The New *Coevolution* of Information Science and Social Science:

From Software Agents to Artificial Societies and Back *or* How More Computing Became Different Computing

Robert L. Axtell[1]
Center on Social and Economic Dynamics
The Brookings Institution
1775 Massachusetts Avenue, NW
Washington, DC 20036

raxtell@brookings.edu

Abstract: The ways in which exponentially increasing IT capabilities are reshaping the social sciences are briefly reviewed. Many of these changes are primarily increases in scale and scope and do not represent new methodology. However, one specific IT-facilitated development—multi-agent systems—holds out the promise of fundamentally altering the ways in which social science models are conceived, built, explored and evaluated. Here the nascent field of multi-agent social science is described and some alternative futures for it are sketched. But it turns out that the road from IT to social science is not a one way street. Increasingly, results from the social sciences are making their way into computer and information science. Specifically, multi-agent systems researchers are progressively utilizing ideas from game theory (e.g., mechanism design), economics (e.g., auction theory), and even sociology (e.g., social networks). Today we are witnessing the beginnings of the *coevolution* of IT and social science, a process that offers to invigorate the social sciences, while simultaneously threatening their very existence as autonomous fields of inquiry.

Keywords: multi-agent systems, agent-based computational social science, coevolution

[1] This work was generously supported by the National Science Foundation and by grants to the Center on Social and Economic Dynamics at Brookings by the John D. and Catherine T. MacArthur Foundation and the Alex C. Walker Foundation.

I Introduction: IT and the Social Sciences

When computer science and social science are mentioned in the same breath, one typically encounters just a few *typical* responses. While it is unclear whether we learn more about the subject or respondents from such answers, let us will review them here nonetheless. Among a certain group of (usually natural) scientists it is assumed that the intersection of these two fields is essentially empty. Since this is an objectively false position we will dismiss it for our purposes here, although the fact that it is a rather widespread view is certainly something of a public relations problem, for one field or another.

A quite different response acknowledges that some social scientists work with significant datasets—indeed, even vast ones—which serve, through information technology, as the basis for much of what passes as empirical social science today. Each step in the process of working with social science data, from collection, to organization, reduction, and ultimate analysis, usually involves IT of one form or another so that the overall process is not really even feasible without modern IT. This is the domain of the applied social scientists, econometricians, survey researchers and so on.

A third type of respondent knows enough about the social sciences to appreciate that in addition to empirical uses of IT, a smaller but no less important enterprise involves model building and analysis. Such efforts come in a variety of flavors, some more computational than others. For example, *systems dynamics* is an inherently computational approach, while game theory is analytical in the main, with some computing done on the margins.

A yet different response has it that among social theorists there is often the need to resort to classical numerical methods in order to solve models. Here the techniques employed look not unlike those commonly encountered in the natural sciences (Judd [1999]).

Perhaps the final canonical response to questions about the intersection of computing and the social sciences mentions the connection between cognitive models and social phenomena. Here invariably the name Herbert Simon comes to the fore, perhaps with co-workers.[2] As a father of artificial intelligence in computer science and the most articulate defender of 'bounded rationality' in the social sciences, he made contributions in both subject areas. But he also conjoined these research enterprises in various ways, for example, through the 'behavioral theory of the firm' research group at Carnegie Tech in the early 1960s (Cyert and March [1963]), which utilized computational models to explore the dynamics of firms.

While each of these five views is true in its own way—indicating mostly the methodological heterogeneity of the social sciences—none of them grasps the whole picture. For the rapid development of information technology has sparked concomitant developments in social science methodology, and obversely, computer science is being invigorated today by major research areas in the social sciences—e.g., market processes, auction design, social networks—in a way that was unimaginable even a short time ago. Thus it is meaningful to speak of the *coevolution* of computer science and the social sciences, whereby advances in one field lead to progress in the other, nucleating further improvements in the original field, and so on. This catalytic process has already given rise to a new subfield in the social sciences, so-called *agent-based computational social science* (sometimes called ABSS), also known as agent-based computational economics (ACE) in its economic variant. Today the face of artificial intelligence—primarily *distributed artificial intelligence* (DAI) and *multi-agent sytems* (MAS)—is being altered by incorporation of ideas of strategic behavior from game theory and economics. These research communities, ABSS and ACE on the one hand and DAI and MAS on the other, are in a state of coevolution.

