Social Life *In Silico*: the Science of Artificial Societies*

Damon M. Centola and Michael W. Macy

Cornell University

To appear in

*Handbook of Group Research and Practice*

Susan Whelan, editor
Sage Publications

August 26, 2004

---

* The research on which this study is based was supported by a National Science Foundation IGERT Fellowship for the study of nonlinear dynamics to the first author and a National Science Foundation grant (SES0241657) to the second author. We also thank Arnout Van de Rijt for helpful comments on an earlier draft.

Computational modeling has become well established as an essential methodology in the biological and physical sciences (Strogatz, 1994), and has recently begun a migration into the social sciences (Epstein & Axtell,1996; Axtell, 2000; Axelrod, 1997; Gilbert & Troitzsch, 1999). Interest in modeling the microfoundations of macrosocial patterns has led to advances in evolutionary and cognitive game theory (Axelrod, 1984, 1997; Roth & Erev, 1995), organizational ecology, (Lomi & Larson, 1998), social networks (Carley, 2003), artificial societies (Epstein & Axtell, 1996), cultural differentiation and diffusion (Mark, 1998), and collective action (Chwe, 1999; Marwell & Oliver, 1993; Macy, 1991).

The growing interest in this tool for theoretical research can be attributed in part to the lowering of technological barriers. Most social scientists now have access to machines capable of running very complex and computationally intensive modeling programs. Moreover, graphical interfaces facilitate highly intuitive representations of the results.

The power of inexpensive desktop computers explains the widening access to computational modeling but not the lure. In the physical and life sciences, the attraction centers on the ability to use numerical integration to solve problems that would otherwise be mathematically intractable. In physics, systems with non-linear dynamics and sensitive dependence on initial conditions (so-called "complex systems") have motivated the use of a wide variety of computational techniques (Strogatz,1994). And in biology, large systems that exhibit stochasticity and periodic dynamics have motivated the adoption of many of the computational techniques developed by physicists (Strogatz, 2003; May 1974, 1976).

Similarly, social scientists have begun to appreciate that the complexity of social systems cannot be understood using traditional analytical techniques. In particular, contrary to traditional intuitions, the emergence of macrosocial patterns out of microsocial interactions exhibits a high degree of non-linearity. Thus, it should not be surprising that interest in computational modeling has migrated from "artificial life" to "artificial societies" (Epstein & Axtell, 1996), in which investigators attempt to grow complex macrosocial patterns from the "bottom up."

Increasingly, these bottom-up computational techniques use agents embedded in a network to study how local interactions can generate system level dynamics. These "agent-based models" (hereafter ABM) allow researchers to study how local decision heuristics and network topology interact to produce highly complex and often surprising global patterns.

*Why not just use game theory instead?*

Agent-based computational modeling extends the traditional microsocial approach of classical game theory. Game theory is not a computational technique, but a formal mathematical tool that allows the behavior of a population of two or more players to be analytically derived from the utility maximizing behavior of every member, based on the solution concept proposed by John Nash (1950). Nash proved that every general sum noncooperative game has at least one Nash equilibrium in mixed strategies. A Nash equilibrium is a configuration of strategies such that no player has a incentive to deviate unilaterally. Hence, Nash equilibria are self-enforcing; no contract is necessary to guarantee compliance. This allows for the possibility that social order can self-organize, even in the absence of the Leviathan.

Although this result is clearly of enormous significance across all the social sciences, there are important limitations to the classical analytical approach that have spurred the development of evolutionary and learning-theoretic extensions. Nash equilibrium analysis tells us if there are any strategic configurations that are stable, and if so, how they are characterized. Knowing that a configuration is an equilibrium means that if this state should obtain, the system will remain there forever. However, this does not imply that this state will ever be reached. Where multiple equilibria exist, the normal form[1] Nash solution cannot tell us which equilibria will be selected or with what relative probability. Nor does it tell us how such states might be achieved, or what will happen if an equilibrium should be perturbed. In short, the static focus on the identification of equilibrium fails to capture the out-of-equilibrium dynamics that guide a population of players from some initial condition to a stable social arrangement or from one equilibrium (when perturbed) to another.

