SOCIAL DYNAMICS: THEORY AND APPLICATIONS*

H. Peyton Young

JOHNS HOPKINS UNIVERSITY AND UNIVERSITY OF OXFORD

CSED Working Paper No. 39

Social Dynamics: Theory and Applications*

H. Peyton Young

Johns Hopkins University and University of Oxford

CSED Working Paper No. 39

April 2005

Abstract

Agent-based models typically involve large numbers of interacting individuals with widely differing characteristics, rules of behavior, and sources of information. The dynamics of such systems can be extremely complex due to their high dimensionality. This chapter discusses a general method for rigorously analyzing the long-run behavior of such systems using the theory of large deviations in Markov chains. The theory highlights certain qualitative features that distinguish agent-based models from more conventional types of equilibrium analysis. Among these distinguishing features are: local conformity versus global diversity, punctuated equilibrium, and the persistence of particular states in the presence of random shocks. These ideas are illustrated through a variety of examples, including competition between technologies, models of sorting and segregation, and the evolution of contractual customs.

1. Adaptive dynamics

Many forms of social and economic behavior evolve from the bottom up: they crystallize from the behavior and beliefs of disparate individuals interacting with each other over time. Language, codes of dress, forms of money and credit, patterns of courtship and marriage, standards of evidence, rules of the road, and economic contracts all have this feature. For the most part no one dictated the form that they have; they emerged through a process of experimentation, historical accident, and the accumulation of precedent. Agent-based models are particularly well suited to studying the dynamics of such processes, since by their nature they involve large numbers of dispersed, heterogeneous actors. In this chapter I shall outline a general framework for analyzing such systems based on theoretical results on large Markov chains, and then show how to apply the theory to concrete situations. Importantly, the theory can be applied without compromising the inherent complexity of the system: agents can be endowed with different characteristics, different levels of rationality, different amounts of information, and different locations.

My starting point is the assumption that agents are boundedly rational but purposeful. They look around them, they gather information, and they act fairly sensibly on the basis of that information.[1] I shall also assume that their choices are not entirely deterministic and predictable, but may be buffeted by random perturbations in the environment, errors of perception, and idiosyncracies in behavior. Whatever the source, these perturbations play a role similar to mutations in biology by injecting variability into agents' behaviors. Moreover,

---

[1] For a discussion of learning models see the chapters in this volume by Thomas Brenner and John Duffy.

of concrete examples, including competing technologies, neighborhood segregation, and the emergence of contractual norms.[2]

To set the stage, let us consider a classical example: the emergence of money as a medium of exchange.[3] History records the great variety of goods that societies have adopted as money: some used gold or silver, some copper or bronze, others used beads, still others favored cattle. In the early stages of economic development, we can conceive of the choice of currency as growing out of individual decisions that gradually converge on some norm. Once enough people in a society have adopted a particular currency, everyone else wants to follow suit.

At the individual level, this sort of decision problem can be cast as a coordination game. Suppose that there are two choices of currency: gold and silver. At the beginning of a period, each person must decide which currency to carry. During the period, each person meets various other people in the society at random, and they can trade only if they are both carrying the same currency. Thus the decision problem at the beginning of the period is to choose the currency that one believes will be chosen by a majority of the others.

Schematically we can model the dynamics as follows. Let p^t be the proportion in the population choosing gold at time t, and let $1 - p^t$ be the proportion choosing silver. In period $t + 1$, one person is drawn at random to reconsider his decision. He or she selects a random sample of s other individuals to determine what they are currently doing. Let $\hat{e}^t$ be the sample proportion of those using gold. Assume for the moment that the properties of gold and silver make them equally desirable as currencies. Then the decision maker chooses gold in period $t + 1$ if $\hat{e}^t > .5$ and chooses silver

---

[2] There is a large literature on the evolution of norms, some of which is related to the approach described here. See in particular Ullman-Margalit (1977), Sugden (1986, 1989), Bendor and Swistak (2001), Hechter and Karl-Dieter (2001), Skyrms (2004), and Bicchieri (forthcoming).

[3] See for example Menger (1871) and Marimon, McGratten and Sargent (1989).

if $\hat{e}^t < .5$.   (If $\hat{e}^t = .5$ we shall assume the agent chooses randomly.) All of this happens with high probability, say $1 - \varepsilon$. But with probability $\varepsilon > 0$ a person chooses gold or silver at random, that is, for reasons external to the model.

Qualitatively this process evolves in the following manner. After an initial shakeout, the process converges quite rapidly to a situation in which most people are carrying the same currency -- say gold. This norm will very likely stay in place for a considerable period of time. Eventually, however, an accumulation of random shocks will "tip" the process into the silver norm. These tipping incidents are infrequent compared to the periods in which one or the other norm is in place. Moreover, once a tipping incident occurs, the process will tend to adjust quite rapidly to the new norm. This pattern -- long periods of stasis punctuated by sudden changes of regime -- is known in biology as the *punctuated equilibrium effect*. The model described above predicts a similar phenomenon in the evolution of economic and social norms. Figure 1 illustrates this idea for the currency game when the two currencies have equal payoffs.

Figure 1. The currency game with equal payoffs, sample size s = 10, and ε = .50.
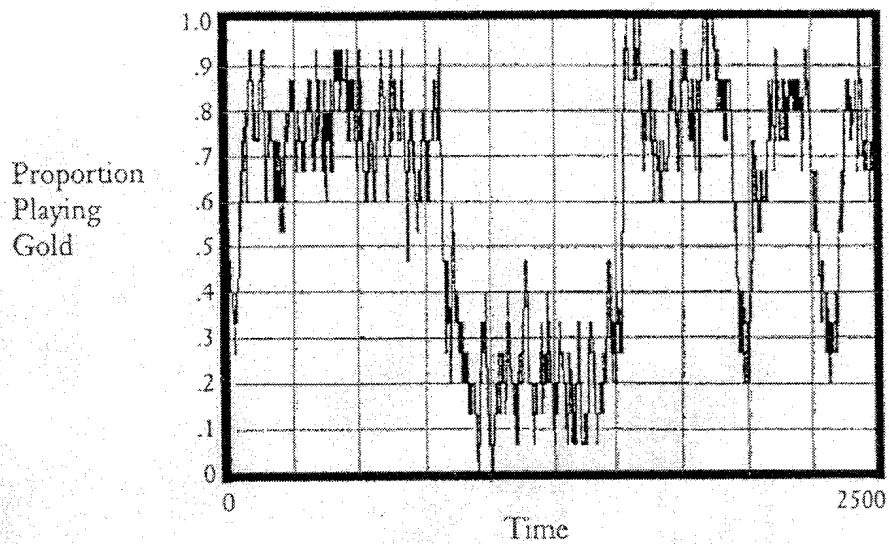
Now let us ask what happens when one currency is inherently better than the other. Suppose, for example, that gold is somewhat preferred because it does not tarnish as easily as silver. Then the decision problem at the individual level is to choose gold if $\hat{e}^t > \alpha$, and to choose silver if $\hat{e}^t < \alpha$, where α is a fraction strictly less than one-half. Now the process follows a path that looks like figure 2. Over the long run there is a bias toward gold, that is, at any given time the society is more likely to have adopted the gold standard than the silver standard. Moreover, the bias becomes *larger* the *smaller* the random perturbations are. Figure 3 shows a characteristic sample path when the noise level is reduced by a factor of ten. Notice that the process is at or near the gold standard a larger fraction of the time, and shifts of regime are more infrequent. These features become more pronounced as the noise level becomes smaller, a fact that can be verified analytically using methods to be discussed in the next section.

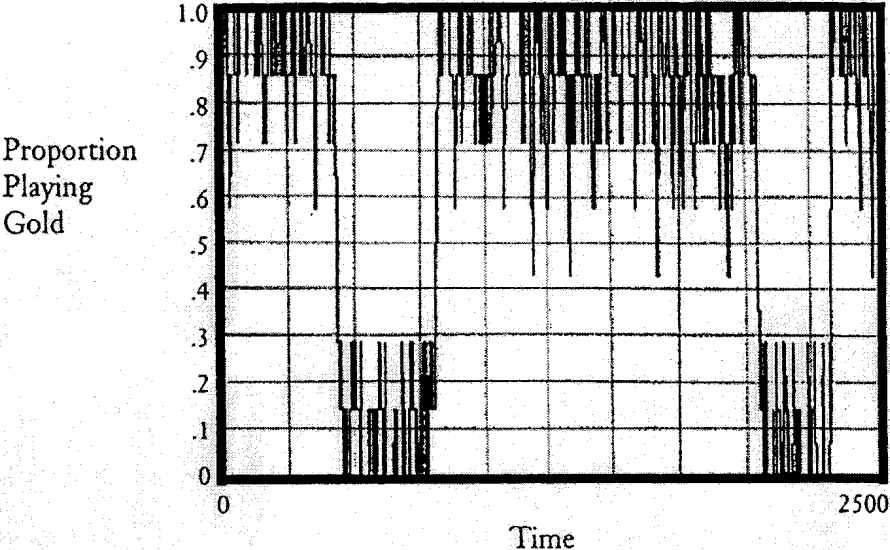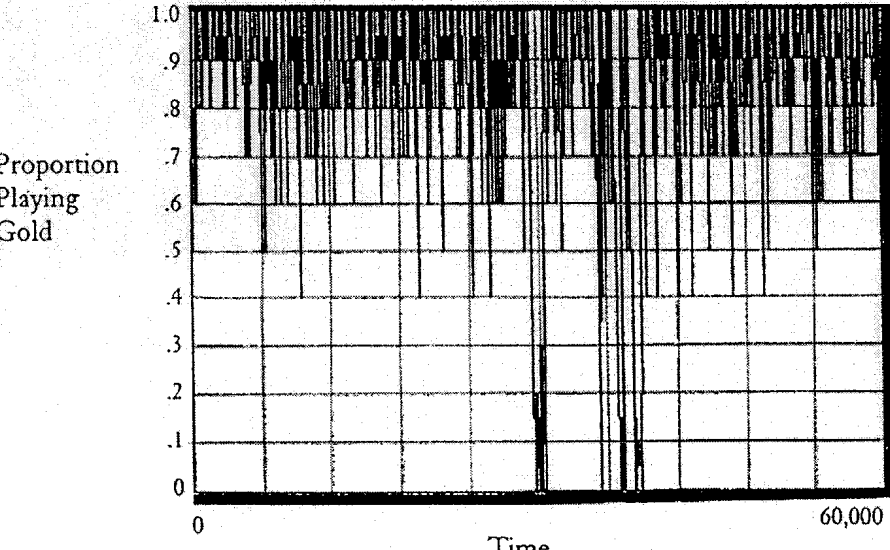Figure 2. The currency game with asymmetric payoffs, $\alpha = 1/3$, s = 10, and $\varepsilon = .50$.