Where this coevolution will lead is anyone's guess and is probably not even fathomable more than about one generation ahead. For such coevolutionary systems have the capacity to fundamentally alter one another and their environment in novel, creative ways. In what follows we offer, after recapitulating recent developments, likely near term advances. But the medium term and long term remain wide open. Speculations on such time scales may not differ substantially from science fiction.

II Social Science Models: More of the Same or More is Different?

Much of social science research today involves building models. This is true of both theoretical and applied work. In the former category, models are often completely symbolic and it is their mathematical or qualitative properties that are studied. In the latter case, models are based on some measured data.

Most models, both historically and today, are mathematical in character. The predominant way computation has been invoked in social science modeling is as a way to

---

[2] Interestingly, most of these people are much better known within either computer science or one of the social sciences. Only Simon himself is revered in each domain.

solve mathematical equations, primarily via numerical methods, although recent advances in symbolic analysis have been significant (Wolfram [1991]). Beginning in the late 1980s, particularly as microcomputer and workstation capabilities grew, DAI and artificial life (AL) modeling emerged in computer science. These approaches treated individuals—people, ants, viruses—as distinct data structures—eventually objects—within computer memory. This approach defeated two classical problems of conventional mathematical modeling in the social sciences: aggregation and equilibrium. Since social systems are typically composed of a large number of individuals, mathematical models in the social sciences have, in order to maintain mathematical tractability, essentially always been one of two types: (1) aggregate models where the heterogeneity of the actual population is either assumed away (e.g., representative agent models) or averaged away by only looking at mean behavior (e.g., systems dynamics models); (2) models written at the level of individuals in which 'solution' of the models involves all agents engaging only in equilibrium behavior (e.g., Nash equilibria in game theory) and all dynamic paths by which such equilibria might be achieved are neglected. It is clear how the agent approach fixes (1), by fully-representing individuals. The way it remedies (2) is by letting the agents interact directly, which usually amounts to out-of-equilibrium interactions, with equilibrium obtaining only if a path to it is realized from initial conditions.

DAI eventually grew into MAS in the mid 1990s, and these developments combined with AL gave rise to agent-based approaches in the social sciences (ABSS, ACE). As computer hardware increased in capacity exponentially, more sophisticated agent models could be built, utilizing either more cognitively complex agents or greater numbers of simple agents, or both. Thus it became possible to build more than 'toy' models, and soon large agent populations were realized in practice[3] leading naturally to the metaphor of an artificial society.

Summarizing, the unprecedented growth in computing power has led not just to more powerful numerical techniques, but has given rise to a new kind of computing based on autonomous agents.

Artificial Societies as a Research Paradigm

In artificial society modeling, a population of objects is instantiated and permitted to interact. Typically, each object represents one individual. These objects have internal data fields that store the specific characteristics of the individuals, things like preferences, endowments, goals and aspirations. The objects also have methods that both modify their internal data as well as describe how they interact. Each object also has some way to assess its own self-interest, i.e., it can rank the value to itself of alternative actions. This self-interestedness or purposefulness makes the objects into agents.

There are four main ways in which agent-based computing adds value to the social science modeling project. Conventional mathematical models in the social sciences rely heavily on a suite of heroic assumptions that are certainly false empirically and arguably do more harm than good as benchmarks. First, mainstream economics makes much of a 'representative agent', conceiving the entire economy as simply a scaled up version of a single decision-maker, or perhaps two agents interacting game theoretically. This specification is easy to relax computationally by assigning random values to some

[3] Models with millions of agents have been realized (Nagel and Pacuzski 1995, Axtell 1999), .

agent fields in accord with empirical data, when available, yielding a heterogeneous population. Second, it is the norm in economics to consider only rational agent behavior, whereby they are able to deduce the optimal behavior *not only for themselves but for all other agents as well*. Not surprisingly, in multi-agent systems having much complexity at all, such non-procedural specifications are essentially unimplementable in practice. Thus resort to bounded rationality is common and usually necessary. Third, modeling norms also dictate that agents do not interact directly with other individuals, but rather either indirectly through aggregate variables or perhaps through some idealized interaction topology (e.g., random graph, lattice). In agent computing, however, any topology, including empirically significant ones, can be easily implemented to serve as the basis for agent interactions. Finally, equilibrium serves as the focal point for all solution concepts in the social sciences. Whether equilibrium obtains or not in an agent system, the dynamics matter and are fully modeled.

