

LEARNING, LOCAL INTERACTION, AND COORDINATION

BY GLENN ELLISON¹

This paper discusses the dynamic implications of learning in a large population coordination game, focusing on the structure of the matching process which describes how players meet. As in Kandori, Mailath, and Rob (1993) a combination of experimentation and myopia creates "evolutionary" forces which lead players to coordinate on the risk dominant equilibrium. To describe play with finite time horizons it is necessary to consider the rates at which the dynamic systems converge. In large populations with uniform matching, play is determined largely by historical factors. In contrast, when players interact with small sets of neighbors it is more reasonable to assume that evolutionary forces may determine the outcome.

KEYWORDS: Learning, neighbors, coordination, rates of convergence.

1. INTRODUCTION

EVEN THE SIMPLEST GAME THEORETIC MODELS all too often have multiple equilibria. A typical example is the coordination game which arises when two players must work together in order to achieve a commonly desired outcome, but in which neither player will benefit from his efforts if his partner does not do his part. We regard the players working together as the "good" equilibrium and speak of coordination failure if it does not occur. In trying to understand play in such games we are led to ask why we should expect players to coordinate on an equilibrium and whether there is any reason to believe that one equilibrium is more likely than the other.

Recent models of Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993) have derived surprisingly strong predictions by exploring the disequilibrium process by which players learn their opponents' play and adjust their strategies over time. For the case of 2×2 coordination games, Kandori, Mailath, and Rob (KMR) (1993) show that the simple combination of random experimentation or mutations and the myopic attempts of players to coordinate with those around them creates powerful dynamic forces which influence the evolution of play over time. In analyzing the long run limit of this dynamic process, they show not only that players will achieve coordination on an equilibrium, but that one particular equilibrium, the "risk dominant equilibrium" will be selected.²

While providing an elegant characterization of the long run influence of evolutionary forces, the KMR analysis is incomplete as a description of

¹I would like to thank Drew Fudenberg, Sara Fisher Ellison, Michi Kandori, Jeff Zwiebel, an editor, and two anonymous referees for their comments. Financial support from the Sloan Foundation is acknowledged.

²See Young (1993) and Kandori and Rob (1992) for a discussion of play in broader classes of games under similar behavioral rules.

economic systems in that it does not consider whether the evolutionary forces would be felt within a reasonable time horizon. In fact, the dynamics justifying their solution are implausible for the large populations we would want to consider to discuss the origins of conventions or focal point in society. In this paper, I show that to assess the relevance of evolutionary forces it is necessary to consider the rate at which play converges to its long run limit. Such an analysis yields two main conclusions. First, the nature (local or global) of the interactions within a population are a crucial determinant of play. Second, when interactions are local, the evolutionary arguments of KMR may be reasonably applied to large populations.

To examine the play of coordination games in large populations I adopt a framework similar to that of KMR. In each period of a dynamic model the players are randomly matched and each pair plays a 2×2 coordination game. The behavioral assumptions incorporate noise and myopic responses by boundedly rational players. The model departs from that of KMR in that it allows for different matching processes within the population.³ I focus on two extremes among the possible matching rules which I will describe as uniform and local. The uniform matching rule is that used in KMR. In contrast, I will describe as local a matching rule in which players interact with a small group of close friends, neighbors, or colleagues.

This paper also departs from KMR in that I consider the rate at which each dynamic process converges. If, as will be the case for one model, a dynamic system takes 10^{100} periods to approach its limit, the limit is not a good prediction for what we will see when the game is repeated a few hundred times. While it is very hard to draw a dividing line and say exactly how fast a system must converge for the limit to be relevant, the models of this paper will exhibit such extreme contrasts that meaningful conclusions are possible. When a system adjusts very slowly, I will conclude that whatever historical factors determine the initial play will continue to determine play long into the future. The model of KMR with a large population is one such example. On the other hand, when a system approaches its limit quickly, I will conclude that we should expect to see the limiting behavior.

The paper is structured as follows. The model is described in Section 2. Section 3 contains some simple examples of the dynamics of the learning model. Section 4 contains the main theoretical results on both limiting distributions of play and rates of convergence. Section 5 discusses numerical simulations of a broader range of specifications.

³ In independent work, Blume (1992) develops a related continuous time model in which players are spatially distributed and interact with a finite set of neighbors. This model is discussed further in Section 4. The importance of particular population structures has long been discussed in evolutionary biology. S. Wright (1931) first argued that small isolated subpopulations allowed more genetic drift and were hence particularly amenable to evolution. See Ewens (1979), especially Sections 1.6, 3.9, and 10.3 for a discussion of more recent work.

2. THE MODEL

2.1. Coordination Games with Bounded Rationality

The model described here has two classes of assumptions: those concerning the nature of the game being played and those describing the particular behavioral rules which players follow.

The basic model is of a repeated game played in periods $t = 1, 2, 3, \dots$. There is a large population of N players (perhaps a few hundred for typical applications). In each period, player i chooses one of two possible actions $a_{it} \in \{A, B\}$. The payoff to player i is given by

$$u_i(a_{it}, a_{-it}) = \sum_{j \neq i} \pi_{ij} g(a_{it}, a_{jt})$$

where the payoffs g are those of the 2×2 coordination game in illustration A. Formally, it is required that $a > d$ and $b > c$ so that (A, A) and (B, B) are both Nash equilibria. In addition, I assume that $(a - d) > (b - c)$ so that (A, A) is the "risk dominant" equilibrium as defined in Harsanyi and Selten (1988). Note that when the strategies have equal security levels ($c = d$), (A, A) is also the Pareto optimum.

In many applications, we will envision the players to be playing a random matching game in which case the weights π_{ij} will represent the probability that players i and j are matched in a given period, and $g(a_{it}, a_{jt})$ gives the payoff to player i when he is matched with player j .

Rather than assuming complete rationality, I simply specify behavioral rules which are a special case of those described in KMR. The rules are intended to capture the intuitive notion that players usually react myopically to their environment. In particular, I assume that in period t player i chooses

$$a_{it} \in \arg \max_{a_i} u_i(a_i, a_{-it-1})$$

with probability $1 - 2\varepsilon$. It is important that player i is reacting to the distribution of play in period $t - 1$, not to the action of one particular opponent with

	A	B
A	a, a	c, d
B	d, c	b, b

ILLUSTRATION A

whom he was matched. Hence, player i must observe his potential opponents' actions, but be fairly naive in predicting how they will play in period t . With probability 2ϵ player i chooses an action at random with 50–50 probability. This randomization is meant to capture the cumulative effect of noise in the form of deliberate experimentation, trembles in strategy choices and the play of new players unfamiliar with the history of the game.

2.2. Local and Uniform Matching Rules

Within the basic framework described above, I will contrast two extreme specifications of the matching process. I term the two types of matching rules uniform and local. The *uniform* matching rule is given by

$$\pi_{ij} = \frac{1}{N-1} \quad \forall j \neq i.$$

The assumption of uniform matching expresses the idea that each player has no information about with whom he would like to coordinate until after he has chosen his action. This assumption might be appropriate, for example, in modeling the interaction of merchants at a large trade fair where each participant would know little about the others. With this rule, a myopic player will choose his period t strategy considering only the fraction of the population playing each strategy at time $t-1$, not the identities of the players using each strategy.

In contrast, I will use the term *local* matching as an informal description of several matching rules in which each player is likely to be matched only with a small fixed subset of the population. For simplicity, I will usually envision the players as being spatially distributed around a circle. In the most extreme local matching rule, each player is only ever matched with one of his two immediate neighbors, i.e.,

$$\pi_{ij} = \begin{cases} \frac{1}{2} & \text{if } i-j \equiv \pm 1 \pmod{N}, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, for any $k \geq 1$ we can define a rule where each player has $2k$ neighbors by

$$\pi_{ij} = \begin{cases} \frac{1}{2k} & \text{if } i-j \equiv \pm 1, \pm 2, \dots, \pm k \pmod{N}, \\ 0 & \text{otherwise.} \end{cases}$$

We could also assign positive probability to any match with the probabilities declining with distance so a player is usually matched with someone nearby, e.g. for N even

$$(1) \quad \pi_{ij} = \begin{cases} \frac{3}{\pi^2} \frac{1}{d^2} & \text{for } d = \min\{|i-j|, N-|i-j|\} \neq \frac{N}{2}, \\ 1 - \frac{3}{\pi^2} \sum_{|i-j| \neq N/2} \frac{1}{d^2} & \text{otherwise.} \end{cases}$$

Local matching rules are appropriate to describe situations where players interact not with the population as a whole, but rather with a few close friends or colleagues. For example, such a rule might describe the interactions at a college reunion where each participant knows in advance who he or she wishes to see. The contacts among a group of firms or economists might also be of this type.