Figure 3. The currency game with asymmetric payoffs, $\alpha = 1/3$, s = 10, and $\varepsilon = .05$.

2. Stochastic stability

Many agent-based models can be represented as Markov processes of very large dimensionality. A *state* of the system is specified by the location, information, and beliefs of the various actors. The *transition probabilities* are specified by the interaction probabilities among agents and the rules by which they adapt their choices and beliefs to perceived conditions (the learning rules). Let Z denote the set of possible states of such a system, which, though finite, may be extremely large. For every pair of states $z, z' \in Z$, let P be a $|Z| \times |Z|$ matrix such that the component $P_{zz'}$ is the *probability* of moving from state z to state z' in one period. P is the *transition probability matrix* of a finite Markov process. We shall always restrict ourselves to processes that are time-homogeneous, that is, the transition probabilities do not change from one period to the next.

Suppose that the initial state is z^0. For every time $t > 0$, let the random variable $f^t(z \mid z^0)$ denote the *empirical frequency* with which state z is visited during the first t periods. It can be shown that, as t goes to infinity, $f^t(z \mid z^0)$ converges almost surely to a limiting frequency distribution. If this distribution depends on the initial state z^0, or on chance events that occur along the way, we shall say that the process is *non-ergodic* or *path-dependent*. If the limiting distribution is uniquely determined independently of z^0, the process is *ergodic*.

There is a simple structural criterion that allows us to say whether or not a process is ergodic.[4] Say that state z' is *accessible* from state z, written $z \to z'$, if there is a positive probability of moving from z to z' in a finite number of periods (including no periods, i.e., z is accessible from z). States z and z' *communicate*, written $z \sim z'$, if

_____

[4] For a discussion of ergodicity in Markov chains see Karlin and Taylor (1975, Chapter 2).

7

each is accessible from the other. Clearly ~ is an equivalence relation, so it partitions the space Z into equivalence classes, which are known as *communication classes*. A *recurrence class* of P is a communication class such that no state outside the class is accessible from any state inside it. It is straightforward to show that every finite Markov chain has at least one recurrence class. A state is *recurrent* if it is contained in one of the recurrence classes; otherwise it is *transient*. In particular, a state is recurrent if and only if, once the process has entered it, the probability of returning to it is one.

A basic result on finite Markov chains is that ergodicity holds if and only if the process has a unique recurrence class. Equivalently, such a process is ergodic if the states can be divided into two disjoint classes A and B such that: there is a positive probability of moving from any state in A to some state in B; there is a positive probability of moving from any state in B to any other state in B; there is zero probability of moving from any state in B to any state in A. A particular instance occurs when A is empty and B constitutes the entire state space; in this case the process is said to be *irreducible*.

The standard approach to analyzing the asymptotic behavior of a Markov chain is to solve for the stationary distribution algebraically. Specifically, let μ be a probability distribution on Z written out as a row vector and consider the system of linear equations

$$\mu P = \mu, \text{ where } \mu \geq 0 \text{ and } \sum_{z \in Z} \mu(z) = 1. \tag{1}$$

This system always has at least one solution μ, called a *stationary distribution* of the process P. The solution is unique if and only if P has a unique recurrence class, that

8

is, if and only if P is ergodic. In this event the empirical frequency distribution converges almost surely to μ independently of the initial conditions:

$$\lim_{t \to \infty} f^t (z \mid z^0) = \mu(z). \tag{2}$$

By contrast, if P has more than one recurrence class, the process is path-dependent, and the initial position -- as well as chance events along the way -- can influence its long-run behavior.

Most of the models we shall consider are ergodic; in fact they have another property that allows us to make even sharper statements about their asymptotic behavior. Given a finite Markov process P and a state z, let N_z be the set of all positive integers n such that there is a positive probability of moving from z to z in *exactly* n periods. The process P is *aperiodic* if, for every z, the greatest common divisor of N_z is unity. If P is aperiodic and ergodic, not only does its *average* behavior converge to the unique stationary distribution μ, so does its probabilistic behavior *at each point in time* t when t is sufficiently large. More precisely, P^t be the t-fold product of P. If the process starts in an arbitrary state y, then in t periods the probability of being in state z is P^t_{yz}. It can be shown that, if P is ergodic and aperiodic, then with probability one

$$\forall \; y, z \in Z, \qquad\qquad \lim_{t \to \infty} P^t_{yz} = \mu(z). \tag{3}$$

In particular, the probability of being in a given state z at a given time t is essentially the same as the probability $f^t (z \mid z^0)$ of being in state z *up through time* t provided that t is large; furthermore both converge to the stationary distribution $\mu(z)$ independently of the initial state.

9

When the state space is very large -- as is usually the case with agent-based models -- the stationarity equation (1) is much too cumbersome to solve explicitly. Fortunately there is an alternative approach, based on the theory of large deviations, that often permits a good approximation of the stationary distribution without having to solve equation (1).

Suppose that the Markov process P can be split into two parts: a basic process P^0, on which is superimposed small trembles or perturbations. An example would be a model in which agents change their behaviors according to a choice rule that has a small probabilistic component. In this case the basic process is given by the probabilities of interaction among the agents, combined with their *expected* change in behaviors; the perturbations correspond to idiosyncratic aspects of individual-level changes in behavior. (We shall consider a number of concrete examples below.) Under certain regularity conditions, one can identify the states that have high probability when the perturbations are small without solving for the stationary distribution explicitly. These are known as *stochastically stable states*, and correspond to the equilibria that have the greatest persistence or robustness in the presence of random perturbations [Foster and Young (1990)].

3. Technology adoption

We shall first illustrate the approach using a model of technology choice with network externalities, which is similar to the currency model discussed earlier. Consider a population of n individuals. At each point in time every individual owns one of two technologies, A or B, hence the system has 2^n possible states. Both technologies generate positive externalities – the payoff from a given choice increases with the proportion of others who make the same choice. A contemporary example is personal computers: if most people own PCs it is advantageous to own a PC; if most people own Macs it is more desirable to own a

Mac. The reason is that the more popular a given model is, the more software will be created for it, and the easier it is to share programs with others.[5]

In each period one individual is chosen at random to make a new choice -- say because her current model wears out. She makes her decision by asking s randomly selected people what choices they made, and then choosing a perturbed best response. The payoffs are as follows: if in the random sample k people have chosen A and s - k have chosen B, the payoff to adopting A is as and the payoff to adopting B is $b(s - k)$. This is equivalent to playing a game against the field in which the row player's payoffs are given by

$$
\begin{array}{ccc}
 & A & B \\
A & a & 0 \\
B & 0 & b
\end{array}
$$

Let us assume that players choose a best response with high probability, but not with certainty. Specifically let us suppose that an individual chooses a best response (given the sample evidence) with probability $1 - \varepsilon$, and chooses an action at random with probability ε. Thus, with low probability the individual does not deliberate about her decision, whereas with high probability she does.

This is a simple example of a perturbed dynamical process. There is a finite (but large) number of states, and there are well-defined transition probabilities from any state to any other state. Unless the population is very small, however, it is extremely cumbersome to write down the transition matrix and to solve the stationarity equation algebraically. Instead we exploit the fact that the process is

---

[5] For other models of network externalities see Katz and Shapiro (1985, 1986), David (1985), and Arthur (1989).

perturbed due to the idiosyncratic choices of agents.