Reaseach in the social sciences—particularly economics and to lesser extents in sociology and political science—involves (at least) two modes of inquiry. In economics these conventionally are called 'positive' and 'normative'; the former purports to describe economic phenomena while the latter offers suggestions for how the performance of economic mechanisms and institutions might be improved. Of course, the dichotomy is not so clear as this in practice, with much of what passes today as economic research first taking an apparently positive approach but quickly descending into normative speculation. This 'linear combination' of research modes cuts short all purely positive ambitions which in turn hamstrings its normative value. Research that utilizes empirically-false behavioral models, i.e., utility maximization, and then quickly draws 'policy conclusions' is quite close to what the physicist Fenyman [1983] called 'cargo cult science': efforts that have the look and feel of real science but which do not follow the basic principles of science. In the next section we review recent research in the social sciences that utilizes agent computing.

Early Results with Artificial Societies

Within each of the social sciences there exist more or less active research program using agent computing. While the nature of these applications is each somewhat idiosyncratic, they are unified methodologically in the search for agent specifications that yield empirically-observed (or at least empirically-plausible) social behavior.

In *anthropology* there are a variety of research groups active with agents. Many of these efforts are described in the recent Santa Fe Institute volume edited by Kohler and Gumerman [2000]. In this work the main idea is to utilize the typically very extensive empirical data that exist for specific archaeologically-significant regions as the target of multi-agent modeling. One then postulates and refines agent specifications that yield simulated historical trajectories that are 'close' to the data. This research program has been pushed farthest, perhaps, in the context of the American Southwest [Diamond 2002], where empirical data, especially on environmental histories, is particularly complete by virtue of progress in tree ring dating techniques. While models of such primitive societies would seem at first blush to have only passing relevance to advanced industrial societies, they serve at least two wider purposes: (a) as necessary precursors to models of more complex societies; (b) as significant for modern underdeveloped societies where people live very close to the land and environmental fluctuations are often responsible for major demographic shifts. Institutionally, within the American Association of Anthropologists there is a special interest group on computational modeling, with a strong representation of agent modelers.

In *geography* there has been rapidly growing interest in geographical information systems (GIS), another IT-facilitated technology. Much work with GIS to date simply amounts to animating recent or historical data on diverse geographical representations. However, recently there has been a move to link multi-agent systems type models to GIS, as exemplified by the recent Santa Fe Institute volume of Gimblett [2001]. The skill set of GIS researchers is quite similar to those required for agent computing, so there should result rapid co-evolution and perhaps coalescence of these research areas over the coming years. Indeed, with the American Society of Geographers there are now sessions at their annual meeting dedicated to just this intersection of techniques.

In *social psychology* the work of Latane *et al.* [1994], Nowak *et al.* [2000] and Kennedy *et al.* [2001] has amply demonstrated the power of the agent-based approach to social situations typically studied by psychologists. The latter takes a unique 'swarm' approach to cognition that is quite different from reigning approaches to learning.

In *sociology* a variety of more or less classical problems involving the tension between individual incentives and group behavior have been the subject of agent-based computational models. Many of these problems fall under the general rubric of *social dilemmas* (e.g., Macy and Flache [2002]. This is an instance of a general problem that arises across multi-agent systems, that of the micro-macro link. That is, if one knows the agent specification fully, what are the resulting societal properties? This problem is difficult in general but progress has been made in specific environments. A recent review article summarizing much of the multi-agent work in sociology is Macy and Willers [2002]. Early work includes Gilbert and Doran [1994] and Gilbert and Conte [1995].

