

The Emergence of Classes in a Multi-Agent Bargaining Model

Robert Axtell, Joshua M. Epstein, and H. Peyton Young

Center on Social and Economic Dynamics
Working Paper No. 9
February 2000

ABSTRACT

The essential idea is to show how norms can emerge spontaneously at the social level from the decentralized interactions of many individuals that cumulate over time into a set of social expectations. Due to the self-reinforcing nature of the process, these expectations tend to perpetuate themselves for long periods of time, even though they may have arisen from purely random events and have no *a priori* justification. We show that social expectations gravitate to one of three conditions: i) an equity norm in which property is shared equally among claimants, and there are no “class” distinctions; ii) a discriminatory norm in which the claimants get different amounts based on observable characteristics that have become socially salient (but are fundamentally irrelevant); and iii) fractious states in which norms of distribution have failed to coalesce, resulting in constant disputes and missed opportunities.

For valuable comments the authors thank Sam Bowles, Jeff Carpenter, Steve Durlauf, Nienke Oomes, John Roemer, John Rust, Thomas Schelling, John Steinbruner, Leigh Tesfatsion, Frank Thompson, Erik Olin Wright, and seminar participants at Brookings, Davis, Michigan, the Santa Fe Institute, Stanford, and the University of Massachusetts (Amherst). Steven McCarroll's research assistance was invaluable. Support from the National Science Foundation under grant IRI-9725302 is gratefully acknowledged. Additional support was provided to the Center on Social and Economic Dynamics by the MacArthur Foundation.

1 Introduction

Norms are self-enforcing patterns of behavior: it is in everyone's interest to conform given the expectation that others are going to conform. Many spheres of social interaction are governed by norms: dress codes, table manners, forms of deference, modes of communication, reciprocity in exchange, and so forth. In this paper we are interested in norms that govern the distribution of property. In particular, we are concerned with the contrast between *discriminatory norms*, which allocate different shares of the pie according to gender, race, ethnicity, age, etc., and *equity norms*, which do not so discriminate. An example of a discriminatory norm is the practice of passing on inherited property to the eldest son (primogeniture). Another is the custom, once common in the southern United States, that blacks should sit in the back of the bus. A third is the notion that certain categories of people (e.g., women, blacks) should receive lower compensation than others doing the same job, and in other cases that they not be given the job at all. These kinds of discriminatory norms can lead to significant differences in economic class, that is, long-lived differences in property rights based on characteristics that are viewed as socially salient.¹

In this paper we study the question of how such classes can emerge and persist, given a norm-free, classless world initially. The framework combines concepts from evolutionary game theory on the one hand and agent-based computational modeling on the other. The essential idea is to show how norms can emerge spontaneously at the social level from the decentralized interactions of many individuals that cumulate over time into a set of social expectations. Due to the self-reinforcing nature of the process, these expectations tend to perpetuate themselves for long periods of time, even though they may have arisen from purely random events and have no *a priori* justification. We show that social expectations gravitate to one of three conditions: i) an equity norm in which property is shared equally among claimants, and there are no "class" distinctions; ii) a discriminatory norm in which the claimants get different amounts based on observable characteristics that have become socially salient (but are fundamentally irrelevant); and iii) fractious states in which norms of distribution have failed to coalesce, resulting in constant disputes and missed opportunities. In both the first and second case, society functions efficiently in the sense that no property is wasted. There is no equity-efficiency tradeoff, just a difference in the way property rights are distributed. The third case, by contrast, is highly inefficient and may involve substantial inequality as well.

¹ For other models of classes see Roemer [1982] and Cole *et al.* [1998]. The present paper differs from these by focusing on the dynamic process by which classes emerge, rather than on the equilibrium conditions that sustain them.

The long run probability of being in these three different regimes can be computed using techniques from stochastic dynamical systems theory (Freidlin and Wentzell [1984]; Foster and Young [1990]; Young [1993a, 1998]; Kandori, Mailath and Rob [1993]). But these methods are less helpful in characterizing the short and intermediate run behavior of these processes. Here agent-based computational techniques can play a central role, by identifying regimes that are long-lived on intermediate time scales, though not necessarily stable over very long time scales (Epstein and Axtell [1996], Axtell and Epstein [1999]).

