

# Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement

Robert L. Axtell and Joshua M. Epstein<sup>†</sup>

Center on Social and Economic Dynamics  
Working Paper No. 1  
May 1999

## ABSTRACT

A model of retirement decision-making is described in which a relatively small number of agents are rational. Such agents behave as if they compute the optimal age at which to retire. A small proportion of agents retire at random. A majority of the population engages in imitative behavior. An imitative agent retires once a certain fraction of the agents in its social network have retired. The model is analyzed by agent-based computational techniques. It is demonstrated that high levels of optimal behavior can result in the aggregate despite low levels of individual rationality.

© 1999 Brookings Institution

---

<sup>†</sup> Robert L. Axtell is a fellow and Joshua M. Epstein is a senior fellow in Economic Studies at the Brookings Institution. Each is also a member of the Santa Fe Institute. They thank George Akerlof, Chris Carroll, Bob Hall, Peyton Young, and participants in the Brookings Work-In-Progress seminar. Research assistance from Trisha Brandon and David Hines is gratefully acknowledged. This research was partially supported by the National Science Foundation under grant IRI-9725302.

## Introduction

Though motivated by a policy question, this work has theoretical dimensions. There are two related theoretical issues. One is the connection between individual rationality and aggregate efficiency—between optimization by individuals and optimality in the aggregate. The second is the role of social interactions, and social networks in individual decision-making and in determining macroscopic outcomes and dynamics. Regarding the first, much of mathematical social science assumes that aggregate efficiency requires individual optimization. Perhaps this is why bounded rationality is disturbing to some economists: they implicitly believe that if the individual is not sufficiently rational it must follow that decentralized behavior is doomed to produce inefficiency. The invisible hand requires rational fingers, if you will.

Experimental economics and psychology have now produced strong empirical support for the view that framing effects, as well as contextual and other psychological factors put a large gap between *homo economicus* and *homo sapiens* (see the recent review of Rabin [1998], for instance). Individual rationality is bounded. The question we pose here is: Does that matter? How does it matter?


To answer these questions, we have developed a model in which imitation in social networks can ultimately yield high aggregate levels of optimal behavior despite extremely low levels of individual rationality. Now, the *fraction* of agents who are rational in such an imitative system will definitely affect the *rate* at which a steady-state sets in. But the eventual (asymptotic) attainment *per se* of such a state need not depend on the extent to which rationality is bounded. Perhaps the main issue then is not how much rationality there is (at the micro level), but how little is enough to generate macro-level patterns in which most agents are behaving “as if” they were rational, and how various social networks affect the dynamics of such patterns. Of particular concern here are the puzzling dynamics of retirement.

### *The Retirement Age Puzzle*

In 1961, Congress reduced—from 65 to 62—the minimum age at which workers could claim social security benefits. By any measure, this was a major policy shift. Yet it took nearly three decades for the modal retirement age to fall from 65 to 62. While various explanations are possible, we shall suggest that imitative behavior and social interactions—factors absent from traditional economic models—may be fundamental in explaining the sluggish response to policy.

For modeling purposes, one can represent retirement decision-making—and perhaps a range of other problems—in the following stylized terms: First, there is an initial state of the world in which the individually optimal age at which to take some action is  $Y$ . Suddenly, a policy is instituted exogenously. Given this policy, the individually optimal age at which to take the action becomes  $Y^*$ . What we observe, however, is not the instantaneous shift from  $Y$  to  $Y^*$ , as would be predicted assuming universal fully-informed rational behavior. Rather a long process of patchy social adjustments transpires, in which different clusters of individuals migrate to  $Y^*$  at different rates, with some groups perhaps not getting there at all.

In our model, the action in question is individual retirement, the exogenously instituted policy is Congress's 1961 reduction in the Social Security eligibility age, and 65 and 62 are  $Y$  and  $Y^*$ , respectively. The actual data are plotted in figure 6-1.<sup>1</sup>



---

<sup>1</sup> We thank Gary Burtless [1998] for supplying these data.




Figure 6-1: Male Retirement Rate by Age, 1960, 1970, and 1995-1996

As noted above, it took nearly three decades for the response—a downward shift in the modal retirement age from 65 to 62—to manifest itself. We develop a relatively general model—involving imitation in social networks—that generates such patchy and sluggish dynamics. It is not the only approach possible.

One body of research has sought to explain the data with aggregate models in which a representative agent solves some life cycle optimization (dynamic programming) problem (e.g., Rust and Phelan [1997], Laibson *et al.*, [1998]). If the goal is simply to fit the data, it is not unreasonable to attribute to agents the capacity to explicitly formulate and solve dynamic programming problems. However, there is strong empirical evidence that humans do not perform well on problems whose solution involves backward induction (Camerer [1997]). For this reason, these models fail to provide a realistic microeconomic—individualist—account of the phenomenon. We would like to provide such an account.

The model we describe will not invoke a representative agent, but will posit a heterogeneous population of individuals. Some of these will behave “as if” they were fully informed optimizers, while others—indeed most—will not. Social networks and social interactions—clearly absent from the prevailing literature—will play an explicit central role. We now turn to the specifics of the model.

## Retirement Age Norms: A Model

The agents in our model fall into three categories. Members of one minority group adopt the (presumably) optimal policy by a process we do not model. Another minority group is composed of randomly behaving agents, who retire with a fixed probability once they reach retirement age. The majority group consists of imitators, who mimic members of their social networks.

