
Markov Economies and Stochastic Dynamical Systems

God does not play dice
with the Universe.

Albert Einstein

Time-discrete stochastic processes are powerful tools for characterizing some dynamical systems. The prerequisites include an understanding of Markov processes (§13.1). Time-discrete systems behave quite differently from dynamical systems based on systems of ordinary differential equations. This chapter presents a Markov model of adaptive learning that illustrates the concept of stochastic stability, as developed in Young (1998). After developing some of the theoretical results, we provide an agent-based model.

13.1 Markov Processes

A *finite Markov process* is a dynamical system that in each time period $t = 0, 1, \dots$ can be any one of n states, such that if the system is in state i in one time period, there is a probability p_{ij} that the system will be in state j in the next time period. Thus, for each i , we must have $\sum_j p_{ij} = 1$, because the system must go somewhere in each period. We call the $n \times n$ matrix $P = \{p_{ij}\}$ the *transition probability matrix* of the Markov process, and each p_{ij} is called a *transition probability*.

Many games can be viewed as Markov processes. Here are some examples:

- a. Suppose two gamblers have wealth k_1 and k_2 dollars, respectively, and in each period they play a game in which each has an equal chance of winning one dollar. The game continues until one player has no more wealth. Here the state of the system is the wealth w of player 1, $p_{w,w+1} = p_{w,w-1} = 1/2$ for $0 < w < k_1 + k_2$, and all other transition probabilities are zero.

- b. Suppose n agents play a game in which they are randomly paired in each period, and the stage game is a prisoner's dilemma. Players can remember the last k moves of their various partners. Players are also given one of r strategies, which determine their next move, depending on their current histories. When a player dies, which occurs with a certain probability, it is replaced by a new player who is a clone of a successful player. We can consider this a Markov process in which the state of the system is the history, strategy, and score of each player, and the transition probabilities are just the probabilities of moving from one such state to another, given the players' strategies (§13.12).
- c. Suppose n agents play a game in which they are randomly paired in each period to trade. Each agent has an inventory of goods to trade and a strategy indicating which goods the agent is willing to trade for which other goods. After trading, agents consume some of their inventory and produce more goods for their inventory, according to some consumption and production strategy. When an agent dies, it is replaced by a new agent with the same strategy and an empty inventory. If there is a maximum-size inventory and all goods are indivisible, we can consider this a finite Markov process in which the state of the system is the strategy and inventory of each player and the transition probabilities are determined accordingly.
- d. In a population of beetles, females have k offspring in each period with probability f_k , and beetles live for n periods. The state of the system is the fraction of males and females of each age. This is a denumerable Markov process, where the transition probabilities are calculated from the birth and death rates of the beetles.

13.2 Long-run Behavior

We are interested in the long-run behavior of Markov processes, which is the relevant equilibrium concept for a stochastic dynamical system. Let $p_{ij}^{(m)}$ be the probability of being in state j in m periods if the Markov process is currently in state i . Thus, if we start in state i at period 1, the probability of being in state j at period 2 is just $p_{ij}^{(1)} = p_{ij}$. To be in state j in period 3 starting from state i in period 1, the system must move from state i to some state k in period 2, and then from k to j in period 3. This happens with probability $p_{ik} p_{kj}$. Adding up over all k , the probability of

being in state j in period 3 is

$$p_{ij}^{(2)} = \sum_k p_{ik} p_{kj}.$$

Using matrix notation, this means the matrix of two-period transitions is given by

$$P^{(2)} = \{p_{ij}^{(2)} | i, j = 1, 2, \dots\} = P^2.$$

Generalizing, we see that the k -period transition matrix is simply P^k . What we are looking for, then, is the limit of P^k as $k \rightarrow \infty$. Let us call this limit

$$P^* = \{p_{ij}^*\}.$$

We calculate P^* as follows. Suppose the $n \times n$ matrix M of left eigenvectors of P is nonsingular. Then if D is the $n \times n$ diagonal matrix with the eigenvalues of P along the diagonal, we have

$$MPM^{-1} = D. \tag{13.1}$$

To see this, note that the i^{th} row of M is the i^{th} eigenvector of P , so the i^{th} row of M times P equals i^{th} left eigenvalue of P times the i^{th} row of M . Thus we have $MP = DM$, and since M is invertible, we have $MPM^{-1} = D$. This equation allows us to calculate P^* rather easily because $P = M^{-1}DM$, so $P^2 = M^{-1}DMM^{-1}DM = M^{-1}D^2M$. Similarly, for all k , $P^k = M^{-1}D^kM$.