Overview of the Model

Our model of class formation is based on Young's evolutionary model of bargaining (Young [1993b]). The model is bottom-up in the sense that norms emerge spontaneously from the decentralized interactions of self-interested agents.² In each time period two randomly chosen agents interact, bargaining over shares of available property. Their behavior, and their expectations about others' behavior, evolve endogenously based on prior experiences. These expectations may be conditioned on certain visible characteristics or "tags" that serve to differentiate people. These tags have no *inherent* social or economic significance—they are merely distinguishing features, such as dark or light skin, or brown or blue eyes. Over time, however, they can acquire social significance due to path dependency effects. It might happen, for example, that blue-eyed people get a larger share of the pie than brown-eyed people due to a series of chance coincidences. The existence of these precedents causes the expectation to develop that blue-eyed people generally get more than brown-eyed people, and a *discriminatory norm* emerges. Alternatively, an *equity norm* can develop in which the tags have no significance, and both sides get equal shares.

It can be shown that, asymptotically, the equity norm is more stable than any discriminatory norm. In other words, starting from arbitrary initial conditions, society is more likely to be at or near an equal sharing regime than an unequal or discriminatory one if we wait long enough. Nevertheless, metastable regimes can emerge that are discriminatory and inequitable, yet persist for substantial periods of time. These inequitable regimes correspond, roughly speaking, to situations where a discriminatory inter-group norm divides society into upper and lower classes, while a different, intra-group norm causes dissension within one (or both) of the classes. Based on many realizations of the agent-based computational model, we estimate the time it takes to exit from these discriminatory regimes as a function of the number of agents, the length of agents' memory, and the level of background noise. In

² We use the term "emergent" as defined in Epstein and Axtell [1996] to mean simply "arising from the local interactions of agents." The term and its history are discussed at length in Epstein [1999].

this case, the waiting time increases *exponentially in memory length and the number of agents*, and can be immense even for relatively modest values of the parameters. The contrast between long-run and intermediate-run behavior illustrates how analytical and computational methods complement one another in studying a given social dynamic.

2 Bargaining

We begin by modelling a bargaining process between individual agents. Consider two players, A and B, each of whom demands some portion of a "pie," which we take as a metaphor for a piece of available property. The exact nature of the property need not concern us here. For simplicity, however, we shall suppose that the property is divisible, and that both parties have an equal claim to it a priori.³ A posteriori differences in claim will emerge endogenously from the process itself.

To specify the process, we must first delineate how agents solve the one-shot bargaining problem. A standard way of modelling this situation is the *Nash demand game*: each party gets his demand if the *sum* of the two demands is not more than 100 percent of the pie, otherwise each gets nothing. For instance, if employers and employees demand more than 100 percent of total revenues, negotiations break down.

To simplify the analysis, we shall suppose that each agent can make just three possible demands: Low (30 percent of the pie), Medium (50 percent), and High (70 percent).⁴ For example, if row demands H and column plays M, their demands sum to 120 and each gets nothing. The payoffs (in percentage share) from all combinations of demands are shown in Table 1.

	H	M	L
H	0,0	0,0	70,30
M	0,0	50,50	50,30
L	30,70	30,50	30,30

Table 1: The Nash demand game

This yields a coordination game in which there are exactly three pure-strategy Nash equilibria, shown in bold: (L, H), (M, M), and (H, L). While various theories have been advanced that identify a particular equilibrium as being most plausible a priori (e.g., Harsanyi and Selten [1988]), we do not find these

³ Indivisible forms of property, such as a bus seat, can be made divisible by giving the claimants equal a priori chances at being the occupant.

⁴ The more general case is considered in Young [1993b].

equilibrium selection theories to be especially compelling. Instead of assuming equilibrium, we wish to explore the process by which equilibrium emerges (if indeed it does) at the aggregate level, from the repeated, decentralized interactions of individuals.

3 The Model with One Agent Type

We begin by studying this question for a population of agents who are indistinguishable from one another, but who have different experiences (life-histories) that condition their beliefs. Then we consider a population consisting of two distinct types of agents, who are differentiated by a visible "tag" (dark or light skin, brown or blue eyes) that has no intrinsic economic significance, but on which agents may condition their behavior. In the latter case, long-lived discriminatory norms can develop purely by historical chance, while this does not happen in the case of homogeneous agents. But in both situations, fractious regimes can emerge in which society fails to develop any coherent norm for long periods of time.