### Agents and Cohorts

The agent population is divided into age cohorts ranging from age 20 to 100. Thus, there are 81 cohorts. Each contains  $C$  agents for a total of  $81C = A$  agents. Each agent is assigned a random death age drawn from  $U[60, 100]$ .<sup>2</sup> The average death age is thus 80. When an agent dies it is replaced by a 20 year old agent.<sup>3</sup> Each time period each agent is activated exactly once and, if it is eligible to retire but not yet retired, decides whether or not to retire.<sup>4</sup>

### Social Networks

Agents are heterogeneous by social network; each has its own. A social network is simply a list of other agents, specified randomly and fixed over the agent's lifetime. The *number* of other agents is set by drawing a random network size,  $S$ , from  $U[a, b]$ . Some of these agents may be younger or older than the agent in question. This *extent*,  $E$ , represents how far, in the cohort dimension, the agent's social network extends above and below its own cohort;  $E$  is drawn from  $U[0, c]$ . So, one agent might have a social network of 17 other agents ranging in age from 5 years younger to 5 years older than itself, while another agent might have a social network consisting of 13 other agents, all within a year of its own age. Any two networks may or may not overlap, that is, have agents in common. At any time the set of all social networks constitutes a single random graph, with the agents as nodes and the network relations as (directed) edges. Figure 6-2 below shows a variety of social networks. Each rectangle represents an agent, with agents in the same cohort going across the page, and progressively older cohorts going down the page (colors will be interpreted below). Three social networks are shown, one each for the agents who are colored black in the age 60, 77, and 94 cohorts. For the age 60 agent, each of the 24 members of its social network are shown with

---

<sup>2</sup> Certain variables in our model are assigned random values. In all cases below the random variables are assumed to be uniformly distributed. The uniform (i.e., rectangular) distribution on the interval  $[a, b]$  is denoted  $U[a, b]$ .

<sup>3</sup> The number of cohorts, number of agents per cohort, and the death age distribution are all easily modified in the software that we have created for this model.

<sup>4</sup> In the computational implementation of the model the order of agent activation is randomized within cohorts each period. Such randomization is commonly held to be necessary in order to suppress the production of so-called simulation "artifacts", that is, spurious correlation in the agent population.

an 'X' in the figure. This agent's network includes 13 younger agents and 11 older agents.

### Agent Types

As noted earlier, there are three broad types of agents. For lack of better terminology, we designate them "rationals," "randoms," and "imitators." Rational agents retire at the earliest possible age allowed by government policy. Random agents retire with probability  $p$  each period once they reach the retirement eligibility age.

Imitator agents are the most heterogeneous and interesting. Each imitator has a unique social network. Within this individual network, there is some fraction  $f$  of eligibles who have actually retired. At each instant this is heterogeneous across agents since the size and composition of networks are agent-specific. Agents are assigned an *imitation threshold*,  $\theta$ . Each agent's behavioral rule then simply amounts to comparing  $f$  with  $\theta$ .<sup>5</sup> If  $f > \theta$ , the agent retires; otherwise, it continues working until the following period when it reevaluates its decision.

Notionally, the imitator agents are playing a simple coordination game within their social networks.<sup>6</sup> That is, agents derive utility from coordinating their behavior with the members of their social network. At every instant each agent in the population is either working or retired. Since  $A$  is the number of agents, call  $x = \{\text{working, retired}\}^A$  the state of the population, and  $x_i$  agent  $i$ 's state. Note agent  $i$ 's social network by  $N_i$ . Then the utility that  $i$  derives from interacting with the members of its social network in state  $x$ ,  $U_i(x)$ , can be written

$$U_i(x) = \frac{1}{|N_i|} \sum_{j \in N_i} u(x_i, x_j),$$

where  $u(x_i, x_j)$  is the utility of  $i$ 's interaction with  $j$ . The function  $u$  can be thought of as the payoff function of a  $2 \times 2$  symmetric game, as given by figure 6-3.

---

<sup>5</sup> It makes a difference to the numerical results whether an agent considers all agents in its social network, or only those who are eligible to retire. However, the qualitative character of the results described below do not depend on this distinction.

<sup>6</sup> This development closely follows Young [1998: 3-4].

|        |        |        |
|--------|--------|--------|
|        | work   | retire |
| work   | $w, w$ | $0, 0$ |
| retire | $0, 0$ | $r, r$ |

Figure 6-3: Retirement as a coordination game

$U_i$  is then the payoff function of the social network game. Note that  $U_i$  can be expressed in payoff terms. When an agent is young, and none of its social peers are retired,  $f = 0$ , and the agent derives maximum utility from working. However, as its friends begin to retire ( $f > 0$ ), the utility from retiring rises to  $rf$ , and the utility to working falls from  $w$  to  $w(1-f)$ . The agent decides to retire if  $f$  rises to a level such that  $rf = w(1-f)$ , or equivalently,  $f = w/(r+w)$ , which expresses the agent's imitation threshold in terms of payoffs as  $w/(w+r)$ .