Before moving on, I would like to emphasize two essential features of the local matching rules described above. First, each player assigns a large weight to a small subset of the population. Second, there is considerable overlap in the groups of neighbors so that a player's neighbors' neighbors are likely to be his neighbors as well. The combination of these features allows for the existence of small clusters within the population, each member of which is matched with another member with probability at least $\frac{1}{2}$. The possibility of a new strategy gaining a foothold within one of these clusters allows for a relatively rapid transition to the risk dominant equilibrium.

3. MODEL DYNAMICS

For the remainder of this paper, I discuss the dynamic pattern of play in the model described above. The approach of the paper is as follows. I assume that at some point in the past, arbitrary historical factors determined the initial strategies of the players. The behavioral rules then generate a dynamic system which describes the evolution of players' strategy choices over time. I will formally discuss both the limit of this system and the rates at which the limit is approached. First though, I describe the dynamic evolution of play in a few simple cases in order to motivate subsequent results.

The dynamics of the model with uniform matching are virtually identical to those described in KMR. Let q_i be the fraction of player i 's opponents who played A in period $t-1$. Note that

$$(2) \quad u_i(A, a_{-it-1}) \geq u_i(B, a_{-it-1}) \Leftrightarrow q_i a + (1 - q_i)c \geq q_i d + (1 - q_i)b \\ \Leftrightarrow q_i \geq \frac{b - c}{(a - d) + (b - c)} \equiv q^*.$$

Hence, player i will play A in period t if and only if $q_i \geq q^*$ (assuming player i chooses A when he is indifferent). The assumption that (A, A) is the risk dominant equilibrium implies that $q^* < \frac{1}{2}$. I will frequently discuss the behavior of the model with payoffs $a = 2$, $b = 1$, and $c = d = 0$ so that $q^* = \frac{1}{3}$ and player i will play A if at least $\frac{1}{3}$ of his opponents did so.

At time t , we describe the state s_t of the system by an N -tuple $(a_1, a_2, \dots, a_N) \in S = \{A, B\}^N$ indicating the strategy used by each player. I will write \bar{A} as shorthand for the state (A, A, \dots, A) and \bar{B} for (B, B, \dots, B) . Because all players have nearly identical sets of possible opponents, the dynamics are largely determined by the total number of players playing A , which will be denoted by $A(s_t)$.

If $A(s_t) < [q^*(N-1)]$, then $q_i < q^*$ for all i so all players play B in period $t+1$ and $s_{t+1} = \vec{B}$. ($[x]$ indicates the smallest integer greater than or equal to x .)

If $A(s_t) > [q^*(N-1)]$, then $q_i > q^*$ for all i so all players play A in period $t+1$ and $s_{t+1} = \vec{A}$.

If $A(s_t) = [q^*(N-1)]$, we have a knife edge case where $q_i < q^*$ if player i played A in period t , and $q_i \geq q^*$ if player i played B in period t . The result is that $A(s_{t+1}) = N - [q^*(N-1)]$ as only those players who played B in period t play A in period $t+1$.

While unsightly, this transition does not play a significant role in the subsequent analysis so the reader should not be too troubled by it.

What is important is that the model with uniform matching and no noise has two steady states \vec{A} and \vec{B} corresponding to the Nash equilibria where all players coordinate on one of the strategies. Further, if play starts close to either steady state it immediately jumps to that equilibrium. When noise is introduced the transitions are governed by a Markov process, but once play approaches either equilibrium it will likely remain nearby for a long period of time. To see this, suppose that most players played B in period t so that $A(s_t) < [q^*(N-1)]$. Each player then plays A in period $t+1$ with probability ε . These randomizations are independent and unless $[q^*(N-1)] \varepsilon$ -probability events occur everyone will again have B as a best response in period $t+2$. There is no gradual evolution from one equilibrium to the other, only the possibility of a jump caused by coincident randomizations. When ε is much smaller than q^* we will have to wait a long time for this to occur.

In models of local interaction, the dynamics are dependent on the locations of the players using each strategy in addition to the aggregate frequencies. To illustrate the dynamics of such models, I discuss a typical case where N players are arranged around a circle and each places equal weight on being matched with his eight closest neighbors. Let the payoffs be such that $q^* = \frac{1}{3}$ so that each player has A as his best response whenever at least three of his eight neighbors play A .

Consider first the model with no noise. Clearly there are at least two steady states, \vec{A} and \vec{B} . Each of these steady states has a nontrivial basin of attraction. If all but one or two of the players are playing B at time t , then each player has at least six neighbors playing B and hence will play B in period $t+1$. We may write two such transitions as

$$(A, B, B, \dots, B) \rightarrow \vec{B} \quad (A, A, B, \dots, B) \rightarrow \vec{B}.$$

Similarly, for period t states sufficiently close to \vec{A} we have $s_{t+1} = \vec{A}$.

An important feature of the dynamics is that the basin of attraction of \vec{B} is relatively small. In particular, the existence of a small cluster of players playing A is sufficient to ensure that the dynamic process will eventually lead all players to play A . Suppose the period t state is $(A, A, A, A, B, \dots, B)$ so that players 1 through 4 played A . Players 1 through 6 and players N and $N-1$ all have at least three neighbors playing A . Those eight players will play A in period $t+1$. In period $t+2$, players $N-2$, $N-3$, 7, and 8 will switch to playing A . The cluster of players playing A will grow until eventually the state \vec{A} is reached.

In contrast, any relatively small cluster of players playing B will disappear over time. It is easy to verify that

$$(B, B, B, B, A, A, A, \dots, A) \rightarrow \vec{A},$$

$$(B, B, B, B, B, B, B, A, \dots, A) \rightarrow (A, A, B, B, B, A, \dots, A) \rightarrow \vec{A}.$$

The basin of attraction of \vec{A} is much larger than that of \vec{B} .

It is the differing sizes of these basins of attraction which causes the relatively rapid convergence of play to a limit concentrated around \vec{A} once noise is introduced. From the dynamics above, it should be clear that we usually only need to wait for four well placed randomizations to create a cluster of players playing A and lead us away from an initial condition where everyone is playing B . When the number of players is large, seeing four adjacent randomizations is far more likely than seeing the $[(N-1)/3]$ simultaneous randomizations required to shift play in the model with uniform matching.

The extreme local matching rule in which each player has only two neighbors is neither an apt description of any economic systems, nor does it have particularly compelling dynamic behavior. Nonetheless, it is the easiest model of local interaction to analyze, and hence reappears throughout this paper. I briefly discuss its dynamics here.

First, note that regardless of the payoffs, the assumptions that (A, A) is the risk dominant equilibrium entails that each player will have A as his best response whenever at least one of his two neighbors plays A . In a model with no noise, we have two steady states, \vec{A} and \vec{B} . There is also one stable cycle when N is even,

$$\begin{aligned} (A, B, A, B, \dots, A, B) &\rightarrow (B, A, B, A, \dots, B, A) \\ &\rightarrow (A, B, A, B, \dots, A, B), \end{aligned}$$

because players in period t myopically respond to their opponents' play in period $t-1$, not to a forecast of their period t play. I will write \vec{AB} for the state $(A, B, A, B, \dots, A, B)$ and \vec{BA} for the state $(B, A, B, A, \dots, B, A)$. The most important aspect of the dynamics is that the steady state \vec{B} now has no other states in its basin of attraction. If at least one player plays A in period t , then at least two players (his neighbors) will play A in period $t+1$. Any state which contains a cluster of two adjacent players playing A lies in the basin of attraction of \vec{A} . Once noise is introduced, this leads to rapid convergence to a steady state concentrated around \vec{A} .

4. LIMITS AND RATES OF CONVERGES

In this section, I discuss the principal theoretical results of the paper. As mentioned above, the motivation for the analysis here is the assumption that at some point the initial actions of the players were determined by historical factors and that for some subsequent period of time play has evolved according to the behavioral rules specified above. The fundamental problem is then to

determine how historical and evolutionary forces combine to determine the play we observe. To this end, I first discuss the limiting behavior of these systems as the number of periods of evolution grows to infinity. Subsequently, I discuss the rates at which the limits are approached in order to assess whether the limits are meaningful given that the economic systems modeled involve only some reasonable finite repetition of play.