Now D^k is just the diagonal matrix with the k^{th} power of the eigenvalues down the diagonal, and P^* exists only if $D^* = \lim_{k \rightarrow \infty} D^k$ exists. We then have

$$P^* = M^{-1}D^*M. \tag{13.2}$$

If D^* exists, then $\lim_{k \rightarrow \infty} \lambda^k$ must exist for all eigenvalues λ of P . Thus we must have $|\lambda| \leq 1$, where $|\lambda|$ is the modulus of λ (i.e., if $\lambda = a + b\sqrt{-1}$, for a and b real, then $|\lambda| = \sqrt{a^2 + b^2}$). Now if $|\lambda| < 1$, then clearly $\lim_{k \rightarrow \infty} \lambda^k = 0$. However when $|\lambda| = 1$, λ^k has a limit only if $\lambda = 1$. Otherwise λ is a root of unity, so λ^k cycles indefinitely. For instance, $\lambda = -1$ is a second root of unity, and λ^k alternates between 1 and -1 . Similarly, $\lambda = \sqrt{-1}$ is a fourth root of unity, and λ^k cycles $\lambda, -1, -\lambda, 1, \lambda$, and so on.

The existence of P^* thus comes down to the question as to when P has a unit eigenvalue, and all its non-unit eigenvalues have modulus strictly less

that unity. It turns out that this is the case when the Markov process is irreducible and aperiodic, terms we are about to explain.

We say a state j of a Markov process can be **reached** from a state i if $p_{ij}^{(m)} > 0$ for some positive integer m . We say two states, i and j , **communicate** if each is reached from the other; i.e., if $p_{ij}^{(m)} > 0$ for some integer m , and $p_{ji}^{(m)} > 0$ for some, possibly different, integer m . We say a Markov process is **irreducible** if every two states communicate.

We say state i in a Markov process is **periodic** with period k if there is some integer $k > 1$ such that $p_{ii}^{(k)} > 0$ and $p_{ii}^{(m)} > 0$ if and only if m is a multiple of k . If no state of the Markov process M is periodic, we say M is **aperiodic**.

An irreducible, aperiodic finite Markov process is called *ergodic*. We have the following *ergodic theorem* for Markov processes, the proof of which can be found in Feller (1950).

THEOREM 13.1 *For any finite ergodic Markov process with transition matrix P , the following equations hold with all u_j unique and strictly positive.*

$$u_j = \lim_{m \rightarrow \infty} p_{ij}^{(m)} \quad \text{for } i = 1, \dots, n \quad (13.3)$$

$$u_j = \sum_i u_i p_{ij} \quad (13.4)$$

$$1 = \sum_k u_k, \quad (13.5)$$

for $j = 1, \dots, n$.

Note that (13.4) can be written in matrix notation as $u = uP$, so u is a left eigenvector of P with an associated unit eigenvalue. The first equation says that P^* exists and all its rows are the same and equal to the eigenvector u ; i.e., u_j is the limit probability of being in state j starting from any state. The fact that such a u is unique implies that P has no other unit eigenvalue, so the other $n - 1$ eigenvalues have modulus strictly less than one, which implies that the limit matrix D^* is all zeros except for a single unit entry along the diagonal which we assume is in the first row, first column, of D^* .

The second equation says that the probability of being in state j is the probability of moving from some state i to state j , which is $u_i p_{ij}$, summed over all states i . The equation states that u_i is the probability of being in

state i in the long run, so the probability of being in state j in the long run is just $u_i p_{ij}$ summed over all states i . The Markov process thus eventually spends a fraction of time u_j in state j , for each j , no matter where it started. It is in this sense that u is the “stationary distribution” of the Markov process.

The final equation affirms that u is indeed a probability distribution over the states of the Markov process. The *recursion equations* (13.4) and (13.5) are often sufficient to determine u , which we call the *invariant distribution* or *stationary distribution* of the Markov process.