If there were no perturbations ($\varepsilon = 0$), the transition probabilities would be calculated as follows. Let the current state consist of m users of A and n – m users of B. At the start of the next period, choose one agent at random and let her draw a sample of size s from the remaining agents. Assume that she chooses a best response to the distribution of A-users and B-users in her sample. The combination of these events determines the probability of transiting to every possible successor state at the end of the period. (Note that the process can only transit to a state that differs from the current state in at most one coordinate, because only one agent reconsiders in each period.) Let P^0 denote the transition probability matrix of the resulting *unperturbed process*. Define a separate process Q in which one agent is drawn at random each period and chooses A or B with equal probability. We can then represent the *perturbed process* (with noise level ε) by the transition matrix $P^\varepsilon = (1 - \varepsilon)P + \varepsilon Q$.

The stationary distribution may now be calculated as follows. First we identify the recurrence classes of P^0. One such class is the absorbing state in which everyone plays A; another is the absorbing state in which everyone plays B. Call these states z^A and z^B respectively. It can be checked that these are the only recurrence classes: from any state the probability is one of eventually landing in one of these two states. Now compute the "path of least resistance" from z^B and z^A and vice versa. Starting from z^B, consider a series of A adoptions (due to perturbations) that lead to a critical or "tipping" state z^*, from which the process can transit to z^A with no further perturbations. This tipping point occurs when there are k^* choices of A, where k^* satisfies the condition $ak^* \geq b(s - k^*)$, that is, $k^* \geq bs/(a + b)$. (An agent who draws these k^* individuals in her sample will choose A instead of B.) The probability of this tipping event is approximately $(\varepsilon/2)^{\lceil bs/(a + b) \rceil}$, where in general

$\lceil x \rceil$ denotes the least integer greater than or equal to x. Define the *resistance* of the transition $z^B \to z^A$ to be the exponent on ε, that is,

$$r(z^B \to z^A) = \lceil bs/(a + b) \rceil.$$

Similarly, the resistance of the transition $z^A \to z^B$ is

$$r(z^A \to z^B) = \lceil as/(a + b) \rceil.$$

The smaller of these numbers determines the shape of the stationary distribution when ε is small. Specifically, if $r(z^A \to z^B) < r(z^B \to z^A)$ then the stationary distribution puts probability close to 1 on the state z^B. If $r(z^A \to z^B) > r(z^B \to z^A)$, the stationary distribution puts probability close to 1 on the state z^A. It follows that, when the sample size s is sufficiently large, the Pareto efficient technology is favored in the long run: if $a > b$, society is much more likely to have a large number of A-users than a large number of B-users, and vice versa.

4. Characterizing the stochastically stable states

We now show how this framework can be generalized to a wide variety of agent-based models. Consider a process such that the size of the perturbations can be indexed by a scalar $\varepsilon > 0$, and let P^ε be the associated transition probability matrix. P^ε is called a *regular perturbed Markov process* if P^ε is ergodic for all sufficiently small $\varepsilon > 0$ and P^ε approaches P^0 at an exponentially smooth rate [Young (1993a)]. Specifically, the latter condition means that

$$\forall \ z, z' \in Z, \qquad \lim_{\varepsilon \to 0^+} P^\varepsilon_{zz'} = P^0_{zz'} \ ,$$

13

and $\qquad P^\varepsilon_{zz'} > 0$ for some $\varepsilon > 0$ implies $0 < \lim_{\varepsilon \to 0^+} P^\varepsilon_{zz'} / \varepsilon^{r(z \to z')} < \infty$,

for some nonnegative real number $r(z \to z')$, which is called the *resistance* of the transition $z \to z'$.

Let P^0 denote the unperturbed process and let its recurrence classes be denoted by $E_1, E_2, \ldots, E_N$. For each pair of distinct recurrence classes E_i and E_j, $i \neq j$, an *ij-path* is defined to be a sequence of distinct states $\zeta = (z_1 \to z_2 \to \ldots \to z_n)$ such that $z_1 \in E_i$ and $z_n \in E_j$. The *resistance* of this path is the sum of the resistances of its edges, that is, $r(\zeta) = r(z_1 \to z_2) + r(z_2 \to z_3) + \ldots + r(z_{n-1} \to z_n)$. Let $\rho_{ij} = \min r(\zeta)$ be the *least resistance* over all ij-paths ζ. Note that ρ_{ij} must be positive for all distinct i and j, because there exists no path of zero resistance between distinct recurrence classes.

Now construct a complete directed graph with N vertices, one for each recurrence class. The vertex corresponding to class E_j will be called "j". The *weight* on the directed edge $i \to j$ is ρ_{ij}. A tree rooted at vertex j, or j-*tree*, is a set of N - 1 directed edges such that, from every vertex different from j, there is a unique directed path in the tree to j. The *resistance* of a rooted tree T is the sum of the resistances ρ_{ij} on the N - 1 edges that compose it. The *stochastic potential* γ_j of the recurrence class E_j is defined to be the minimum resistance over all trees rooted at j. The following theorem gives a simple criterion for determining the stochastically stable states [Young (1993a, Theorem 4)].

THEOREM 1. *Let P^ε be a regular perturbed Markov process and for each $\varepsilon > 0$ let μ^ε be the unique stationary distribution of P^ε. Then $\lim_{\varepsilon \to 0} \mu^\varepsilon$ exists and the limiting distribution μ^0 is a stationary distribution of P^0. The stochastically stable states (the support of μ^0) are precisely those states contained in the recurrence classes with minimum stochastic potential.*

We shall illustrate this result with the preceding example. In this situation there are two recurrence classes, $\{z^A\}$ and $\{z^B\}$, and exactly two rooted trees, as shown in Figure 4.

Figure 4. The currency game with two recurrence classes.

The tree with least resistance points toward the Pareto dominant equilibrium, and confirms our earlier calculation that this is the stochastically stable outcome.

A more complex example is the following. Consider a technology choice game in which there are three choices of technology -- A, B, C -- and the payoffs from networking are

	A	B	C
A	5	0	0
B	0	4	0
C	0	0	3

In this case there are three recurrence classes, one for each of the absorbing states z^A, z^B, z^C, and there are nine trees, as shown in figure 5.

15

Figure 5. An example with nine rooted trees and three recurrence classes.
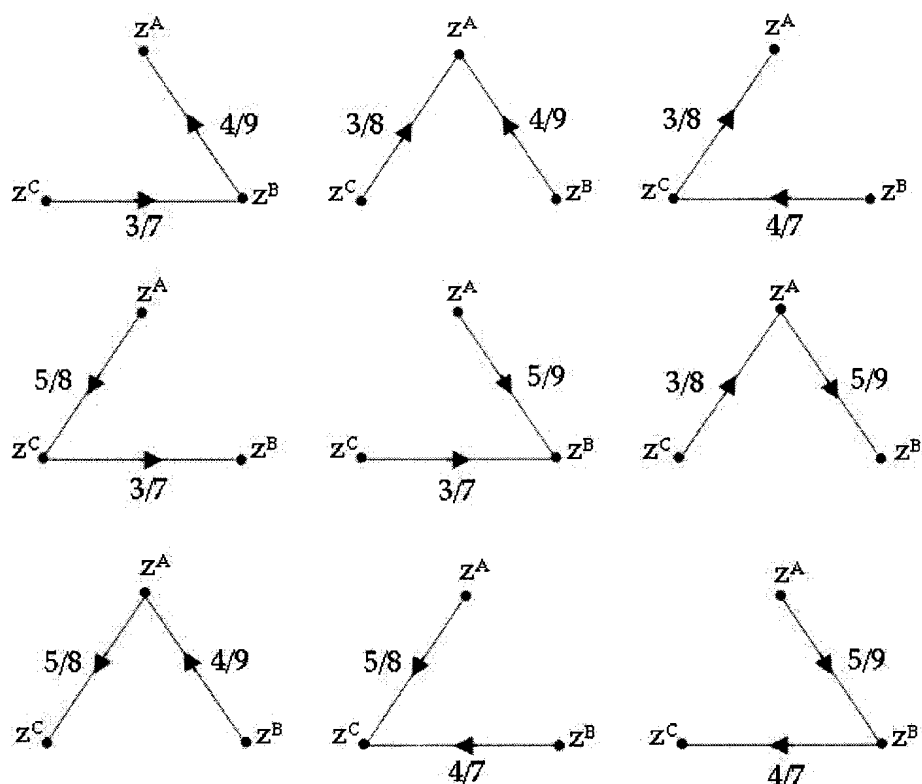
The sum of the resistances is minimized for the middle tree in the top row (3/8 + 4/9 is the smallest sum among the nine trees). Hence the root of this tree, which corresponds to the state in which everyone adopts technology A, is the stochastically stable state.

5. Efficiency versus stochastic stability

The preceding examples should not lull the reader into believing that evolution invariably selects efficient norms or standards. On the contrary, this state of affairs is quite exceptional, and hinges on the form of the payoff matrix. In this section we discuss the connection between efficiency and stochastic stability when there are two alternatives; a more extended discussion may be found in Young (1993a, 1998).

