

Effects of Interaction Topology and Activation Regime in Several Multi-Agent Systems

Robert Axtell
Center on Social and Economic Dynamics
The Brookings Institution
1775 Massachusetts Avenue, NW
Washington, D.C. 20036 USA
raxtell@brook.edu

Abstract

The effects of distinct agent interaction and activation structures are compared and contrasted in several multi-agent models of social phenomena. Random graphs and lattices represent two limiting kinds of agent interaction networks studied, with so-called 'small-world' networks being an intermediate form between these two extremes. A model of retirement behavior is studied with each network type, resulting in important differences in key model outputs. Then, in the context of a model of firm formation, in which multi-agent structures (firms) are emergent, it is demonstrated that the medium of interaction—whether through individual agents or through firms—affects the qualitative character of the results. Finally, alternative agent activation 'schedules' are studied. In particular, two activation modes are compared: (1) all agents being active exactly once each period, and (2) each agent having a random number of activations in every period with mean 1. In many circumstances these two regimes produce indistinguishable results at the aggregate level, but in certain cases the differences between them are significant.

I Introduction¹

One class of multi-agent systems (MAS) consists of a relatively small number of agents, each of whom has relatively sophisticated behavior (e.g., a rich cognitive model, perhaps for dealing with a complex task environment [18]). A different type of MAS involves relatively large numbers of behaviorally simple agents. This second family of multi-agent systems is of significant interest as the basis for empirically-relevant models of human social and economic phenomena. Such models typically involve the use of aggregate social or economic data to estimate parameters of a MAS in which agents have heterogeneous internal states (e.g., preferences) but a common repertoire of behaviors (e.g., economic exchange).

One reason for the elevated attention given to simple agents is that the prevailing norm in the mathematical

social sciences is to build models that abstract from the details of cognition.² Stated differently, the focus of economists and other quantitative social scientists on behaviorally simple models is a symptom of the lack today of anything like a universal model of cognition. A second reason for differential interest in models composed of moderate or large numbers of simple agents is that such systems are quite capable of complex aggregate behavior, involving, for example, the spontaneous emergence of behavioral norms (e.g., [19]) or the self-organization of multi-agent coalitions (e.g., [3]). Understanding the origin of these complex patterns of emerged behavior is often a significant challenge, and would be even more difficult if individual agents were complex in their own right—if individual decisions were also emergent.

Given the relative simplicity of individual agents in such systems, it is almost certainly true that model specifications beyond the individual level play a somewhat more important role in such models than in MAS involving few agents. In particular, the interconnections between agents—the interaction topology—and the relative amount of individual activity in the agent population—the agent activation regime—surely must matter in wide varieties of models, especially to the extent that such models have empirical ends.

In this paper it is demonstrated that these factors—interaction topology and activation regime—can be crucially important in multi-agent systems, illustrated through a variety of empirically-oriented models of social phenomena. Specifically, when structures of interaction and activation are systematically altered, the aggregate statistics produced by such models can vary substantially. The next section addresses interaction networks in two multi-agent models. First, an existing multi-agent model of retirement dynamics, in which social networks are dynamic, random graphs, is modified to have lattice-type networks in the space of age cohorts. This is shown to systematically alter the overall behavior of the agent society as measured by the time required for establishment of a social norm in

¹ Thanks are due Michael Cohen for stimulating my interest in this subject, and to workshop participants at the Program for the Study of Complex Systems at the University of Michigan, for many useful comments.

² There are many reasons for this state of affairs. The fact that cognitive psychology has focussed little on economic behavior [15] is certainly one reason, remediable in practice. That most economists have little training in cognitive science represents a larger challenge.

retirement age. Second, in a model of endogenous firm formation, where agents learn of alternative employment opportunities through social networks, a key statistic of the model is shown to depend crucially on the structure of these networks. Then, in section three we study the effect of two asynchronous agent activation regimes: *uniform activation*, in which all agents are active *exactly* once each period, and *random activation*, in which agents are active once a period *on average*. The effect of changing regimes is described in multi-agent models of trade, cultural transmission, and firm formation. A final section draws conclusions.

