliver. 1993. *The Critical Mass in Collective Action: A Micro-Social Theory*. Cambridge: Cambridge University Press.

McAdam, D. 1988. *Freedom Summer*. New York: Oxford University Press.

Meadows, D. L., W.W. Behrens III, D. H. Meadows, R. F. Naill, J. Randers, and E. K. Zahn. 1974. *The Dynamics of Growth in a Finite World*. Cambridge, MA: MIT Press.

Meeker and Leik. 1997. "Use of Computer Simulation for Theory Development: An Evolvin Component of Sociological Research Programs." Pp. 47-70 in J. Szmatko, J. Skvoretz, and J. Berger, eds. *Status, Network, and Structure: Theory Development in Group Processes.* Palo Alto, CA: Stanford University Press.

Nowak, A. and R.R. Vallacher. 1998. "Toward Computational Social Psychology: Cellular Automata and Neural Network Models of Interpersonal Dynamics." In *Connectionist Models of Social Reasoning and Social Behavior*. Eds. S.J. Read and L.C. Miller. Pp.277-311.

Orbell, J., L. Zeng, and M. Mulford. 1996. "Individual Experience and the Fragmentation of Societies." *American Sociological Review*. 61:1018-1032.

Pedone, R. and Parisi, D. 1997. “In What Kind of Social Groups can ‘Altruistic’ Behaviors Evolve?” In *Simulating Social Phenomena*. Conte, R., Hegselmann, R. and P. Terna (eds.) Berlin: Springer-Verlag. pp. 195-201.

Rapoport, A. and A.M. Chammah. 1965. *Prisoner’s Dilemma: A Study in Conflict and Cooperation*. Ann Arbor: University of Michigan Press.

Resnick, Mitchell. 1997. *Termites Turtles and Traffic Jams.* Cambridge, MA: MIT Press.

Reynolds, Craig W. 1987. “Flocks, Herds, and Schools: A Distributed Behavioral Model.” *Computer Graphics*, 21: 25-34.

Rosenkopf, L. and E. Abrahamson. 1999. “Modeling Reputational and Informational Influences in Threshold Models of Bandwagon Innovation Diffusion.” *Computational and Mathematical Organization Theory*. 5(4): 361-384.

Rummelhart, D. and J. McClelland. 1988. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.

Saam, N.J. and A. Harrer. 1999. “Simulating Norms, Social Inequality, and Functional Change in Artificial Societies.” *Journal of Artificial Societies and Social Simulation.* 2(1).

Saam, Nicole. 1999. “Simulating the Micro-Macro Link: New Approaches to an Old Problem and an Application to Military Coups.” *Sociological Methodology* 29:43-79.

Sawyer, K. 2001. “Artificial Societies and the Micro-Macro Link in Sociological Theory.” Unpublished manuscript.

Schelling, T. 1971. “Dynamic Models of Segregation.” *Journal of Mathematical Sociology.* 1:143-186.

Schumpeter, Joseph A. 1909. "On the Concept of Social Value." *Quarterly Journal of Economics*, 23:213-232.

Shibutani, Tomatsu. 1978. *Derelicts of Company K : A Sociological Study of Demoralization*. Berkeley, CA: University of California Press.

Simon, H. 1998. *The Sciences of the Artificial*. Cambridge, MA: MIT Press.

Skvoretz, J. and T. Fararo. 1995. “The Evolution of Systems of Interaction.” *Current Perspectuives in Social Theory*. 15:275-299.

Smith, T. S. and G. T. Stevens. 1999. “The Architecture of Small Networks: Strong Interaction and Dynamic Organization in Small Social Systems.” *American Sociological Review.* 64:403-420.

Strang, D. and M. Macy. 2001. “’In Search of Excellence:’ Fads, Success Stories, and Adaptive Emulation.” *American Journal of Sociology*, forthcoming, vol. 106.

Takahashi, Noboyuki. 2000. “The Emergence of Generalized Exchange.” *American Journal of Sociology*. 105(4):1105-34.

Watts, Duncan J. 1999. “Networks, Dynamics, and the Small-World Phenomenon.” *American Journal of Sociology*. 105(2):493-527.

Willis, Paul. 1977. *Learning to Labor.* New York: Columbia University Press.

Vakas-Duong, Deborah, and Kevin Reilley. 1995. "A System of IAC Neural Networks as the Basis for Self-Organization in a Sociological Dynamical System Simulation." *Behavioral Science.* 40: 275-303.

Figure 1: Typology of Agent-Based Models

Article	Substantive Problem	Network[1]	Elective Ties[2]	Number Agents	Adaptative Mechanism[3]	Adaptive Criteria	Manipula-tion[4]
Lomi-Larsen ’98	Differentiatn	Spatial	N	10000	Reproduction	Frequency	Macro
Mark ’98	Differentiatn	Social	Y	6-100	Imitation	Familiarity	Macro
Axelrod ‘97a	Differentiatn	Spatial	Y	10^2-10^4	Imitation	Familiarity	Macro
Orbell et al. ’96	Differentiatn	Social	Y	1000	Learning	Observation	Macro
Bullheimer-Zeller ’98	Diffusion	Global	N	10	Imitn & learn.	Success	Micro
Rosenkopf-Abr. ’99	Diffusion	Global	N	20	Imitation	Frequency	Macro
Strang-Macy ’01	Diffusion	Global	N	100	Imitn & learn.	Success	Both
Eshel et al. ’00	Social Order	Spatial	N	1000	Imitation	Success	Both
Flache-Hegs. ’99	Social Order	Spatial	Y	315	Learning	Success	Micro
Chwe '99	Social Order	Spatial	N	30	Imitation	Frequency	Macro
Saam-Harrer ’99	Social Order	Spatial	Y	50	Reproduction	Success	Macro
Cohen et al. ’01	Social Order	Spatial	Y	256	Imitation	Success	Macro
Takahashi ’00	Social Order	Gl & Sp	N	20-100	Reproduction	Success	Macro
Axelrod ‘97b	Social Order	Global	N	20	Reproduction	Success	None
Kim-Bearman ’97	Social Order	Social	N	100	Imitation	Success	Macro
De Vos et al. ’01	Social Order	Social	Y	10-50	Reproduction	Success	Macro
Macy-Skvoretz ‘98	Social Order	Social	Y	1000	Reproduction	Success	Macro
Kitts-Macy-Flac. ‘98	Social Order	Social	Y	3-15	Learning	Success	Both

1. Spatial: restricted by physical distance; social: restricted by social distance; global: not restricted by distance.
2. Is interaction forced or voluntary, based on an option to move, exit a relationship, or assortative mating?
3. Reproduction: Successful agents replace or convert unsuccessful ones; imitation: agents copy observed behavior (but not the underlying rule); learning: agents change behavior based on experience.
4. Macro: manipulation of global parameter; micro: manipulation of agent parameter.