Let the population consist of N agents. Each time period consists of $\lfloor N/2 \rfloor$ "matches." In each match, one pair of agents is drawn at random from the population, and they play the game in Table 1.⁵ Each agent's data about its world—its beliefs—are based on experience from previous plays. In particular, every agent remembers the demands—H, M, or L—played by each of her last m opponents, where m is *memory length*.⁶ The concatenation of all agent memories defines the current *state* of the society. Behaviorally, each agent forms an expectation about her opponent's demands. She assumes that the probability of the current opponent demanding L, M, or H is equal to the relative frequency with which her previous opponents made these demands in the last m interactions. But with some relatively small probability, ε , she selects her demand randomly. Her behavior is thus a kind of 'noisy best reply' to her past experience:

- o With probability $1 - \varepsilon$ an agent makes a demand that maximizes her expected payoff given her expectations about the opponent's behavior. If several demands maximize expected payoff, they are chosen with equal probability.
- o With probability ε the agent does not optimize, but chooses one of the three demands, H, M, or L, at random.

These rules for matching, belief formation, and behavior define a particular *social dynamic* as a function of the population size N , memory length m , and error rate ε . Notice that it is a Markov process, because there is a well-defined

⁵ Some agents may be active more than once in a particular period, while others are inactive. On average, agents are active once per period.

⁶ Some agents may have larger memories than others, that is, m may be a random variable in the agent population.

probability of moving from any given state s to any other state s' in the next period.

In this model, agents' beliefs evolve according to their particular experiences. Thus, at any given time, the beliefs can be highly heterogeneous because agents will have had different histories of interactions with others. Importantly, moreover, these beliefs may be inconsistent with the actual state of the world. A given agent's experiences may not be representative of behavior in the whole population. For example, one agent, say A , might by chance have been matched against opponents who demanded H in each of the last m periods. Thus A will believe that the next opponent is likely to demand H , so she is very likely to demand L (which is a best reply to H). But another agent, say B , may have been matched against opponents who always demanded M ; for this agent it makes sense to demand M . The reality, however, could be that most people in the population actually plan to demand L , in which case the beliefs of both A and B are at variance with the facts. Moreover, if A is matched against B in the next round, they will make the demands (L, M) with high probability, which is not an equilibrium of the one-shot game bargaining game.

A *social norm* is a self-perpetuating state in which players' memories, and hence their best replies, are unchanging. In other words, it is a rest point or equilibrium of the dynamical system when the error term $\varepsilon = 0$. Consider, for example, the state in which everyone's experience is that opponents always demand M . Then everyone believes that her next opponent will play M . Given these beliefs, M is a best response. Assuming there are no errors ($\varepsilon = 0$), both sides demand M in the next period. Thus, agents' beliefs about opponents turn out to be correct, and this situation perpetuates itself from one period to the next. This is the *equity norm* in which everyone *expects* the other to demand one-half, and as a result everyone *does in fact* demand one-half. Note that this social norm involves no tradeoff between equity and efficiency: the solution is equitable because both sides get equal shares of each pie, and it is efficient because there is no rearrangement of shares that makes all agents better off. It can be verified that, when there are no observable differences among agents, the equity norm is the unique equilibrium of the Nash demand game, and is the unique rest point of the unperturbed social dynamic.

Simplex representation of agent states

We represent the state of the agent population on a simplex with three differently shaded regions, as shown in figure 1. At each time, every agent occupies a position on the simplex that is determined by the content of her memory. For example, an agent who has encountered only agents playing L is located at the lower right vertex of the simplex (labeled 'low'). The shading

within the simplex represents the best reply strategy *given* the agent's memory. That is, since each agent best replies to her memories, an agent's location on the simplex can be thought of as representing her *expectation* about her opponents' play. In the white region L is the best reply since memory configurations here are dominated by H. In the dark gray zone the opposite occurs—memories are dominated by Hs—so L is the best reply. Agents in the light gray zone have memories for opponents playing M, so it is best for them to play M as well.




Figure 1: Memory simplex for one agent type

Starting from different initial states, we can examine various realizations of the process.⁷ Suppose for example that $N = 100$, $m = 10$ and $\epsilon = 0.2$, and the initial state is random about the point of indifference between the three strategies. After 80 periods the process can evolve to the situation shown in Figure 2.⁸ In this new state, all agents have encountered frequent demands of M in the past, and thus they expect their opponents to play M in the next period. Given this expectation, M is the best response. Hence most agents play M next period, which reinforces the expectation of M. However, by a process we do not model explicitly, agents occasionally deviate from best reply and play either H or L. This may occur due to random errors, conscious experimentation, simple imitation or for any number of other reasons. This is analogous to mutation in biological models and serves to create variety in the population.⁹

⁷ A working version of this model is available on the world-wide web. It is written in Java and can be found at the URL: www.brookings.edu/es/dynamics/papers/classes.