Moreover, Nash equilibria cannot explain social stability among interacting agents who are themselves dynamically changing their strategies. Put differently, a strong limitation of the Nash equilibrium concept is the assumption that population stability is predicated upon the stability of the individual members – the unwillingness of anyone to change strategies unilaterally. Agent-based models extend equilibrium analysis to cases in which social stability arises

1 The problems with equilibrium selection and out-of-equilibrium dynamics apply mainly to normal form games (those in which players must choose simultaneously). The Nash solution for extensive form games with complete information allows for subgame perfect equilibrium selection and for identification of the path to that equilibrium via backward induction.

even though individual strategies are constantly changing.[2] A dynamic equilibrium obtains when the forces pushing the system in one direction are precisely balanced by countervailing forces, as happens in a market when changes in consumer behavior prompt changes in the distribution of production strategies, which in turn drive the consumer behavior back into balance. Empirical studies of market behavior (Kirman, 1989, 1992) show that agent-based models with dynamic agents yield better approximations of market equilibria than traditional game theoretic models with all players in Nash equilibrium.

The Nash solution concept has another important limitation – the mathematical intractability of heterogeneity of preferences among the actors. It is difficult to derive equilibria for a population of players with different preferences and incomplete information about the preferences of other players.

[2] Young (1993, 1998) has proposed an analytical solution concept -- *stochastic stability* -- that overcomes this limitation of the Nash solution concept. Other recent extensions to classical game theory also provide analytical tools for exploring boundedly rational agents, heterogeneity, and simple network structures. However, agent based models have the advantage that model parameters can be much more easily explored, allowing researchers to understand a system's behavior under a wider range of parameter settings than would be possible (or mathematically tractable) using an analytical apparatus.

ABMs show how agents with different preferences can interact dynamically to produce macrosocial outcomes.

The tractability of classical of game theoretic solutions also requires that players have unlimited cognitive capacity, complete and perfect information, and common knowledge.[3] However, the analytical power is purchased at a very high price. Game theory becomes the theory of games played by game theorists (Macy & Flache, 2002).

ABMs allow us to relax the cognitive and informational demands of the classical approach and to explore the behavior of a lay population. In place of the standard *homo œconomicus* of classical game theory, ABMs allow investigators to explore how different cognitive and behavioral mechanisms affect aggregate dynamics. These mechanisms include imitation, heuristic decision, stochastic learning, Bayesian updating, best reply with finite memory, and local optimization.

We conclude this litany of limitations by pointing to what we regard as the most important of all – the assumption that the players interact in a fully

3 Perfect information means that each player knows the entire history of moves and outcomes; complete information means there is no private information; and common knowledge means that all players know that all of the other players know what they know, know that they know it, and so on.

connected (or complete) network or in randomized subsets. Although game theoretic analysis has been applied to network interactions (Buskens 2002), mathematically tractability imposes severe constraints on the exploration of structural parameters. Yet social scientists increasingly appreciate the importance of social structure in the study of population dynamics. The assumption of random or fully connected ties among players in a population is an abstraction that ignores a crucial vector of explanation in the social sciences, the structure of social ties. ABMs facilitate exploration of the effects of network topology on aggregate dynamics.

These extensions of traditional game theory can be stated very succinctly: *Agent-based models allow researchers to systematically explore (independently and in combination) the effects of heterogeneity, bounded rationality, and network structure on the dynamics and stability of social systems* (Axtell, 2000). This short summary implies an indefinitely long research agenda, with unexpected branches and turns that remain to be pursued.

A simple but highly illuminating example comes from recent efforts to model social cooperation in a game called the Prisoner's Dilemma (PD). The PD formalizes the paradox in which individually rational players produce a collectively irrational outcome. In PD, "defection" (or cheating) is the dominant strategy, that is, each player is better of defecting, no matter what the partner chooses. Yet the payoff for mutual defection is less than that for mutual cooperation.

An early computational refinement of this result, called the “replicator dynamics,” looked at changes over time in the distribution of “Cooperate” and “Defect” strategies in a population of players. While the Nash solution for the PD is static, the replicator dynamics assumes a heterogeneous population of agents, and evaluates differential reproduction rates based on their local success in the PD game by using numerical integration to iterate over many generations. Interestingly, the replicator dynamic solution is identical to the Nash solution: global defection. Indeed, Nash noted the convergence between the static equilibrium that bears his name and what he called the “mass action” equilibrium based on reasoning that anticipated the replicator dynamics.