In this social network game, then, how does the behavior of interest—the shift to earlier retirement—diffuse through coupled heterogeneous networks? And, how do the dynamics vary with key parameters, such as the number of rational agents, the distribution of imitation thresholds, and the probability that a random player will retire when eligible? We will resolve these questions quantitatively by appeal to an agent-based computational model.<sup>7</sup> Before delving into detailed analysis of model runs, perhaps a brief introduction to the general approach is in order.

## Agent-Based Computational Models<sup>8</sup>

Compactly, in agent-based computational models a population of data structures representing individual agents is instantiated and permitted to interact. One then looks for systematic regularities, often at the macro-level, to emerge, that is, arise from the local interactions of the agents. The shorthand for this is that macroscopic regularities “grow” from the bottom-up. No equations governing the overall social structure are stipulated in multi-agent computational modeling, thus avoiding any aggregation or misspecification bias. Typically, the only equations present are those used by individual agents for decision-making. Different agents may have different decision rules and different information; usually, no agents have global information, and the behavioral rules involve bounded computational capacities—the agents are

---

<sup>7</sup> Coordination games on *fixed* social networks have been studied by Blume [1995] and Young [1998]. Because the networks here are transient, these analytical results do not apply.

<sup>8</sup> For extended discussions of the agent-based computational approach, see Epstein and Axtell [1996] or Axelrod [1997].

“simple”. This relatively new methodology facilitates modeling agent heterogeneity, boundedly rational behavior, non-equilibrium dynamics, and spatial processes.<sup>9</sup> A particularly natural way to implement agent-based models is through so-called object-oriented programming. Our object-oriented implementation of the present model is described in the appendix.

## **Realizations of the Model: Establishment of an Age 65 Norm**

We begin our analysis by describing in detail two particular realizations of the model, one with a relatively large fraction of rational agents and the other with relatively few. Because the model involves stochastic elements, each realization is essentially unique, even for fixed numerical values of all parameters. Eventually we will characterize large numbers of realizations statistically, but first we focus on individual realizations in order to build up some intuition about how the model works.

In all runs of the model to be described below each cohort consists of  $C=100$  agents. Therefore, the population size,  $A$ , is 8100. In the first realization 15% of the agents are rational, 75% are imitators, and 5% are random agents. The size of each individual's social network is set by drawing a random number from  $U[10, 25]$ . Each agent's network extends up to 5 age cohorts above and below. Imitating agents have a homogeneous imitation threshold,  $\theta$ , of 0.5, meaning that 50% of the members of an agent's social network must be retired before that agent will retire. Random agents retire with probability  $p = 0.5$  each period, once they are eligible. Government retirement eligibility age is 65, and there is no forced retirement age.

Animation 6-1 portrays the evolution of retirement in our agent society, and conveys a sense of how imitation propagates the retirement decision through social networks.<sup>10</sup> As in figure 6-2, each agent is a rectangle. Agents are arrayed across the page by cohort, down the page by increasing age. Retired agents are shown in red and dead agents are colored white. Among the unretired agents, pink ones are rational, the blue are imitators, while the few yellow agents are random.

It is worthwhile to spell out exactly how to “read” an animation. At the start, there are 100 agents in each of 81 age cohorts, of which the eldest 46

---


<sup>9</sup> For more on the comparative advantages of this modeling technique, see Epstein and Axtell [1996].

<sup>10</sup> Readers with access to the World Wide Web are invited to the URL <http://www.brook.edu/es/dynamics/papers/retirement>, where QuickTime™ movies of this and subsequent animations are stored and available to be downloaded.



are displayed. So, the top row represents 100 agents of age 55. Let's call the upper left hand agent Tom. At time  $t = 1$ , Tom is cell (1, 1) in matrix notation. At  $t = 2$ , Tom is the cell immediately beneath this one, cell (2, 1). In general, at time  $t$ , Tom's status appears in cell  $(t, 1)$ . A shift in color over time indicates that an agent has either retired or died.


Observing animation 6-1, notice that a uniform retirement age of 65 quickly sets in, despite the fact that only a fairly small minority (15%) of the population arrives at this decision rationally. Figure 6-4 gives a time series plot of the fraction of agents eligible for retirement who are actually retired.



**Figure 6-4:** Fraction of eligible agents actually retired over time in a population with 20% rational agents, typical realization

Note that this trajectory is essentially monotone. Within the first 6 periods essentially all of the eligible population has retired.

For the next realization, only the mix of agent types is changed: now there are only 5% rationals and 90% imitators. Animation 6-2 is a typical result. Note that the older cohorts show extensive fluctuation in retirement levels before the system converges to full retirement at age 65. It is as if retirement 'percolates up' from older to younger agents. Figure 6-5 gives the time series of the fraction of agents eligible for retirement who are retired.



**Figure 6-5:** Fraction of eligible agents actually retired over time in a population with 5% rational agents, typical realization

It takes a long time for the absorbing state to be achieved in this case. Notice that now the trajectory is not monotone.

### *Some Sensitivity Analysis*

Each of the realizations described above yielded interesting qualitative information about the model. However, in order to quantitatively characterize the model's overall behavior it is necessary to make many realizations for a particular set of parameters and progressively build-up a statistical portrait of the solution space computationally. That is, the intrinsic stochasticity of the model can be approximately characterized through a sufficiently large number of realizations. Once this is done for a particular configuration, the effect of varying the model parameters can be studied.