⁸ There are less than 100 dots shown in the figures because some agents have the same memory state.

⁹ Each matched agent chooses randomly with probability $\epsilon = 0.20$. However, there is a one-third chance that the random choice will in fact be the best reply, hence the probability that an "error" is realized is 0.1333...




Figure 2: Convergence to the equity norm

If the process is allowed to continue from the state shown in Figure 2, the probability is high that most agents will remain in the light gray region for quite a long period of time. This is because the equity norm has a large basin of attraction, and even substantial deviations caused by random ‘mutations’ in individual behavior may not be enough to tip society into a fundamentally different regime. Nevertheless such tipping events will eventually occur, and they can lead to regimes that have a fundamentally different character.

Such inequitable regimes may also emerge right away when we start from a different initial state. Figure 3 illustrates this for one realization of the process, showing the state after 150 periods.




Figure 3: Emergence of a fractious state

In this *fractious* state, people at each instant are either aggressive or passive; they have not learned to compromise. If, in one's experience, a sufficient proportion of one's opponents are aggressive (demand H), then it is better to submit (play L) than to offer to share equally, and conversely. (It can be checked, in fact, that M is *never* a best response for someone who has never experienced an opponent who played M.) This fractious state persists in excess of 10^9 time periods, though it is neither equitable nor efficient. There is frequent miscoordination in which the players either demand too much (both play H) or they demand too little (both play L) and end up leaving part of the pie on the table. In the state shown in Figure 3, the average share of pie per person in each period is only about one-quarter, or about half the expected share under the equity norm. But while this is an inefficient state, it does not exhibit classes, because agents frequently migrate between zones, sometimes demanding H and sometimes demanding L.

Transitions between regimes

Using asymptotic methods, it can be shown that when both m and N/m are sufficiently large, the probability of being in the equity region is substantially higher than being in the fractious region if one waits long enough and the error rate ε is small. In the terminology of evolutionary game theory, the equity norm is *stochastically stable* (Foster and Young [1990]). The intuitive reason is that it takes much longer to undo the equity norm once it is established, than to undo the fractious regime once it is in place. However, the *inertia* of the system—the waiting time to reach the stochastically stable regime—can be very large indeed. Suppose that we start the agent society off in the fractious regime with $N = 10$, $\varepsilon = 0.10$, and compute the expected number of periods to transit to a neighborhood of the equity norm (namely, to a state where all agents have at least $(1 - \varepsilon)m$ instances of M in their memories). As Figure 4 shows, the waiting time increases exponentially in memory length. For example, when $m = 13$ it takes in excess of 10^5 periods on average for the fractious regime to be displaced in favor of the equity norm.




Figure 4: Transition time between regimes as a function of memory length, $N = 10$, various ϵ

Similarly the transit time increases exponentially with population size, as shown in Figure 5.




Figure 5: Transition time between regimes as a function of population size, $m = 10$, various ϵ

Hence, although the equity norm is stochastically stable, the agent-based computational model reveals that—depending on the number of agents and

the memory length—the waiting time to transit from the fractious regime to the equity norm may be astronomically long.¹⁰

Broken ergodicity¹¹

In figure 4, for $m = 18$, the expected number of time periods the society must wait in order to move from the split regime to the equity norm is $O(10^6)$. In human societies, a million interactions per agent is not realizable. So how are we to interpret such large interaction requirements?

Dynamical systems which are formally ergodic but which possess subregions of the state space that confine the system with high probability over a long time scale are said to display *broken ergodicity* with respect to that time scale.¹² Call $R_{\text{trans}}(m, N, \epsilon)$, the rate of transition from the split state to the equity norm. For example, from figure 4, for $m = 18$, this rate would be approximately 10^{-6} . Now, say that the lifetime of the society is $T \ll 1/R_{\text{trans}}(m, N, \epsilon)$. Then, to a first approximation the probability of regime transition $Pr_{\text{trans}}(T, m, N, \epsilon) = TR_{\text{trans}}(m, N, \epsilon)$. A system has *effective* broken ergodicity if $Pr_{\text{trans}}(T, m, N, \epsilon) < p_0$, where p_0 is some small level of significance, say 0.001. Clearly, the exponential dependence of transition times on memory length and population size implies that our model society displays broken ergodicity.