In *political science* the early work by Axelrod [1984] can be seen as having kept the early flame of multi-agent systems modeling alive in the social sciences during the 1980s. Within the past five years or so there has been an explosion of interest, with a great range of models appearing, from state evolution through political party dynamics all the way down to game theoretic models of social dilemmas. In the related field of *political economy* there have been several papers utilizing agents; see Kollman *et al.* [1997] for a survey.

Economics has probably been the most active of the social sciences to date in applying agent computing. This probably is primarily due to the fact that methodologically individualism—writing models in terms of individual agents—has long been the norm among economics researchers, and because some economists have reasonably well-developed computing skills. The general areas of economics in which these models have been applied include to traditional markets (Kirman and Vriend [2002]) as well as financial ones (see next paragraph), to consumer behavior (Allen and Carroll [2001]), to social norms and conventions (Axtell and Epstein [1999], Axtell, Epstein and Young [2001], Young [1998]), labor markets (Tesfatsion [2001]), to the formation and evolution of firms (Axtell [1999], Luna [2000]), to public goods problems (Kollman *et al.* [1998]), to macroeconomics (Bullard and Duffy [2001]) and to international economics (Arifovic [1996]). Tesfatsion maintains an ACE website (www.econ.iastate.edu/tesfatsi/ace.htm).

Finance has been a particularly fertile area for the application of multi-agent systems modeling. In this area there are extensive datasets that are extremely useful in estimating and calibrating models. Pioneering work has been done by LeBaron [2001] and Lux [1998], among others. These models involve heterogeneous agents who make forecasts about future prices and take speculative positions in order to profit. These models have had significant success in explaining empirical data. Finance is also the area where the so-called *econophysicists* have begun to work, and there has been some convergence between their efforts and agent modeling (Farmer and Lo [1999]) The physicists wont is to work with highly idealized but tractable models, so they have made use of agents precisely at the interface of analytical intractability.

In *organizational science* the recent volume of Prietula, Carley and Gasser [2000] is a pioneering effort to bring agent computing to a field otherwise dominated by case studies and management science/operations research-type models. Here the basic approach is to take some organizational form or topology as given and populate each node of the organizational chart with adaptive agents. These agents glean information from their environment and pass messages up and down the organizational graph, resulting in organizational performance that is quantifiable and to some extent comparable to real-world organizations. These models have both positive and normative uses, since elaboration of such models usually leads to recommendations for how actual organizations can be improved.

In *business* there has been a variety of both academic and commercial work, with individual projects too numerous to list here comprehensively. Painting with a broad brush, there have been outright agent simulation projects of particular firms' operations or client bases, such as the highly detailed NASDAQ simulation (Darley and Outkin [2003]) and life insurance policy-holder behavior (Shumrak *et al.* [1999]). Agents have also been employed in more of a normative way, to better design logistical functions, such as cargo routing and job scheduling. The third way that agents have been utilized is commercially is as the basis for evolutionary tools that help engineers design better systems, whether physical (e.g., vehicles) or financial (e.g., risk management) ones. The recent articles by Bonabeau and Meyer [2001] and Bonabeau [2002] describe many of these efforts. In the near term it would appear that there will continue to be strong commercial demand for agent systems.

In the arena of *public policy*, where even simple computational models (e.g., spreadsheets) have revolutionized the practice of public administration, there has been limited adoption of agent computing to date. Recent projects include environmental resource management (Janssen [2002]) and a model of a large urban school district (Saunders-Newton [2002]). One interesting project at the local level, and funded with philanthropically, is the comprehensive agent-based planning model of Trewell, Vermont (Bernard). The software upon which this model is based has been generalized and attempts to commercialize it as a tool for town/city planning are underway. At the state and provincial level there are also some initial forays into agent modeling, including regional water quality management (Moss *et al* [2001]) and aspects of technological policy. At the national level there are a variety of agent models that claim some policy relevance. For example, several non-DOD US government agencies have adopted agent computing to better understand their client bases—typically some sub-group of U.S. citizens. Here agent computing typically competes with some more conventional modeling approach, such as statistics/econometrics. Many of the early uses of agents so far have involved circumstances where, for one reason or another, these tried an true methodologies are either weak or altogether inappropriate, such as when a regime change has occurred and little or no historical data are available to estimate empirical models. Thinks tanks—primarily my own—have proffered agent models as the basis for improving public policy (e.g., Axtell and Epsten [1999]). Despite the preliminary character of many of these efforts, there seems to be significant appetite for this highly visual and intuitive modeling technique in policy communities (Bourges [2002]).