The replicator dynamics are not agent based, however. An early contribution of agent-based modeling to the study of the PD was presented by Nowak & May (1993). They used a spatial lattice to show that local clusters of cooperation could be sustained indefinitely in a heterogeneous population constrained by non-random interactions. The success of cooperation in the Nowak and May model is due to the topology of interaction. Agents interacting in a spatial lattice have more contact with neighbors who are similar to them than to those who differ; thus cooperators will fare better (on average) than defectors, and be able to reproduce at competitive rates. This result shows how "bottom-up" aggregation, from micro interaction to population behavior, can be very different when the system is studied as an ABM than when studied under either the replicator dynamics or the Nash formalization.

*What are we modeling anyway?*

Despite the advantages of ABMs over traditional game theoretic methods, skepticism remains about the validity and robustness of results obtained from "computer simulations." The criticisms are similar to those directed at laboratory research in social psychology: How can experiments under artificial conditions be used to make predictions about behavior outside the lab?

These criticisms reflect a misunderstanding of experimental methodology that applies to both laboratory and computational research. Experiments are used to test predictions about outcomes under controlled conditions, not in natural settings where predictions are likely to be confounded by unmeasured effects. Should the predictions be supported in the lab, it may be useful to apply the theory to "real world" conditions that violate its scope, so as to generate intuitions about how the theory might be extended or elaborated.

With agent-based modeling, the misunderstanding may also reflect confusion created by the widely misused term "simulation," whose meaning was shaped by earlier generations of computational techniques. Classical simulation models (e.g., Cyert & March, 1963) used computers to simulate dynamical systems, such as control and feedback processes in organizations, industries, cities, and even global populations. The value of these models depends entirely on predictive accuracy and realism.

In contrast, artificial societies are highly abstract "thought experiments."[4] Contrary to the critics, making these models more "realistic" would add complexity that is likely to undermine their usefulness. The priority in building ABMs is to keep them as simple and abstract as possible. From this perspective, the artificiality of ABMs is a virtue, not a vice. 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 (Axelrod, 1997, p. 5). "But if the goal is to deepen our understanding of some fundamental process," Axelrod continues, "then simplicity of the

4 For a more detailed overview of the history of social simulation, see Gilbert and Troitzsch 1999.

assumptions is important and realistic representation of all the details of a particular setting is not."

This principle extends to the behavioral assumptions about the cognitive complexity of the agents. The principle of emergence suggests that the complexity of social life need not be reducible to the cognitive complexity of individuals. Although agents may follow simple rules, the interactions can produce global patterns that may not be at all obvious and are very difficult to understand. Modeling artificial worlds allows us to explore the complexity of the social environment by removing the cognitive complexity (and idiosyncrasy) of constituent individuals. The intent behind these models is to understand 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.

*Don't the results simply depend on the assumptions?*

Another concern is that the global patterns generated by ABMs are artifacts of arbitrary assumptions about local behavior or interaction. In one important sense, this should always be the case, since the explanatory strategy in agent modeling, as in mathematical proofs, is to demonstrate that the *explanandum* follows from the *explanans*. Nevertheless, it can also happen that two similar but non-identical models generate similar but non-identical results. Which is to be believed?

The answer can be found through a modeling technique called "alignment" (Axtell et al., 1996). Alignment is a way of integrating two models such that each becomes a special case of a more general model, in order to formalize the logical implications of the different model specifications. A recent example of alignment involved two prominent and very similar models of social dilemma[5] games played by adaptive agents, the Bush-Mosteller stochastic learning model (Bush and Mosteller, 1955) and the Roth-Erev payoff-matching model (Roth and Erev, 1995). The two models share three behavioral assumptions that relax the cognitive demands of the Nash approach – experiential induction (vs. logical deduction), reward and punishment (vs. utility), and melioration (vs. optimization). Both models identify two new solution concepts for the problem of cooperation in social dilemmas. One solution is a socially deficient self-correcting equilibrium (or social trap) and the other is a self-reinforcing equilibrium that is usually (but not always) socially