We can summarize these results as follows. *Occasional random choices create noise in the system, which implies that no state is perfectly absorbing. However, there are two regions of the state space—one equitable, the other fractious—that are very persistent: once the process enters such a region, it tends to stay there for a long period of time. A particular implication is that, while there is only one pure equilibrium of the game (corresponding to the equity norm), it may be difficult for decentralized decision makers to discover this equilibrium from certain initial conditions. Put differently, the*

¹⁰ It is important to note that the expected waiting time depends crucially on the geometry of the interaction structure. In this model we have assumed that agents are paired at random from the whole society. In reality, agents interact in social networks in which there are both local (neighborhood) and global (long range) interactions. The existence of such neighborhood structures can greatly reduce the dependence of the social learning process on population size (Ellison [1993], Young [1998]). Intuitively, the reason is that a local switch in regime—say from fractious to equitable—may be relatively easy because it involves only a small number of agents (the local population size is small). Agent behavior in local interaction models is, however, quite different than in the model described here. Agents repeatedly interact with the same agents in such models, and memory plays no essential role, i.e., interactions are not anonymous.

¹¹ The authors thank Kai Nagel and Maya Paczuski for suggesting the relevance of this concept to our results.

¹² For a review article on broken ergodicity see Palmer [1989].

computation of the equity norm by a decentralized society of agents is "hard" in the sense that it takes exponential time to achieve it from some states.¹³

4 Two Agent Types: The "tag" Model

Thus far agents have been indistinguishable from one another. Even though they have different experiences that lead them to act differently, they look the same to others. Let us now suppose that agents carry a distinguishing tag (e.g., light or dark).¹⁴ The tag is completely meaningless in that agents are identical in competence; for example they have the same amount of memory and follow the same behavioral rule (conditional on experience). However, the presence of the tag allows agents to condition their behavior on the tag of their opponents. To be specific, assume that each agent records in his memory the tag of his opponent and the demand that he made. Faced with a new dark opponent, the agent demands an amount that maximizes the expected payoff against his remembered distribution of dark opponents.¹⁵ Faced with a light opponent, the agent plays a best reply against his remembered distribution of light opponents. All of this happens with high probability, but with some small probability $\varepsilon > 0$ agents make random demands. In this model, the social possibilities are richer than before, since equity or fractiousness can prevail both between and within types.

To fix ideas, assume for the moment that there is no noise in the agents' strategy choice ($\varepsilon = 0$). Define an *intergroup equilibrium* as a state in which each agent in the light group demands x against members in the dark group, each agent in the dark group demands $1 - x$ against each opponent from the light group, and this is true for every previous encounter that each agent remembers. An *intragroup equilibrium* is a state in which everyone demands one-half against members of his own group, and this is true for every previous encounter that each agent remembers.

Using methods from perturbed Markov process theory (Young, [1993a]), it can be shown that when m and N/m are sufficiently large, then the unique stochastically stable state corresponds to the particular case where $x = 1/2$; that is, equity prevails both between and within groups. When ε is sufficiently small, this state or something close to it will be observed with very high probability in the long run. But, as before, there exist fractious states and inequitable norms that have considerable staying power. Furthermore, the

¹³ The view of social systems as distributed computational devices, and the associated characterization of various social problems as computationally hard, are developed more fully in Epstein [1999] and Epstein and Axtell [1996]; see also Shoham and Tennenholtz [1996] and DeCanio and Watkins [1998].

¹⁴ For different uses of tags and tag-like devices in agent-based models see Epstein and Axtell [1996], Holland [1996], and Axelrod [1997].

¹⁵ In the event that an agent has no memory of Blue opponents it picks a random strategy.

dynamics governing the emergence (and dissolution) of inter-group norms differs from that governing intra-group norms.

To study these dynamics computationally, we shall represent events on two simplexes: the one on the right corresponds to agent memory states when playing agents of the *opposite* type—it depicts the inter-group dynamics—while the one on the left displays agent memories for playing agents of the same type—the intra-group dynamics. Black dots refer to dark agents, gray dots to light ones. In each run, there are 100 agents in total, 50 of each type. All agents have memory length 20 and the noise level $\varepsilon = 0.2$. The initial state differs between the runs in order to illustrate the effects of path dependency.

Figure 6 illustrates our first case. Starting from random initial conditions, it depicts the state of the system at time $t = 150$.




Figure 6: Equity between and within types

At this point the process has reached a state where something close to the equity norm prevails both between and within groups. In particular, the process is in the basin of attraction of the equity norm for dark against dark, light against light, and light against dark. Average payoffs in this regime are high, because most agents succeed in dividing the pie rather than fighting over it.