We begin this analysis by defining a 'base case' configuration of the model in table 6-1.

| Parameter                | Value     |
|--------------------------|-----------|
| Agents/cohort, $C$       | 100       |
| rational agents          | 10%       |
| imitative agents         | 85%       |
| imitation threshold,     | 0.50      |
| social network size, $S$ | U[10, 25] |
| network age extent, $E$  | U[0, 5]   |
| random agents            | 5%        |

|     |      |
|-----|------|
| $p$ | 0.50 |
|-----|------|

Table 6-1: Base case parameterization of the model

We will study the effect of each of these parameters on the time required for a the age 65 retirement norm to emerge, the *transition time*. The first parameter,  $C$ , the number of agents per cohort, was found to have no effect on the average transition time. So we begin exploration of the model by varying the relative proportions of the three agent types—rationals, imitators, and randoms—with all other parameters as in the base case. Fifty realizations were made for each configuration of the model, and mean transition times were estimated along with standard deviations. Figure 6-6 shows the average transition times for three levels of randomly-behaving agents, as a function of the fraction of rationals (and hence imitators). Note that the ordinate is in logarithmic coordinates; error bars are  $\pm 1$  standard deviation, and are asymmetrical due to the logarithmic scale.

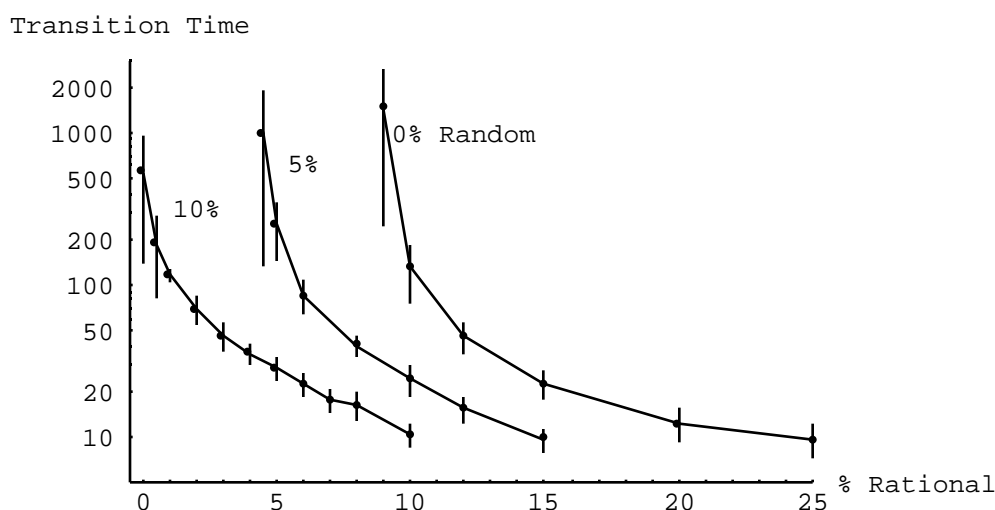



Figure 6-6: Transition time to the age 65 retirement norm, as a function of the fraction of rational agents, parameterized by the number of randomly-behaving agents

Reducing the proportion of rationals, while holding constant the proportion of randoms, increases transition time. When randoms comprise 0 percent or 5 percent of the population, certain minimum proportions of the population must be rational for a retirement age norm to arise. For a given fraction of rationals, the transition time decreases as the proportion of randoms increases. Notice that the variances increase rapidly with transition times.

The effect of the imitation threshold,  $\theta$ , on transition times is the next sensitivity analysis described. Now, since social networks are composed of individuals, the fraction of agents in one's network engaged in some behavior can only take on certain discrete values. That is, small changes in

may have no effect on agent decision making, and thus no effect on transition times. For example, imagine that all agents have social networks of size ten. Then, clearly, increasing  $\theta$  from 0.55 to 0.58 has no effect; agents either have 5 or fewer retired agents in their network, or 6 or more. Only when  $\theta$  is moved across a discrete boundary—say from 0.58 to 0.62 in our example—does it have an effect.


Therefore, instead of studying the dependence of transition times on the average imitation threshold—surely a very 'lumpy' dependence—we investigate the effect of making the threshold progressively more heterogeneous in the agent population while holding the average value of constant. Figure 6-7 shows how transition times depend on the standard deviation in the imitation threshold, with the average threshold fixed at 0.50. Once again, the ordinate is in logarithmic coordinates.



**Figure 6-7:** Transition time to the age 65 retirement norm, as a function of the standard deviation in agent imitation threshold


Increasing the variance in the threshold decreases the average transition time. The reason is that in high variance populations there are relatively more agents having low thresholds, and these agents quickly retire, leading the rest of the population to retire quickly as well. Note that for low variance in the threshold there is significantly more variance in the transition time.

In figures 6-8a, 6-8b and 6-8c the dependence of transition time on the size of agent networks is shown, for two levels of random agents. The separate effects of changing the average size and the size variance are treated in the first two figures, while the overall (opposite) effects are combined in the third figure. In particular, figure 6-8a describes the effect of increasing network size, with constant variance.



**Figure 6-8a:** Transition time to the age 65 retirement norm, as a function of the average size of agent social networks


Note that the time required to transit to a uniform retirement age increases very rapidly with increasing social network size; in large networks it is difficult for a new norm to establish itself. Figure 6-8b gives the dependence of the transition time on the dispersion (the population standard deviation) in the social network size, holding the average size constant.