As noted above, we may view the time t strategy profiles as the states s_t of a Markov process. We may represent the time t probability distribution over the states by an 1×2^N vector v_t . The evolution of the process is governed by

$$v_{t+1} = v_t P(\varepsilon)$$

where $P(\varepsilon)$ is the transition matrix whose elements are given by

$$p_{ij}(\varepsilon) = \text{Prob}\{s_{t+1} = j | s_t = i\}.$$

Write $P^u(\varepsilon)$ for the transition matrix with uniform matching and $P^{2k}(\varepsilon)$ for that of the model with $2k$ neighbor matching. For example, we have

$$p_{ij}^u(\varepsilon) = \varepsilon^{A(j)}(1 - \varepsilon)^{N-A(j)}$$

whenever $i < [q^*(N - 1)]$ as each player's best response is B so state j arises when a particular $A(j)$ ε -probability randomizations occur. Note that each transition matrix $P(\varepsilon)$ is strictly positive for $\varepsilon > 0$ so by standard results on Markov processes there is an unique steady-state distribution $\mu(\varepsilon)$ such that

$$\mu(\varepsilon) = \mu(\varepsilon)P(\varepsilon).^4$$

The steady-state distribution reflects the long-run behavior of the process. I write $\mu^u(\varepsilon)$ and $\mu^{2k}(\varepsilon)$ for the steady states with uniform and $2k$ neighbor matching, and $\mu_s(\varepsilon)$ or $\mu(\varepsilon)(s)$ for the probability assigned to state s by the distribution $\mu(\varepsilon)$.

Because we are interested in small ε , we shall consider the asymptotic behavior of $\mu(\varepsilon)$ as $\varepsilon \rightarrow 0$. We write $f(x) = O(g(x))$ ($x \rightarrow 0$) if there exists a constant C such that $|f(x)| \leq C|g(x)|$ for sufficiently small x . The formula $f(x) = O(g(x))$ is said to be best possible if in addition there exists a constant $c > 0$ such that $c|g(x)| \leq |f(x)| \leq C|g(x)|$. In each of the theorems of this section the O -approximations are meant in this sense.

The first result compares the steady state distributions of the uniform and $2k$ neighbor models. KMR show that the evolutionary forces in a model virtually identical to the uniform model yield a steady state limit in which the risk dominant equilibrium (A, A) is played with very high probability. The statement that $\mu_A^u(\varepsilon) \rightarrow 1$ and $\mu_A^{2k}(\varepsilon) \rightarrow 1$ verifies this both for the uniform model I have defined and for the model with $2k$ neighbor matching. The second part of the theorem discusses the asymptotics of $\mu(\varepsilon)$ in order to give a more precise comparison of the relative probabilities with which the entire population coordinates on the equilibrium (B, B) . For sufficiently small ε , this equilibrium is even

⁴ See Karlin and Taylor (1975, p. 85).

less common in the model with two neighbor matching than it is under uniform matching (although it is extremely rare in both models).

THEOREM 1: *Let $\mu^u(\varepsilon)$ and $\mu^{2k}(\varepsilon)$ be the steady state distributions of the general model of Section 2 under the uniform and $2k$ neighbor matching rules, respectively. Then, for N sufficiently large*

- (a) $\lim_{\varepsilon \rightarrow 0} \mu_A^u(\varepsilon) = 1,$
 $\lim_{\varepsilon \rightarrow 0} \mu_A^{2k}(\varepsilon) = 1;$
- (b) $\mu_B^u(\varepsilon) = O(\varepsilon^{N-2\lfloor q^*(N-1)\rfloor+1}),$
 $\mu_B^{2k}(\varepsilon) = \begin{cases} O(\varepsilon^{N-2}) & \text{for } N \text{ even,} \\ O(\varepsilon^{N-1}) & \text{for } N \text{ odd.} \end{cases}$

PROOF: Both the statement and the proofs of the results for the uniform model are virtually identical to those given in KMR, and hence I will omit the proofs. N must be sufficiently large that $\lfloor q^*(N-1) \rfloor < N/2$.

The proof for the $2k$ neighbor model relies on a characterization of the steady state used by Foster and Young (1990) and KMR. The reader may refer to those papers or to Freidlin and Wentzell (1984) for an exposition of the background material. An x -tree t on S is a function $t : S \rightarrow S$ such that $t(x) = x$ and such that for all $s \neq x$ there exists m with $t^m(s) = x$. We may think of an x -tree as a set of arrows connecting elements of S in which every element has an unique successor and all paths eventually lead to x . The steady state distribution $\mu^{2k}(\varepsilon)$ is characterized by

$$\mu_x^{2k}(\varepsilon) = c(\varepsilon) \sum_{t \in H_x} \prod_{i \neq x} p_{it(i)}^{2k}(\varepsilon)$$

where H_x is the set of x -trees on S . Note that $p_{ij}^{2k}(\varepsilon)$ is a polynomial in ε whose constant term is nonzero if and only if the transition $i \rightarrow j$ occurs in the model with no noise ($\varepsilon = 0$). For any state x , the expression above allows us to express the quantity $\mu_x^{2k}(\varepsilon)/\mu_A^{2k}(\varepsilon)$ as a ratio of polynomials in ε .

To get the result in (a), it suffices to show that $\mu_x^{2k}(\varepsilon)/\mu_A^{2k}(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$ for all $x \neq \vec{A}$. This will follow if we demonstrate that for $N > (k+1)(k+2)$ and any x -tree t ($x \neq \vec{A}$) we have $\prod_{i \neq x} p_{it(i)}^{2k}(\varepsilon) / \sum_{t' \in H_{\vec{A}}} \prod_{i \neq \vec{A}} p_{it'(i)}^{2k}(\varepsilon) \rightarrow 0$. This in turn follows if we show that there exists an \vec{A} -tree t' such that $\prod_{i \neq x} p_{it(i)}^{2k}(\varepsilon) / \prod_{i \neq \vec{A}} p_{it'(i)}^{2k}(\varepsilon) \rightarrow 0$. Write $s(z)$ for the successor of state z in the model with no noise and $D(\vec{A}) = \{z | s^m(z) = \vec{A} \text{ for some } m\}$ for the basin of attraction of A . I show that the latter ratio converges to zero by considering two cases:

Case 1: $x \in D(\vec{A})$. Define t' by

$$t'(z) = \begin{cases} s(z) & \text{if } z \in D(\vec{A}), \\ t(z) & \text{otherwise.} \end{cases}$$

t' is an \vec{A} -tree because for any state z the path $z \rightarrow t'(z) \rightarrow t'^2(z) \rightarrow \dots$ initially coincides with t and hence eventually enters $D(\vec{A})$. From the first point at which $t'^m(z) \in D(\vec{A})$, the tree maps every point to its successor in the no noise model and hence reaches \vec{A} . The ratio is

$$\frac{\prod_{i \neq x} p_{it(i)}^{2k}(\varepsilon)}{\prod_{i \neq \vec{A}} p_{it'(i)}^{2k}(\varepsilon)} = \frac{p_{\vec{A}t(\vec{A})}^{2k}(\varepsilon)}{p_{xst(x)}^{2k}(\varepsilon)} \frac{\prod_{i \in D(\vec{A}) - \{\vec{A}, x\}} p_{it(i)}^{2k}(\varepsilon)}{\prod_{i \in D(\vec{A}) - \{\vec{A}, x\}} p_{is(i)}^{2k}(\varepsilon)}.$$

The right side converges to 0 as $\varepsilon \rightarrow 0$ because $p_{it(i)}^{2k}(\varepsilon)/p_{is(i)}^{2k}(\varepsilon)$ is bounded, $p_{\vec{A}t(\vec{A})}^{2k}(\varepsilon) \rightarrow 0$, and $p_{xst(x)}^{2k}(\varepsilon) \rightarrow 1$.

Case 2: $x \notin D(\vec{A})$. Define t' by

$$t'(z) = \begin{cases} s(z) & \text{if } z \in D(\vec{A}), \\ t(z) & \text{if } z \notin D(\vec{A}), z \neq x, \\ y & \text{if } z = x, \end{cases}$$

where y is an element of $D(\vec{A})$ such that $p_{xy}^{2k}(\varepsilon)$ is of minimum order. t' is again an \vec{A} -tree. The ratio is

$$\frac{\prod_{i \neq x} p_{it(i)}^{2k}(\varepsilon)}{\prod_{i \neq \vec{A}} p_{it'(i)}^{2k}(\varepsilon)} = \frac{\prod_{i \in D(\vec{A})} p_{it(i)}^{2k}(\varepsilon)}{p_{xy}^{2k}(\varepsilon) \prod_{i \in D(\vec{A}) - \{\vec{A}\}} p_{is(i)}^{2k}(\varepsilon)}.$$

The denominator is of ε -order at most $k+1$ because $p_{is(i)}^{2k} \rightarrow 1$ and the transition $x \rightarrow y$ requires at most $k+1$ randomizations. (Any state with $k+1$ adjacent players playing \vec{A} is in $D(\vec{A})$.) To see that the numerator is of higher order consider the product $\prod_{j=0}^{m-1} p_{t^j(\vec{A})t^{j+1}(\vec{A})}^{2k}(\varepsilon)$ where m is the smallest integer such that $t^m(\vec{A}) \notin D(\vec{A})$. For $i=0, 1, \dots, k+1$, let $r(i)$ be the first time period such that $t^{r(i)}(\vec{A})$ does not have all players in the set $\{i(k+1)+1, \dots, i(k+1)+k+1\}$ playing A . The transition $t^{r(i)-1}(\vec{A}) \rightarrow t^{r(i)}(\vec{A})$ involves a mutation among the players in that set because each of them had k neighbors playing A in period $r(i)-1$. Hence, there are at least $k+2$ distinct ε -probability events in this path and the numerator is of order at least $k+2$.