Figure 7 tells a different story. Starting from a different random initial conditions, it shows the system at $t = 150$. Internally, the darks (black dots) have come close to the equity norm while the lights (gray dots) are still in a fractious state. However, something close to the equity norm prevails *between* the lights and the darks.




Figure 7: Equity between, but not within, types

Classes

Yet another history unfolds in figure 8. In this case the process evolves fairly rapidly (after 225 periods) to a state in which the equity norm holds within each group, whereas a discriminatory norm governs relations between the two groups. When agents meet others of their own type, most of them expect to divide the pie in half. But when a dark agent meets a light agent, the darks act aggressively and the lights act passively. The result is that, on average, the payoff to dark agents (70) is over twice as high as it is to light agents (30). In other words, *class distinctions have emerged endogenously*. Once established, such class structures can persist for very long periods of time. The reason is that lights have come to expect that darks will be very demanding, so it is rational to submit to their demands. Similarly, darks have come to expect that lights will submit, so it is rational to take advantage of them.




Figure 8: Equity within, but not between, types

The final case is to us the most interesting. Starting from a different random initial state, society evolves after 260 periods to the state shown in Figure 9. As evident in the right (inter-type) simplex, the darks dominate the lights. However, from the left simplex, it is clear that the equity norm prevails within the dominant darks while the lights are a fractious society. This, then, is the picture of a *divided underclass oppressed by a unified elite*. This result seems particularly disturbing in that every individual is behaving rationally—playing the best reply strategy—and yet the social outcome is far from optimal. Even though this regime does not correspond to a coordination equilibrium of the bargaining game (unlike Figure 8), it may nevertheless persist for long periods of time.




Figure 9: Equity above, division below

Transition Dynamics

Figures 4 and 5 above show how the transition time from the fractious configuration to the equity state depends on the population size and memory length, for various values of ϵ . A similar analysis is possible for the class-like configurations displayed in figure 9. That is, we can start the system off in a configuration with classes and measure how long it takes to transit to the equity norm, as a function of the model parameters. We have not executed such analyses for a simple reason: even for model configurations that should be hospitable to such transitions (e.g., 10 agents of each type, $m = 10$, and $\epsilon = 0.1$), these events are very rare, and thus difficult to systematically investigate. This is in sharp contrast to earlier results where $O(10^3)$ periods were sufficient on average for such equity transitions to occur. The 'basin of attraction' of the class-like configuration is much deeper than the fractious outcome, and the transition times are correspondingly longer. It is an open problem to estimate analytically the expected duration of these transient regimes as a function of the parameters of the process.

5 Summary

Although class systems can certainly arise through outright coercion (Wright [1985]), we have argued that various kinds of social orders—including segregated, discriminatory, and class systems—can also arise through the decentralized interactions of many agents in which accidents of history become reinforced over time. In these path-dependent dynamics, society may self-organize around distinctions that are quite arbitrary from an a priori standpoint. As we have shown, initially meaningless "tags" can acquire socially organizing salience: tag-based classes emerge. While equity norms have an advantage over discriminatory norms in the very long run, computational analysis indicates that long-lived regimes may emerge that are far from equitable, and may be highly inefficient as well.

Appendix: State Space of the Multi-Agent Bargaining Game

This model admits a Markovian formulation. Briefly, call ξ the set of all possible individual memory configurations—each one a string of length m (the memory length) recording the demands (H, M, or L) made by an agent's opponents in the most recent m periods played. In a population of N agents, the state space Z of this process is the set of all possible N -tuples of ξ . The random matching and strategy choice rules then determine a Markov chain with fixed transition probabilities—that is, a $|Z| \times |Z|$ transition matrix, dependent on N , m , and the noise level ε .

The origin of the broken ergodicity displayed by this model for seemingly modest configurations—10 to 100 agents, each of whom has memory length $O(10)$ —arises from the enormous dimension of the state space, Z . For memory length m and three strategies, the number of distinct memory configurations is 3^m . Generally, for S strategies there are S^m memory configurations. For N agents, since individual memories are independent, $|Z| = 3^{Nm}$; S^{Nm} generally. Therefore, the $|Z| \times |Z|$ transition matrix will have 3^{2Nm} entries, S^{2Nm} generally. However, because any individual's memory configuration can only be converted into 9 others in a single interaction (S^2 others generally), the transition matrix is sparse—there are only 3^{2N} transitions possible for each state, thus only $3^{2N} \times 3^{Nm} = 3^{N(m+2)}$ entries in the transition matrix are non-zero; generally, $S^{N(m+2)}$. Table A.1 below gives numerical values for these various quantities as a function of m , for a population of 10 agents ($N = 10$).