**Figure 6-8b:** Transition time to the age 65 retirement norm, as a function of the standard deviation in size of agent social networks

Here we see that as the variance increases the transition time decreases, although this is a relatively weak effect. The reason for this is that the small networks catalyze the transition to a new norm, and as the variance increases there are more small networks. Finally, both of these effects are combined in figure 6-8c, where the abscissa, call it  $\bar{S}$ , represents the maximum size any

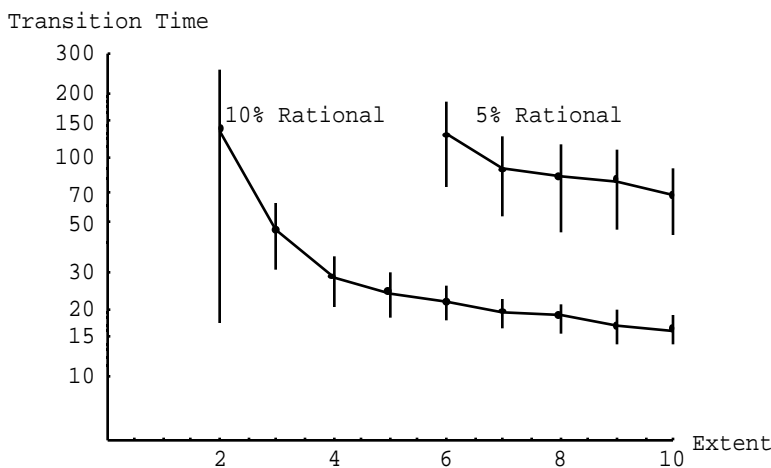
social network, i.e., the size,  $S$ , of an agent's network is set by drawing a random number from  $U[10, \bar{S}]$ .



**Figure 6-8c:** Transition time to the age 65 retirement norm, as a function of the maximum size of agent social networks

As  $\bar{S}$  increases, the average social network size rises as does the variance, and the two competing effects on transition time given in figures 6-8a and 6-8b play out, yielding figure 6-8c. Overall, the general effect is that the transition time increases very rapidly with  $\bar{S}$ .

Next, consider the effect of extending not the size of agent social networks, but rather their extent in the age cohort dimension. For all of the above the maximum extent has been 5, i.e., up to 5 cohorts above and below an agent's own cohort. Here we vary this, with the ordinate in figure 6-9 representing the extent in each direction.



**Figure 6-9:** Transition time to the age 65 retirement norm, as a function of the extent of agents' social network (by age cohorts)

Note that the effect of increasing the extent (in the age dimension) of agent social networks is to decrease the transition times. The reason for this is that networks having greater extent include older agents, who are more likely to be retired.


*Dynamics and “As If”*

Notice that in figure 6-6 the only variable affected by the fraction of rationals is the transition time. The attainment *per se* of the age 65 retirement norm is compatible with any rationality fraction above a critical level. So while in establishing the social norm the system does behave “as if” all agents are rational it also behaves “as if” none are! However in taking a long time to achieve the norm it does not behave “as if” all agents are rational; indeed, it behaves as if most are not.

**Response to Policy Change: Retirement Age Moves from 65 to 62**

So far, our model permitted agents to retire when they wished; no mandatory retirement age was in effect. Now we require that all agents retire at age 70. This increases the speed at which the age 65 retirement norm is established. We wish to investigate the effect of policy intervention on retirement age norms. Thus, once the age 65 norm is established, we throw a 'policy switch' and lower the retirement age from 65 to 62, mimicking Congress's 1961 policy change. In our model, this switch means only that rationals claim benefits at age 62 and that randoms and imitators *may* receive benefits at age 62. We then measure how long it takes for a new retirement age norm to become established. Keep in mind that after the eligibility age for social security benefits was lowered from 65 to 62, it took around 35 years for a new norm to emerge (see figure 6-1). Animation 6-3, which uses the same parameters as did animation 6-2, shows that a new norm indeed emerges after twenty to thirty periods. In short, the model replicates the sluggish adjustment that occurred, at least qualitatively.

Many realizations of this model have been made, varying the number of rationals in the population. The results are shown in figure 6-10.



**Figure 6-10:** Transition time to the age 62 retirement norm, as a function of the fraction of rational agents in the population

Note that the transition time to the age 62 norm falls as the fraction of rational agents increases. Based on this parameterization of the model, a new norm is instituted in about 35 periods if between 1 and 4 percent of the population responds rationally—that is to say, immediately—to the new policy. The sensitivity analyses described in figures 6-7 through 6-9 indicate how the speed of adjustment depends on other parameters in the model. In particular, we expect the time required to adjust to a policy shock to rise for larger values of the imitation threshold,  $\theta$ , and for increases in the average size and extent of social networks.

This use of the agent-based computational model—as a kind of laboratory in which alternative policies can be studied—seems to us a fertile application of this technology, and one that has not been systematically exploited.