Assume then that there are two competing technologies, A and B. In the preceding section we assumed that there were gains only from networking with the same technology (the payoff matrix has zeroes off the diagonal). In general, however, there may be positive payoffs from networking with different technologies, and there may also be payoffs that arise from using the technology independently of networking effects. (For example, in the case of computer software there is a payoff from ease of file-sharing with other users, but there is also a payoff from the convenience of the software itself.) To be concrete, suppose that A-A interactions yield a payoff of 4 to each user, A-B interactions yield a payoff of 1 to each user, and the use of A yields a payoff of 1 to the user *in addition to* the networking payoffs. Similarly, suppose that B-B interactions yield a payoff of 1 to each user, B-A interactions also yield a payoff of 1, while using B yields a payoff of 3 in addition to the networking payoffs. The combination of these effects leads to the following total payoff matrix:

	networking		*own use*		*total payoff*	
	A	B	A	B	A	B
A	4, 4	1, 1	1, 1	1, 3	5, 5	2, 4
			+		=	
B	1, 1	1, 1	3, 1	3, 3	4, 2	4, 4

We claim that the efficient outcome is for everyone to adopt A, but the stochastically stable outcome is for everyone to use B. To see why this is so, we need to compute the two resistances $r(z^B \to z^A)$ and $r(z^A \to z^B)$. This involves finding the smallest number, k^*, of mistakes or mutations that are needed to tip the process from z^B to z^A. This is the least integer satisfying the inequality

$$5k^* + 2(s - k^*) \geq 4k^* + 4(s - k^*).$$

Subject to rounding this leads to the estimate $r(z^B \to z^A) \approx 2s/3$. Similarly we find that $r(z^A \to z^B) \approx s/3$. Since the latter is smaller, it follows from theorem 1 that (when s is sufficiently large) the stochastically stable state is all-B, which of course is not efficient.

Suppose, more generally, that the payoff matrix is of form

	A	B
A	a	c
B	d	b

When $a > d$ and $b > c$, this is a symmetric coordination game with coordination equilibria (A, A) and (B, B). We say that alternative A is *strictly risk-dominant* if $a - d > b - c$. Similarly, B is *strictly risk-dominant* if the reverse inequality holds. Note that risk dominance is not the same as efficiency, which is determined by the larger of a and b. One implication of the preceding analysis is the following.

Theorem 2. *Let G be a 2 x 2 coordination game with a strictly risk-dominant equilibrium. If G is played by a population of n players using samples of size s, then for all sufficiently large s and n (s < n/2) the unique stochastically stable state is the one in which everyone plays the risk-dominant alternative.*

This result has an interesting implication for the relative "fitness" of competing technologies. Consider again the situation in which each individual's payoff can be decomposed into a payoff from networking and a payoff from own use. We can write this in the following general form:

18

	networking		own use		total payoff	
	A	B	A	B	A	B
A	a	c	a'	a'	$a + a'$	$c + a'$
B	c	b	b'	b'	$c + b'$	$b + b'$

with $+$ between the networking and own use tables and $=$ between the own use and total payoff tables. (5)

Assume that $a > c$ and $b > c$, so that both A and B are coordination equilibria. By definition, risk dominance is determined by the inequality

$$(a + a') - (c + b') > (b + b') - (c + a'),$$

that is,

$$a + 2a' > b + 2b'. \qquad (6)$$

This has the following implication for the producers of A and B. Suppose that one of the firms – say the A-producer – is contemplating whether to invest in improvements that lead to greater networking transparency with other As, or to greater ease of use. Where should the money be invested to maximize the chance that A will take over the market? The answer is that investment in networking should be chosen only if it increases each user's utility at least *twice* as much as a similar investment in non-networking improvements. For example, suppose that A and B represent two types of cellphones. Suppose that, for a given expenditure, the firm producing A can either improve the clarity of the signal with other A-users, or improve the ease of reading the monitor independently of other users. Say that the first improvement increases the payoff to each A-user by Δ per call made to other A's, whereas the second increases it by Δ' per call made to anyone. If everyone in the population were using A, the firm would simply evaluate which is larger: Δ or Δ'. But in a competition for acceptance, the relevant criterion is the larger of Δ or $2\Delta'$. The reason is that Δ results from externalities with other A-users, whereas Δ'

does not. To my knowledge this point has not been previously recognized in the literature on network externalities.

6. Application to Schelling's segregation model

We turn now to a more complex example that illustrates the power of the analytical method discussed above. One of the earliest agent-based models in the social science literature is Schelling's illustration of how segregated neighborhoods can emerge spontaneously from decisions by individuals who would in fact prefer to live in integrated settings [Schelling (1971, 1978)].   Here we shall present a variant of Schelling's model that lends itself to the stochastic analysis discussed above; for an extension of the analysis to more complex environments see Zhang (2004a, 2004b).

Assume that the population consists of n individuals, who are of two types: A and B. They cannot change their type, but they can choose where to live. Suppose for simplicity that they are located around a circle as shown in figure 6. We shall say that an individual is *discontent* if his two immediate neighbors are unlike himself; otherwise he is *content*.   An *equilibrium* is a state in which no two individuals want to trade places.   In other words, there is no pair of agents such that one (or both) is currently discontent, and both would be content after they trade locations. (If only one agent is discontent beforehand, we can imagine that he compensates the other to move, so that both are better off after the move than they were before.)

Figure 6. A disequilibrium state.

A = ●
B = ○
d = discontent
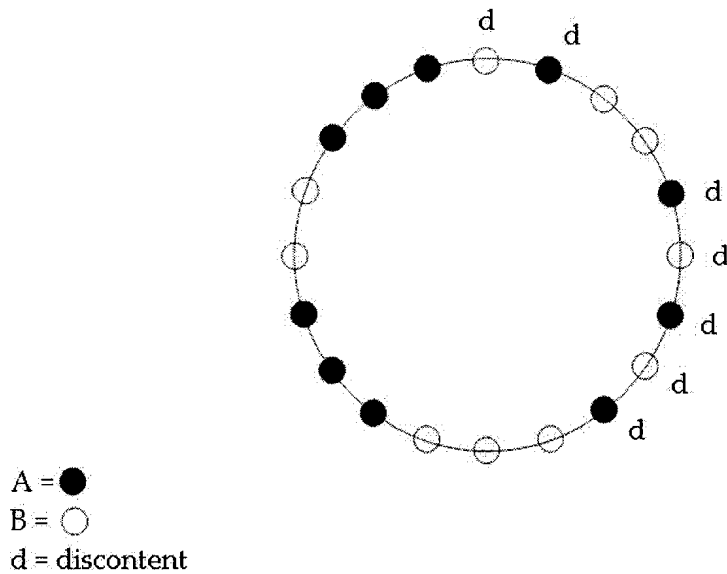
We claim that, if there are at least two agents of each type, then in equilibrium no one is discontent. To see why this is so, suppose to the contrary that an A is surrounded by two B's: . . .BAB . . .. Moving clockwise around the circle, let B* be the last B-type in the string of Bs who follows this A, and let A* be the agent who follows B*:

$$. . . BAB . . . BB^*A^* . . .$$

Since there are at least two agents of each type, we can be sure that A* differs from the original A. But then the original discontent A could switch with B* (who is content), and both would be content afterwards. Thus we see that the equilibrium configurations consist of those arrangements in which everyone lives next to at least one person of his own type. No one is "isolated." In general there are many different kinds of equilibrium states: some consist of small enclaves of A's and B's scattered around the landscape, while others exhibit full segregation with the A's

21

living on one side of the circle and the B's on the other.

Consider the following adjustment dynamic. In each discrete time period a pair of individuals is selected at random, where all pairs are equally likely to be chosen. Consider such a pair of individuals, say i and j. We shall assume that the probability that they trade depends on their prospective gains from trade. Let us assume that every trade involves moving costs. Thus there can be positive gains from trade only if the partners are of opposite types and at least one of them (say i) was discontent before and is content afterwards. This means that, before the trade, i was surrounded by people of the opposite type, so in fact both i and j are content afterwards. (We shall assume that if j is content before and after the trade, i can compensate j for his moving costs and still leave both better off.) Such Pareto improving trades are said to be *advantageous;* all other trades are *disadvantageous.*

Assume that each advantageous trade occurs with high probability, and that each disadvantageous trade occurs with low probability. Specifically, let us suppose that there exist real numbers $0 < a < b < c$ such that the probability of a disadvantageous trade is ε^a if neither partner's degree of contentment changes (so the losses involve only moving costs), the probability is ε^b if both partners were content before and one is discontent after, and it is ε^c if both were content before and both are discontent after. (These are the only possibilities.) Advantageous trades are assumed to occur with probabilities that approach one as $\varepsilon \to 0$; beyond this we need not specify the probabilities exactly. The resulting perturbed Markov process P^ε is ergodic for every $\varepsilon > 0$, and regular in the sense defined earlier.

To apply the theory, we first need to identify the recurrence classes of the unperturbed process P^0. These obviously include the absorbing (equilibrium) states. We claim that these are in fact the only recurrence classes of P^0. To prove this, consider a state that is not absorbing. It contains at least one discontent individual,

22

say i; without loss of generality we may assume that i is of type A. Going clockwise around the circle, let i' be the next individual of type A. (Recall that there are at least two individuals of each type.) The individual just before i' must be of type B. Call this individual j. If i and j trade places, both will be content afterwards. In any given period there is a positive probability that this pair will in fact be drawn, and that they will trade. The resulting state has fewer discontent individuals. Continuing in this manner, we see that from any non-absorbing state there is a positive probability of transiting to an absorbing state within a finite number of periods. Hence the absorbing states are the only recurrent states.