	3^m	3^{Nm}	$3^{N(m+2)}$	$S^{Nm}; S = 5$
$m = 2$	9	$\approx 3 \times 10^9$	$\approx 1 \times 10^{19}$	$\approx 1 \times 10^{14}$
$m = 7$	2187	$\approx 3 \times 10^{33}$	$\approx 9 \times 10^{42}$	$\approx 8 \times 10^{48}$
$m = 10$	59,049	$\approx 5 \times 10^{47}$	$\approx 2 \times 10^{57}$	$\approx 8 \times 10^{69}$
$m = 20$	3,486,784,401	$\approx 3 \times 10^{95}$	$\approx 9 \times 10^{104}$	$\approx 6 \times 10^{139}$

Table A.1: Number of memory configurations, dimension of the state space and number of entries in the sparse transition matrix, for $N = 10$, various m

Even for this relatively small population size, most of these quantities are enormous. As a practical matter, a state-of-the-art workstation is not even capable of holding the $m = 2$ state vector in memory, since this would require some 6 GB of RAM at two bytes per entry, a conceivable although untypically large quantity of memory (c. 1999). Furthermore, the corresponding (sparse) transition matrix is so large that it could not be stored by conventional means—its entries would therefore have to be computed as needed.

The situation is vastly worse for a population size of 100. Table A.2 gives the number of memory configurations, dimension of the state space, and the size of the sparse transition matrix, this time for $N = 100$.

	3^m	3^{Nm}	$3^{N(m+2)}$	$S^{Nm}; S = 5$
$m = 2$	9	$\approx 3 \times 10^{95}$	$\approx 7 \times 10^{190}$	$\approx 6 \times 10^{139}$
$m = 7$	2187	$\approx 1 \times 10^{334}$	$\approx 3 \times 10^{429}$	$\approx 2 \times 10^{489}$
$m = 10$	59,049	$\approx 1 \times 10^{477}$	$\approx 4 \times 10^{572}$	$\approx 9 \times 10^{698}$
$m = 20$	3,486,784,401	$\approx 2 \times 10^{954}$	$\approx 5 \times 10^{1049}$	$\approx 9 \times 10^{1397}$

Table A.2: Number of memory configurations, dimension of the state space and number of entries in the transition matrix, for $N = 100$, various m

These quantities are unimaginably large. However, it turns out that it is possible to shrink these sizes significantly. This is because the best reply (BR) rule of the type employed here does not use any information on the order in which past opponents' strategies were encountered. That is, for $m = 6$, memory string (H, H, H, L, L, L) is equivalent to (L, L, L, H, H, H) for purposes of BR; in each the frequency of L and H is 0.5. Because the order of an agent's memories is unimportant—at least to this variant of BR—the number of BR-distinct memory configurations is much smaller than S^m . This permits significant reduction in sizes of the state space and transition matrix of the overall Markov process. Let us call Z , where $|Z| = 3^{Nm}$, the naive state space. We explore this smaller (aggregated) state space, Z' , presently.¹⁶

Call n_L, n_M, n_H the number of low, medium and high memories, respectively, that some particular agent possesses. Because these must sum to m , n_H can be written as $m - n_L - n_M$. Thus, the pair (n_L, n_M) gives all information needed by the agent in order to execute BR. Now, since each $n_{(\cdot)} \in [0, 1, \dots, m]$, the number of distinct memory states is simply $(m+1) + m + (m-1) + \dots + 1 = (m+1)(m+2)/2 = (m^2 + 3m + 2)/2$; for $m = 10$ the total is 66. So, $|Z'| = [(m+1)(m+2)]^N / 2^N$; for $N = m = 10$ the state space has $66^{10} \approx 1.6 \times 10^{18}$ dimensions, which is smaller than the naive state space from Table A.1 of 5×10^{47} by approximately 3×10^{29} . A dense transition matrix for a state space of this size is $[(m+1)(m+2)]^{2N} / 4^N$ in size. But for the problem at hand this is yet sparse—each state can be converted into only 9 others—and thus only $3^{2N} [(m+1)(m+2)]^N / 2^N$ entries need to be stored; for $N = m = 10$, there are some 5×10^{27} non-zero entries, which is a massive reduction from the 3×10^{52} entries of the transition matrix associated with the naive state space.