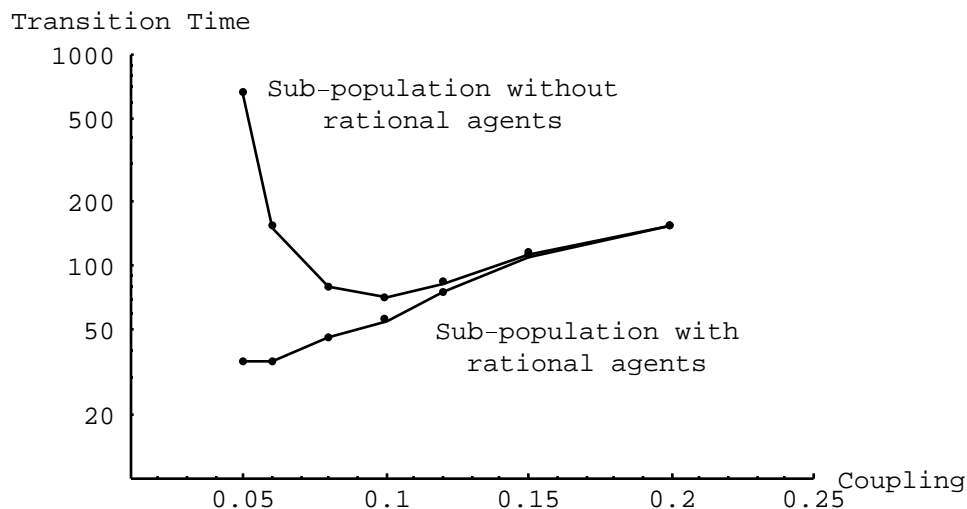
## Two Sub-Populations, Loosely Coupled Through Social Networks

Different subgroups in society may be better informed and educated than others. And such differences can affect the relative rates at which the communities adopt various norms. In animation 6-4 the agents have been broken into 2 distinct sub-populations. The 50 agents on the left do not include any rational agents, while those on the right include 10% rationals. However, the populations are coupled through social networks as follows: 10% of each agent's network belongs to the other sub-population, with the remainder being members of its own group. We term this quantity—10 percent—the *coupling* between sub-populations. Even this rather loose



coupling is sufficient for the group containing some rationals to pull the other into conformity with its retirement norm, as shown in animation 6-4.

We studied this general effect by systematically varying the extent of coupling between sub-populations and measuring the time required for each group to reach a retirement norm of age 65 from an initially unretired state. The results are shown in figure 6-11; each point is an average over 50 realizations.



**Figure 6-11:** Transition time to the age 65 retirement norm, as a function of the amount of social network overlap between sub-populations

Note that very little coupling is needed for the non-rational sub-population to be pulled into conformity with the more rational sub-population.

## Conclusions

With social network interactions and imitative dynamics, very little individual rationality may be needed for society as a whole to ultimately exhibit optimal behavior. More pointedly, there is a large literature, experimental and theoretical, devoted to the question: how rational are individual humans? From the perspective of network imitation, it may not matter. Second, the non-equilibrium dynamics and the social patchiness of a response to policy will both depend on the size and structure of networks. It is not clear how one would adapt the representative agent approach to study either of these dependencies. They are naturally explored within the agent-based computational framework.

This paper has barely scratched the surface of a rich and promising area. Many fruitful avenues for future research suggest themselves, both analytical

and computational. On the analytical side, it would be extremely useful to have—for the transient networks described here—theorems analogous to those of Blume [1995] and Young [1998], which give conditions under which social norms will be established eventually for static networks. Furthermore, it would be desirable to have formal expressions for the way in which transition time distributions depend on model parameters, like the fraction of imitators and the size of social networks.

Computationally, it would be useful to extend this retirement age norms model to include income shocks and imitative consumption behavior over the life cycle, as in the agent-based model of Carroll and Allen [1997]. We suspect that doing so would add more heterogeneity to the outcomes observed in our model. Furthermore, such a model would provide a useful laboratory in which to explore new theoretical ideas, like the effect of hyperbolic discounting, as well as to experiment with policy alternatives, like increasing the retirement age or privatizing social security.

#### *Retirement as an Instance of More General Social Phenomena*

While we have interpreted this model as applying to retirement, it could be applied to a wide range of settings in which social interactions mediate purely rational behavior. Obvious candidates include contagion behavior in markets, migration to different health plans, adoption of various social norms, or the diffusion of technological innovations. In reality, these phenomena occur in social networks, while most existing models treat them as occurring either in 'perfectly mixed' environments, where each agent knows what every other is doing, or in local interaction models on regular lattices or other highly specialized topologies. The agent-based computational approach is well-suited to studying such processes with any topology of interactions.

### **Appendix: Implementation of the Model—Agents as Objects**

There are many ways to computationally implement agent-based models. This can be done in any modern programming language, or with a number of mathematical or simulation software packages. However, since the model is stated in terms of individual agents, there is one idea from modern computer science that renders the implementation both transparent and efficacious. This is the notion of object-oriented programming.

Objects are contiguous blocks of memory that contain both data—so-called *instance variables*—and functions for modifying this data—the object's so-called *methods*. This ability of objects to hold both data and functions

operating on data is called *encapsulation*. Agent-based models are very naturally implemented using objects by interpreting an object's data as an agent's state information, while the object's functions become the agent's rules of behavior.<sup>11</sup> A population of agents that have the same behavioral repertoire but local state information is then conveniently implemented as multiple instantiations of a single agent object type or class.<sup>12</sup>

The model described above has been implemented using object-oriented programming. Not only are individual agents objects, but cohorts are objects too, albeit of a different class than agent objects. In fact, it has proven to be convenient for the population of cohorts (and thus of agents) as a whole to be an object as well.