Denote the set of all absorbing states by Z^0. For any two states z and z' in Z^0, let r(z, z') denote the *least* resistance among all paths from z to z'. The stochastic potential of $z \in Z^0$ is defined to be the resistance of the minimum resistance z-tree on the set of nodes Z^0. By theorem 1, the stochastically stable states are those with minimum stochastic potential. We claim that these are precisely the *segregated absorbing states*, that is, states in which all the A's are lined up on one side of the circle and all the B's are on the other.

To prove this claim, let $Z^0 = Z^s \cup Z^{ns}$ where Z^s is the set of segregated absorbing states and Z^{ns} is the set of non-segregated absorbing states. We claim that (i) for every $z \in Z^{ns}$, every z-tree has at least one edge with resistance b or c (which by assumption are greater than a); and (ii) for every $z \in Z^s$, there exists a z-tree in which every edge has resistance exactly equal to a. Assume for the moment that (i) and (ii) have been established. In any z-tree there are exactly $|Z^0| - 1$ edges, and the resistance of each edge is *at least* a. It follows from (i) and (ii) that the stochastic potential of every segregated state equals $a|Z^0| - a$, while the stochastic potential of every non-segregated state is at least $a|Z^0| - 2a + b$, which is strictly larger. Theorem 1 therefore implies that the segregated states are precisely the

stochastically stable states.

To establish (i), let $z \in Z^{ns}$ be a non-segregated absorbing state. Given any z-tree T, there exists at least one edge in T that is directed from a segregated absorbing state z^s to a non-segregated absorbing state z^{ns}. We claim that any such edge has resistance at least b. The reason is that any trade that breaks up a segregated state must create at least one discontent individual, hence the probability of such a trade is either ε^b or ε^c (see figure 7). Thus the resistance of the edge from z^s to z^{ns} must be at least b, which establishes (i).

Figure 7. Single disadvantageous trade leading from an equilibrium to a disequilibrium state.
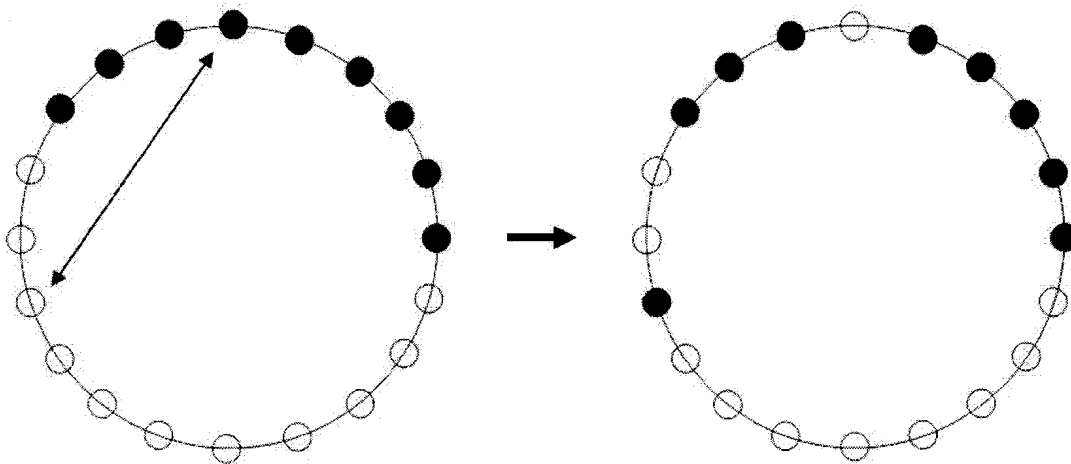
To establish (ii), let $z \in Z^s$ be a segregated absorbing state. From each state $z' \neq z$ we shall construct a sequence of absorbing states $z' = z^1 \to z^2 \to \ldots \to z^k = z$ such that $r(z^{j-1} \to z^j) = a$ for $1 < j \leq k$. Call this a z'z-*path* . We shall carry out the construction so that the union of all of the directed edges on all of these paths forms a z-tree. Since each edge has a resistance of a and the tree has $|Z^0| - 1$ edges, the total resistance of the tree is $a|Z^0| - a$ as claimed in (ii).

24

Suppose first that z' is also segregated, that is, z' consists of a single contiguous A-group and a complementary contiguous B-group. Label the positions on the circle 1, 2, . . ., n, in the clockwise sense. Let the first member of the A-group trade places with the first member of the B-group. Since both were content before and after, this trade has probability ε^a. It also results in a new absorbing state, which shifts the A-group and the B-group by one position clockwise around the circle. Hence within n steps we can reach any absorbing state, and in particular we can reach z. Thus we have constructed a sequence of absorbing states that leads from z' to z, where the resistance of each successive pair in the sequence equals a.

Suppose alternatively that z' is not segregated. Moving clockwise from position 1, let **A** denote the first complete group of contiguous As. Let **B** be the next group of Bs, and **A'** the next group of As. Since z' is absorbing, each of these groups contains at least two members. Let the first player in **A** trade places with the first player in **B** (in the clockwise labeling). Since both players were content before and after the trade, its probability is ε^a. This trade also shifts group **A** one position clockwise and reduces by one the number of B players between **A** and **A'**. It either results in a new absorbing state, or else a single B player remains between **A** and **A'**. In the latter case this B player can then trade with the first player in group **A**, and this trade has zero resistance. The result is an absorbing state with fewer distinct groups of As and Bs.

Repeat the process described in the preceding paragraph until all the As are contiguous and all the Bs are contiguous. Then continue as in the earlier part of the argument until we reach the target state z. This construction yields a sequence of absorbing states that begins at z' and ends at z, where the resistance between each successive pair of states is a. The path contains no cycles because the number of distinct groups never increases; indeed with each transition one of the groups

25

shrinks until it is eliminated. Thus the union of these paths forms a z-tree whose total resistance is $a \mid Z^0 \mid - a$. This concludes the proof that the stochastically stable states are precisely the segregated ones.

7. Local interaction models

Schelling's model is an example of a situation in which agents adapt their behaviors to the actions of their near neighbors. We can easily imagine that the same issue could arise in a model of technological adoption. What happens if people adopt practices or technologies based only on the choices of their immediate neighbors, as opposed to a sample drawn from the population at large? In this section we show how to address this problem using methods from statistical mechanics, an approach pioneered by Blume (1993, 1995).

Consider a group of n agents who are located in a social or geographic space that allows us to talk about their proximity. A very general model of this sort is to suppose that each agent lives at the vertex of a graph. The edges of the graph have weights that indicated the degree of proximity or influence that pertains to each pair of agents. To be specific, let V denote the set of vertices, and let $i \in V$ denote a particular vertex (which is identified with an agent whom we shall also call i). Let $w_{ij} \geq 0$ be a weight that measures the proximity of agents i and j in a geographical (or social) sense. We shall assume that this is a symmetric relation, that is, $w_{ij} = w_{ji}$.

Let X be a finite set of available options or choices. The *state* of the process at time t specifies the choice of each agent at that time. A state can therefore be represented as an n-dimensional vector $x^t \in X^n$, where x_i^t is i's choice at time t. Each individual i gets to reconsider his choice at random times governed by a Poisson random variable ω_i. We shall assume that the random variables ω_i are independent and identically distributed among agents, and that time is scaled so that, on average,

26

there is one revision opportunity per time period at each location. (Allowing differences in the rates of revision opportunities does not change the analysis in any fundamental way.)

In line with our earlier discussion, we shall decompose the utility of each agent i into two parts: the utility of the choice itself (without externalities), and the positive externality from doing what "the Joneses" do. Specifically, let $w_{ij} u(x, y)$ be the *externality payoff* from choosing x at location i when one's neighbor at location j chooses y. Thus $e_i(\mathbf{x}) = \sum_{j \neq i} w_{ij} u(x_i, x_j)$ denotes i's externality payoff in state $\mathbf{x}$. Let $v_i(\mathbf{x})$ denote the utility that i derives from x_i itself without regard to externalities. Assume that i's utility in state $\mathbf{x}$ at time t is given by

$$U_i(\mathbf{x}^t) = v_i(x_i^t) + e_i(\mathbf{x}^t) + \varepsilon_i^t, \tag{7}$$

where ε_i^t is an unobserved utility shock. It is analytically convenient to assume that the ε_i^t are independent and identically distributed according to the extreme value distribution.[6] Suppose that i chooses x_i^t to maximize U_i given that the others' choices at time t are fixed. It can be shown that, from the observer's point of view, i chooses $x_i^t \in X$ according to the logistic distribution

$$P(x_i^t \mid x_{-i}^t) = \exp \beta[v_i(x_i^t) + e_i(\mathbf{x}^t)] / \sum_{y_i \in \mathbf{X}} \exp \beta[v_i(y_i) + e_i(y_i, x_{-i}^t)] \tag{8}$$

---

[6]The random variable z is extreme value distributed if its cumulative distribution function F(z) takes the form $\ln F(z) = -e^{-\beta z}$. This distribution is analytically convenient because it yields a simple closed-form solution for the stationary distribution of the adjustment process; moreover it is standard as a model of discrete choice [McFadden (1974), Blume (1993, 1995), McKelvey and Palfrey (1995), Durlauf (1997), Brock and Durlauf (2001). Alternative error distributions can be analyzed using the methods discussed in section 2.