The problem with (13.2) for calculating the stationary distribution of a Markov process is that it is difficult to calculate and invert the matrix of eigenvectors. It is often easier to solve the recursion equations (13.3)-(13.5) directly, as we now describe.

13.3 Solving for the Stationary Distribution

Consider first the n -state Markov process called the *random walk on a circle*, in which there are n states, and from any state $t = 2, \dots, n - 1$ the system moves with equal probability to the previous or the next state, from state n it moves with equal probability to state 1 or state $n - 1$, and from state 1 it moves with equal probability to state 2 and to state n . In the long run, it is intuitively clear that the system will be all states with equal probability $1/n$. To derive this from the recursion equations, note that the probability transition matrix for this Markov process is given by

$$P = \begin{bmatrix} 0 & 1/2 & 0 & \dots & 0 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & \dots & 0 & 0 & 0 \\ & & & \vdots & & & \\ 0 & 0 & 0 & \dots & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 0 & \dots & 0 & 1/2 & 0 \end{bmatrix}.$$

Clearly this Markov process is irreducible, and for odd n , it can be shown to be aperiodic, so the Ergodic Theorem holds.¹ The equations governing

¹The aperiodicity of this Markov process for odd n appears not to be completely trivial. You can check that each row of P^k has a pair of adjacent non-zero entries for the first time when $k = (n - 1)/2$, and each successive power of P increases the length of this series of adjacent non-zero entries by two, until when $k = n$, all entries are non-zero. This shows that the process is aperiodic.

this system are thus given by

$$\begin{aligned} u_1 &= \frac{1}{2}u_n + \frac{1}{2}u_2 \\ u_i &= \frac{1}{2}u_{i-1} + \frac{1}{2}u_{i+1} \quad i = 2, \dots, n-1 \\ u_n &= \frac{1}{2}u_1 + \frac{1}{2}u_{n-1} \\ \sum_{i=1}^n u_i &= 1. \end{aligned}$$

This set of equations has solution $u_i = 1/n$ for $i = 1, \dots, n$. The Ergodic Theorem asserts that this solution is unique, but you can prove this directly for yourself in this case. This result conforms to our intuition.

However, this calculation holds whether or not n is odd, but the resulting u is a stationary distribution only for odd n . Note that we did not use (13.3) in calculating u , and indeed (13.3) is guaranteed to hold only if n is odd.

Consider next a closely related n -state Markov process called the *random walk on the line with reflecting barriers*, in which from any state $2, \dots, n-1$ the system moves with equal probability to the previous or the next state, but from state 1 it moves to state 2 with probability 1, and from state n it moves to state $n-1$ with probability 1. Intuition in this case is a bit more complicated, because states 1 and n behave differently from the other states. The probability transition matrix for the problem is given by

$$P = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & \dots & 0 & 0 & 0 \\ & & & \vdots & & & \\ 0 & 0 & 0 & \dots & 1/2 & 0 & 1/2 \\ 0 & 0 & 0 & \dots & 0 & 1 & 0 \end{bmatrix}.$$

The recursion equations for this system are given by

$$\begin{aligned} u_1 &= u_2/2 \\ u_i &= u_{i-1}/2 + u_{i+1}/2 \quad i = 2, \dots, n-1 \\ u_n &= u_{n-1}/2 \end{aligned}$$

$$\sum_{i=1}^n u_i = 1.$$

It is easy to check directly that $u_i = 1/(n - 1)$ for $i = 2, \dots, n - 1$, and $u_1 = u_n = 1/2(n - 1)$.

This answer, however, is wrong! The problem is that this Markov process, while irreducible, is periodic. Indeed, if $p_{ij}^{(k)} > 0$, then $p_{ij}^{(k+1)} = 0$ and $p_{ij}^{(k+2)} > 0$. Thus, all states have period 2, so the Markov process is not ergodic. We can calculate the eigenvalues of P assuming a particular value for n . For instance, assuming $n = 7$, we get