II Effect of Agent Interaction Topologies

Social networks play a critical role as the medium within which human beings are socially situated and through which interactions between individuals occur. Therefore, it is hardly surprising that positive models of human social processes should include such networks. What is perhaps surprising is the extent to which relatively small changes in network structure can lead to large changes in macro-social outcomes. Experiments with human subjects involving systematic changes to social networks are difficult to perform, of course. Here we utilize multiagent systems in their role as social science laboratories, since such models allow us to methodically alter such networks and to then discover, by spinning the models forward in time, the overall effects of such alterations. In this section, two models—one on the dynamics of retirement norms and the other a multi-agent model of firm formation—have their social networks systematically altered. In the first case an extant model employing random graph networks is morphed to a lattice configuration, and then to a so-called ‘small world’ graph.³

Recently, procedures for creating ‘small world’ graphs have been introduced in [17]. Analysis of such graphs demonstrate that they possess a form intermediate between regular and random graphs, and have many properties of real-world social networks. Specifically, start with a d -dimensional lattice having a fixed number of edges. Then, systematically break each edge with some probability, p , and reattach each broken edge to a random node. There results a graph having a well-defined sense of ‘localness’ as in a regular graph (i.e., agents who know each other have significant overlap in their individual social networks), as well as relatively short paths between any two agents in the graph, a characteristic feature of random graphs. These ‘small world’ graphs are thought to be relatively good models of certain social networks (cf. [11], [16]).

Coordination in Transient Social Networks: A Model of the Timing of Retirement

In [5] a model of retirement behavior is described in which there is a population of agents of various ages, with each agent having to decide when to retire. The model

abstracts from economic factors in attempting to explain a certain puzzle in the evolution of the modal age of retirement—namely the long lag between the last big change in benefits policy and the systematic change in the overall behavior of retirees. Specifically, in 1961 an ‘early retirement’ age of 62 was instituted by the U.S. Social Security Administration. Benefits received by age 62 retirees were reduced in comparison with age 65 recipients, by an ‘actuarially neutral’ amount—typical people would be, it was thought, indifferent between retiring at any age between age 62 and 65. The puzzle is that in 1961 the modal retirement age was 65, and this remained so up through 1990. Only by 1995 had this shifted, rather abruptly, to age 62. Standard rational actor accounts of retirement decision-making have difficulty explaining the long lag between changes in benefits and responses in overall behavior.

Our model of the retirement process employs a heterogeneous population in which there are three kinds of agents:

1. ‘rationals’ retire at the earliest age permitted by law;
2. imitators play a coordination game in their social networks—if the fraction of agents who are retired among those eligible for retirement within an agent’s network exceeds a threshold, t , then the agent too retires, else it continues working;
3. randomly-behaving agents retire with some fixed probability, p , once they are eligible to do so.

This simple model is capable of reproducing certain features of retirement data from the United States, particularly a relatively long lag time from the change in earliest retirement age until its establishment as a behavioral norm in the population.

During each period of model execution, each agent ages one year and gets to decide whether or not to retire (there is also some chance of dying). The behavior of the first and third agent types is straightforward, while that of the second type, the imitators, can be characterized using game theoretic notions.⁴ Consider a population of A agents, in which the state of agent i is $x_i \in \{\text{working}, \text{retired}\}$. Then the state of society is given by $x \in \{\text{working}, \text{retired}\}^A$. Agent i has a social network, consisting of a set, N_i , of other agents. Overall, the utility that agent i derives, U_i , from interacting with members of its network is given by

$$U_i(x) = \hat{A} \sum_{j \in N_i} u(x_i, x_j)$$

where the $u(x_i, x_j)$ can be thought of as payoffs in a 2 x 2 game, as in table 1.

	working	retired
working	w, w	$0, 0$
retired	$0, 0$	r, r

Table 1: Imitation as a coordination game

An agent’s imitation threshold, t , can be stated in terms of these payoffs, given the graph weights in its social

³ An illustration of the importance of social network structure in several game theoretic models can be found in [13].

