Diffusion in Social Networks

H. Peyton Young†

Center on Social and Economic Dynamics
Working Paper No. 2
May 1999

ABSTRACT

We consider processes in which norms of behavior are transmitted through social or geographic networks. Agents adopt behaviors based on a combination of their inherent payoff and their local popularity (the number of neighbors who have adopted them) subject to some random error. Extending work of Blume (1993, 1995), Ellison (1993), and Morris (1997), we characterize the long-run dynamics of such processes in terms of the *geometry* of the network, but without placing a priori restrictions on the network structure. We show first that the relative likelihood of different states can be described in terms of a potential function that is inversely related to the length of the boundary between regions where norms of behavior differ. As in a variety of other evolutionary models, the most likely state is the one in which everyone is coordinated on the risk-dominant equilibrium. We then show that, if agents interact in sufficiently small, close-knit groups, the expected waiting time until almost everyone is playing the risk-dominant equilibrium is bounded above independently of the number of agents and independently of the initial state. Simulation results indicate that convergence is surprisingly rapid even in very large networks, provided they are close-knit.

---

† H. Peyton Young is a senior fellow and co-director of the Center on Social and Economic Dynamics at the Brookings Institution, and a professor of economics at Johns Hopkins University. This research was supported by grant SBR 9601743 from the National Science Foundation. The paper benefited from constructive comments by Dean Foster and Stephen Morris.

1. Introduction

New ideas and ways of doing things do not necessarily take hold all at once, but often spread gradually through social and geographic networks. In a classic study, Coleman et al. (1966) documented how doctors' willingness to prescribe the new antibiotic tetracycline diffused through personal and professional contacts. A similar phenomenon has been observed in the adoption of new agricultural techniques, family planning methods, and various other innovations (Rogers, 1983). In the first stage a few innovators adopt, which makes it more likely that their neighbors adopt, then their neighbors' neighbors, and so forth, until the innovation spreads to the entire social group.

Such processes can be used to model not only technological innovation, but also the diffusion of customs, laws, and norms of behavior. Consider the history of rules of the road in Europe. While driving on the right is currently the norm in continental European countries, it was not always so. Prior to World War I, for example, driving on the left was common practice in Austria, Hungary, Bohemia, Sweden, Portugal, and parts of Italy, and probably other European countries as well. Looking back still further into the eighteenth century, we find that driving on the left was customary in France. Under the *ancien regime* horse-drawn carriages kept to the left, and pedestrians walked on the right in order to face the oncoming traffic and be prepared to jump out of the way if necessary. Keeping to the right was therefore a custom of the average citizen, while keeping to the left was a practice of the privileged classes. After the French Revolution, the right-hand rule was adopted as a symbol of the new, more democratic order. It was subsequently instituted by Napoleon for his armies, and probably took hold in some of the countries he occupied. From this point onward the tide shifted gradually as one country after another opted for the right-hand rule, Sweden being the last to switch in 1967. Thus a single, idiosyncratic event in one country triggered a change in custom that gradually diffused throughout Europe.[1]

The network structure that describes which European countries are neighbors has probably not changed very much over the past couple of hundred years. One can therefore study the specific properties of this network and how innovations would diffuse

---

[1]For historical accounts of rules of the road see Hopper (1982), Hamer (1986), and Lay (1992).

within it. Such cases are quite exceptional however. In general, it is very difficult to pin down the precise structure of a social network empirically; moreover, these structures are somewhat fluid as relationships develop and dissolve. Thus, in characterizing diffusion processes in such networks, we would like to have results that depend on certain qualitative features of the network, without assuming that they take specific geometric forms, such as lines, lattices, circles, and the like.

Here we investigate this problem using the following simple framework. Agents are located at the vertices of a graph, which represents their social network. In each period one agent is chosen at random and he or she adopts one of two behaviors depending on: i) its intrinsic payoff to the agent; and ii) the proportion of the agent's neighbors who have adopted the behavior. Agents are assumed to maximize their expected payoff subject to some small tremble or error that represents unobservable shocks. These kinds of local best reply models were considered by Blume (1993, 1995), and Ellison (1993); see also Anderlini and Ianni (1996), Ely (1995), Valente (995, 1996), Goyal and Janssen (1996, 1997), Glaeser, Sacerdote, and Scheinkman (1996), and Glaeser and Scheinkman (1998). Another approach is to suppose that agents react to some aggregate statistic about the distribution of behaviors in the general population (Brock and Durlauf, 1995; Durlauf, 1997).[2] While these models differ in various ways, we remark that the global case can be treated as a special case of the local one by assuming that every agent interacts with everyone else in society (the network structure is a complete graph).

In this paper we examine the following two questions. What patterns of behavior are most likely to emerge from arbitrary initial conditions, given the structure of the network and the payoffs from the interactions? And how long does it take society to reach or come close to this most likely state? The answers to these questions exploit the notion of a potential function, an approach pioneered by Blume (1993). Using this technique, we show that if the network has the property that agents interact in sufficiently small, close-knit groups, then the expected waiting time until the process comes close to its asymptotic distribution is bounded independently of the number of agents and independently of the initial state. Simulation results show that convergence is surprisingly rapid even in very large networks, provided they are close-knit.

---

[2]These are called "threshold models" in sociology .

2. The model

A social network will be represented by a graph Γ consisting of a finite set V of vertices and a set E of undirected edges. Each vertex i represents an agent in the system, and an edge {i, j} joins two agents i and j if and only if they are "neighbors," that is, their actions influence one another. Each interaction {i, j} is assumed to have an *importance weight* $w_{ij} = w_{ji} > 0$. Thus $w_{ij} > w_{ik}$ means that i attaches more importance to j's actions than to k's.

Each agent has two alternative behaviors, A or B. The desirability of choosing A or B depends on the prevalence of these actions among those with whom one interacts. For example, the agents might be countries, and A might be a statute that mandates driving on the left, while B mandates driving on the right. The utility of choosing A or B depends on the practices of neighboring countries; moreover, the practices of some neighbors may be more important than others. For example, the weight of the interaction {i, j} might depend on the amount of cross-border traffic between these two countries. Many other examples can be imagined, e.g., whether or not to have a child out of wedlock, whether or not to go to college, or whether to commit a certain kind of crime, as a function of the choices of their friends, peers, or neighbors (Crane, 1991; Glaeser, Sacerdote, and Scheinkman, 1996, Akerlof, 1997).