$$1, -1, \frac{\sqrt{3}}{2}, \frac{\sqrt{3}}{2}, -\frac{1}{2}, \frac{1}{2}, 0.$$

The diagonal matrix D^* thus does not exist, as for large k , the diagonal of D^k alternates between D_1 with diagonal $1, 1, \epsilon_3, \dots, \epsilon_n$ and D_2 with diagonal $1, -1, \epsilon'_3, \dots, \epsilon'_n$, where the ϵ_i and ϵ'_i are very small and go to zero as $k \rightarrow \infty$. We can then calculate that the long-run behavior of the system is to alternate between $M^{-1}D_1M$ and $M^{-1}D_2M$. These two matrices, for $n = 7$ are given by

$$\begin{bmatrix} 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \end{bmatrix} \quad (13.6)$$

and

$$\begin{bmatrix} 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \\ 0 & 1/3 & 0 & 1/3 & 0 & 1/3 & 0 \\ 1/6 & 0 & 1/3 & 0 & 1/3 & 0 & 1/6 \end{bmatrix} \quad (13.7)$$

We will address later (§13.5) how to deal with nonergodic finite Markov processes.

In the present case, it is obvious that we can split the Markov process into two processes, one for the odd periods and the other for the even. We form P^2 , the two-period transition matrix, and drop the odd-numbered states, we get the transition matrix for starting in an even-numbered state.

$$\begin{bmatrix} \frac{3}{4} & \frac{3}{4} & 0 \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ 0 & \frac{1}{4} & \frac{3}{4} \end{bmatrix}$$

This Markov process is ergodic, and it is easy to check that its stationary distribution is $u = (1/3, 1/3, 1/3)$.

If we drop the even-numbered states in P^2 , we get the transition matrix for starting in an odd-numbered state:

$$\begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} & 0 \\ 0 & \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

This Markov process is also ergodic, and it is easy to check that its stationary distribution is $u = (1/6, 1/3, 1/3, 1/6)$.

Now we can reassemble the two es' stationary distributions, getting (13.7).

For another example, consider the Markov process with transition matrix

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

You can check that the eigenvalues of P are $\{\omega, \omega^2, 1, 0\}$ where ω is the complex cube root of unity, given by $\omega = -(-1)^{1/3}$. This Markov process is thus not ergodic, and indeed you can check that all states are periodic with period 3. Thus $P^* = \lim_{k \rightarrow \infty} P^k$ does not exist. Indeed, you can check that

$$P^2 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad P^3 = \begin{bmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{bmatrix},$$

and $P^4 = P$, so if you start in some state i , in the long run the Markov process will be the i^{th} row of P in states of the form $3t$, the i^{th} row of P^2 in states of the form $3t + 1$, and the i^{th} row of P^2 in states of the form $3t + 2$. If you look at P^3 , you see that state 1 communicates only with state 4, and if we drop states 2 and 3, we get an ergodic subprocess with stationary distribution $\{1/2, 1/2\}$. States 2 and 3 are absorbing in P^3 , so we can reconstruct the behavior of the Markov process as before.

13.4 Solving Linear Recursion Problems

In analyzing the stationary distribution of a Markov process, we commonly encounter an equation of the form

$$u_k = a_1 u_{k-1} + a_2 u_{k-2} + \dots + a_r u_{k-r}, \tag{13.8}$$

along with some boundary conditions, including $u_i \geq 0$ for all i and $\sum_i u_i = 1$. Note that this recursion equation is *linear* in the sense that if $u_k = g_i(k)$ for $i = 1, \dots, m$ are m solutions, then so are all the weighted sums of the form $u_k = \sum_{j=1}^m b_j g(j)$ for arbitrary weights b_1, \dots, b_m .

A general approach to solving such equations is presented by Elaydi (1999) in the general context of difference equations. We here present a short introduction to the subject, especially suited to analyzing Markov processes. First, form the associated k -degree *characteristic equation*

$$x^r = a_1 x^{r-1} + a_2 x^{r-2} + \dots + a_{r-1} x + a_r. \tag{13.9}$$

The general solution to (13.8) is the weighted sum, with arbitrary coefficients, of solutions of the following form. Let λ be a root of (13.9) of multiplicity m . Then $u_l = l^j \lambda^l$ are independent solutions for $j = 0, \dots, m-1$. Now, choose the weights of the various terms to satisfy the system's boundary conditions.