5 A social dilemma is a noncooperative game with at least one Pareto deficient Nash equilibrium, such as Prisoner's Dilemma or Stag Hunt (in which mutual defection is deficient), or Chicken (probabilistic cooperation is deficient).

efficient. The models also identify the mechanism by which players can escape the social trap – stochastic collusion, based on a random walk in which both players wander far enough out of equilibrium that they escape its "gravitational" pull. Random walk, in turn, implies that a principle obstacle to escape is the coordination complexity of stochastic collusion. Thus, these learning models direct attention to conditions that reduce coordination complexity, including small-world networks (Watts, 1999) which minimize the average number of partners for each player yet still permit locally successful strategies to propagate. Temporal Schelling points are another potential solution, such as the holidays that coordinated locally self-organizing truces in the trenches of World War I. Coordination complexity can also be reduced by high learning rates which reduce the number of steps that must be coordinated in the random walk.

It is here – the effect of learning rates on stochastic collusion – that the two learning models diverge. The divergence might easily go unnoticed, given the theoretical isomorphism of two learning theoretic models based on the same three fundamental behavioral principles. Yet each model implements these principles in different ways, and with different results.

In order to identify the differences, Flache and Macy (2002) aligned and then integrated the two models as special cases of a general reinforcement learning model. The integration uncovered a key hidden assumption, the "Power Law of Learning." This is the curious but plausible tendency for learning to diminish with success and intensify with failure, which they labeled "fixation." Fixation, in turn, impacts the effective learning rate, and through that, the probability of stochastic collusion. This exercise showed how the integration of alternative models can uncover underlying principles and lead to a more general theory.

*Is "computational experiment" an oxymoron?*

Social psychologists interested in the subtleties of motivation may find ABMs too abstract and coarse-grained to be useful for studying highly idiosyncratic individual behavior. However, we believe it is just the opposite. ABMs are ideally suited for experimental social psychology precisely because of their abstractness. Computational experiments using ABMs can provide a propaedeutic that lays the groundwork for laboratory experiments. ABMs serve a proto-experimental function by testing the internal validity of proposed explanatory mechanisms. By assigning agents simple rules that isolate a mechanism, an experimental researcher can use the agent model as a "test run" to see whether the proposed mechanism can, in principle, account for the aggregate patterns of behavior that are of interest. Computational results can help to focus and refine directions of empirical research, and also assist in framing the design of experimental tests. Computational experiments may even show that the behavioral mechanism being studied can produce unanticipated consequences, thus leading to new experimental hypotheses.

For example, suppose we want to know whether residential segregation can be diminished by promoting multiculturalism? One of the earliest ABMs, implemented by Thomas Schelling (1971) on a large checkerboard, suggests that a better strategy might be to promote ethnic "color blindness," rather than appreciation of diversity. His "tipping" model revealed a surprisingly strong tendency toward neighborhood segregation, so long as individuals take into account ethnic composition, even when they have a preference for diversity.

ABMs can also identify social psychological mechanisms that might account for enigmatic behaviors. These mechanisms can then be tested under controlled conditions in the laboratory to see if the predicted outcomes obtain. For example, a recent application (Centola, Willer & Macy, *forthcoming*) addressed the curious tendency for people to pretend to believe something they know not to be true, and to disparage those who disagree, a behavioral pattern that is mocked in Hans Christian Andersen's story of the "Emperor's New

Clothes." Everyday examples include college students who celebrate and encourage intoxication through drinking games, yet who privately express discomfort with excessive consumption (Prentice & Miller, 1993). The foible is not limited to students. We all know "naked" scholars whose incomprehensible writings are applauded by those who pretend to understand and appreciate every word, and who disparage critics for being intellectually shallow. The pattern extends to social snobbery by posturing elitists who are privately bored by Bavarian opera, and to citizens in totalitarian regimes who denounce their neighbors so as to affirm their loyalty to the regime (Kuran, 1995).