A *state* of the process is a choice of behavior (A or B) by each individual, that is, it is a vector $\mathbf{x} \in \{A, B\}^V$, where x_i denotes i's choice in state **x**. The utility to i of choosing A or B depends on two factors: the inherent utility of these choices to i, and the number of i's neighbors who are choosing A or B. Denote the set of i's neighbors by N_i, that is, the set of agents j $\neq$ i such that {i, j} is an edge in G. The utility to i in state **x** is

$$U_i(x) \sum_{j \in N_i} w_{ij} u(x_i, x_j) + v_i(x_i). \tag{1}$$

Here $v_i(x_i)$ represents the inherent utility of x_i to agent i, and $w_{ij}u(x_i, x_j)$ is the utility from i's interaction with j weighted by its importance. We can regard u as the payoff function of a 2 x 2 symmetric game in which the strategies are A and B. Further, we can regard U_i as the payoff function of an n-person game played on the graph Γ, where each player has the two strategies A and B. This is the *spatial game* induced by Γ.[3]

Consider the example of left/right driving mentioned above. In this case the agents are countries, and two agents are neighbors if they share a common border. Moreover, the game is perfectly symmetric: coordinating on A has the same payoff as coordinating on B. A more interesting situation arises if one technology is inherently better than another. (An example would be different rail gauges, where wider rails have greater stability but also cost more.) Suppose, for instance, that the payoff matrix and the network are as shown in Figure 1. For simplicity we shall assume that all edges have the same importance weight (all $w_{ij} = 1$), and that all agents have the same inherent preference for the two technologies, that is, $v_i(A) = v_i(B) = 0$ for all i.

	A	B
A	3, 3	0, 0
B	0, 0	2, 2

(2)

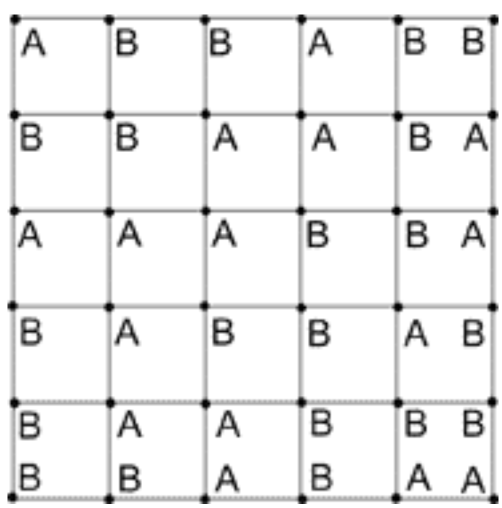

Figure 1. A two-dimensional lattice of social interactions with unit weights.

---

[3]Blume (1993) considered the case where the graph is a d-dimensional lattice and called the corresponding object a *lattice game.*

An *equilibrium* of the spatial game is a state **x** such that no agent strictly prefers to switch actions given the actions of his neighbors. The most prominent equilibria are the *homogeneous* states in which everyone plays A or everyone plays B. But there also exist many *heterogeneous* equilibria in which one subgroup of the population is coordinated on A while the other is coordinated on B. Specifically, define an *enclave* to be a rectangle of agents who are coordinated on the same action (see Figure 2). Consider a state in which every two A-enclaves are at least three edges apart. Then no B-player is adjacent to more than one A-player, and no A-player is adjacent to more than two B-players. Furthermore, no B-player has more than one-third of her interactions with A-players, and no A-player has more than one-half of her interactions with B-players. Hence the state is an equilibrium. In contrast, a B-enclave surrounded by As is not an equilibrium, because an agent at a corner of such an enclave must be adjacent to two A-players and two B-players, so this person would prefer to play A. The only B-enclaves that are in equilibrium are *strips* that run from one end of the lattice to the other, as shown in figure 3.[4]

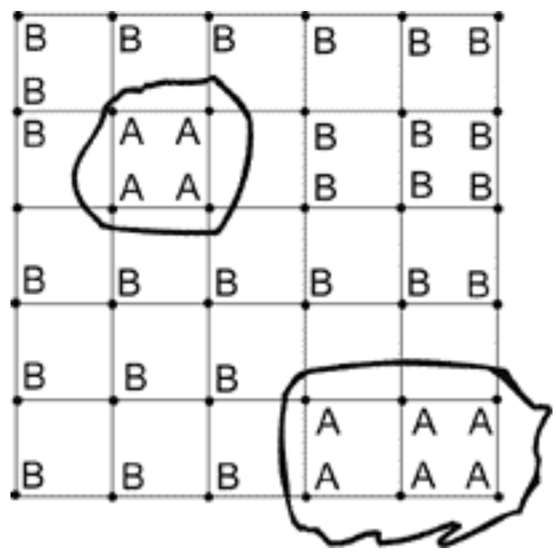

Figure 2. Equilibrium state consisting of two A-enclaves.

[4]If the lattice is embedded on the surface of a torus, the strip wraps entirely around the torus. If the lattice is planar as above, there is also an equilibrium in which the boundary vertices are all B, and the interior is all A.

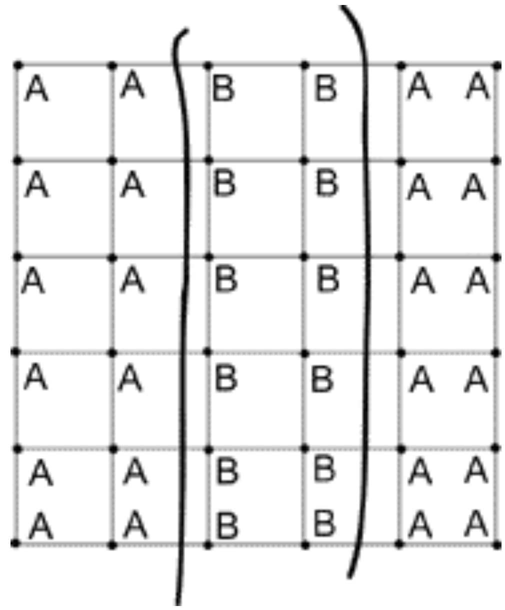

Figure 3. Equilibrium state consisting of two A-strips and one B-strip.