For part (b), we find the minimum order \vec{A} - and \vec{B} -trees explicitly. For example, for two neighbor matching and $N \geq 4$ even, one can show that the minimum order \vec{A} -tree has $t(\vec{B}) = (A, B, B, \dots, B)$, $t(\vec{A}\vec{B}) = (B, A, A, A, B, A, \dots, B, A)$, and all other states mapped to their successors in the no noise model. The minimum order \vec{B} -tree has $t(\vec{A}) = \vec{A}\vec{B}$ and $t(\vec{A}\vec{B}) = \vec{B}$. These trees are of order 2 and N respectively. Q.E.D.

REMARKS: 1. The fact that models with local and uniform matching both have steady state distribution concentrated on \vec{A} is more general than the hypotheses of the theorem. The proof does not rely on the fact that the matching

distribution has finite support. It would suffice for the set of neighbors with whom a player is matched with probability $2q^*$ to be finite or grow slowly as $N \rightarrow \infty$. More precisely, it suffices that there exist $k'(N)$ such that $(k'(N) + 1)(k'(N) + 2) < N$ and $\sum_{j=i-k'(N)}^{i+k'(N)} \pi_{ij} > q^*$ for all $i' \in \{i - k'(N), \dots, i + k'(N)\}$. The matching rule given in (1) is one example. Further extensions are no doubt possible.

2. The steady-state distribution is *not* concentrated on \vec{A} if the matching rule is too concentrated. If $\pi_{ij} > 1 - q^*$ the steady-state probability of the cycle in which players i and j alternately play (A, B) and (B, A) while all other players play A cannot vanish relative to the probability of \vec{A} , because a single ε -probability event leads from \vec{A} to the cycle and vice-versa.

3. The long-run outcome of models with uniform and $2k$ neighbor matching may differ once we move beyond 2×2 games. For example, Young (1993) notes that in a similar model the long-run outcome in a 3×3 game is a complicated function of the payoffs. Modifying his example, it can be shown that with uniform matching the steady-state distribution of the game on the left in illustration B is concentrated on \vec{B} while that of the game on the right is concentrated on \vec{C} . Intuitively, the minimum order \vec{B} - and \vec{C} -trees are of nearly equal order in the two games, because while C is pairwise risk dominant B has a greater advantage against A . Because the orders are nearly equal, it is possible to reverse their ranking with only a small change in the payoffs. In contrast, with two neighbor matching the dynamics (and hence the steady state) are completely determined by the best responses to the six possible configurations of a player's neighbors: $\{A, A\}, \{A, B\}, \{A, C\}, \{B, B\}, \{B, C\}, \{C, C\}$. Each of these best responses are identical across the two games, so the outcome with two neighbor matching cannot resemble that with uniform matching in both games.

Theorem 1 implies that if the coordination games we have described are repeated enough times, we expect eventually to see the risk dominant equilibrium played almost all the time. It remains to be seen, however, whether this 'eventually' is relevant. If, for example, we are modeling weekly interactions, evolutionary forces will only be felt within a few decades if they affect play in

	A	B	C
A	6, 6	0, 5	0, 0
B	5, 0	7, 7	5, 5
C	0, 0	5, 5	8, 8

	A	B	C
A	6, 6	0, 5	0, 1
B	5, 0	7, 7	5, 5
C	1, 0	5, 5	8, 8

ILLUSTRATION B

the first few thousand periods. In applying evolutionary arguments to games which are not repeated infinitely often, we must ask whether the number of repetitions is sufficiently large for play in the game to resemble the steady-state limit we have identified. I begin with some theoretical results on the rates of convergence of models with uniform and local matching rules.⁵

Let ρ be a probability distribution describing play in the initial period of the game. For the Markov process with transition matrix $P(\varepsilon)$, the distribution of period t play is given by $\rho P(\varepsilon)^t$. The steady-state distribution reflects long run behavior in the sense that

$$\rho P(\varepsilon)^t \rightarrow \mu(\varepsilon)$$

for any ρ as $t \rightarrow \infty$. For any two distributions μ and ν define

$$\|\mu - \nu\| \equiv \max_{s \in S} |\mu_s - \nu_s|.$$

The quantity $\|\rho P(\varepsilon)^t - \mu(\varepsilon)\|$ then measures the difference between the distribution of play in period t and the steady state. In discussing the rate of convergence we are trying to say when this difference is small.

To do so, I characterize the asymptotic behavior of $\|\rho P(\varepsilon)^t - \mu(\varepsilon)\|$ as $t \rightarrow \infty$. For finite state Markov processes, convergence is always at an exponential rate. (See Seneta (1973, p. 8).) Loosely speaking, we have $\|\rho P(\varepsilon)^t - \mu(\varepsilon)\| = O(r^t)$ for some $r < 1$. The distinction I draw between fast and slow convergence is simply a quantitative one comparing different values of r . Meaningful comparisons are nonetheless possible, because the values of r turn out to be very different.⁶

The principal implication of Theorem 2 is that when ε is small, r will be much closer to one in the model with uniform matching than it is in the model with two neighbor matching. To see what kind of effect this has on convergence, consider the following numerical example. Suppose there are 100 players with payoffs such that $[q^*(N-1)] = 33$. Suppose also that the randomization probability ε is small and consider the effect of reducing ε to $\varepsilon/2$. In the two neighbor model, the model will take about twice as long to converge. In the model with uniform matching, Theorem 2 tells us that $1 - r^u(\varepsilon/2) \approx 2^{-33}(1 - r^u(\varepsilon))$. If $1 - r^u(\varepsilon)$ is small we have the first order approximation

$$r^u\left(\frac{\varepsilon}{2}\right)^{2^{33}} \approx \left(1 - \frac{1 - r^u(\varepsilon)}{2^{33}}\right)^{2^{33}} \approx r^u(\varepsilon).$$

Hence, convergence in the uniform model will take not twice as many periods but rather 2^{33} or over 8 billion times as many.

⁵ See Karlin and Taylor (1975, Ch. 3 and Appendix 2), or Seneta (1973, Sections 1.1 and 4.2) for formal statements of the theorems mentioned below.

⁶ The idea of such a distinction may be puzzling, especially to econometricians who are used to seeing distributions converge at rate $1/\sqrt{t}$ and think of exponential convergence as rapid. What is important to note is that exponential convergence can be quite slow in a practical sense if r is extremely close to one. For example, if $r = 0.9999999$, then $r^{1,000,000} \approx 0.9$, so with every million periods play gets only 10% closer to the steady state.

THEOREM 2: Let $P^u(\varepsilon)$ and $P^2(\varepsilon)$ be the transition matrices for the uniform and two neighbor models and let $\mu^u(\varepsilon)$ and $\mu^2(\varepsilon)$ be the associated steady state distributions. Assume $|q^*(N-1)| < N/2$. For Δ the set of probability distributions on S define

$$r^u(\varepsilon) \equiv \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho P^u(\varepsilon)^t - \mu^u(\varepsilon)\|^{1/t},$$

$$r^2(\varepsilon) \equiv \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho P^2(\varepsilon)^t - \mu^2(\varepsilon)\|^{1/t}.$$

Then,

$$1 - r^u(\varepsilon) = O(\varepsilon^{\lfloor q^*(N-1) \rfloor}),$$

$$1 - r^2(\varepsilon) = O(\varepsilon),$$

as $\varepsilon \rightarrow 0$.