⁴ Formal results for games of this type played on static graphs are developed in [20].

network. In the case of uniform weighting—i.e., the behavior of all neighbors considered equally— $t = w/(r+w)$; for $r = w$, $t = 1/2$. Furthermore, if the payoffs are made heterogeneous in the population then the imitation thresholds become heterogeneous as well.

For the imitators in the model, their networks are random graphs in [5].⁵ Note that although each agent has a fixed network, the overlapping generations character of the population, in which old agents die and new, younger agents are born, renders the overall networks in society as transient. Here we study the effect of first moving to regular graphs of a certain type and then to small world graphs. Since human social networks tend to be highly correlated with age—i.e., people’s designations of their friends are often dominated by people within a few years of their own age—the regular graph we employ is localized in the cohort space. That is, each agent has friends who are near it in age.

The key statistical output of this model is a measure of how long it takes a social norm of age 65 retirement to establish itself from an initial condition of no retirement age norm. In [5] the way in which this measure depends on various parameters of random graphs—e.g., their size, their extent in the cohort dimension, their heterogeneity—is studied. Here we investigate how altering the network structure to lattice social networks in the cohort dimension modifies the time required for establishment of this social norm. When populations are homogeneous and interactions occur on regular graphs, it is known analytically that the transition time between equilibria can be sped up through local interactions [7]. But for heterogeneous populations on transient graphs little is known analytically about waiting times, the lifetimes of transients, and so on.

In order to facilitate comparisons, the size and extent of agent networks is kept fixed as we vary network type. These and other parameters are described in table 2. The agents are parametrically homogeneous here, but each agent has a unique social network and therefore the population of agents is heterogeneous. These parameters differ somewhat from the ‘base case’ described in [5], involving larger social networks that have less extent in the cohort dimension (‘network age extent’ of 2 means that an agent’s network includes agents who are within 2 years of its own age).

Parameter	Value
Agents/cohort	100
Imitation threshold, t	50%
Social network size	24
Network age extent	2
Random agents	5%
p	0.50

Table 2: Configuration of the retirement model

This set of parameters yields a model in which the typical time required to establish an age 65 retirement norm is somewhat greater than in the base case in [5], due to the

larger social networks. The rightmost line in figure 1 below gives these times as a function of the fraction of rational agents in the population, holding the number of randomly behaving agents constant at the level of table 2. Note that for approximately 15% rational agents, the average transition time is around 100 years. The bars above and below the average values in the figure represent ± 1 standard deviation; these are asymmetrical because the ordinate is in logarithmic coordinates.

Running this model demonstrates that the state of being ‘retired’ percolates upward from below—from older cohorts—as it were.⁶ The model behaves as if retirement were *diffusing through* or *infecting* the population, from older to younger individuals. In effect, the agents are an *excitable media* through which retirement behavior can spread, with agent interaction serving to speed the adoption of retirement while at the same time the aging of the population and the inherently transient nature (through die-off) of all social networks acts in the opposite direction, to limit adoption. Overall there is continual ebb and flow of the retirement state through social networks, until eventually it takes hold among essentially all agents capable of being in such a state. This uniform adoption of retirement behavior is best understood as the emergence of a social norm of uniform retirement age. While such norms, once established, can be destabilized by chance events, in general they are quite robust.




Figure 1: Time until establishment of age 65 retirement norm, as a function of social network configuration

These results can be compared with those that obtain for lattice-type regular social networks. These are also depicted in figure 1, as the leftmost line. For a given fraction of the population acting rationally, lattice networks clearly require significantly less time for social norms to arise than random graph networks. The effect is dramatic at small levels of rationality, where random graphs would essentially take forever to lock in to a social norm, but lattice networks may require only a few decades. Equivalently, for a specified transition time a much smaller fraction of the population needs to be rational when social networks are lattice-like.

⁵ For analysis of the properties of random graphs see [6].

⁶ A typical realization of this model can be found in QuickTime™ format at www.brook.edu/es/dynamics/papers/interaction.

Typical model realizations⁷ reveal that this model has a different dynamic for the establishment of retirement norms. In contrast to retirement percolating from older to younger cohorts, with lattice social networks retirement behavior starts among age 65 agents, first in a small group, and grows outward through that cohort, then down through the population as it ages. Visually it is clear that this is a very different mechanism than with random graph networks.