3. Best reply dynamics

The preceding analysis shows that equilibrium patterns can be surprisingly varied and complex. What happens when society begins in an out-of-equilibrium state? Consider the following simple adjustment process (Blume, 1995). Each individual updates his strategy at random times that are governed by a Poisson arrival process with unit expectation. In other words, each person updates once, on average, per unit time interval. We assume that these updating processes are independent among the individuals, so the probability is negligible that more than one person updates at any given time. When an individual updates, she chooses a best reply given the current choices of her neighbors. (If the two actions A and B are tied as best replies, the agent stays with her current choice out of inertia.)

Proposition 1. *Given any symmetric 2 x 2 coordination game played on a finite graph* Γ*, the probability is one that the process will reach an equilibrium state in finite time from any initial state.*

The proof is a straightforward extension of an argument in Blume (1995). Consider a symmetric 2 x 2 coordination game with payoff matrix

$$
\begin{array}{ccc}
 & A & B \\
A & a & c \\
B & d & b
\end{array}
\qquad a > d,\ b > c.
$$

(3)

For each state **x**, let A(**x**) be the sum of the weights on all edges {i, j} such that $x_i = x_j = A$. Similarly, let B(**x**) be the sum of the weights on all edges {i, j} such that $x_i = x_j = B$. Finally, let $v(\mathbf{x}) = \sum v_i(x_i)$ be the sum of the inherent payoffs of the choices in state **x**. Now define the potential of **x** to be

$$\rho(x) = (a - d)A(\mathbf{x}) + (b - c)B(\mathbf{x}) + v(\mathbf{x}).$$

(4)

It can be verified that whenever an agent chooses a best reply to what his neighbors are doing, the potential increases by at least some amount $\delta > 0$. Suppose that the process is in some nonabsorbing state at time t. By the end of the next period, the potential will have increased by at least d with probability 1/n, and with certainty it will not have decreased. Thus the probability is one that the process will reach an absorbing (i.e., equilibrium) state in finite time.

We remark that a similar argument applies to any finite potential game. Let G be a symmetric two-person game with finite strategy space X and payoff function $u: X \times X \to R$, where u(x, y) is the payoff to the agent playing x and u(y, x) is the payoff to the agent playing y. G is a *potential game* if there exists a *symmetric* function $\rho* : X \times X \to R$ such that $u(x, y) - u(x', y) = \rho*(x, y) - \rho*(x', y)$ for all choices of x, x', and y. If G is a potential game, then it is easy to show that the associated spatial game on the graph Γ is also a potential game with potential function

$$\rho(x) = \sum_{\{i,j\} \in E} w_{ij}\rho*(x_i, x_j) + \sum_{i \in V} v_i(x_i).$$

(5)

Proposition 1 and its proof extend immediately to this case.

4. Seeding new behaviors

Suppose that society is trapped in a suboptimal equilibrium, such as all-B in the technology adoption game. How can it be tipped into the more desirable outcome all-A? One approach would be to intervene and forcibly change the behavior of some of the agents to A, thus triggering a series of adjustments that result in the all-A equilibrium. Another possibility is to imagine that some agents are innovators and simply change on their own. If there are enough innovators they may be able to trigger a movement to all-A. What then is the minimum number of agents that suffice to trigger such a change, and which agents are they?

The answer to this question is provided by the following framework due to Morris (1997). For each vertex i let d_i be the degree of i, that is, the number of i's neighbors. Given a real number $r \in [0,1]$, a subset of agents S is said to be r-*cohesive* if everyone in S has at least r of his or her interactions with other members of S, that is, $|S\ N_i| \geq rd_i$. Suppose now that we force everyone in S to play A initially, while everyone else plays B, and then let the best reply process operate. For the behavior A to spread to the whole society with positive probability, it is both necessary and sufficient that there exist a sequencing of the individuals in V - S, say $i_1, i_2, \ldots, i_h$, such that

$$S \cup \{i_1, i_2, \ldots, i_j\} \text{ is r-cohesive for } 1 \leq j \leq h \text{ for some } r \geq r*,$$

where

$$r* = (b-c)/(a-d+b-c). \tag{6}$$

In this case we say that S *seeds* behavior A.

For example, let the players be arranged along a line, as shown in Figure 4. In the game with payoff matrix (2), it suffices to plant just one A anywhere along the line, and society will eventually tip with probability one into the equilibrium in which everyone plays A.

Figure 4. A linear social network.

This linear structure is quite exceptional however: in more general interaction structures it may require a very large number of agents, spread throughout the whole society, to tip it into the all-A equilibrium. Consider, for example, a lattice embedded on the surface of a torus, where each individual has exactly four neighbors. Let the process begin in a state where everyone in some subset S plays A, while everyone in the complementary set $\overline{S}$ plays B. Suppose further that S does not intersect some two adjacent lines in the lattice. Then A can never spread to these lines, because everyone in them will be adjacent at all times to at least three other B-players, and thus have at least 3/4 of their interactions with B-players. Hence no such agent will ever be motivated to switch to A.

More generally, consider any equilibrium **x** of the spatial game (other than the all-A equilibrium), and let $S_B(\mathbf{x})$ be the subset of agents that choose B in **x**. We shall call $S_B(\mathbf{x})$ the B-*support* of **x**. Since **x** is an equilibrium, $S_B(\mathbf{x})$ cannot seed A. Thus a simple necessary condition for S to seed A is that *S meet the B-support of every equilibrium of the spatial game whose* B-*support is nonempty.*

5. Perturbed best reply dynamics

We now consider a variation of the above process in which people sometimes deviate from best-reply behavior. We can think of these perturbations as resulting from unobservable shocks to individual utilities, or from deliberate experimentation, or as just plain randomness. A particularly convenient formulation is to suppose that each individual chooses a nonbest reply with a probability that decreases exponentially with the prospective loss in payoff (Blume, 1993). Specifically, we suppose that there exists a real number $\beta > 0$ such that

$$\Pr\{\text{i chooses A in state x}\} = e^{\beta U_i(A|x_{-i})} \Big/ \left[e^{\beta U_i(A|x_{-i})} + e^{\beta U_i(B|x_{-i})} \right]$$
$$\Pr\{\text{i chooses B in state x}\} = e^{\beta U_i(B|x_{-i})} \Big/ \left[e^{\beta U_i(A|x_{-i})} + e^{\beta U_i(B|x_{-i})} \right] \tag{7}$$