The resulting stochastic adjustment process can be represented as a finite Markov chain. This process has a unique recurrence class (namely the whole state space) because the choice model implies that any choice will be made with positive probability whenever an agent reconsiders. Hence the process is ergodic and has a unique stationary distribution μ^β on the set of states $\mathbf{X}$. For each $\mathbf{x} \in \mathbf{X}$, $\mu^\beta(\mathbf{x})$ represents the long-run relative frequency with which state $\mathbf{x}$ is visited starting from any initial state.

A noteworthy feature of this set-up is that the stationary distribution can be expressed in a simple closed form. Specifically, define the *potential* of state $\mathbf{x}$ to be

$$\rho(\mathbf{x}) = \sum_{i=1}^{n} v_i(x_i) + (1/2) \sum_{i=1}^{n} e_i(\mathbf{x})]. \tag{9}$$

Thus the potential of a state equals the nonexternality payoffs generated individuals' choices, plus one-half the externalities generated by social interactions. It can be shown that the long-run distribution of the process has the following simple form, known as a Gibbs representation:

$$\tag{10}$$

$$\mu(\mathbf{x}) = e^{\beta\rho(\mathbf{x})} \Big/ \sum_y e^{\beta\rho(y)}$$

It follows that, when β is large, the probability is close to one that the process will be in a state that maximizes potential, that is, the stochastically stable states are precisely those that maximize $\rho(\mathbf{x})$.

Theorem 3. *Starting from an arbitrary initial state, the long-run probability of being in any given state $\mathbf{x}$ is proportional to $e^{\beta\rho(\mathbf{x})}$. When β is large, the probability is close to one that the process is in a state $\mathbf{x}$ that maximizes $\rho(\mathbf{x})$.*

We remark that this model can be applied to the technology adoption problem discussed in section 5. Recall that in this case the choice set consists of just two options, A and B, and the utilities are given by the payoff matrix

	networking		own use			total payoff	
	A	B	A	B		A	B
A	a	c	a'	a'		$a + a'$	$c + a'$
		+			=		
B	c	b	b'	b'		$c + b'$	$b + b'$

Now let us suppose a social structure among the agents that determines who interacts with whom. Specifically, assume that each agent is joined by an edge to s other agents (the graph is regular of degree s), and that the weight on each edge is 1/s.

Consider any state x, and let $n_{AA}(x)$ be the total number of edges such that the agents at both ends of the edge choose A. Similarly, let $n_{BB}(x)$ be the total number of edges such that the agents at both ends choose B, and let $n_{AB}(x)$ be the total number of edges such that the agent at one end chooses B and the agent at the other end chooses A. Next, let $n_A(x)$ be the number of agents who choose A and let $n_B(x)$ be the number who choose B. Note that $n_A(x) + n_B(x) = n$ and $n_{AA}(x) + n_{BB}(x) + n_{AB}(x) = n/2$. The potential function in (9) can then be written as follows:

$$\rho(x) = a'n_A(x) + b'n_B(x) + (1/s)(an_{AA}(x) + bn_{BB}(x) + cn_{AB}(x)). \tag{11}$$

This is maximized either by the all-A state x^A or the all-B state x^B. Thus we wish to evaluate which is larger:

$$\rho(x^A) = a'n + an/2 \quad \text{or} \quad \rho(x^B) = b'n + bn/2 \ . \tag{12}$$

This amounts to finding the larger of $a + 2a'$ and $b + 2b'$, which is exactly the risk dominance criterion (see the derivation of (6)). It can be shown, in fact, that risk dominance is the relevant criterion of stochastic stability in a wide variety of binary choice situations [Kandori, Mailath, and Rob (1993), Blume (2003)], though this is not true when more than two choices are available [Young (1993a)].

8. Contractual norms

The framework outlined above has potential application to any situation in which social norms influence individual agents' decisions. Cases in which this possibility has been discussed include the use of addictive substances, dropping out of school, and criminal behavior [Case and Katz (1991, Crane (1991), Glaeser, Sacerdote and Scheinkman (1996)]. In this section we apply the theory to yet another domain, the role of social norms in shaping the terms of economic contracts. In particular we show how it can illuminate the pattern of cropsharing contracts found in contemporary U.S. agriculture [Young and Burke (2000)].[7]

A *share contract* is an arrangement in which a landowner and a tenant farmer split the gross proceeds of the harvest in fixed proportions or shares. The logic of such a contract is that it shares the risk of an uncertain outcome while offering the tenant a rough-and-ready incentive to increase the expected value of that outcome. When

---

[7] Applications of the theory to the evolution of bargaining norms may be found in Young (1993b) and Young (1998, chapter 9).

contracts are competitively negotiated, one would expect the size of the share to vary in accordance with the mean (and variance) of the expected returns, the risk aversion of the parties, the agent's quality, and other relevant factors. In practice, however, shares seem to cluster around "usual and customary" levels even when there is substantial heterogeneity among principal-agent pairs, and substantial and readily observed differences in the quality of different parcels of land. These contractual customs are pinned to psychologically prominent focal points, such as 1/2-1/2, though other shares -- such as 1/3-2/3 and 2/5-3/5 -- are also common, with the larger share going to the tenant.

A striking feature of the Illinois data is that the above three divisions account for over 98% of all share contracts in the survey, which involved several thousand farms in all parts of the state. An equally striking feature is that the predominant or customary shares differ by region: in the northern part of the state the overwhelming majority of share contracts specify 1/2-1/2, whereas in the southern part of the state the most common shares are 1/3-2/3 and 2/5-3/5 [Illinois Cooperative Extension Service (1995)].[8] Thus, on the one hand, uniformity *within* each region exists in spite of the fact that there are substantial and easily observed differences in the soil characteristics and productivities of farms within the region. On the other hand, large differences exist *between* the regions in spite of the fact that there are many farms in both regions that have essentially the same soil productivity, so in principle they should be using the same (or similar) shares. The local interaction model discussed in section 7 can help us to understand these apparent anomalies.

---

[8] This north-south division corresponds roughly to the southern boundary of the last major glaciation. In both regions, farming techniques are similar and the same crops are grown -- mainly corn, soybeans, and wheat. In the north the land tends to be flatter and more productive than in the south, though there is substantial variability within each of the regions.

Let us identify each farm i with the vertex of a graph. Each vertex is joined by edges to its immediate geographical neighbors. For ease of exposition we shall assume that the social influence weights on the edges are all the same. The *soil productivity index* on farm i, s_i, is a number that gives the expected output per acre, measured in dollars, of the soils on that particular farm. (For example, $s_i = 80$ means that total net income on farm i is, on average, $80 per acre.) The contract on farm i specifies a share x_i for the tenant, and $1 - x_i$ for the landlord, where x_i is a number between zero and one. The tenant's expected income on farm i is therefore $x_i s_i$ times the number of acres on the farm. For expositional convenience let us assume that all farms have the same size, which we may as well suppose is unity. (This does not affect the analysis in any important way.)

Assume that renegotiations occur on each farm according to i.i.d. Poisson random variables, as described in the preceding section. When the time comes to renegotiate on a particular farm, say i, the landlord makes an offer, say x_i. The tenant accepts if and only if his expected return $x_i s_i$ is at least w_i, where w_i is the reservation wage at location i. The expected monetary return to the landlord from such a deal is $v_i(x_i) = (1 - x_i)s_i$.

To model the impact of local custom, suppose that each of i's neighbors exerts the same degree of social influence on i. Specifically, for each state x, let $\delta_{ij}(x) = 1$ if i and j are neighbors and $x_i = x_j$; otherwise let $\delta_{ij}(x) = 0$. We assume that i's utility in state x is $(1 - x_i)s_i + \gamma \Sigma_j \delta_{ij}(x)$, where γ is a *conformity parameter*. The idea is that, if a landlord offers his tenant a contract that differs from the practices of the neighbors, the tenant will be offended and may retaliate with poorer performance. Hence the landlord's utility for different contracts is affected by the choices of his neighbors. The resulting potential function is

$$\Sigma_i(1 - x_i)s_i + (\gamma/2) \Sigma_{i,j} \delta_{ij}(x). \tag{14}$$

Note that $\Sigma_i(1 - x_i)s_i$ represents the total *rent to land*, which we shall abbreviate by $r(x)$. The expression $(1/2) \Sigma_{i,j} \delta_{ij}(x)$ represents the total number of edges (neighbor-pairs) that are coordinated on the same contract in state x, which we shall abbreviate by $e(x)$. We can therefore write

$$\rho(x) = r(x) + \gamma e(x). \tag{15}$$

As in (10) it follows that the stationary distribution, $\mu(x)$, has the Gibbs form

$$\mu(x) \propto e^{\beta[r(x) + \gamma c(x)]}. \tag{16}$$

It follows that *the log probability of each state x is a linear function of the total rent to land plus the degree of local conformity*. Given specific values of the conformity parameter γ and the response parameter β, we can compute the relative probability of various states of the process, and from this deduce the likelihood of different geographic distributions of contracts. In fact, one can say a fair amount about the qualitative behavior of the process even when one does not know specific values of the parameters.