There are a variety of applications of agents beyond the edges of conventional social science, For example, in *transportation science and policy*, the agent-based approach has become dominant methodological paradigm for modeling traffic flow (Nagel and Paczuski [1996]). The highways of entire cities (e.g., Albuquerque, Dallas-Fort Worth and Portland) have been simulated using the TRANSIMS code. Further, the view of transportation systems as complex adaptive systems, coupling mobility, air quality, and energy considerations has recently been articulated (Simon *et al.* [2003]) and this perspective leads naturally to agent computing. In *public health/epidemiology* the use of agents to more accurately model disease dynamics—taking into account realistic social networks for instance—has proven important in AIDS models, among others (Wayner [1996]) Here the limitations of conventional mathematical representations are clearly very severe and apparently the only way to make progress is through agents. Because transmission in urban environments is the dominant concern, there is even a version of the TRANSIMS code called EPISIMS for conducting large-scale modeling experiments. In *demography* the thesis that highly localized social norms of fertility importantly influence a population's overall demographics has recently received much attention and seems to be broadly confirmed empirically (Kohler [2001]). Finally, the *military* has been an early adopter of agent computing. A decade ago *equation-based* modeling of combat situations was essentially the only method in use, with vector supercomputer time being extensively utilized to study alternative rules of engagement, tactics, hardware, and so on. Since then there has been a transition to agent computing techniques for modeling such 'blue vs. red' interactions, with this process accelerating in the past five years particularly. These war-gaming environments, based solely by autonomous software agent players, can also be configured to support human players. A related kind of war-gaming is economic instead of lethal, and here there has also emerged a suite of agent-based tools to accomplish this (www.SEASLLC.com).

While in none of these fields is agent computing anything like a dominant paradigm today, in toto these many successful applications speak to the breadth and potential of the approach.

III Multi-Agent Systems: From Engineering to Social Science

My subject so far as been the ways in which agent computing is changing the practice of the social sciences by serving as a tool for social scientists. Here we describe how the social sciences are altering the face of AI/DAI/MAS.

Early in the evolution of DAI from AI, there emerged the need to decentralize agent systems in a meaningful way. Given the focus on decentralization within microeconomcs and general equilibrium theory, ideas from these fields were early on incorporated into DAI under the rubric 'market-oriented programming' (Wellman [1995]).

Since then there has been rapid realization that the rational model of human behavior could serve as a possible basis for building agent systems. Subsequently, game theoretic ideas have been systematically incorporated into MAS, and their importance has grown dramatically in recent years until perhaps as many has half of all papers in this area are based on these ideas (e.g., Sandholm [1999], Sen [2002], Shoham and Tennenholtz [1997]).

However, the ideas from economics and game theory that have been utilized by MAS researchers are not those emanating from the agent-based computing community in the social sciences, but rather older, largely rational and equilibrium ideas. For example, *mechanism design* refers to an approach to the synthesis of interaction environments in which the desired performance of a mechanism is specified and one then figures out what incentives to give the agents such that the (e.g., Nash) equilibria that are individually rational and incentive compatible achieve the objective. This formalism was developed largely in the 1980s and is today viewed by some as a viable way to design MAS (e.g., Parkes [2001]). Unfortunately, because such mechanisms are not credible behaviorally for humans—real people are not fully rational!—the utility of mechanism design principles is quite opaque. In lieu of using mechanism design ideas to create perfect societies, researchers in ABSS and ACE are much more likely to build behaviorally realistic systems and then study the performance of the artificial society wrt alternative policies in the environment in question. This seems both more practical and more like the way actual policy is formulated in the real world. Presumably MAS researchers can make good use of new ABSS and ACE research, and the latter can surely benefit from better understanding of what is being achieved by the former. But today these research communities are only tenuously joined by the interdisciplinary interests of a few individuals.