The larger β is, the higher is the probability that i chooses a best reply in any given state. This is known as a *log-linear response model with response parameter* β, or simply a β*-response model*.[5]

The associated spatial process is ergodic, and has a unique stationary distribution $\mu^{\beta}(\mathbf{x})$. Starting from an arbitrary initial state, $\mu^{\beta}(\mathbf{x})$ is almost surely the frequency with which state **x** is visited over the long run. In the present case, the stationary distribution takes the following special form, known as a *Gibbs state* (Liggett, 1985):

$$\mu(\mathbf{x}) = e^{\beta\rho(\mathbf{x})} \Big/ \sum_{\mathbf{y}\in\Xi} e^{\beta\rho(\mathbf{y})}. \tag{8}$$

where ρ is the potential function defined in (5). It follows that the most likely state **x** is the one that maximizes the potential function $\rho(\mathbf{x})$; moreover as β becomes large the probability is close to one that the process is in a state that maximizes potential. Such a state is said to be *stochastically stable* (Foster and Young, 1990).

The stochastically stable state does not necessarily correspond to a social optimum. Consider, for example, the 2 x 2 game with payoff matrix (3). Suppose that all agents are inherently indifferent between A and B, so that v(**x**) is a constant for all **x** (say v(**x**) = 0). Then the potential function takes the form $\rho(\mathbf{x}) = (a-d)A(\mathbf{x}) + (b-c)B(\mathbf{x})$. If

---

[5] Variants of this model have been used to explain how subjects make choices in experimental games. See for example Suppes and Atkinson, 1960; Roth and Erev, 1995; Camerer and Ho, 1997.

a - d > b - c and a < b, then A is risk dominant while B is Pareto dominant. It is evident from the potential function, however, that **A** is stochastically stable and **B** is not. A similar result holds in a variety of other evolutionary models (e.g., Kandori, Mailath, and Rob (1993), Ellison (1993), and Young (1993a).)[6]

Expression (8) allows us not only to find the stochastically stable states, but to estimate the *relative likelihood* of various equilibrium and nonequilibrium states. In particular, it illustrates the importance of boundaries in destabilizing a given regime. Given a state **x**, define the *boundary length* L(**x**) to be the total number of interactions (the number of edges in the graph) such that one end is coordinated on A and the other is coordinated on B. In other words, L(**x**) = $|E|$ - A(**x**) - B(**x**), where $|E|$ is the total number of edges in the graph. Assume that the situation is a priori symmetric in A and B, that is, $v_i(A) = v_i(B)$ for all i, and a = b, c = d. Then the potential function takes the form $-\lambda L(\mathbf{x}) + \nu$ for some constants $\lambda > 0$ and ν. It follows that, *when* A *and* B *are equally desirable, the log odds of any given state is inversely proportional to the length of the boundary, that is,* $\mu(\mathbf{x}) \propto Ce^{-\beta\lambda L(\mathbf{x})}$ *for some positive constant* C.

6. Social inertia.

The preceding analysis allows us to predict the asymptotic behavior of the social diffusion process in the long run, but how long is the long run? If it takes eons to reach the stochastically stable states predicted by theory, then the theory may not usefully describe real-world phenomena. In this section we examine the "inertia" of the dynamical process, that is, the expected waiting time until the stochastically stable state (or something close to it) first becomes established, starting from an arbitrary initial state. As we shall see, the waiting time depends crucially on the geometry of social interaction, that is, on who interacts with whom. When people interact mainly with small groups of neighbors, the inertia is dramatically lower than when everyone interacts with everyone else. Ellison (1993) was the first to demonstrate this point in a model where

[6]There exist adaptive models in which the Pareto optimal outcome is selected in the long run, but they typically differ from perturbed best reply models by relying on imitation, satisficing, or other type of choice behavior (Ely, 1995; Robson and Vega-Redondo, 1996).

people interact with their neighbors on a circle. Here we shall extend the analysis to general interaction structures.

Consider a symmetric two-person game with two actions A and B, and suppose that A is the risk dominant action. (Action A may or may not be Pareto optimal as well.) For simplicity we assume throughout this section that idiosyncratic payoffs are zero: $v_i(x_i) = 0$ for all i and all x_i. Let Γ be the graph that defines who interacts with whom. For a given response parameter β let $W(\Gamma, \beta, \delta)$ be the expected waiting time until the first time that at least 1 - δ of the population is playing action A, given that the process started in state **B**. The parameter δ is the degree of *precision* with which the process approximates state **A**. Since there is a positive probability that any given individual will choose B even when all of his neighbors choose A, the probability is very small that the process will actually be in state **A** at any given time when the number of individuals is large. Thus we need to choose δ (for a given level of β) so that with high probability the process stays within δ of state **A** once the process comes within δ of **A**.

We claim that, for certain natural kinds of interaction structures (of which lattices are a special case), the expected waiting time $W(\Gamma, \beta, \delta)$ is bounded *independently of the number of vertices in* Γ. Surprisingly, this result does not hinge on the connectivity of the interaction structure, but on the degree to which people interact in small, close-knit groups, as we shall show below.

Before stating this result formally, let us consider an example. Assume that players are arranged on a square lattice that is embedded on the surface of a torus, and they play the coordination game in (2). Using Monte Carlo simulations, we can compute the average waiting time until the first time that at least 95% of the agents are playing A, starting from the state where all played B. The results are shown in Figure 5 for two values of β. The waiting time is obviously sensitive to the choice of β, because this determines the probability that people make mistakes, and thus the probability that the system can exit from the all-B state. The more interesting point is that, for a given value of β and increasing sizes of grids, the waiting time plateaus fairly quickly; indeed it is more or less independent of the population size when the latter is large.