As has been argued briefly above, neither random nor regular graphs well-represent empirically-significant social networks. In random graphs there is no sense of location, while in lattices the 'social distance' between two agents can be very large. Rather, real-world social networks seem to have the 'small world' property, i.e., networks are localized yet the path length between any two individuals is not large. From [17] we know there is a well-defined sense in which we can move from lattices to random graphs by randomly selecting links to be broken and randomly reattached. Such a process yields small world graphs. Figure 1 also reports results for the dynamics of retirement in such small world social networks, with two different values of the probability of breaking and randomly reattaching a link, p . For $p = 10\%$, the social networks retain their lattice look, but now there exist much shorter paths between any pair of agents. The establishment of retirement norms looks much like the lattice social network model in this case, although the transition times are somewhat longer. In the case of $p = 25\%$, the paths between arbitrary pairs of agents are even shorter. Here, the fraction of rational agents necessary to achieve a specified transit time to a retirement age norm is about halfway between the pure lattice and random graph social network cases. Clearly, these small world graphs behave as an intermediate form between regular graphs and random ones.

The Emergence of Firms

In [3] a model is developed in which heterogeneous, self-interested agents form groups in a team production environment with increasing returns. Each agent has preferences for income, gained from working, and for leisure (all time not spent working). Nash equilibrium effort levels in a group are Pareto-dominated by higher effort levels, although these are not individually-rational. The main analytical result for this strategic situation is that there is a size beyond which any group is unstable. That is, there is a maximum stable group size for any distribution of preferences. As groups exceed their maximum size, agents are permitted to join other groups or to start up new groups if it is welfare-improving to do so. It turns out that meta-stable groups—temporary coalitions—of agents can survive in this model out of equilibrium for significant lengths of time. We have studied such transient groups via a multi-agent system.

Interestingly, many of the statistical features of these groups closely resemble what is known about the

population of firms. In particular, the size distribution of such groups is highly skewed, approximating a power law (Pareto distribution) in the upper tail, a well-known property of firms in industrial countries. Second, the growth rate distributions in the model are closely related to those found empirically for U.S. firms—essentially, non-normal distributions with heavy tails. Third, the way in which growth rate variance depends on firm size in the model can be made almost identical to the data (more on this later). Fourth, it is something of an empirical puzzle that wages tend to increase with firm size. This phenomenon can be found in the model, but only for particular interaction topologies. To understand how this is so, it is necessary to consider the details of agent decision-making.

When an agent is activated in this model, it assesses how its utility could be increased by altering its effort level. Perhaps several new agents have joined the group since it last re-evaluated its effort, or maybe other agents have systematically altered their effort contributions to production. This assessment by the agent could involve utility maximization, taking as given other agents' behaviors or taking into account the reactions of others to its own change in effort level. Alternatively, it could be simply utility-improving through a process of groping for better effort levels. However potential utility increases are determined, the agent stores these new efforts and utilities for comparison with other options, including joining other firms as well as starting up a new firm on its own. For each of these options the agent determines effort levels that improve its utility. But which extant firms do agents consider? There are many ways to do this. First, an agent could simply pick a firm at random in the population of firms. Alternatively, it could pick an agent at random from the overall population and consider joining its firm. Similarly, it might carry around with it a social network of 'friends' and each period consider joining the firms of its friends. Or, perhaps most realistically, firms that can profit most from hiring could post 'ads' in a virtual newspaper in order to attract potentially interested agents.

Clearly, these are quite varied ways of selecting prospective employers, and there is no obvious reason why they should yield the same results, especially given the highly skewed distribution of firm sizes. That is, due to the size distribution skewness, sampling a random firm is very different from selecting an agent at random. The former process produces a small firm with high probability—median firm size is under 10, the mode is 1 or 2—while the latter more frequently samples larger firms. Interestingly, it turns out that most of the empirical features of this model are robust, qualitatively, to such variations, i.e., size distributions remain skewed, growth rate distributions have fat tails, and growth rate variance scales with size. However, one empirical feature is sensitive to the structure of interactions, the wage-size effect. In figure 2 the dependence of wages on size is shown for two kinds of networks, one in which random agents are selected (lower line), and one based on choosing new firms (upper line). It is clear from this figure that there is little wage-size effect in the former

⁷ Also available at www.brook.edu/es/dynamics/papers/interaction.

case (i.e., random agents), while wages increase with size in the latter case. Empirically wages are $\propto \text{size}^{0.10}$, while the upper line in figure 2 describes an almost identical relationship.