This out-of-step evolution between MAS proper and ABSS and ACE is not uncharacteristic of co-evolutionary environments in their early stages, and should soon be replaced by a more synchronous co-evolution.

IV Agents: Evolution or Revolution?

I am acutely aware that by portraying agent computing as something really new and different for the social sciences I am exposing these ideas to grave risks. Firstly, when outlining the features that distinguish a new methodology there is the risk that these distinguishing features will become, in essence, walls that serve to 'ghetto-ize' the new technique, effectively limiting the interactions between the old and new ways of working. Secondly, in touting the apparent effectiveness of a new methodology there is always the risk that hype will overtake substance, and expectations inflate beyond what is realistic. In caricature the method then looks like a panacea when it is unlikely to be one.

The alternative approach is to paint an evolutionary picture, where today's new methodology is seen as simply the logical extension of adequate but dated conventional methodology, representing natural progress instead of abrupt change. This view is more easily 'sold' to existing research communities and has an easier time insinuating itself into conventional discourse.

Evolution or revolution? Continuous or abrupt change? Smooth transition or phase transition? One is tempted to invoke Kuhn [1962] at this point and delve into the sociology of science. However, it is possible to abstract from such abstruse considerations in arguing on behalf of the latter interpretation—agents as revolutionary—by merely pointing out that the technical skills employed by those who foment the agent revolution today are quite different from those administered to today's graduate students, and thus those possessed by today's faculty members. A very small subset of social science researchers knows enough about computer science and information technology to actually perform agent-based modeling within their area of domain expertise. This fact is simultaneously the major barrier to the systematic adoption of these new techniques as well as the ultimate evidence that agents constitute a discontinuous advance.

Be that as it may, many open research questions remain, and much fertile ground between computer and information science and the social sciences remains to be ploughed. For example, an essentially virgin territory involves incorporating results from experiments with human subjects into agent-based models. Similarly, how best to statistically estimate agent models from aggregate data is today largely an open question. The ways in which *emergence* is dealt with between the two fields is quite different (see Axtell 2003), as are aspects of model verification and validation. Finally, aspects of model inter-comparision and robustness are just now coming to the fore and need to be better understood if this new paradigm is going to have significant staying power (Hales [2002]).

V Coevolution: Rise or Demise of the Social Sciences?

Assuming that Moore's law will continue to be realized for the next generation or so—say 20-30 years—agent computing will double in capabilities every 18-24 months.

From the social science side, this ostensibly exogenous technological revolution will permit the construction ever larger models involving ever greater numbers of more complex agents. When one contemplates the possible desktop hardware of 2020—just over 16 years from now, at least 8 doubling times and some $2^8 = 256$ times greater power—one imagines 500 gigabytes of ultra-fast RAM, clock speeds of at least 500 gigahertz, bus speeds of perhaps 100 gigahertz, and hard disks having capacity of 25 terabytes.

This continuing revolution in IT will fundamentally alter the kinds of social science models that can be built. It will also alter the practice of the social sciences, as equations give way fully to agents, empirically-tested cognitive models arise, and as decision models grounded in neuroscience emerge.

It will too forever change the face of computer science as seen through the guise of MAS, as more powerful social science models reflect back to MAS. Overall two distinct scenarios of coevolution seem plausible from my vantage point.

Should social science methodology adapt as agent computing evolves, one imagines the journals of 2025 filled with papers describing multi-agent models in which the theorems that can be proved in special cases relegated to appendices. However, if MAS methodology becomes so pervasive (through commercial deployment) that it substitutes for traditional private (e.g., commercial) or public (e.g., governmental) mechanisms and institutions, then analysis of such social structures will, of necessity, be analysis of MAS. This could lead to the displacement of traditional economics and other social sciences, turning them into a branch of computer and information science, in much the same way as library science has been suddenly transformed over the past decade into information science. How all this plays out is anyone's guess, but surely, in the style of the Chinese proverb, this is an exciting time to be alive.

References

Allen, T.W. and C.D. Carroll. 2001. Individual learning about consumption. *Macroeconomic Dynamics, 5*.