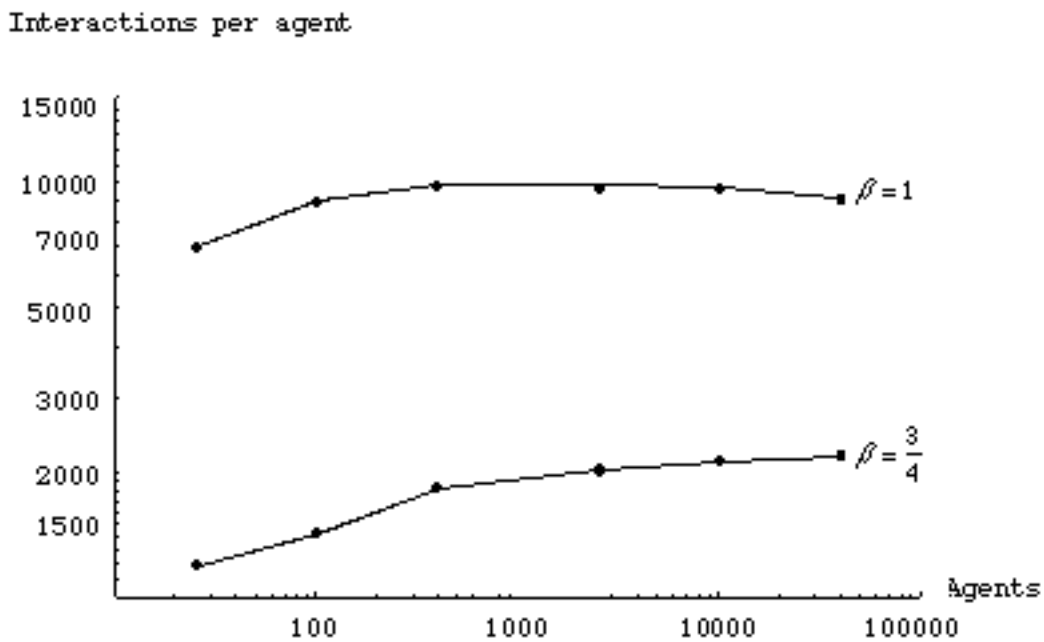

Figure 5. Average waiting time, in number of interactions per agent, until at least 95% of population first plays A, conditional on all playing B initially. Each point estimate is the mean of fifty trials.

This is a special property of the system that holds whenever the interaction structure is sufficiently "close-knit." The intuitive idea is that, for the new behavior A to spread rapidly in a society that was previously playing B, it must be able to take hold locally in small groups of individuals. In other words, the new behavior must be able to establish local "beachheads" that do not fall apart when pitted against the neighbors who have not yet converted.

To formulate this concept precisely, consider any two nonempty subsets S', S'' of vertices, and let $e(S', S'')$ denote the number of edges such that one end is in S' and the other is in S''. The *degree* of vertex i is the number of edges that meet vertex i. The vertex is *isolated* if $d_i = 0$. Let Γ be a graph with no isolated vertices. A nonempty subset S of vertices is r-*close-knit* for some real number $0 \leq r \leq 1/2$ if

$$\forall S' \subseteq S, S' \neq \emptyset \qquad e(S', S) \Big/ \sum_{i \in S'} d_i \geq r. \tag{9}$$

In other words, for every subset S' of S, the number of interactions between members of S' and members of S must be at least r times the sum of the degrees in S'. In particular, for every vertex $i \in S$, at least r of i's interactions must be with other members of S. It follows that if S is r-close-knit, then S is r-cohesive. The converse is not true however. For example, if each member of S has *exactly* r of its interactions with other members of S, then S is only (r/2)-close knit. The reason is that, with $S' = S$ in (9), every interaction between members of S is counted once in the numerator but twice in the denominator. Thus r-cohesiveness only implies r/2-close-knittedness.

Why is condition (9) required to hold for all subsets S' of S, instead of just individuals? The answer is that, even if each individual in S is impervious to conversion by outsiders, there could nevertheless exist a subgroup of S that is convertible, because collectively this subgroup has too many interactions with outsiders compared to interactions with insiders. Condition (9) states that S *and all of its subgroups* are sufficiently impervious to outside influences.

Given a positive integer k, we say that the graph Γ is (r, k)-*close-knit* if every person belongs to some group of size at most k that is at least r-close-knit. A family F of graphs is *close-knit* if for every 0 < r < 1/2 there exists an integer k (possibly depending on r) such that every graph in the class is (r, k)-close-knit.

As an example, consider the class of all polygons. In a polygon, the degree of every vertex is two. Each subset S of k consecutive vertices contains k - 1 edges, so $e(S,S) / \sum_{i \in S} d_i = (k-1)/2k$. It is easy to check that in fact $e(S',S) / \sum_{i \in S'} d_i = (k-1)/2k$ for every nonempty subset S' of S, hence every subset of k consecutive vertices is (1/2 - 1/2k)-close-knit. Since every vertex is contained in such a set, the class of polygons is close-knit. For a square lattice embedded on the surface of a torus, it can be verified that every subsquare of side h is $(1/2 - 1/2h, h^2)$-close-knit. It follows that the family of square lattices is close-knit.

Our principal result is that the inertia of a social diffusion process is bounded *independently of the number of vertices* provided that the interaction structures form a close-knit family.

Theorem 1. *Let G be a symmetric 2 x 2 coordination game with a strictly risk-dominant equilibrium, and let F be a close-knit family of graphs. Given any precision $\delta > 0$, there exists a number β_{δ} such that for each fixed $\beta \geq \beta_{\delta}$ the waiting time $W(\Gamma, \beta, \delta)$ bounded independently of the number of vertices in Γ.*

The formal argument is given in the Appendix. Intuitively, however, the idea of the proof is straightforward. Given a symmetric 2 x 2 game G with payoff matrix (3), choose $r* < r < 1/2$ and let k be such that all members of F are (r, k)-close knit. Fix a specific graph $\Gamma \in F$. Each individual is contained in an r-close-knit group S of size k or less. The probability that such an individual chooses a nonbest reply in a unit interval of time is bounded below by a positive number that depends on β, k, r, and the payoff matrix, but does not depend on the particular graph Γ. Since these parameters are fixed, the expected waiting time is bounded until the first time that all k members of S play A simultaneously. Once this happens, the close-knittedness of S assures that at every subsequent time, everyone in S plays action A with high probability irrespective of what the players outside of S are doing (assuming that β is sufficiently large). Since every individual is in such a group S, and the process is running simultaneously for all individuals, the waiting time is bounded until the first time that a high *proportion* of the individuals in the graph are playing A.