Figure 2: Effect of firm size on wages, search networks based on new firms (upper) and random agents (lower line)

This clear difference in wage-size effects that results from changing the networks in which search is conducted is due to the different kinds of prospective employers yielded by the sampling process, as alluded to above. Due to increasing returns, growing firms in this model have high output per employee (productivity) and create high utility for their workers. When job-seeking agents can 'see' such firms, as with firm-based search networks, then output per worker grows with size. Alternatively, if such firms are infrequently sampled, as with random agent networks, then the 'arbitrage' in marginal utility that takes place in this model as agents migrate between firms yields constant output per agent across firms.

III Effect of Agent Activation Regime

By agent activation we refer to the order and frequency of agent action. Because real social processes are rarely synchronous, modeling them by permitting each agent to update its states exactly once each period, using last period's state information, is usually inappropriate.⁸ Two types of asynchronous activation regimes are commonly employed in large population, simple agent MAS described in this paper. In *uniform activation* a period is defined as the time during which all agents are active once and only once. In *random activation* each agent has an equal *a priori* probability of being active each period, but stochastic variations lead some agents to be more active than others. We study these in turn below.⁹ In general, we can expect these alternative execution regimes to yield different individual agent histories, and possibly different macro-social outcomes as will be illustrated below. Indeed, even for a given type of activation it is only under very

specialized conditions that distinct agent updating histories will yield invariant histories of agent states [9].

Uniform Activation

In uniform activation, each agent is activated once per period. A possible problem with this execution regime is that if agent i always acts immediately before agent j over the course of many periods then there exists the possibility that correlation between these agents will develop that is unrelated to the agent behavioral rules, but is rather a direct consequence of the activation structure. We call this *spurious* agent-agent correlation and suggest that in most cases it should be considered a programming artifact—an unintended consequence or side effect—and avoided. Happily, by randomizing the *order* of agent activation from period to period it is usually possible to remove all artifacts of this type and thus avoid any spurious correlation. Stated differently, with uniform activation it is crucial to periodically randomize the order of agent updating.

But how much randomization is appropriate? That is, given the order in which agents were serially activated last period, how many agents need to be repositioned in this sequence so that in the next period most of the agents either precede or follow a different agent—i.e., most agents have one or more new neighbors in the execution sequence? We have tried to answer this question analytically, without success.¹⁰ To get some feeling for the difficulty of this question, consider the simple case of a single pair of randomly selected agents who are swapped in the agent activation list. How many agents will have 1 new neighbor? How many will have 2 new neighbors? There are three cases to consider:

1. The agents selected to be swapped are neither neighbors nor have any neighbors in common, so rearranging them will give each of their 4 previous neighbors 1 new neighbor. Therefore, the fraction of the population that has exactly 1 new neighbor is $4/A$. Furthermore, each of the agents who moved have 2 new neighbors.
2. The agents to be swapped are immediate neighbors, so the process of rearranging them yields 4 agents who have exactly 1 new neighbor, and no agents with 2 new neighbors.
3. The agents to be moved have 1 neighbor in common, so the rearrangement process also produces 4 agents with exactly 1 new neighbor, and no agents with 2 new neighbors.

Now, to figure out the probabilities of having exactly 1 or 2 new neighbors it is necessary to determine the relative probabilities of these 3 cases. Then, the case of rearranging 2 agent pairs (4 agents total) involves 3×3 cases: the first agent pair generates the three cases above, and for each one of these the next agent pair creates 3 more cases. For an arbitrary number of rearrangements, this quickly becomes a

⁸ An important exception to this statement are multi-agent models of traffic. In [10] a critique of the use of synchronous updating in certain game theoretic models is rendered, demonstrating that results from such models are typically not robust to asynchronous updating.