Anderson, P.W. 1972. "More is Different." *Science*.

Arifovic, J. 1996. "The Behavior of Exchange Rates in the Genetic Algorithm and Experimental Economies." *Journal of Political Economy, 104* (3): 510-541.

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

Axtell, R.L. 2003a. "Economics as Distributed Computation." In K. Takadama *et al.* eds. XXX. Springer-Verlag: Tokyo.

———. 2003b. "A Positive Theory of Emergence for Multi-Agent Systems." Working paper. Center on Social and Economic Dynamics. The Brookings Institution: Washington, D.C.

———. 2000. "Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences."

Axtell, R.L. 1999. "The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions." Working paper, Center on Social and Economic Dynamics, The Brookings Institution.

Axtell, R.L. and J.M. Epstein. 1999. "Coordination in Transient Social Networks: An Agent-Based Model of the Timing of Retirement." In H. Aaron, ed. *Behavioral Dimensions of Retirement Economics*. Brookings Institution Press: Washington, DC

Axtell, R.L., J.M. Epstein and H.P. Young. 2001. "Emergence of Class Norms in a Multi-Agent Model of Bargaining." In S. Durlauf and H.P. Young, eds., *Social Dynamics*. MIT Press: Cambridge, Mass.

Bankes, S. 2002. "Agent-Based Modeling: A Revolution?." *Proceedings of the National Academy of Sciences USA, 99*: 7199-7200.

Bernard, R. No date. "Policy Simulator: A Decision Support System for Local Government." Working paper. CommunityViz.

Bonabeau, E. 2002. "Agent-Based Modeling: Methods and Techniques for Simulating Human Systems." *Proceedings of the National Academy of Sciences, 99*: 7280-7287.

Bonabeau, E. and C. Meyer. 2001. “Swarm Intelligence: A Whole New Way to Think About Business.” *Harvard Business Review*.

Bourges, C. 2002. "Computer-Based Artificial Societies May Create Real Policy." *Washington Times* (May 12).

Bullard, J. and J. Duffy. 2001. Learning and excess volatility. *Macroeconomic Dynamics, 5*: 272-302.

Cyert, R.M. and J. March. 1963. A Behavioral Theory of the Firm. Prentice-Hall: Englewood Cliffs, NJ

Critchon, M. 2002. *Prey*. Harper Collins: New York.

Darley, V. and S. Outkin. 2003. Bios technical report. Santa Fe, NM.

Diamond, J.M. 2002. “Life with the Artificial Anasazi.” *Nature*, vol 419: 567-569.

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

Farmer, J.D. and A.W. Lo. 1999. “Frontiers of Finance: Evolution and Efficient Markets.” *Proceedings of the National Academy of Sciences USA, 96*: 9991-9992.

Feynmann, R. 1983. *Surely You’re Joking, Mr. Feynman.*

Gilbert, N. and J. Doran, eds. 1994. *Simulating Societies: The Computer Simulation of Social Phenomena*. UCL Press: London.

Gilbert, N. and R. Conte, eds. 1995. *Artificial Societies: The Computer Simulation of Social Life*. UCL Press: London.

Gimblett, R. ed. 2001. *Integrating Geographic Information Systems and Agent-Based Modeling Techniques for Understanding Social and Ecological Processes*. Oxford University Press: New York, NY.

Hales, D. *et al*. 2002. Model to Model Workshop (cfpm.org/m2m).

Janssen, M., ed. 2002. *Complexity and Ecosystem Management: the theory and Practice of Multi-Agent Systems*. Elgar: London.

Judd, K. 1999. *Numerical Methods in Economics*. MIT Press: Cambridge, Mass.

Kennedy, J., R.C. Eberhardt and Y. Shi.. 2001. *Swarm Intelligence*. Morgan Kaufmann.

Kirman, A.P. and N.J. Vriend. 2000. Learning to be loyal: A study of the Marseille Fish Market. In D. Delli gatti, M. Gallegati and A.P. Kirman, eds., *Interaction and Market Structure: Essays on Heterogeneity in Economics*. Springer: Berlin.