For instance, consider equations (13.6). We can write the recursion equation as

$$u_{i+1} = 2u_i - u_{i-1}.$$

The corresponding characteristic equation is $x^2 = 2x - 1$, which has a double root $x = 1$. Thus the solutions are of the form $u_k = a k^0 1^k + b k^1 1^k = a + bk$ for unknowns a and b , which we evaluate using the special conditions for u_1 and u_n . We have

$$u_1 = a + b = u_n/2 + u_2/2 = (a + bn)/2 + (a + 2b)/2,$$

which implies $b = 0$. Then

$$u_n = a + bn = u_1/2 + u_{n-1}/2 = (a + b)/2 + (a + b(n - 1))/2.$$

is satisfied for any a , but $\sum_{i=1}^n u_i = 1$ requires $a = 1/n$.

Sometimes the recursion equations have an *inhomogeneous* part, as in

$$u_i = u_{i-1}p_{i-1,i} + u_i p_{ii} + u_{i+1}p_{i+1,i} + g(i) \quad (13.10)$$

There is no general rule for finding the solution to the inhomogeneous part, but generally trying low-degree polynomials works.

For instance, consider the finite random walk, between points $-w$ and w , starting at k , with $0 < k < w$. We assume the end points are absorbing, so we may think of this as a gambler's wealth, where he is equally likely to win, lose, or draw in each period, until he is bankrupt or has reached wealth w . The recursion equations for the mean time to absorption into state $-w$ or w are then given by

$$\begin{aligned} m_{-w} &= 0 \\ m_w &= 0 \\ m_n &= m_n/3 + m_{n-1}/3 + m_{n+1}/3 + 1 \quad -w < n < w. \end{aligned}$$

We can rewrite the recursion equation as

$$m_{n+1} = 2m_n - m_{n-1} - 3.$$

The associated characteristic equation is $x^2 = 2x - 1$, with double root $x = 1$, so $m_n = a + nb$. To deal with the inhomogeneous part (-3) , we try adding a quadratic term, so $m_n = a + bn + cn^2$. We then have

$$a + b(n+1) + c(n^2 + 2n + 1) = 2(a + bn + cn^2) - (a + b(n-1) + c(n-1)^2) - 3$$

which simplifies to $c = -3/2$. To solve for a and b , we use the boundary conditions $m_{-w} = m_w = 0$, getting

$$m_n = \frac{3}{2}(w^2 - n^2).$$

We can use similar equations to calculate the probability p_n of being absorbed at $-w$ if one starts at n . In this case, we have

$$\begin{aligned} p_{-w} &= 1 \\ p_w &= 0 \\ p_n &= p_n/3 + p_{n-1}/3 + p_{n+1}/3 \quad -w < n < w. \end{aligned}$$

We now have $p_i = a + bi$ for constants a and b . Now, $p_{-w} = 1$ means $a - bw = 1$, and $p_w = 0$ means $a + bw = 0$, so

$$p_i = \frac{1}{2} \left(1 - \frac{i}{w} \right).$$

Note that the random walk is “fair” in the sense that the expecting payoff if you start with wealth i is equal to $w(1 - p_i) - wp_i = i$.

13.5 Analyzing Nonergodic Finite Markov Processes

We say a state i of a Markov process is *transient* if, whenever in state i , the probability of returning to i in the future is less than one. With probability one, a transient state will only be entered a finite number of times. To see this, suppose when state i occurs, the probability that it occurs again is $p < 1$. Then the probability that it reoccurs exactly k times is $p^k(1 - p)$, and

$$\sum_{k=0}^{\infty} p^k(1 - p) = (1 - p) \sum_{k=0}^{\infty} p^k = 1,$$

so the probability of an infinite number of returns is zero. In a finite Markov process, then, after a finite number of periods, with probability one, no transient states will reappear.

Consider, for instance, the Markov process depicted in Figure 13.1, where the arrows indicate strictly positive transition probabilities. Clearly the states in the outer circle are transient, while the inner five states, considered alone, form an ergodic Markov process. With probability one each transient state eventually is “absorbed” in the inner circle of states, from which it never escapes. Of course, we cannot say when that will happen, but we can calculate the expected time before a transition to the inner circle as a function of where we start in the outer circle.