Centola, Willer and Macy (*forthcoming*) wanted to know if the fear of exposure as an imposter was sufficient to sustain a cascade that might then trap a population into thinking a norm was highly popular, when in fact almost everyone shared their strong (but private) wish that the norm would go away. More precisely, they wanted to see if this could happen even in the absence of any explicit pressure to enforce the norm. In addition, they wanted to identify the conditions that might make a population highly vulnerable to such a cascade.

To find out, they modeled a population of agents who are willing to comply with a norm they privately dislike in order to win social approval or avoid disapprobation, but only if the pressure to comply is sufficiently strong.

Having complied with the norm, they may also enforce compliance by others for a similar reason – out of fear of exposure as an opportunistic imposter who does not really believe in the norm. The motivating question was whether popular enforcement of an unpopular norm could spread without any explicit social pressure to enforce the norm.

The "Emperor's Dilemma" is a cascade model in which agents are assigned thresholds for compliance and enforcement that reflect the strength of the agent's antipathy toward the norm. The enforcement threshold is always higher than the compliance threshold, reflecting the additional costs of sanctioning others, over and above the costs of compliance. Agents comply and enforce when the social pressure for and against the norm exceeds the compliance and enforcement thresholds, respectively. This pressure comes not from seeing one's neighbors comply (as in a model of herd behavior), but only from seeing them enforce compliance (as in a model of norm enforcement).

The model is initialized with a population that contains only a tiny fraction of vigilant "true believers" (1 percent), who truly support the norm and will always apply pressure if they spot any deviance at all. The rest of the population opposes the norm, and in the absence of enforcement, the skeptics refuse to comply. The model assumes that agents live in a clustered network, but with nontrivial overlap between the clusters, such that non-neighbors who have one neighbor in common are likely to have more than one.

The model demonstrated that it is indeed possible for widespread enforcement of an unpopular norm to emerge even in the absence of any pressure to enforce the norm, so long as the fear of exposure as an imposter is sufficient to motivate enforcement as a way to signal the sincerity of compliance. This suggests the need to test whether this fear of exposure can lead to enforcement behavior under controlled conditions in the laboratory.

A skeptic might counter that one could have arrived at the computational result by mere intuition, and then gone directly to the lab. Not so. The computational model also identifies some surprising conditions that are necessary for these cascades to emerge. For example, it turns out that cascades are not possible if there are too many true believers. For another, the authors discovered that unpopular norms tend to collapse if agents eventually change their private beliefs to conform to their public behavior, thereby making the norm *less* unpopular over time. Allowing agents with weak convictions to eventually switch sides caused a reverse cascade that undermined support for the unpopular norm. The norm collapses not when a child laughs at the naked

Emperor, but when the adults begin to actually believe they can see the beautiful new clothes!

Having observed this unexpected result, the authors dug deeper to find out how this had happened, and found a straightforward explanation that might easily have been missed had they relied solely on intuition to predict the macrosocial consequences of the microsocial postulates. An agent who comes to actually believe the norm no longer worries about exposure as an imposter and is thus no longer compelled to falsely enforce. Thus, converted agents conform to the norm, but they no longer enforce, given the weakness of their convictions. When these converts stop enforcing, there is insufficient social pressure to induce compliance by those of their neighbors with strong private convictions. They in turn stop complying, and pressuring others to do so, which further reduces local pressure in the neighborhood. The social stability of an unpopular norm is thus undermined by increased popularity of the norm!

We also found that changing the network structure dramatically changed the results. In particular, adding "bridge ties" between otherwise distant clusters tends to inhibit the cascades. Why this sensitivity to network topology? These "shortcuts" across the network tend to dissipate social pressure below the local critical mass needed to fuel the cascade. Contrary to an extensive literature on small worlds and "the strength of weak ties," bridge ties across clusters actually reduce the propagation of unpopular norms.

These results deepened our understanding of the dynamics of a puzzling social phenomenon that has perplexed social scientists for years. Further, the results helped to elaborate the experimental design in a laboratory study of the effects of social structure and changing preferences on the enforcement decisions of individuals.

*Conclusion*

We conclude with a brief set of principles of agent-based modeling. First, the researcher must specify the rules that govern the agent's behavior, and whether they are deterministic or probabilistic. These rules define the agents as instrumental (agents respond to payoffs relative to aspiration levels), emotional (agents respond to behavior by others), or normative (agents act on internalized obligations). These agent specifications need not be fixed. On the contrary, by manipulating agent knowledge and ability, the researcher can explore the macrosocial consequences of changes in individual behavior. Agent's rules can also change endogenously over time, through processes of learning or evolution.