Unfortunately, the vast reduction in the size of the state space in going from Z to Z' does not make the problem tractable computationally. In particular, consider the case of $S = 3, N = m = 10$. In this instance there are only two recurrent communication classes (see Young [1993a: 68] for definition), one in which all agents are in state (M, M, M, M, M, M, M, M, M, M), the unperturbed equity norm—call it H_1 , $|H_1| = 1$ —and one in which each agent

¹⁶ For more on aggregating Markov processes, see Howard [1971].

has some combination of (only) L's and H's in memory—call this H_2 , and note that there are at most $2^{100} \approx 1.3 \times 10^{30}$ of these states in the naive state space, while in Z' , $|H_2| = 11^{10} \approx 2.6 \times 10^{10}$. Since $|Z'| = 66^{10}$, the number of states outside of both H_1 and H_2 is $66^{10} - 11^{10} - 1 \approx 66^{10} \approx 1.6 \times 10^{18}$. Finding the path with least total resistance between H_1 and H_2 is indeed a shortest path problem [Young 1993a: 69], but in 1.6×10^{18} vertices with approximately $9^{10} \approx 3.5 \times 10^9$ times that many edges, i.e., 5.5×10^{27} total. Now, shortest path problems can be solved in an amount of time linear in the number of vertices + edges (cf., Bertsekas and Tsitsiklis [1989]), but a problem of this magnitude is far beyond the scope of conventional computation.

References

- Axelrod, R. 1997. The Dissemination of Culture: A Model with Local Convergence and Global Polarization. **Journal of Conflict Resolution** 41: 203-226.
- Axtell, R.L. and J.M. Epstein. 1999. Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement. In H. Aaron, ed. **Behavioral Dimensions of Retirement Economics**. Brookings Institution Press: Washington, D.C.
- Bertsekas, D.P. and J.N. Tsitsiklis. 1989. **Parallel and Distributed Computation**. Prentice-Hall: Englewood Cliffs, New Jersey.
- Cole, H.L., G.J. Mailath and A. Postlewaite. 1998. Class Systems and the Enforcement of Social Norms. **Journal of Public Economics** 70: 5-35.
- DeCanio, S.J. and W.E. Watkins. 1998. Information Processing and Organizational Structure. **Journal of Economic Behavior and Organization** 36: 275-294.
- Ellison, G. 1993. Learning, Local Interaction, and Coordination. **Econometrica** 61: 1047-1071.
- Epstein, J.M. 1999. Agent-Based Computational Models and Generative Social Science. **Complexity**.
- Epstein, J.M. and R. Axtell, 1996. **Growing Artificial Societies: Social Science from the Bottom Up**. MIT Press and Brookings Institution Press: Cambridge, Mass. and Washington, D.C.
- Foster, D.P., and H.P. Young, 1990. Stochastic Evolutionary Game Dynamics. **Theoretical Population Biology** 38:219-232.
- Freidlin, M., and A. Wentzell. 1984. **Stochastic Dynamical Systems Theory**. New York: Springer Verlag.
- Harsanyi, J. and R. Selten. 1988. **A General Theory of Equilibrium Selection in Games**. MIT Press: Cambridge, Mass.
- Holland, J.H. 1996. **Hidden Order: How Adaptation Builds Complexity**. Perseus Press.
- Howard, R. 1971. **Dynamic Probabilistic Systems**. Wiley: New York.
- Kandori, M., G. Mailath, and R. Rob, 1993. Learning, Mutation, and Long-Run Equilibria in Games. **Econometrica** 61: 29-56.

- Nash, J. 1950. The Bargaining Problem. **Econometrica** 21: 128-140.
- Palmer, R. 1989. Broken Ergodicity. In **Lectures in the Sciences of Complexity**, D. Stein, ed. Volume 1 in the Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley Longman: Reading, Mass.
- Roemer, J.E. 1982. **A General Theory of Exploitation and Class**. Harvard University Press: Cambridge, Mass.
- Shoham, Y. and M. Tennenholtz. 1996. On Social Laws for Artificial Agent Societies: Off-Line Design. In P.E. Agre and S.J. Rosenschein, eds., **Computational Theories of Interaction and Agency**. MIT Press: Cambridge, Mass.
- Wright, E.O. 1985. **Classes**. Verso: London.
- Young, H.P. 1993a. The Evolution of Conventions. **Econometrica** 61 (1): 57-84.
- . 1993b. An Evolutionary Model of Bargaining. **Journal of Economic Theory** 59 (1): 145-168.
- . 1998. **Individual Strategy and Social Structure**. Princeton University Press: Princeton, New Jersey.