PROOF: To begin we simplify the right hand side of the expressions defining $r^u(\varepsilon)$ and $r^2(\varepsilon)$ using standard matrix theorems. Let P be any strictly positive transition matrix and let μ be its unique steady state. The first result is that if we order the eigenvalues of P so that $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_N|$, then $\lambda_1 = 1$ and $|\lambda_2| < 1$. From this follows a result directly applicable to our problem, namely that

$$\sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho P^t - \mu\|^{1/t} = |\lambda_2|.$$

The proof of this second result for a diagonalizable P is quite simple and is sketched below.⁷ Suppose

$$P = \Phi \Lambda \Phi^{-1},$$

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & \lambda_N \end{bmatrix}.$$

For any fixed ρ ,

$$\rho P^t = \rho \Phi \Lambda^t \Phi^{-1}.$$

Hence,

$$\rho P^t - \rho \Phi \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix} \Phi^{-1} = \rho \Phi \begin{bmatrix} 0 & 0 & \dots & 0 \\ 0 & \lambda_2^t & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_N^t \end{bmatrix} \Phi^{-1}.$$

⁷ For a nondiagonalizable P a similar argument can be constructed using the Jordan canonical form. See Karlin and Taylor (1975, pp. 542-551) or Seneta (1973, pp. 1-8) for more thorough discussions and alternate proofs.

The right-hand side of this expression converges to zero. Writing E_{ij} for the matrix whose ij th element is 1 and with all other elements equal to zero, we have from the uniqueness of the steady state that

$$\rho \Phi E_{11} \Phi^{-1} = \mu$$

for all distributions ρ . Hence,

$$\begin{aligned} \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho P^t - \mu\|^{1/t} &= \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho P^t - \rho \Phi E_{11} \Phi^{-1}\|^{1/t} \\ &= \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \left\| \lambda_2^t \rho \Phi \left(\frac{\Lambda^t - E_{11}}{\lambda_2^t} \right) \Phi^{-1} \right\|^{1/t} \\ &= |\lambda_2| \sup_{\rho \in \Delta} \limsup_{t \rightarrow \infty} \|\rho \Phi E_{22} \Phi^{-1}\|^{1/t} \\ &= |\lambda_2|. \end{aligned}$$

We also see from this calculation that the supremum is in fact achieved for any ρ such that $\rho \Phi E_{22} \Phi^{-1} \neq 0$.

Given this result, the problem of finding $r^u(\varepsilon)$ and $r^2(\varepsilon)$ is reduced to the problem of finding the second largest eigenvalues of the matrices $P^u(\varepsilon)$ and $P^2(\varepsilon)$. The remainder of the proof is the rather lengthy solution to this problem and can be found in the Appendix. *Q.E.D.*

To better illustrate the behavior of models with local and uniform matching I define an alternate measure of the extent to which play resembles its long run limit by

$$W(N, \varepsilon, \alpha) = E\left(\min\{t | A(s_t) \geq (1 - \alpha)N\} | s_0 = \bar{B}\right),$$

i.e. $W(N, \varepsilon, \alpha)$ is the expected waiting time until at least $1 - \alpha$ of the players simultaneously play A given that everyone starts off playing B . Table I lists the value $W^u(N, \varepsilon, 0.25)$ for the uniform model for two sets of payoffs: the first $a = 2, b = 1, c = d = 0$ as before, and the second a more extreme example where the payoff to (A, A) has been increased to 5. (See the proof of Theorem 3 for the computation of $W^u(N, \varepsilon, \alpha)$. It is approximately $1/(1 - r^u(\varepsilon))$ for ε small or N large.) Table II presents Monte Carlo estimates of the corresponding waiting times $W^2(N, \varepsilon, 0.25)$ for the model with two neighbor matching. Comparing the tables, it is clear that the nature of the matching rule has a tremendous effect on the pattern of play. Moreover, this basic insight is quite robust with extreme contrasts for values of ε which are far from infinitesimal. In the uniform model, it takes an extremely long time for play to shift from one equilibrium to the other, so for economically reasonable horizons the evolutionarily determined limits do not apply. This is not just a negative result. We can predict that play should exhibit great inertia with a historically determined equilibrium repeated over and over again. In the two neighbor model play shifts to the risk dominant equilibrium within six to fourteen periods, suggesting that evolutionary forces

TABLE I
EXPECTED WAITING TIMES WITH UNIFORM MATCHING^a

	Expected Wait ($\alpha = 2$)		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
	$N = 10$	611	89
$N = 50$	3.83×10^{14}	6.54×10^9	2.63×10^5
$N = 100$	1.30×10^{27}	8.13×10^{17}	3.09×10^9
$N = 1000$	1.09×10^{266}	1.96×10^{173}	4.46×10^{88}

	Expected Wait ($\alpha = 5$)		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
	$N = 10$	42	13
$N = 50$	2.65×10^5	1324	19
$N = 100$	1.86×10^9	1.06×10^5	49
$N = 1000$	1.61×10^{82}	1.82×10^{41}	2.16×10^{10}

^a N = population size, ϵ = randomization probability, α = payoff to (A, A) . Other payoffs are $u(B, B) = 1$, $u(A, B) = u(B, A) = 0$. Waiting times ($W^u(N, \epsilon, 0.25)$) are until 3/4 of the players play A simultaneously.

TABLE II
EXPECTED WAITING TIMES WITH TWO NEIGHBOR MATCHING^a

	$W^2(N, \epsilon, 0.25)$		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
	$N = 10$	14.5	9.0
$N = 50$	11.0	8.1	6.3
$N = 100$	11.1	8.2	6.4
$N = 1000$	11.0	8.1	6.2

^a Waiting times are independent of payoffs for $q^* \in (0, 1/2)$. Estimated standard errors are 0.1 or less for all estimates.

will be a powerful determinant of play. I comment later on the robustness of this result.

I noted in the introduction that models with local interaction are of particular interest if they allow evolutionary arguments to be used in discussing the development of conventions in large societies. I therefore ask now how the rate at which play converges depends on the size of the population. Theorem 3 characterizes play in terms of the expected waiting time until a transition to the risk dominant equilibrium occurs. In the uniform model, $W^u(N, \epsilon, \alpha)$ increases dramatically in N reflecting our earlier intuition that the simultaneous randomizations which produce a transition become extremely unlikely in large populations. In contrast, with a local matching rule play approaches the risk dominant equilibrium within a period of time bounded above by a constant independent of N . Intuitively, the proof formalizes the idea that players who are sufficiently far from player l cannot significantly slow his adoption of strategy A by considering the wait until player l plays A in a constrained model where all players m or more units away play B in every period. Convergence times in the

constrained model are independent of N and are shown to provide an upper bound on convergence times in the standard local matching model.

THEOREM 3: For ε sufficiently small we have

$$W^u(N, \varepsilon, \alpha) = O(\sqrt{N} e^{((q^* - \varepsilon)/\varepsilon(1 - \varepsilon))^N}),$$

$$W^{2k}(N, \varepsilon, \alpha) = O(1),$$

as $N \rightarrow \infty$.

PROOF: In the model with $2k$ neighbor matching it suffices to show that for $\varepsilon < \bar{\varepsilon}$ there exist constants \underline{N} , η , and T , such that for all $N \geq \underline{N}$

$$\min_s \text{prob}(A(s_{t+T}) \geq (1 - \alpha)N | s_t = s) > \eta.$$

From this it immediately follows that $W^{2k}(N, \varepsilon, \alpha) < T/\eta$ which provides the desired bound.

To begin, consider a constrained model with $2m + 1$ players on a line of whom players $-m, -m + 1, \dots, -m + k - 1$ and $m - k + 1, \dots, m$ play B in every period while the rest of the players play a best response to their $2k$ closest neighbors with probability $1 - \varepsilon$. Write $P^{2k'}$ for the transition matrix, μ' for the steady-state distribution, and \vec{A}' for the state where all unconstrained players play A in this model. For $2m + 1 > (k + 1)(k + 2) + 2k$ a proof identical to that of Theorem 1(a) shows $\mu'_{\vec{A}'}(\varepsilon) \rightarrow 1$ as $\varepsilon \rightarrow 0$. Pick $\bar{\varepsilon}$ such that $\varepsilon < \bar{\varepsilon}$ implies $\mu'_{\vec{A}'}(\varepsilon) > 1 - \alpha/2$. Now fix $\varepsilon < \bar{\varepsilon}$. Because $P^{2k'}$ is ergodic we can choose T such that $e_{\vec{B}'} P^{2k'}(\varepsilon)^T(\vec{A}') > 1 - 3\alpha/4$ for $e_{\vec{B}'}$ an initial distribution putting mass 1 on \vec{B}' .