Note that this argument does not depend on the idea that actions spread by diffusion (which they may indeed do). For example, if one person switches to action A, it becomes more likely that nearby individuals will also adopt action A, which makes it more likely that *their* neighbors will adopt action A, and so forth. Clearly such a process further speeds up the waiting time, but it also requires a certain degree of connectivity in the interaction structure. The result assumes nothing about connectivity, however. For example, it applies equally well to graphs that consist of many distinct connected components, each of size k. The driving force behind the result is local reinforcement: if people interact mainly within a small group, any switch by the group to the risk-dominant equilibrium takes a long time to undo, and before that happens most of the other groups will have switched to the risk-dominant equilibrium too.

7. Conclusion

In this paper we have argued that, when agents interact in sufficiently small, close-knit groups, the time it takes for an adaptive process to come close to its stochastically stable state can be surprisingly short even for very large populations of agents. We demonstrated this point for symmetric 2 x 2 games played by agents using a log linear response rule. In this case the stochastically stable state occurs when everyone plays the risk-dominant equilibrium. How general is this result? A distinct advantage of the log linear response rule is that it leads to an analytically tractable formula for the long run distribution of the process (a Gibbs representation). While it seems reasonable to conjecture that an analog of Theorem 1 holds for some other classes of perturbed best reply models, establishing this fact will require somewhat different methods of proof. Another generalization of the model would be to suppose that agents pay attention both to what their neighbors are doing, and also to what people in the population at large are doing. One might assume, for example, that agents choose a perturbed best reply to some weighted combination of the distribution of choices by their neighbors, and the distribution in the population at large. It seems reasonable to conjecture that, when the weight placed on the local distribution is sufficiently large relative to the weight placed on the aggregate, then the process behaves (in terms of waiting times) very much like a purely local best reply process.

Finally, we have made the simplifying assumption that the graph representing the interaction structure is fixed. Yet there are many situations in which both the structure and the strength of interactions is to some degree endogenous (for example, bilateral trading relationships). The question then becomes: what structures are most stable over time, and how does their evolution interact with the evolution of play itself? This issue can also be explored using the concept of stochastic stability (e.g., Jackson and Watts, 1998); characterizing the waiting times in these models remains an open problem.

Appendix

Proof of Theorem 1. Let F be a close-knit family of graphs, and let G be a 2 x 2 symmetric coordination game with payoff matrix

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

Assume that a - d > b - c, so that (A, A) is the strictly risk dominant equilibrium. Let $r* = (b-c)/((a-d)+(b-c)) < 1/2$ and fix $r \in (r*, 1/2)$. Since F is close-knit, there exists an integer k such that every graph in F is (r, k)-close-knit. The parameters r and k will be fixed for the remainder of the proof. We are going to show that, given any $\delta \in (0,1)$, there exists a β_δ such that for each $\beta \geq \beta_\delta$, the p-inertia of the process is bounded above for all $\Gamma \in F$.

Choose $\Gamma \in F$ having vertex set V, and let S be an r-close-knit subset of size k. The adaptive process on G will be denoted by $P^{\Gamma,\beta}$. Now consider the following modification of P: whenever agents in S updates, they do so according to the log-linear response process with parameter β, but the agents in $\bar{S}$ always choose B. Denote this *restricted process* by $P^{\Gamma,S,\beta}$. States of the restricted process will be denoted by **y**, and states of the *unrestricted process* $P^{\Gamma,\beta}$ will be denoted by **x**. Let Ξ_S denote the set of restricted states, and Ξ the set of all states.

By assumption every agent is indifferent between the inherent properties of A and B, and thus we can drop the term v(**x**) and write the potential function as

$$\rho(\mathbf{x}) = (a-d)A(\mathbf{x}) + (b-c)B(\mathbf{x}).$$

By (8) we know that the stationary distribution $\mu^{\Gamma,\beta}(\mathbf{x})$ of $P^{\Gamma,\beta}$ satisfies $\mu^{\Gamma,\beta}(\mathbf{y}) \propto e^{\beta\rho(\mathbf{x})}$ for all $\mathbf{x} \in \Xi$. Similarly it can be shown that the stationary distribution $\mu^{\Gamma,S,\beta}(\mathbf{y})$ of $P^{\Gamma,S,\beta}$ satisfies $\mu^{\Gamma,S,\beta}(\mathbf{y}) \propto e^{\beta\rho(\mathbf{y})}$ for all $\mathbf{y} \in \Xi_S$.

Let $\mathbf{A}_S$ denote the state in Ξ_S such that everyone in S chooses action A, and everyone in $\overline{S}$ chooses action B. We claim that $\mathbf{A}_S$ uniquely maximizes $\rho(\mathbf{y})$ among all restricted states **y**. To see this, consider any restricted state **y** and let $S' = \{i \in S : y_i = B\}$. Then

$$\rho(\mathbf{y}) = (a-d)e(S-S',S-S') + (b-c)\left[e(S',S') + e(S',\overline{S}) + e(\overline{S},\overline{S})\right],$$

and

$$\rho(\mathbf{A}_S) = (a-d)e(S,S) + (b-c)(\overline{S},\overline{S}).$$

It follows that

$$\begin{aligned}\rho(\mathbf{A}_S) - \rho(\mathbf{y}) &= (a-d)e(S',S) - (b-c)\left[e(S',S') + e(S',\overline{S})\right] \\ &= (a-d)e(S',S) - (b-c)\left[e(S',S) + e(S',V) - e(S',S)\right].\end{aligned}$$

Thus $\rho(\mathbf{A}_S) - \rho(\mathbf{y}) > 0$ if and only if

$$\left[(a-d) + (b-c)\right]e(S',S) > (b-c)\left[e(S',V) + e(S',S')\right]. \tag{10}$$

The latter holds because by assumption

$$e(S',S)\big/\left[e(S',V)+e(S',S')\right] > r* = (b-c)\big/\left[(a-d)-(b-c)\right].$$

(11)

Note that $e(S',V)+e(S',S') > 0$, because by assumption there are no isolated vertices in any of the graphs we shall consider. Thus $\mathbf{y} = \mathbf{A}_S$ uniquely maximizes $\rho(\mathbf{y})$ as claimed. It follows that $\mu^{\Gamma,S,\beta}$ puts arbitrarily high probability on the state $\mathbf{A}_S$whenever β is sufficiently large.