Kohler, H.P. 2001. *Fertility and Social Interactions*. Oxford University Press: New York.

Kohler, T.A. and G.J. Gumerman. 2000. *Dynamics of Human and Primate Societies: Agent-Based Modeling of Social and Spatial Processes*. Oxford University Press: New York, NY.

Kollman, K., J. Miller, and S. Page. 1997. “Computational Political Economy.” In Arthur *et al.*, eds.

Kollman, K., J. Miller, and S. Page. 1998. “Political Institutions and Sorting in a Tiebout Model” *American Economic Review*.

Kuhn. T. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press: Chicago.

Latane, B., A. Nowak and J.H. Liu. 1994. “Measuring Emergent Social Phenomena: Dynamicsm, Polarization and Clustering as Order Parameters of Social Systems.” *Behavioral Science,* vol. 39: 1-24.

LeBaron, B. 2001. “Evolution of Time Horizons in an Agent-Based Stock Market.” *Macroeconomic Dynamcis, 5*: 225-254.

Luna, F. 2000. “Induction and Firm Creation.” In F. Luna and B. Steffanson, eds. *Multi-Agent Simulation in Economics and Finance using SWARM*. Kluwer Academic.

Lux, T. 1998. The Socio-Economic Dynamics of Speculative Markets: Interacting Agents, Chaos, and the Fat Tails of return Distributions. *Journal of Economic Behavior and Organization, 33*: 143-165.

Macy, M.W. and A. Flache. 2002. “Learning Dynamics in Social Dilemmas.” *Proceedings of the National Academy of Sciences USA, 99* (10): 7229-7236,

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

Moss, S. *et al*. 2001. FIRMA: Freshwater Integrated Resource Management with Agents. Working paper. Manchester Metropolitan University: Manchester, UK.

Nagel, E. and M. Paczuski. 1996. "Emergent Traffic Jams." *Physical Review E*...

Nowak, A., R.R. Vallacher, A. Tesser, and W. Borkowski. 2000. “Society of self: The emergence of collective properties in self-structure.” *Psychological Review*, vol. 107, 39-61.

Parkes, D. 2001. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. Department of Computer and Information Science, University of Pennsylvania: Philadelphia, Penn.

Prietula, M., K. Carley and L. Gasser, eds. 1998. *Simulating Organizations*. MIT Press: Cambridge, Mass.

Rauch, J. 2002. "Seeing Around Corners." *The Atlantic Monthly* (April).

Sandholm, T. 1999. In Weiss, ed. *Multi-Agent Systems*.

Saunders-Newton, D. 2002. Working paper. University of Southern California.

Sen, S. 2002. "Believing Others: Pros and Cons." *Artificial Intelligence, 142* (2): 179-203.

Shoham, Y. and M. Tennenholtz. 1997. "On the emergence of social conventions: modeling, analysis and simulations. *Artificial Intelligence, 94* (1): 139-166.

Shumrak, M., M. Greenbaum, V. Darley and R. Axtell. 1999. Modeling Annuity Policyholder Behavior using Behavioral Economics and Complexity Science. Canadian Institute of Actuaries Annual Meeting.

Simon, C.P. 2003. Working draft. University of Michigan. Ann Arbor, Mich.

Simon, H.A. 1997. *The Sciences of the Artificial*. 3rd ed. MIT Press: Cambridge, Mass.

Tesfatsion, L. www.iastate.edu.

Tesfatsion, L. 2001. "Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search." *Journal of Economic Dynamics and Control, 25*: 419-457.

Wayner, P. 1996. "Computer Simulations: New-Media Tools for Online Journalism." Online at www.wayner.org/texts/aids.

Weiß, G. 2000. *Multi-Agent Systems*. MIT Press: Cambridge, Mass.

Wellman, M. 1995. "Market-Oriented Programming: Some Early Lessons." In S. Clearwater, ed., *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific.

Wolfram, S. 1991. *Mathematica: A System for doing Mathematics by Computer*. 2nd ed. Addison-Wesley: New York.

Young, H.P. 1998. *Individual Strategy and Social Structure*. Princeton University Press: Princeton, NJ. 60