⁹ The case in which agents have differential incentive to be active is treated in [12].

¹⁰ This problem is very similar to the card shuffling problem—how much shuffling is sufficient to satisfactorily "mix" a deck?

Figure 4: Probability that a particular agent in a population of size A has been inactive over T periods when k agents are activated at once

As T increases beyond A , that is, $T \gg A$, the mean number of interactions rises together with the variance in proportion to T , while the skewness coefficient approaches 0; for $A \gg k$ skew vanishes like $1/(2T)$.

Example: *Distribution of the number of interactions per agent in bilateral exchange processes*

Consider an economy in which agents are randomly paired to engage in bilateral exchange, a single pair of agents trades each period ($k = 2$), and the probability that an agent is part of a trading pair is uniform across the population. For 100 agents ($A = 100$) over 1000 periods ($T = 1000$) the probable number of activations per agent is shown below in figure 5.




Figure 5: Probability mass function for the number of agent activations in a population of 100 agents over 1000 periods when agents are paired sequentially with uniform probability

The mean number of activations in figure 5 is exactly 20, although the modal event of 20 activations is only slightly more probable ($p = 0.0897$) than 19 ($p = 0.0896$). The standard deviation of this distribution is 4.43, and it is skewed slightly, with a 'fatter' right tail, as evidenced by a skewness coefficient of 0.217.

Clearly, uniform and random activation represent very different models of agent activity. In a certain sense, uniform activation is the zero variance limit of random activation—the number of activations per agent has no variance in the case of uniform activation. In the next subsection the differential effects of these two activation regimes are compared. But before moving on, a third activation regime will be briefly described.

Imagine that each agent has its own 'Poisson clock' that wakes it up periodically in order to be active, such that the probability of its being active over a single time period is k/A . Then, over time T each agent is active kT/A times on average, just as in random activation, with the variance equal to the mean. Here, however, during any period of time the total number of agent activations is a random variable, and the random activation model described above is a kind of zero variance limit of random activation of the Poisson variety. There is a clear connection between these two types of random activation, since the Poisson distribution closely approximates the binomial for large values of A and T .

Comparison of Activation Regimes

The so-called Sugarscape model [8] is a multi-agent system designed for the study of demography, economics, disease propagation, and cultural transmission, on spatial landscapes. It utilizes uniform activation with execution order being randomized each period. A somewhat more elaborate model for the transmission of culture that uses random activation has been studied by Axelrod [2]. A 'docking experiment' was undertaken in which the culture component of the Sugarscape MAS was extended to incorporate Axelrod's model, in hopes of being able to reproduce the latter's somewhat counter-intuitive results. As described in [4], initial attempts to 'dock' the two distinct multi-agent systems yielded qualitatively similar results despite different activation regimes. However, statistical tests rejected the hypothesis that the data from the two models were the same. Only by altering the Sugarscape execution model to random activation was it possible to quantitatively 'dock' the two models. So this is an example of agent activation regime having a measurable statistical effect, although not altering the qualitative character of the output.

Returning now to the model of firm formation proposed in [3], one of the most striking regularities in the empirical data is that the variance in $\log(\text{growth rates})$ decreases with firm size. Large firms simply have less variation in their growth rates than do small firms. Figure 6 below reproduces a figure in [3] that describes the close quantitative agreement between the model and the data on U.S. firms—essentially, the standard deviation in the logarithm of growth rates is a power law in firm size, with exponent approximately equal to $-1/6$. (The data scatter at large firm sizes is due to small sample sizes.)




Figure 6: Effect of firm size on variance in the logarithm of growth rate with random activation

The basic version of this model uses random activation. In the course of investigating the effect of alternative activation regimes on the model it was discovered that uniform activation (with 50% randomization each period) generates the reverse dependence of growth rates on size—variance *increases* with size. This is shown in figure 7.