Now fix $\delta \in (0,1)$. It follows from the preceding discussion that there exists a finite value $\beta(\Gamma, S, \delta)$ such that $\mu^{\Gamma,S,\beta}(\mathbf{A}_S) \geq 1-\delta^2/2$ for all $\beta \geq \beta(\Gamma,S,\delta)$. Fix such a value β. Consider the restricted process $P^{\Gamma,S,\beta}$ starting in the initial state $\mathbf{y}^0$, and let the random variable $\mathbf{y}^\tau$ denote the state of this process at time τ. The probability that $\mathbf{y}^\tau$ is in any given state **y** approaches the long-run probability of $\mathbf{y}, \mu^{\Gamma,S,\beta}(\mathbf{y})$ as τ goes to infinity. (This follows from the fact that the embedded finite chain is irreducible and aperiodic.) In particular, $\lim_{\tau\to\infty} \Pr\left[\mathbf{y}^\tau = \mathbf{A}_S\right] = \mu^{\Gamma,S,\beta}(\mathbf{A}_S)$. Hence there is a finite time $\tau(\Gamma, S, \delta, \beta)$ such that, from any initial state $\mathbf{y}^0$,

$$\forall\beta \geq \beta(\Gamma,S,\delta), \forall\tau \geq \tau(\Gamma,S,\delta,\beta), \Pr\left[\mathbf{y}^\tau = \mathbf{A}_S\right] \geq 1-\delta^2$$

Observe now that the continuous process $P^{\Gamma,S,\beta}$ depends on Γ and S only through the configuration of internal edges that link vertices of S to other vertices of S, and on the configuration of external edges that link vertices of S to vertices outside of S. Since S is of size k, there is a finite number of internal edges and a finite number of ways in which they can be configured. Since S is of size k and r-close-knit, there is a finite number of external edges, and a finite number of ways in which they can be configured vis-a-vis vertices outside of S. Thus, for a given r and k, there is a finite number of distinct processes $P^{\Gamma,S,\beta}$ up to isomorphism. In particular, we can find

$\beta(r,k,\delta)$ and $\tau(r,k,\delta,\beta)$ such that, among all graphs Γ in F and all r-close-knit subsets S with k vertices, the following holds independently of the initial state:

$$\forall\beta \geq \beta(\Gamma,k,\delta), \forall\tau \geq \tau(r,k,\delta,\beta), \quad \Pr\left[\mathbf{y}^{\tau} = \mathbf{A}_S\right] \geq 1 - \delta^2 \tag{12}$$

For the remainder of the discussion, we shall fix r, k, and δ as in the theorem. Let us also fix $\beta* \geq \beta(r,k,\delta)$ and $\tau* = \tau(r,k,\delta,\beta*)$. (In effect, $\beta(r,k,\delta)$ is the value β_δ claimed in the theorem.)

Let Γ be a graph in F with n vertices, and let S be an r-close-knit subset in Γ of size k. We shall couple the unrestricted process $P^{\Gamma,\beta*}$ and the restricted process $P^{\Gamma,S,\beta*}$ as follows. Create two disjoint isomorphic copies of the graph Γ, say Γ_1 and Γ_2, where the ith vertex in Γ_1 corresponds to the ith vertex in Γ_2. We will define a single process that mimics $P^{\Gamma,\beta*}$ on Γ_1, and mimics $P^{\Gamma,S,\beta*}$ on Γ_2. For each state **x** of the unrestricted process $P^{\Gamma,\beta*}$, let $q_i(A|\mathbf{x})$ denote the probability that i chooses A when i updates, given that the current state is **x**. Similarly, for each state **y** of the restricted process $P^{\Gamma,S,\beta*}$, let $q_i'(A|\mathbf{y})$ denote the probability that i chooses A when i updates, given that the current state is **y**. Note that $q_i'(A|\mathbf{y}) = 0$ for all $i \in \overline{S}$.

The coupled process operates as follows. The *state* at time τ is a pair $(\mathbf{x}^{\tau},\mathbf{y}^{\tau})$ where x_i^{τ} is the choice (A or B) at the ith vertex in Γ_1, and y_i^{τ} is the choice (A or B) at the ith vertex in Γ_2. Each matched pair of vertices in the two graphs is governed by a single Poisson process with unit expectation, and these processes are independent among the n matched pairs. Thus whenever the ith agent in Γ_1 updates the ith agent in Γ_2 updates, and vice versa. Let U be a random variable that is distributed uniformly on the interval [0, 1]. Suppose that the ith pair of individuals updates at time τ. Draw a value of U at random, and denote it by u. The ith individual in Γ_1 chooses A if $u \leq q_i(A|\mathbf{x}^{\tau})$ and chooses B if $u > q_i(A|\mathbf{x}^{\tau})$. Similarly, the ith individual in Γ_2 chooses A if $u \leq q_i'(A|\mathbf{y}^{\tau})$ and chooses B if $u > q_i'(A|\mathbf{y}^{\tau})$.

For every two states $\mathbf{x}$ and $\mathbf{y}$ on Γ_1 and Γ_2 respectively, write $\mathbf{x} \geq_A \mathbf{y}$ if $y_i = A$ implies $x_i = A$ for all i. In other words, if A appears at the ith vertex in $\mathbf{y}$ then A appears at the ith vertex in $\mathbf{x}$. It is evident that $\mathbf{x} \geq_A \mathbf{y}$ implies $q_i(A|\mathbf{x}) \geq q_i'(A|\mathbf{y})$ for all i. By construction of the process, if i chooses A in $\mathbf{y}^\tau$ then necessarily i chooses A in $\mathbf{x}^\tau$. Hence if $\mathbf{x}^\tau \geq_A \mathbf{y}^\tau$ at some time τ, then $\mathbf{x}^{\tau'} \geq_A \mathbf{y}^{\tau'}$ at all subsequent times $\tau' \geq \tau$.