Let Ω be the underlying probability space for the circle model with $\omega \in \Omega$ consisting of a uniform $[0, 1]$ random variable ω_{it} for each player at each point in time such that player i mutates to A in period t if $\omega_{it} > 1 - \varepsilon$ and to B if $\omega_{it} < \varepsilon$. Ω also acts as a probability space for the constrained model with the randomizations of player j in that model corresponding to the randomizations of player $l + j$ in the standard model on the circle. The constrained player 0 plays A in a given period only if player l on the circle does so for the same realization of ω with corresponding initial conditions. Hence,

$$\begin{aligned} E(A(s_{t+T}) | s_t = s) &\geq N \text{prob}(\text{Player 0 plays } A \text{ in period } t + T | s'_t = \vec{B}') \\ &\geq N e_{\vec{B}'} P^{2k'}(\varepsilon)^T(\vec{A}') \geq (1 - 3\alpha/4)N. \end{aligned}$$

Using $A(s_{t+T}) \leq N$ we have the desired inequality

$$\text{prob}(A(s_{t+T}) \geq (1 - \alpha)N | s_t = s) > \frac{1}{4}.$$

In the uniform model, let N be sufficiently large so that $[q^*(N - 1)] < N/2$. The probability that $A(s_{t+1}) > (1 - \alpha)N$ depends only on whether $A(s_t)$ is in $\{0, \dots, [q^*(N - 1)] - 1\}$, $\{[q^*(N - 1)]\}$, or $\{[q^*(N - 1)] + 1, \dots, N\}$. Write η_0 ,

η_1 , and η_2 for these three probabilities. Note that $\eta_0 \rightarrow 0$, $\eta_1 \rightarrow 0$, and $\eta_2 \rightarrow 1$, as $N \rightarrow \infty$ for $\varepsilon < \alpha < q^*$. Write W_0 , W_1 , and W_2 for the expected wait until $A(s_{i+\tau}) > (1 - \alpha)N$ conditional on s_i lying in each of these three sets. ($W_0 = W^u(N, \varepsilon, \alpha)$.) We have

$$\begin{aligned} W_0 &= (1 - (p_1 + p_2))(W_0 + 1) + p_1(W_1 + 1) \\ &\quad + (p_2 - \eta_0)(W_2 + 1) + \eta_0, \\ W_1 &= q_1(W_0 + 1) + (1 - (q_1 + q_2))(W_1 + 1) \\ &\quad + (q_2 - \eta_1)(W_2 + 1) + \eta_1, \\ W_2 &= r_1(W_0 + 1) + r_2(W_1 + 1) + (1 - (r_1 + r_2 + \eta_2))(W_2 + 1) + \eta_2, \end{aligned}$$

with p_i , q_i , and r_i as in (A1). Solving these three equations gives a closed form expression for $W^u(N, \varepsilon, \alpha)$ from which it follows that

$$W_0 = \frac{1}{p_1 + p_2} + o\left(\frac{1}{p_1 + p_2}\right).$$

From the central limit theorem

$$p_1 + p_2 \sim 1 - \Phi\left(\frac{[q^*(N - 1)] - \varepsilon N}{\sqrt{N\varepsilon(1 - \varepsilon)}}\right) \quad (N \rightarrow \infty).$$

The desired result now follows by replacing the right-hand side of the expression above by the approximation (see Ross (1984, p. 162)):

$$1 - \Phi(x) \sim \frac{1}{x\sqrt{2\pi}} e^{-x^2/2} \quad (x \rightarrow \infty). \quad Q.E.D.$$

REMARKS: 1. As with Theorem 1, the proof that $W^{2k}(N, \varepsilon, \alpha) = O(1)$ applies to more general matching rules. (See Remark 1 after Theorem 1.) However, because the value of $\bar{\varepsilon}$ used in the construction must be independent of N , the proof extends only to matching rules such as (1) where $\pi_{ij}^N = f^N(d(i, j))$, $f^N(i) \geq f(i)$ for N large, and f is monotone decreasing with $\sum_{i=1}^\infty f(i) = 1/2$.

2. Blume (1992) specifies a similar model of interactions among spatially distributed players in a way which allows tools from statistical mechanics to be applied. Principally, his model differs in that the population is countable, strategies are updated asynchronously in continuous time, and randomization probabilities are related to payoff differences. While Blume does not directly address the rate at which play converges in his model he is able to draw on an extensive literature to determine whether behavior is ergodic. The obvious conjecture is that play in an infinite population model is ergodic whenever convergence times in suitably defined finite population approximations remain finite as $N \rightarrow \infty$.

Figure 1 illustrates the dependence of waiting times on the population size for three local matching rules. Monte Carlo estimates of $W^{2k}(N, \varepsilon, 0.25)$ are

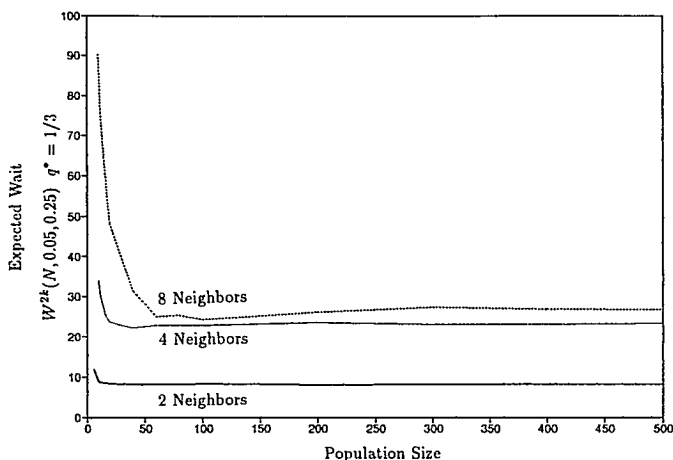


FIGURE 1.—Estimated waiting times for various population sizes.

graphed for $\varepsilon = 0.05$ and $q^* = 1/3$. The waiting times are sharply decreasing for small population sizes and appear remarkably constant for populations of more than about 100 players. It appears that with local matching rules evolutionary forces are not at all diminished in large populations and that any predictions we make will be extremely robust to the size of the population.

5. SIMULATION RESULTS

The contrasting behavior of models with local and uniform matching is not empirically meaningful when convergence is very slow even with local matching. For this reason, I now present two sets of numerical simulations which investigate further the conditions under which local matching rules allow evolutionary forces to determine play. The first set of simulations explores a variety of specifications of the model of Section 2. Although rapid convergence is hard to define, it is possible that evolutionary forces will determine play in a variety of models with local matching. The second set of simulations discusses an extension of the basic model and argues that rapid evolution may also be seen in a much broader range of models.

5.1. Matching Rules

We have already seen that the risk dominant equilibrium arises very quickly with two neighbor matching. I investigate here the extent to which this remains true as the assumptions of the two neighbor rule are gradually relaxed. Typically we have only a rough idea of the matching process which best describes a given population. We therefore hope that the predictions of the model are fairly

TABLE III
 EXPECTED WAITING TIMES WITH VARIOUS MATCHING RULES^a

	$W^{2k}(100, \epsilon, 0.25)$		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
2 Neighbors	11	8	6
4 Neighbors	44	23	12
8 Neighbors	93	25	11
12 Neighbors	522	45	11

^a Payoffs have $q^* = 1/3$. Estimated standard errors are less than 2% of the value shown.

robust to the particular specification chosen. Recall that a transition to the risk dominant equilibrium is possible with very few ϵ -probability events if strategy A is simultaneously adopted by a small cluster of adjacent players. It therefore seems reasonable that the likelihood of a speedy transition will be greatly affected by the factors which allow for the existence of such stable clusters. In particular, I examine here the degree to which the matching rule is concentrated on a few neighbors and to which the groups of neighbors overlap.

Table III describes the behavior of models with less concentrated matching rules. For $k = 1, 2, 4$, and 6 , Monte Carlo estimates of the expected waiting time $W^{2k}(100, \epsilon, 0.25)$ are given. Payoffs with $q^* = 1/3$ (e.g. $a = 2, b = 1, c = 0, d = 0$) are assumed. Note that the effect of the matching rule varies with the frequency of the randomizations. For the smallest value of ϵ shown, waiting times increase significantly when players have more neighbors. In this case, evolution is only likely to be seen when the matching rule is concentrated on a very few neighbors. For the larger values of ϵ , the waiting times are shorter and less dependent on the particular matching rule. Evolution may even be faster for less concentrated matching rules as stable clusters are almost sure to form immediately and the speed at which they spread becomes more important.

The k neighbor matching rules discussed so far are far from general. The assumption that the players are arranged around a circle has been maintained.⁸ Inherent in this assumption is a great overlap of the groups of neighbors in that a player's neighbors' neighbors are likely to be his neighbors as well. More realistic models would usually involve less overlap. For example, if an economist's "neighbors" include colleagues in her own department, others with similar research interests and older friends, then many of the neighbors might not know each other. With less overlap, evolution may be slower because stable clusters must be larger.