We illustrate with a concrete example that is meant to capture some of the key features of the Illinois case. Consider the hypothetical state of Torusota shown in figure 8. In the northern part of the state--above the dashed line--soils are evenly divided between High and Medium quality soils. In the southern part they are evenly divided between Medium and Low quality soils. As in Illinois the soil types are interspersed, but average soil quality is higher in the north than it is in the

33

south. Let n be the number of farms. Each farm is assumed to have exactly eight neighbors, so there are 4n edges altogether. Let us restrict the set of contracts to be in multiples of ten percent: $x = 10\%, 20\%, \ldots, 90\%$. (Contracts in which the tenant receives 0% or 100% are not considered.) For the sake of concreteness, assume that High soils have index 85, Medium soils have index 70, and Low soils have index 60. Let the reservation wage be 32 at all locations.

Figure 8. The hypothetical state of Torusota. Each vertex represents a
farm, and soil qualities are High (H), Medium (M), or Low (L).
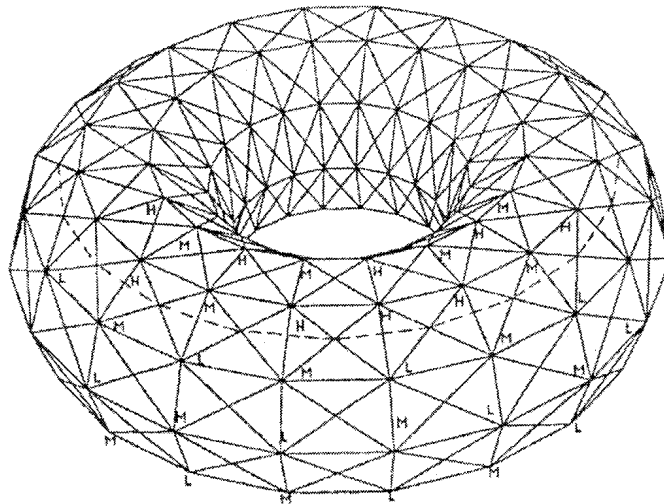
We wish to determine the states of the process that maximize the potential function $\rho(x)$. The answer depends, of course, on the size of γ, that is, on the tradeoff rate between the desire to conform with community norms and the amount of economic payoff one gives up in order to conform.

Consider first the case where $\gamma = 0$, that is, there are no conformity effects. Maximizing potential is then equivalent to maximizing the total rent to land, subject always to the constraint that labor earns at least its reservation wage on each

34

class of soil. The contracts with this property are 40% on High soil, 50% on Medium soil, and 60% on Low soil. The returns to labor under this arrangement are: 34 on H, 35 on M, and 36 on L. Notice that labor actually earns a small premium over the reservation wage (w = 32) on each class of soil. This *quantum premium* is attributable to the discrete nature of the contracts: no landlord can impose a less generous contract (rounded to the nearest 10%) without losing his tenant. Except for the quantum premium, this outcome is the same as would be predicted by a standard market-clearing model, in which labor is paid its reservation wage and all the rent goes to land. We shall call this the *competitive* or *Walrasian* state **w**.

Notice that, in contrast to conventional equilibrium models, our framework actually gives an account of how the state **w** comes about. Suppose that the process begins in some initial state x^0 at time zero. As landlords and tenants renegotiate their contracts, the process gravitates towards the equilibrium state **w** and eventually reaches it with probability one. Moreover, if β is not too small, the process stays close to **w** much of the time, though it will rarely be *exactly* in equilibrium.

These points may be illustrated by simulating the process using an agent-based model. Let there be 100 farms in the North and 100 in the South, and assume a moderate level of noise (β = 0.20). Starting from a random initial seed, the process was simulated for three levels of conformity: γ = 0, 3, and 8. Figure 9 shows a typical distribution of contract shares after 1000 periods have elapsed. When γ = 0 (bottom panel), the contracts are matched quite closely with land quality, and the state is close to the competitive equilibrium. When the level of conformity is somewhat higher (middle panel), the dominant contract in the North is 50%, in the South it is 60%, and there are pockets here and there of other contracts. (This looks quite similar to the Illinois case.) Somewhat surprisingly, however, a further increase in the conformity level (top panel) does not cause the two regional customs

to merge into a single global custom; it merely leads to greater uniformity in each of the two regions.

To understand why this is so, let us suppose for the moment that everyone is using the *same* contract x. Since everyone must be earning their reservation wage, x must be at least 60%. (Otherwise southern tenants on low quality soil would earn less than w = 32.) Moreover, among all such global customs, x = 60% maximizes the total rent to land. Hence the 60% custom, which we shall denote by **y**, maximizes potential among all global customs. But it does not maximize potential among all states. To see why this is so, let **z** be the state in which everyone in the North uses the 50% contract, while everyone in the South everyone uses the 60% contract. State **z**'s potential is almost as high as **y**'s potential, because in state **z** the only negative social externalities are suffered by those who live near the north-south boundary. Let us assume that the number of such agents is on the order of $\sqrt{n}$, where n is the total number of farms. Thus the *proportion* of farms near the boundary can be made as small as we like by choosing n large enough. But **z** offers a higher land rent than **y** to all the northern farms as compared to **y**. To be specific, assume that there are n/2 farms in the north, which are evenly divided between High and Medium soils, and that there are n/2 farms in the south, which are evenly divided between Medium and Low soils. Then the total income difference between **z** and **y** is 7n/4 on the Medium soil farms in the north, and 8.5n/4 on the High soil farms in the north, for a total gain of 31n/8. It follows that, if γ is large enough, then for all sufficiently large n, the regional custom **z** has higher potential than the global custom **y**.[9]

---

[9] A more detailed calculation shows that z uniquely maximizes potential among

αλλ στατεσ ωηενεϖερ γ is sufficiently large and n is sufficiently large relative to γ.

Figure 9. Simulated outcomes of the process for n = 200, β = 0.20.
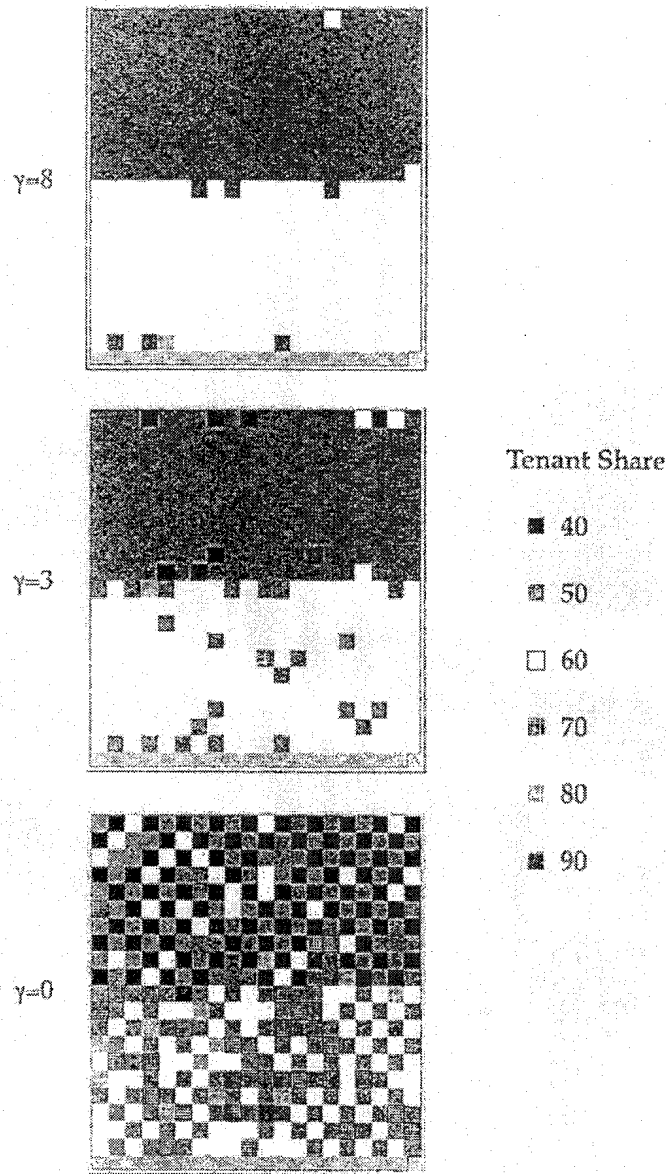
While the details are particular to this example, the logic is quite general. Consider any distribution of soil qualities that is heterogeneous locally, but exhibits substantial shifts in average quality between geographic regions. For intermediate values of conformity γ, it is reasonable to expect that potential will be maximized by a distribution of contracts that is uniform locally, but diverse globally--in other

words the distribution is characterized by *regional customs*. Such a state will typically have higher potential than the competitive equilibrium, because the latter involves substantial losses in social utility when land quality is heterogeneous. Such a state will typically also have higher potential than a global custom, because it allows landlords to capture more rent at relatively little loss in social utility, provided that the boundaries between the regions are not too long (i.e., there are relatively few farms on the boundaries).

In effect, these regional customs form a compromise between completely uniform contracts on the one hand, and fully differentiated, competitive contracts on the other. Given the nature of the model, we should not expect perfect uniformity within any given region, nor should we expect *sharp* changes in custom at the boundary. The model suggests instead that there will be occasional departures from custom within regions (due to idiosyncratic influences), and considerable variation near the boundaries. These features are precisely what we see in the distribution of share contracts in Illinois.