Figure 7: Effect of firm size on variance in the logarithm of growth rate with uniform activation, 50% randomization each period

Here the slope is approximately $+1/6$, just the inverse of random activation. Such data could hardly be in greater disagreement with the facts. The reason why this happens is that when large firms begin to decline—and it is a theorem in the model that all firms eventually decline—the process is hastened by uniform activation. Once a few agents realize there are better opportunities elsewhere, other agents in the firm progressively learn this. Such ‘unraveling’ processes are more diffuse with random activation—those agents inactive during a period of decline in their firm look up from their desks one day to find things in bad shape, and only then do they consider leaving. The random activation model implies more heterogeneous rates of sampling of external job opportunities among the agents within a firm.

IV Conclusions

The effects of agent interaction topology and agent activation regime have been investigated in several multi-agent systems. In moving between regular graphs (lattices) and random graphs, through small world-type graphs, the overall behavior of a model of the timing of retirement changed significantly. Then, altering the qualitative character of social networks in an empirically-accurate model of firm formation caused the wage-size effect in the model to disappear.

Two distinct agent activation schemes were compared and contrasted, uniform and random activation. These produced qualitatively similar but statistically different output models of cultural transmission. In the firm formation model, random activation yields empirically-significant results. Moving to uniform activation generates qualitatively different and unrealistic output.

Clearly, agent interaction and activation structures can play important roles in MAS. Careful consideration of these ‘architectural details’, including systematically studying how model output changes as they are varied, must be a key to robust analysis of multi-agent systems.

References

1. Aaron, H., ed. 1999. **Behavioral Dimensions of Retirement Economics**. Russell Sage Foundation/Brookings Institution Press: N. Y., N.Y./ Wash., D.C.
2. Axelrod, R. 1997. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." **Journal of Conflict Resolution**, **41**: 203-226.
3. 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 03-019-99. Santa Fe Institute: Santa Fe, N.M. Available at www.brook.edu/es/dynamics/papers.
4. Axtell, R.L., R. Axelrod, J.M. Epstein and M.D. Cohen. 1996. "Aligning Simulation Models: A Case Study and Results." **Computational and Mathematical Organization Theory**, **1** (2): 123-41.
5. Axtell, R.L. and J.M. Epstein. 1999. "Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement." In Aaron, ed. [1]. Available in working paper form at www.brook.edu/es/dynamics/papers.
6. Bollobás, B. 1979. **Graph Theory**. Springer-Verlag: N.Y.
7. Ellison, G. 1993. "Learning, Local Interaction, and Coordination." **Econometrica**, **61**: 1047-71.
8. Epstein, J.M. and R. Axtell. 1996. **Growing Artificial Societies: Social Science from the Bottom Up**. MIT Press/Brookings Institution Press: Cambridge, Mass./Washington, D.C.
9. Gács, P. 1997. "Deterministic Computations whose History is Independent of the Order of Asynchronous Updating." Working paper. Boston Univ.: Boston, Mass.
10. Huberman, B.A. and N.S. Glance. 1993. "Evolutionary Games and Computer Simulations." **Proceedings of the National Academy of Sciences USA**, **90**: 7716-18.
11. Kochen, M., ed. 1989. **The Small World**. Ablex Publishing Corporation: Norwood, N.J.
12. Page, S.E. 1997. "On Incentives and Updating in Agent Based Models." **Computational Economics**, **10** (1): 67-87.
13. ———. 1999. "Network Structure Matters." Working paper. Department of Political Science. University of Michigan.
14. Rubinstein, A. 1998. **Modeling Bounded Rationality**. MIT Press: Cambridge, Mass.
15. Simon, H.A. 1998. In Rubinstein [14].
16. Watts, D. 1999. **Small Worlds: The Dynamics of Networks Between Order and Randomness**. Princeton University Press: Princeton, N.J.
17. Watts, D.J. and S.H. Strogatz. 1998. "Collective Dynamics of Small-World Networks." **Nature**, **393**: 440-442.
18. Weiß, G., ed. 1998. **Multi-Agent Systems**. MIT Press: Cambridge, Mass.
19. Young, H.P. 1998. **Individual Strategy and Social Structure**. Princeton University Press: Princeton, N.J.
20. ———. 1999. Diffusion in Social Networks. Center on Social and Economic Dynamics working paper. Brookings Institution: Washington, D.C. Available at www.brook.edu/es/dynamics/papers.