As a proxy for the degree of overlap of groups of neighbors, I explore the dimension of the lattice on which the players are arranged. Besides arranging 400 players around a circle, we can arrange them at the vertices of a 20×20 lattice on the surface of a torus, or on a $4 \times 4 \times 5 \times 5$ lattice in four dimensions.

⁸ Blume (1992) discusses the effects of the lattice used to represent the matching process in a similar model and notes that higher dimensional systems may behave very differently.




FIGURE 2.—Eight neighbor matching in one and two dimensions.

Figure 2 pictures the sets of neighbors which give eight neighbor matching in one and two dimensions. Moving from one to two to four dimensions reduces the maximum number of neighbors which any two players have in common from six to four to two.

For each of the lattices mentioned above, Table IV gives estimates of the expected waiting time $W(N, \epsilon, 0.25)$ again for payoffs with $q^* = 1/3$. For most of the parameter values shown, the expected waiting times are not greatly affected as we reduce the overlap of the neighbor groups by moving to higher dimensional lattices. However, the increased waiting times in the eight neighbor model for $\epsilon = 0.025$ are a reminder that the structure of the matching rule has the potential to greatly affect the behavior of the model.

The general conclusion that local matching leads to relatively fast convergence appears to be fairly robust to the choice of matching rule. However, when randomizations are infrequent, there are reasonable models in which play converges fairly slowly. In addition, whether convergence in say 50 periods is fast enough to be relevant depends on the application. As a result, while local matching allows evolutionary forces to be felt in large populations, we cannot predict that the risk dominant equilibrium will arise as confidently as we are

TABLE IV
EXPECTED WAITING TIMES FOR DIFFERENT GEOMETRIES^a

Lattice	Four Neighbor Matching Expected Wait		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
400×1	46	23	12
20×20	43	21	12
Lattice	Eight Neighbor Matching Expected Wait		
	$\epsilon = 0.025$	$\epsilon = 0.05$	$\epsilon = 0.1$
400×1	70	28	11
20×20	122	21	9
$4 \times 4 \times 5 \times 5$	1740	32	8

^a Payoffs have $q^* = 1/3$. Waiting times are until 3/4 of the players play A . Estimated standard errors are less than 3% of the values shown for all estimates except the $\epsilon = 0.025$ 8 neighbor cases.

TABLE V
WAITING TIMES FOR HETEROGENEOUS POPULATIONS^a

Variance of $u_i(B, B)$	Expected Wait		
	$\varepsilon = 0.025$	$\varepsilon = 0.05$	$\varepsilon = 0.1$
0	522	45	11
0.1	75	19	9
0.2	28	14	7

^a Expected waiting times are for populations of 100 players with 12 neighbor matching, $\alpha = 0.25$, and with the distribution of payoffs matching that generated by independent draws from lognormal distributions with $E(u_i(B, B)) = 1$ and $u_i(A, A) \stackrel{D}{=} (17/7)u_i(B, B)$. Estimated standard errors are less than 3% of the values shown.

able to predict that historical factors determine play in the model with uniform matching.

5.2. Heterogeneity

In a large population it is reasonable to assume that players have different tastes. I suggest here that this heterogeneity may allow rapid convergence for a much broader range of specifications than was previously indicated. Intuitively, in a diverse population there will be a few players who derive great utility from the equilibrium (A, A) and will therefore play A if only a few of their neighbors have done so. In the vicinity of these players, smaller stable clusters of players playing A are possible. Such clusters will therefore arise more quickly. Once they do the fact that a few players prefer (B, B) will do little to slow their spread.

Table V examines the effect of such heterogeneity on the expected waiting time. The players are assumed to have heterogeneous tastes $u_i(A, A)$ and $u_i(B, B)$ with the empirical distribution of the utilities matching that of independent draws from lognormal distributions with $u_i(A, A) \stackrel{D}{=} (17/7)u_i(B, B)$.⁹ A population of 100 players with 12 neighbor matching is assumed. The first line of the table gives the familiar waiting times for a homogeneous population and subsequent lines record the effect of increasing dispersion in the payoff distributions. ($E(u_i(B, B))$ is fixed at one.) To facilitate interpretation of the scale, note that when $\text{var}(u_i(B, B)) = 0.2$, there is one player who plays A if only one neighbor did so, and 2 players who will not play A unless 8 or more neighbors did so. When evolution is already rapid for a homogeneous population, heterogeneity has only a limited effect. More importantly though, when ε is small heterogeneity dramatically increases the rate at which play converges.

⁹ The constant 17/7 is used for comparison with the results of Section 5.1. When $\text{var}(u_i(B, B)) = 0$, $q^* = 7/24$ is in the center of the interval $(3/12, 4/12)$ and the dynamics are identical to those of the last model in Table III.

6. CONCLUSION

In this paper I have discussed a class of coordination game in order to examine the implications of a learning process among a large population of boundedly rational players. Kandori, Mailath, and Rob (1993) introduced such a model and showed how the players' myopic adjustments create evolutionary forces which may select among the equilibria.

The analysis yields several conclusions which appear to be fairly robust. Most generally, the rate at which a dynamic model converges can be an important consideration for economic applications. In the coordination problems discussed here, the nature of the matching rule which describes the interactions in a population helps determine the behavior we will observe. When each individual cares equally about coordinating with a great many opponents, play will exhibit great inertia and reflect arbitrary historical factors for a long period of time. On the other hand, in communities in which players are only likely to be matched with a few close friends or colleagues evolutionary forces may be felt early in the game. I would therefore hope that models with local interactions will allow the further application of evolutionary models to social behavior.

Department of Economics, Harvard University, Cambridge, MA 02138, U.S.A.

Manuscript received December, 1991; final revision received September, 1992.

APPENDIX

PROOF OF THEOREM 2: We begin with a general discussion of second largest eigenvalues applicable to both transition matrices. For v a probability distribution on S , v_s will denote the probability assigned to state s . It will often be convenient to write v as a 2^N -tuple ordered as

$$(v_{\bar{A}}, v_{\bar{B}}, \dots)$$

so that, for example, $(0, 1, 0, \dots, 0)$ represents a distribution which assigns probability 1 to all players playing B . (In writing this vector, some ordering is understood for the $2^N - 2$ states I have not specially named. Each of these states will be assigned probability zero in the 2^N -tuples which appear below.)

Let $CP_{P(\epsilon)}(x)$ be the characteristic polynomial of $P(\epsilon)$. $CP_{P(\epsilon)}$ and $CP_{P(0)}$ are polynomials of the same degree whose coefficients converge as $\epsilon \rightarrow 0$ so the set of roots of $CP_{P(\epsilon)}$ converges to the set of roots of $CP_{P(0)}$ (with multiplicity). If we write the eigenvalues of $P(\epsilon)$ with multiplicity as $1 = \lambda_1(\epsilon) > |\lambda_2(\epsilon)| \geq |\lambda_3(\epsilon)| \geq \dots \geq |\lambda_{2^N}(\epsilon)|$ this implies

$$|\lambda_i(\epsilon)| \rightarrow |\lambda_i(0)|.^{10}$$

Note that the matrices $P(0)$ have a two dimensional space of eigenvectors of eigenvalue 1 spanned by $(1, 0, 0, \dots, 0)$ and $(0, 1, 0, \dots, 0)$. These eigenvectors correspond to the steady states where all players coordinate on A and B respectively. Assuming for now that $P(0)$ has no other steady states or cycles, all other eigenvalues are of magnitude less than one. Therefore, $\lambda_2(\epsilon) \rightarrow 1$ and all other eigenvalues are bounded away from 1 for ϵ near zero.

Let $v(\epsilon)$ be an eigenvector with eigenvalue $\lambda_2(\epsilon)$ normalized so that $\|v(\epsilon)\| = 1$ and $v_{\bar{B}}(\epsilon) \geq 0$. I first show that $v(\epsilon) \rightarrow (-1, 1, 0, \dots, 0) \equiv v$. If not, for some $\eta > 0$ we can find a sequence $\{\epsilon_i\}$ converging to zero for which $\|v(\epsilon_i) - v\| > \eta$. Choosing an appropriate subsequence we may assume

¹⁰ See Franklin (1968, p. 191).

that $v(\epsilon_i)$ converges to a nonzero limit v^* . Note that

$$v^*P(0) = \lim_{i \rightarrow \infty} v(\epsilon_i)P(\epsilon_i) = \lim_{i \rightarrow \infty} \lambda_2(\epsilon_i)v(\epsilon_i) = v^*,$$

so that v^* is an eigenvector of $P(0)$ with eigenvalue 1. This gives $v^* = (x, y, 0, \dots, 0)$ with $y \geq 0$ and $\max(|x|, |y|) = 1$. Further, as $P(\epsilon)$ is a probability transition matrix, the sum of the elements of $v(\epsilon_i)$ is equal to the sum of the elements of $\lambda_2(\epsilon_i)v(\epsilon_i)$. This is only possible with $\lambda_2(\epsilon_i) < 1$ if the sum is zero. This gives the contradiction $v^* = v$.