Second, the researcher must specify a network structure in which the agents will interact. The network structure is the environment, or social space, in which the agents are located and constrains who may interact with whom. Note that this specification does not preclude the possibility that agents decide with whom to interact, creating a dynamic network. Network structure can also be a complete graph in which all agents interact with everyone. Since it is exterior to the agent, the network specification is the purely structural component of the model. Agents can interact in social networks, with spatial constraints, with random pairing, or in specified or dynamic group structures.

With these simple specifications, we can study the emergent (and often surprising) macrosocial consequences of psychological mechanisms and local interactions. These results can then inform experimental designs for testing the theories in the laboratory.

The process can also go the other way. Laboratory experiments can also inform elaboration of agent models. This feedback between agent modeling and laboratory experimentation can lead to a fecund "bottom up" exploration of the self-organization of social life.

*References*


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

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

Axtell, R. (2000). Why Agents: On the Varied Motivations for Agent Computing in the Social Sciences. *CSED Working Paper No. 17*, The Brookings Institution, Washington, DC

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

Bush, R.R. and Mosteller, F. (1955). *Stochastic Models for Learning*. New York: John Wiley and Sons.

Buskens, V. (2002). *Social Networks and Trust.* Kluwer Academic Publishers.

Carley, K. M. (2003). Dynamic Network Analysis. In *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, Breiger, R., Carley, K. & Pattison, P. (Eds.) Committee on Human Factors, National Research Council, National Research Council. (pp.133-145).

Centola, D., Willer, R. & Macy, M. (*forthcoming*). The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms. *American Journal of Sociology*.

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

Cyert, R. & March, J. G. (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.

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

Flache, A. & Macy, M. (2002). Stochastic Collusion and the Power Law of Learning. *Journal of Conflict Resolution*. *46*,629-653.

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

Kirman, A. (1989). The Intrinsic Limits of Modern Economic Theory: The Emperor Has No Clothes. *Economic Journal*, *Vol. 99* (395), 126-39.

Kirman, A. (1992). Whom or What Does the Representative Individual Represent? *Journal of Economic Perspectives*, *Vol. 6* (2), 117-36.

Kuran, T. (1995). *Private Truths, Public Lies: the Social Consequences of Preference Falsification*. Cambridge, MA: Harvard University Press.

Lomi, A. & Larsen, E. R. (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 Flache, A. (2002). Learning Dynamics in Social Dilemmas. *Proceedings of the National Academy of Sciences, 99,* 7229-36.

Mark, N. 1998. Social Differentiation from First Principles. *American Sociological Review*. *63*,309-30.

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

May, R.M. (1974). *Stability and Complexity in Model Ecosystems*. Princeton, NJ: Princeton University Press

May, R.M. (1976). "Simple mathematical model with very complicated dynamics." *Nature*, *26*,457.

Nash, J. (1950). "Equilibrium Points in N-Person Games." Proceedings of the National Academy of Sciences, *36*, 48-49.

Nowak, M. & May, R. (1993). "Evolutionary Games and Spatial Chaos." *Nature*, *359*, 826-929.

Prentice, D.A. & Miller, D.T. (1993). Pluralistic Ignorance and Alcohol Use on Campus: Some Consequences of Misperceiving the Social Norm. *Journal of Personality and Social Psychology 64* (2), 243-56.

Roth, A.E. and Erev, I. (1995). Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. Special Issue: Nobel Symposium, *Games and Economic Behavior*, *8,*164-212.

Schelling, T. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology, 1,* 143-186.

Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering.* Reading, MA: Perseus Books.

Strogatz, S. H. (2003). *Sync: The Emerging Science of Spontaneous Order*. New York: Hyperion.

Watts, D. J. (1999). Networks, Dynamics and the Small-World Phenomenon. *American Journal of Sociology*, *105*, 493-527.

Young, H. P. (1993). The Evolution of Conventions. *Econometrica*, 61, 57-84.

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