The agent object has a variety of state variables and behavioral methods. An agent's state information includes its type (i.e., rational, imitator, or random), its age, its current employment status (i.e., whether working or retired), and the age at which it will die—its 'death age.'<sup>13</sup> All of this information is stored so to speak, locally, in the agent object. Each agent also keeps track of some number of other agents that are identified as its social network. This data is maintained in a social network object, described below. The agent's main decision in the present model is whether or not to retire. This is the agent object's basic *method*. This agent object specification is summarized in pseudo-code block 1.

```
OBJECT agent;  
  type;  
  age;  
  death_age;  
  alive_or_dead;  
  social_network;  
  working_or_retired;  
  next_agent_in_agent_list;  
  FUNCTION initialize;  
  FUNCTION retirement_decision;  
  FUNCTION draw.
```

**Pseudo-code block 1: Agent object**

---

<sup>11</sup> Other features of the object model, including *inheritance* and *polymorphism*, seem to be less relevant to agent-based computational models than encapsulation.

<sup>12</sup> For a discussion on the distinction between object and agent, see Jennings *et al.* [1998].

<sup>13</sup> The agent is assumed not to know this datum.

In practice it makes sense to implement as private some of these data and methods, while others are public, although this is not essential.<sup>14</sup>

Each social network is also conveniently implemented as an object. The size of each social network is data local to that object, as is an array of pointers to (i.e., memory addresses of) the agents who constitute the network. Methods associated with this object include routines for determining how many agents in the network are eligible to retire as well as how many are actually retired. This is summarized in pseudo-code block 2, below.

```
OBJECT social_network;  
  size;  
  array_of_agents;  
  FUNCTION initialize;  
  FUNCTION number_eligible_to_retire;  
  FUNCTION number_retired;  
  FUNCTION fraction_retired_of_eligible;  
  FUNCTION draw.
```

**Pseudo-code block 2: Social network object**

Cohorts are also conveniently implemented as objects. The size of the cohort is kept as local data, as is an array of agents who constitute the cohort. The methods of this object are primarily data gathering and statistical routines, useful in characterizing the behavior of the cohort overall. The cohort object is summarized as pseudo-code block 3.

```
OBJECT cohort;  
  size;  
  array_of_agents;  
  FUNCTION initialize;  
  FUNCTION average_social_network_size_among_agents_in_cohort;  
  FUNCTION number_retired;  
  FUNCTION fraction_retired_of_eligible  
  FUNCTION draw;
```

**Pseudo-code block 3: Cohort object**

The population of cohorts is also an object. Similar to the cohort object, it is merely an array of entities—here cohorts—together with data gathering and statistical methods for discerning the state of the population overall.

Putting all of this together the agent-based computational model amounts to:

- (1) initializing all agents, social networks and cohorts;
- (2) choosing an agent at random and incrementing its age;

---

<sup>14</sup> Private data and methods are accessible only by the agent to whom they belong, unless other objects are given special access privileges.

- (3) checking to see if the agent has achieved its death age; if yes then go to (2); else
- (4) having the agent decide whether to retire;
- (5) repeat (2) through (4) for all agents;
- (6) periodically gather and report statistics on the population.

This algorithm is summarized in pseudo-code block 4.

```
PROGRAM retirement;  
  initialize agents;  
  initialize social networks;  
  initialize cohorts;  
  repeat:  
    select an agent at random  
    increment its age;  
    if age < death_age then  
      do retirement_decision;  
    get statistics on the agents and cohorts;  
  until user terminates.
```

**Pseudo-code block 4:** Pseudo-code for the model overall

The object model is largely responsible for the relatively short description of this code.<sup>15</sup>

## References

- Axelrod, R. 1997. **The Complexity of Cooperation**. Princeton University Press: Princeton, New Jersey.
- Blume, L. 1995. The Statistical Mechanics of Strategic Interaction. **Games and Economic Behavior**, 4: 387-424.
- Burtless, G. 1998. Private communication.
- Camerer, C. 1997. Progress in Behavioral Game Theory. **Journal of Economic Perspectives**, 11 (4): 167-188.
- Carroll, C. and T. Allen. 1997. "Learning About Intertemporal Choice." Presentation at Santa Fe Institute Workshop on Local Interactions Models in Economics.

---

<sup>15</sup> The actual source code is less than 2000 lines of C++ and compiles in the CodeWarrior environment for the Macintosh. A Java implementation is available at <http://www.brook.edu/es/dynamics/papers/retirement>.