First, consider the model with two neighbor matching. Throughout the proof I assume that N is odd. The proof for N even is similar although more complicated because of the cycle $\overrightarrow{AB} \rightarrow \overrightarrow{BA} \rightarrow \overrightarrow{AB}$. A complete proof is available from the author on request.

When N is odd, $P^2(0)$ has only \bar{A} and \bar{B} as steady states and has no cycles so by the reasoning above we know $\lambda_2(\epsilon) \rightarrow 1$ and $v(\epsilon) \rightarrow (-1, 1, 0, \dots, 0)$. Hence, for ϵ sufficiently small, $v_{\bar{B}}(\epsilon) > 0$ so we can write the second largest eigenvalue as

$$\begin{aligned} \lambda_2(\epsilon) &= \frac{\sum_{s \in \bar{S}} p_{s\bar{B}}^2(\epsilon)v_s(\epsilon)}{v_{\bar{B}}(\epsilon)} \\ &= p_{\bar{B}\bar{B}}^2(\epsilon) + \frac{\sum_{s \neq \bar{B}} p_{s\bar{B}}^2(\epsilon)v_s(\epsilon)}{v_{\bar{B}}(\epsilon)}. \end{aligned}$$

We have

$$p_{\bar{B}\bar{B}}^2(\epsilon) = (1 - N\epsilon + O(\epsilon^2)).$$

For $s \neq \bar{B}$, the successor of s in the model with no noise has at least two players playing A , so a transition from s to \bar{B} requires at least two ϵ -probability events. We then have that $p_{s\bar{B}}^2(\epsilon) = O(\epsilon^2)$ and $v_s(\epsilon)$ is bounded. Hence, $p_{s\bar{B}}^2(\epsilon)v_s(\epsilon) = o(\epsilon)$ and

$$\lambda_2(\epsilon) = 1 - N\epsilon + o(\epsilon)$$

as desired.

For the uniform model, we simplify the problem by looking instead at a Markov process on the $N + 1$ element state space $A(S)$. Formally, we define a transition matrix P' on $A(S)$ by

$$p'_{ij}(\epsilon) = \sum_{A(s)=j} p_{r(i)s}^u(\epsilon)$$

where $r(i)$ is any state with $A(r(i)) = i$. Note that we can write P' as a product $P' = DP''E$ where D is an $(N + 1) \times 2^N$ matrix which maps e_i to $e_{r(i)}$ and E is a $2^N \times N + 1$ matrix which classifies states by mapping e_s to $e_{A(s)}$. The matrix P' will be useful because its second largest eigenvalue is identical to that of P'' for ϵ small. To see this, note that the eigenvector of $P''(\epsilon)$ associated with λ_2 satisfies $v(\epsilon) \rightarrow v = (-1, 1, 0, \dots, 0)$. Hence $v(\epsilon)E \neq 0$ for ϵ small, and $v(\epsilon)EP' = v(\epsilon)EDP''E = v(\epsilon)P''E = \lambda_2 v(\epsilon)E$, so that $v(\epsilon)E$ is an eigenvector of P' with eigenvalue λ_2 . Because P' has a unique non-unit eigenvalue with magnitude converging to one, the second largest eigenvalues of P'' and P' must coincide.

We can further reduce the dimensionality of the problem by noting that $p'_{ij}(\epsilon) = p'_{i'j}(\epsilon)$ if states i and i' have the same successor in the model with no noise. Hence, we may write

$$P'(\epsilon) = QR(\epsilon)$$

where

$$Q = \begin{bmatrix} 1 & 0 & 0 \\ \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 1 \end{bmatrix}$$

groups the states into three classes, the precursors of 0, $\{q^*(N-1)\}$, and the precursors of N , and $R(\varepsilon)$ is an $3 \times N + 1$ matrix which gives the probability of each state arising given the class of the previous state. For example,

$$R_{ij}(\varepsilon) = p'_{ij}(\varepsilon) \quad \forall i \in \{0, 1, \dots, [q^*(N-1)] - 1\}.$$

From this decomposition, it is clear that $P'(\varepsilon)$ has exactly three nonzero eigenvalues.

Let $C(\varepsilon) = R(\varepsilon)Q$. $C(\varepsilon)$ is a 3×3 matrix which can be regarded as giving the transitions between the three classes of states. Note that $P'(\varepsilon)^n = QC(\varepsilon)^{n-1}R(\varepsilon)$. Let v be an eigenvector of $P'(\varepsilon)$ with eigenvalue $\lambda > 0$. We have $vQ \neq 0$ and $vQC(\varepsilon) = vQR(\varepsilon)Q = \lambda vQ$ so vQ is an eigenvector of $C(\varepsilon)$ with eigenvalue λ . As $C(\varepsilon)$ has rank 3, this gives a one-to-one correspondence between the nonzero eigenvalues of $P'(\varepsilon)$ and the nonzero eigenvalues of $C(\varepsilon)$.

We can write

$$(A1) \quad C(\varepsilon) = \begin{bmatrix} 1 - (p_1 + p_2) & p_1 & p_2 \\ q_1 & 1 - (q_1 + q_2) & q_2 \\ r_1 & r_2 & 1 - (r_1 + r_2) \end{bmatrix}$$

where, for example, p_1 gives the probability of a transition from any precursor of 0 to the state $\{q^*(N-1)\}$,

$$p_1 = \left([q^*(N-1)] \right) \varepsilon^{[q^*(N-1)]} (1 - \varepsilon)^{N - [q^*(N-1)]}.$$

Look at the characteristic polynomial of $C(\varepsilon)$ as a polynomial in $z \equiv 1 - \lambda$. We find that

$$(A2) \quad \text{Det}(C(\varepsilon) - \lambda I) = z(z^2 - a_1z + a_0)$$

with

$$a_1 = p_1 + p_2 + q_1 + q_2 + r_1 + r_2,$$

$$a_0 = (p_1 + p_2)(q_1 + q_2) + (q_1 + q_2)(r_1 + r_2) + (r_1 + r_2)(p_1 + p_2) \\ - q_2r_2 - p_1q_1 - p_2r_1.$$

The root $z = 0$ corresponds to the eigenvalue $\lambda = 1$. Let $z_1(\varepsilon) \geq z_2(\varepsilon)$ be the other two roots of this equation. As $\varepsilon \rightarrow 0$,

$$z_1(\varepsilon) + z_2(\varepsilon) \rightarrow 1$$

because $q_2(\varepsilon) \rightarrow 1$ when $N - [q^*(N-1)] > [q^*(N-1)]$ and all other terms in the expression for a_1 converge to zero. We also have

$$z_1(\varepsilon)z_2(\varepsilon) = O(\varepsilon^{[q^*(N-1)]})$$

as $q_2(\varepsilon)p_1(\varepsilon) = O(\varepsilon^{[q^*(N-1)]})$ and all other terms in the expression for a_0 are of strictly higher order. Hence, we must have $z_1(\varepsilon) \rightarrow 1$ and $z_2(\varepsilon) = O(\varepsilon^{[q^*(N-1)]})$. Clearly, $1 - z_2(\varepsilon)$ is the second largest eigenvalue in absolute value for sufficiently small ε so the desired result for the uniform model follows from our characterization $r''(\varepsilon) = 1 - z_2(\varepsilon)$.

REFERENCES

- BLUME, L. (1992): "The Statistical Mechanics of Strategic Interaction," Mimeo.
 EWENS, W. (1979): *Mathematical Population Genetics*. New York: Springer Verlag.
 FOSTER, D., AND H. P. YOUNG (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38, 219-232.
 FRANKLIN, J. (1968): *Matrix Theory*. New Jersey: Prentice Hall.
 FREIDLIN, M., AND A. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. New York: Springer Verlag.

- HARSANYI, J., AND R. SELTEN (1988): *A General Theory of Equilibrium in Games*. Cambridge: MIT Press.
- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61, 29–56.
- KANDORI, M., AND R. ROB (1992): "Evolution of Equilibria in the Long Run: A General Theory and Applications," CARESS Working Paper #92-06R.
- KARLIN, S., AND H. TAYLOR (1975): *A First Course in Stochastic Processes*. San Diego: Academic Press.
- ROSS, S. (1984): *A First Course in Probability*. New York: Macmillan.
- SENETA, E. (1973): *Non-negative Matrices*. New York: Wiley.
- WRIGHT, S. (1931): "Evolution in Mendelian Populations," *Genetics*, 16, 97–159.
- YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.