Now let the coupled process begin in the initial state $\mathbf{x}^0$ on Γ_1 and $\mathbf{y}^0$ on Γ_2, where $x_i^0 = y_i^0$ for all $i \in S$, and $y_i^0 = B$ for all $i \in \overline{S}$. Obviously we have $\mathbf{x}^0 \geq_A \mathbf{y}^0$ initially, hence we have $\mathbf{x}^\tau \geq_A \mathbf{y}^\tau$ for all $\tau \geq 0$. From (12) and the choice of $\tau*$ we know that

$$\forall \tau \geq \tau*, \quad \Pr\left[\mathbf{y}^\tau = \mathbf{A}_S\right] \geq 1 - \delta^2$$

hence

$$\forall \tau \geq \tau*, \quad \Pr\left[\mathbf{x}^\tau = \mathbf{A}_S\right] \geq 1 - \delta^2$$

This holds for every r-close-knit set S in Γ. Since every vertex i is, by hypothesis, contained in such a set, it follows that

$$\forall \tau \geq \tau*, \quad \Pr\left[\mathbf{x}_i^\tau = \mathbf{A}_S\right] \geq 1 - \delta^2$$

Letting α^τ be the *proportion* of individuals in Γ_1 playing action A at time τ it follows that

$$\forall \tau \geq \tau*, \quad E\left[\alpha^\tau\right] \geq 1 - \delta^2 \tag{13}$$

We claim that this implies

$$\forall \tau \geq \tau*, \quad \Pr\left[\alpha^\tau \geq 1 - \delta\right] \geq 1 - \delta. \tag{14}$$

If this were false, the probability would be greater than δ that more than δ of the individuals at time τ were playing B. But this would imply that $E\left[\alpha^\tau\right] \leq 1 - \delta^2$, contradicting (13). Thus (14) holds. It follows that the expected waiting time in $P^{\Gamma,\beta*}$

until at least 1 - δ of the individuals are playing action A is bounded above by $\tau * / (1 - \delta)$. Moreover, once such a state is reached, the probability is at least 1 - δ that the process is in such a state at all subsequent times. Since the time τ holds for all graphs Γ in the family F, the δ-inertia of the family of processes $\left\{ \mathrm{P}^{\Gamma,\beta *} \right\}_{\Gamma \in F}$ is bounded as claimed. This concludes the proof of theorem 1.

References

Anderlini, Luc, and Antonella Ianni. 1996. "Path Dependence and Learning from Neighbors." *Games and Economic Behavior* 13:141-77.

Akerlof, George A. 1997. "Social Distance and Social Decisions." *Econometrica* 65:1005-27.

Blume, Larry. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 4: 387-424.

-----------------. 1995. "The Statistical Mechanics of Best-Response Strategy Revision." *Games and Economic Behavior* 11: 111-45.

Brock, William A., and Steven N. Durlauf. 1995. "Discrete Choice with Social Interactions. I: Theory." NBER Working Paper 5291. Cambridge Mass.

Camerer, Colin, and Teck-Hua Ho. 1997. "Experience-Weighted Attraction Learning in Normal Form Games." Mimeo, California Institute of Technology.

Coleman, James S., Elihu Katz, and Herbert Menzel. 1966. *Medical Innovation: A Diffusion Study*. New York: Bobbs Merrill.

Crane, Jonathan. 1991. "The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping Out and Teenage Childbearing." *American Journal of Sociology* 96: 1226-59.

Steven N. Durlauf. 1997. "Statistical Mechanical Approaches to Socioeconomic Behavior." In *The Economy as a Complex Evolving System*. vol. 2, edited by W. Brian Arthur, Steven N. Durlauf, and David Lane. Redwood City, Calif.: Addison-Wesley.

Ellison, Glenn. 1993. "Learning, Social Interaction, and Coordination." *Econometrica* 61: 1047-71.

Ely, Jeffrey. 1995. "Local Conventions." Mimeo, University of California at Berkeley.

Foster, Dean P. and H. Peyton Young, "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38: 219-32.

Glaeser, Edward L., Bruce Sacerdote, and Jose A. Scheinkman. 1996. "Crime and Social Interactions." *Quarterly Journal of Economics* 11: 507-48.

Glaeser, Edward L., and Jose A. Scheinkman. 1998. "Measuring Social Interactions." Forthcoming in *Social Dynamics*, edited by Steven N. Durlauf and H. Peyton Young. The Brookings Institution, Washington, D. C.

Goyal, Sanjeev, and Maarten Janssen. 1996. "Interaction Structure and Social Change." *Journal of Institutional and Theoretical Economics*, 152, 472-495.

----------------. 1997. "Non-Exclusive Conventions and Social Coordination." *Journal of Economic Theory*, 77, 34-57.

Hamer, Mick. 1986. "Left is Right on the Road: The History of Road Traffic Regulations." *New Scientist* 112.

Hooper, R. H. "Left-Right: Why Driving Rules Differ." *Transportation Quarterly* 36: 541-48.

Jackson, Matthew O., and Alison Watts. 1998. "The Evolution of Social and Economic Networks." Working Paper, California Institute of Technology and Vanderbilt University.

Kandori, Michihiro, George Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long-Run Equilibria in Games." *Econometrica* 61: 29-56.

Lay, Maxwell G. 1992. *Ways of the World*. New Brunswick, N.J.: Rutgers University Press.

Morris, Stephen. 1997. "Contagion." Working Paper, Department of Economics, University of Pennsylvania.

Robson, Arthur, and Fernando Vega-Redondo. 1996. "Efficient Equilibrium Selection in Evolutionary Games with Random Matching." *Journal of Economic Theory* 70: 65-92.

Rogers, Everett M. 1983. *Diffusion of Innovations.* 3rd edition. New York: Free Press.

Roth, Alvin E., and Ido Erev. 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior* 8: 164-212.

Suppes, Patrick, and Richard Atkinson. 1960. *Markov Learning Models for Multi-Person Interactions*. Stanford: Stanford University Press.

Valente, Thomas W. 1995. *Network Models of the Diffusion of Innovations*. Creskill, NJ: Hampton Press.

-----------------. 1996. "Social Network Thresholds in the Diffusion of Information." Social Networks 18: 69-89.

Young, H. Peyton, 1993a. "The Evolution of Conventions." *Econometrica* 61: 57-94.

----------------. 1993b. "An Evolutionary Model of Bargaining. " *Journal of Economic Theory* 59: 145-68.

-----------------. 1998. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions.* Princeton: Princeton University Press.