- Epstein, J.M. and R. Axtell. 1996. **Growing Artificial Societies: Social Science from the Bottom Up.** MIT Press and Brookings Institution Press: Cambridge, Massachusetts, and Brookings Press: Washington, D.C.
- Jennings, N.R., K. Sycara, and M. Wooldridge. 1998. A Roadmap of Agent Research and Development. **Autonomous Agents and Multi-Agent Systems**, 1, number 1: 7-38.
- Kagel, J.H., and A.E. Roth, eds. 1995. **The Handbook of Experimental Economics.** Princeton University Press.
- Kochen, M., ed. 1989 **The Small World.** Ablex Publishing Corporation: Norwood, New Jersey.
- Laibson, D.I., A. Repetto and J. Tobacman. 1998. Self-Control and Retirement Savings: Do 401(k)'s Help? Working paper: Harvard University, Cambridge, Massachusetts.
- Rabin, M. 1998. Psychology and Economics. **Journal of Economic Literature**, Vol. XXXVI (March): 11-46.
- Rust, J. and C. Phelan. 1997. How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets. **Econometrica**, 65 (4): 781-831.
- Simon, H.A. 1956. Rational Choice and the Structure of the Environment. **Psychological Review**, 63, 129-138.
- Young, H.P. 1998. Diffusion in Social Networks. Working paper. Brookings Institution: Washington, D.C.






Figure 6-2: Typical agent social networks

Animation 5-1. Rapid Propagation of Retirement Behavior through Social Networks, No Mandatory Retirement Age

Frame 1




Frame 2




Frame 3




Frame 4



Frame 5




Frame 6






*Animation 5-2. Slow Propagation of Retirement Behavior through Social Networks, No Mandatory Retirement Age*


Frame 1




Frame 2




Frame 3




Frame 4




Frame 5




Frame 6



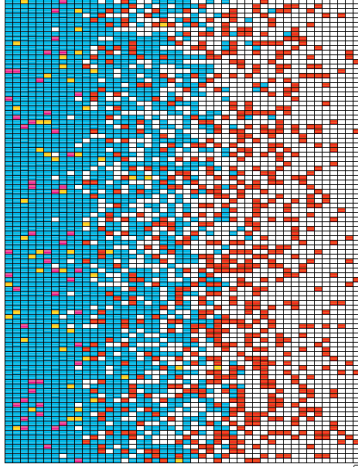
Frame 9



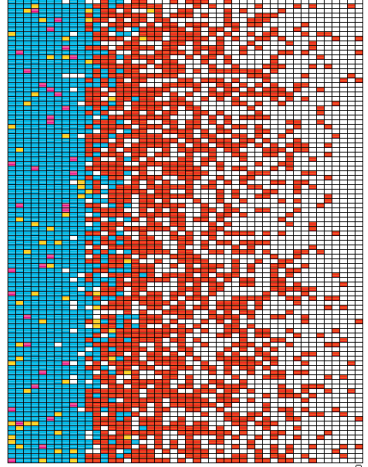
Frame 12




Frame 8




Frame 11



Frame 7




Frame 10




Animation 5-3. *Propagation of Retirement Behavior through Social Networks, Mandatory Retirement Age of 70 and Policy Change from Earliest Retirement Age of 65 to 62*

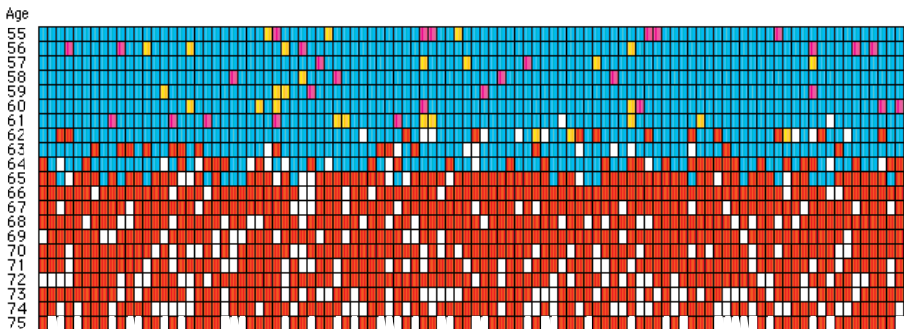
Frame 1



Frame 2




Frame 3




Animation 5-3. (continued)


Frame 4



Frame 5

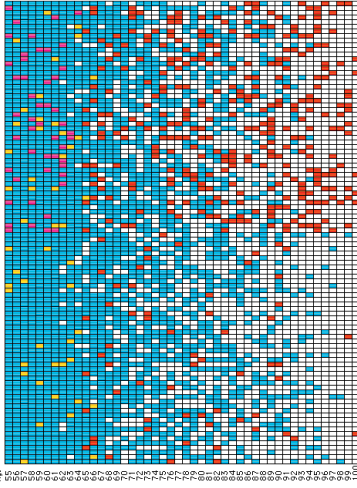


Frame 6




Animation 5-4. Retirement Dynamics in Two Loosely Coupled Subpopulations


Frame 1




Frame 2



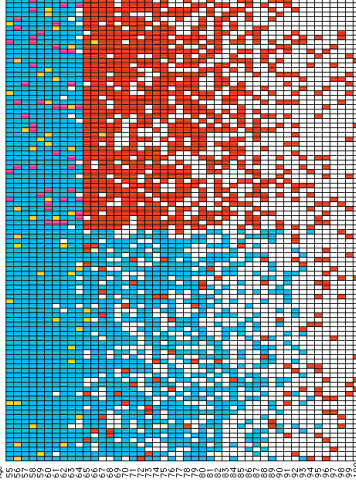
Frame 3




Frame 4



Frame 5




Frame 6







Frame 9




Frame 8




Frame 7



Frame 12



Frame 11



Frame 10