9. Conclusion

In this paper I have described a framework for analyzing the asymptotic behavior of a wide variety of agent-based models. In particular, the theory makes quantitative predictions about the long-run probability of various outcomes and thus avoids the hazards of drawing conclusions solely from simulations. (Of course, simulations can still be extremely helpful in understanding the short and medium run dynamics of a process.) I have shown how the theory plays out in specific contexts, including technology adoption, neighborhood segregation, and the evolution of contractual norms. Perhaps the most important aspect of the theory, however, is that it brings into focus certain qualitative features that are common to many agent-based models, but that one does not tend to find in conventional

equilibrium types of analysis. The three critical features are: i) local conformity vs. global diversity, ii) punctuated equilibrium; iii) persistence of particular states in the presence of stochastic shocks [Young (1998)].

We can illustrate these concepts by imagining a collection of distinct societies whose members do not interact with each other. Over time, each will develop distinctive institutions to cope with various forms of economic and social coordination -- forms of contracts, norms of behavior, property rights, technological standards, and so forth. The solutions that each society finds to these coordination problems will typically take the form of an equilibrium state in an appropriately defined dynamical system. Due to the positive externalities that arise from conforming to the reigning equilibrium, one will tend to find a substantial amount of conformity within a given community. But in separate, non-interacting communities, one may find that the same basic problem is solved in different ways. This is the *local conformity/global diversity effect*. It can apply even within a given society if interactions are sufficiently localized and the externalities are sufficiently strong; Illinois agricultural contracts provide a real-world instance of this phenomenon. A concrete prediction is that two individuals are more likely to exhibit similar behaviors if they come from the same society (or are close in the relevant social network) than if they come from different societies, holding constant all other explanatory variables.

The second qualitative feature of this class of models has to do with the look of the dynamic paths. The theory predicts that the process will tend to exhibit long periods of stasis in which a given equilibrium--or something close to an equilibrium--is in place, punctuated by bursts in which the equilibrium shifts in response to stochastic shocks. In the context of residential segregation Schelling called this the "tipping" phenomenon; here I refer to it as the *punctuated equilibrium effect*.

The third key feature highlighted by the theory is that some equilibrium states are more *persistent* or stable than others. Once established they tend to stay in place for long periods of time because they are robust against stochastic shocks. The methodology outlined allows us to identify these stochastically stable states using the concept of a stochastic potential function. This approach also allows us to make predictions about the long-run behavior of specific dynamical systems, such as segregated outcomes being more stable than integrated ones, and risk dominant technologies being more stable than efficient ones.

I conclude this essay by drawing attention to an important aspect of the theory that we cannot explore in depth here, but that deserves particular recognition, namely, the length of time that it takes for the long-run asymptotic behavior of an evolutionary process to reveal itself. From an empirical point of view it obviously makes a difference if a process takes ten years or a million years to reach its long-run distribution. In the latter case, the short-run dynamics are more important than the long-run asymptotics, and the process may be effectively path dependent even if it is not so from a truly long-run perspective. In practice, however, it is quite difficult to say how long the long run really is. There are several reasons for this. One is that time periods in the model do not correspond to real time intervals; they simply represent markers between distinct events in the model, such as revision decisions by individuals. When the population is large and people interact often, thousands or even millions of such events might be compressed within a short period of real time, such as an hour or a day. Second, the speed of adjustment depends on a number of modeling factors, including the degree of local interaction [Ellison (1993), Young (1998, Chapter 6)], the amount of information that people use to make their decisions, and the extent to which agents' errors are correlated. If agents react only to the behavior of a few neighbors, or they get their information by asking a few friends, or they react similarly to the same conditions, the process

can tip from one equilibrium to another in relatively short order. Thus unless we know quite a lot about the topology of interaction and the agents' decision-making processes, estimates of the speed of adjustment could be off by many orders of magnitude.

The theory discussed above identifies those aspects of evolutionary, agent-based models that are critical to determining the speed with which change occurs. The remaining challenge is to bring these theoretical predictions to bear on the forms of social structure that we see in the real world.

References

Arthur, W. Brian. (1989), "Competing technologies, increasing returns, and lock-in by historical events", Economic Journal 99: 116-31.

Bendor, Jonathan, and Piotr Swistak. (2001), "The evolution of norms", American Journal of Sociology 106: 1493-1545.

Bicchieri, Cristina (forthcoming), The Grammar of Society: The Emergence and Dynamics of Social Norms (Cambridge University Press, Cambridge).

Blume, Lawrence E. (1995), "The statistical mechanics of strategic interaction", Games and Economic Behavior 4: 387-424.

Blume, Lawrence E. (1995), "The statistical mechanics of best-response strategy revision", Games and Economic Behavior 11: 111-45.

Blume, Lawrence E. (2003), "How noise matters", Games and Economic Behavior 44: 251-271.

Brock, William A. and Steven N. Durlauf. (2001), "Discrete choice with social interactions", Review of Economic Studies 58: 313-323.

Case, Anne, and Lawrence Katz. (1991), "The company you keep: the effects of family and neighborhood on disadvantaged families", National Bureau of Economic Research Working Paper No. 3705.

Crane, J. (1991), "The epidemic theory of ghettos and neighborhood effects on dropping out and teenage childbearing", American Journal of Sociology 96: 1226-1259.

David, Paul. (1985), "Clio and the economics of QWERTY", American Economic Review 75: 332-37.

Durlauf, Steven N. (1997), "Statistical mechanics approaches to socioeconomic behavior" in: W. Brian Arthur, Steven N. Durlauf, and David Lane, The Economy as a Complex Evolving System, vol. 2 (Addison-Wesley, Redwood City, California).

Ellison, Glenn. (1993), "Learning, local interaction, and coordination", Econometrica 61: 1047-71.

Foster, Dean, and Peyton Young. (1990), "Stochastic evolutionary game dynamics", Theoretical Population Biology 38: 219-32.

Freidlin, Mark, and Alexander Wentzell. (1984), Random Perturbations of Dynamical Systems (Springer-Verlag, Berlin).

Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman. (1996), "Crime and social interactions", Quarterly Journal of Economics 111: 507-548.

Hechter, Michael, and Karl Dieter. (2001), Social Norms. New York: Russell Sage Foundation.

Illinois Cooperative Extension Service. (1995), "1995 Cooperative Extension Service Farm Leasing Survey", Department of Agricultural and Consumer Economics, Cooperative Extension Service, University of Illinois at Urbana-Champaign.

Kandori, Michihiro, George Mailath, and Rafael Rob. (1993), "Learning, mutation, and long-run equilibria in games", Econometrica 61: 29-56.

Karlin, Samuel, and Howard M. Taylor. (1975), A First Course in Stochastic Processes (Academic Press, New York).

Katz, Michael, and Carl Shapiro. (1985), "Network externalities, competition, and compatibility", American Economic Review 75: 424-440.

Katz, Michael, and Carl Shapiro. (1986), "Technology adoption in the presence of network externalities", Journal of Political Economy 94: 822-841.

Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent. (1990), "Money as a medium of exchange in an economy with artificially intelligent agents", Journal of Economic Dynamics and Control 14: 329-373.

McFadden, Daniel. (1974), "Conditional logit analysis of qualitative choice behavior" in: Paul Zarembka, F rontiers in Econometrics (Academic Press, New York) pp. 105-142.

McKelvey, Richard D. and Thomas R. Palfrey. (1995), "Quantal response equilibria for normal form games", Games and Economic Behavior 10: 6-38.

Menger, Karl. (1871), Grundsätze der Volkswirtschaftslehre (W. Braumüller, Vienna). English translation by James Dingwall and Bert F. Hoselitz, under the title Principles of Economics (Free Press, Glencoe, Illinois).

Schelling, Thomas C. (1971), "Dynamic models of segregation", Journal of_Mathematical Sociology 1:143-86.

Schelling, Thomas C. (1978), Micromotives and Macrobehavior (Norton, New York).

Skyrms, Brian. (2004), The Stag Hunt and the Evolution of Social Culture Princeton, NJ: Princeton University Press.

Sugden, Robert. (1986), The Evolution of Rights, Cooperation, and Welfare (Basil Blackwell, New York).

Sugden, Robert. (1989), "Spontaneous order", Journal of Economic Perspectives 3: 85-97.

Ullman-Margalit, Edna. (1977), The Emergence of Norms (Oxford University Press, Oxford, UK).

Young, H. Peyton. (1993a), "The evolution of conventions", Econometrica 61 (1): 57-84.

Young, H. Peyton. (1993b), "An evolutionary model of bargaining", Journal of Economic Theory 59(1): 145-168.

Young, H. Peyton. (1998), Individual Strategy and Social Structure: An Evolutionary Theory of Institutions (Princeton University Press, Princeton, NJ).

Young, H. Peyton, and Mary A. Burke. (2001), "Competition and custom in economic contracts: a case study of Illinois agriculture", American Economic Review 91: 559-573.

Zhang, Junfu. (2004a), "A dynamic model of residential segregation", Journal of Mathematical Sociology 28: 147-170.

Zhang, Junfu. (2004b), "Residential segregation in an all-integrationist world", Journal of Economic Behavior and Organization 54: 533-550.