A non-transient, or so-called *recurrent* state i , must then satisfy $q_i = 1$, and must recur an infinite number of times with probability one. It is then clear that all the states S that communicate with a recurrent state i form an irreducible Markov process. For if i is recurrent and $p_{ij} > 0$, then $p_{ji} > 0$, or else $q_i < 1$, which is by assumption impossible. Thus S is itself an irreducible Markov process. Although S may be periodic, we understand how to analyze the long-run properties of S : we analyze the aperiodic subprocesses of S using the Ergodic Theorem, and we put the




Figure 13.1. A Markov process with an outer circle of transient states and an inner circle of states that, considered alone, form an ergodic Markov process

subprocesses back together to get a set of long-run transition matrices for S . In general, an aperiodic finite Markov process will have a number other. In S_1, \dots, S_k of such ergodic subprocesses, none of which communicates with any other. In addition, for any transient state i , there will be a probability distribution $\{\sigma_k\}$ over S_1, \dots, S_k such that state i is eventually “absorbed” into the ergodic subprocess S_k with probability σ_k .

13.6 The Ergodic Theorem for Denumerable Markov Processes

If i is a *recurrent* state of a finite Markov process, the mean time of return to state i is finite, and given by (??). If the Markov process is denumerable, then a recurrent state can have an infinite mean recurrence time: when the process is in state i , it will return to state i in the future with probability one, but the expected waiting time can be infinite. We say recurrent state i is *positive recurrent* if the expected number of periods before returning to state i is finite. Every finite Markov process has at least one recurrent state because there are an infinite number of periods $t = 1, 2, \dots$ and the Markov process must be in some state in each period. Indeed, every non-transient state in a finite Markov process is recurrent.

Let M be an irreducible, recurrent, aperiodic denumerable Markov process with transition probabilities $P = \{p_{ij}\}$. Let μ_i be the expected number of states before the Markov process returns to state i , which we

know occurs with probability one because M is recurrent. In a finite Markov process we know μ_i is finite. But in a denumerable Markov process, there is no guarantee that $\mu_i < \infty$.

We call μ_i the *mean recurrence time* of state i . We say a denumerable Markov process is *non-null* if μ_i is finite for all states i . Finally, we say a denumerable Markov process is *ergodic* if it is irreducible, recurrent, non-null, and aperiodic. We then have the following Ergodic Theorem.

THEOREM 13.2 *For any denumerable ergodic Markov process with transition matrix P , the following equations hold with all u_j unique and strictly positive.*

$$u_j = \lim_{m \rightarrow \infty} p_{ij}^{(m)} \quad \text{for } i = 1, \dots, \infty \tag{13.11}$$

$$u_j = \sum_i u_i p_{ij} \tag{13.12}$$

$$1 = \sum_k u_k. \tag{13.13}$$

for $j = 1, \dots, \infty$.

13.7 A Denumerable Markov Process

For an example of a denumerable Markov process, suppose an animal is in state $d_k = k + 1$ if it has a $k + 1$ -day supply of food. The animal forages for food only when $k = 0$, and then he finds a $k + 1$ -day supply of food with probability f_k , for $k = 0, 1, \dots$. This means that the animal surely finds enough food to subsist for at least one day. This is a Markov process with $p_{0k} = f_k$ for all k , and $p_{k,k-1} = 1$ for $k \geq 1$, all other transition probabilities being zero. The recursion equations in this case are

$$u_i = u_0 f_i + u_{i+1}$$

for $i \geq 0$. If we let $r_k = f_k + f_{k+1} + \dots$ for $k \geq 0$ (so r_k is the probability of finding at least a $k + 1$ days' supply of food when foraging), it is easy to see that $u_k = r_k u_0$ satisfies the recursion equations; that is,

$$r_i u_0 = u_0 f_i + r_{i+1} u_0.$$

The requirement that $\sum_i u_i = 1$ becomes $u_0 = 1/\mu$, where $\mu = \sum_{k=0}^{\infty} r_k$. To see that μ is the expected value of the random variable d , note that

$$\begin{aligned} \mathbf{E}d &= 1f_0 + 2f_1 + 3f_2 + 4f_3 + 5f_4 + \dots \\ &= r_0 + f_1 + 2f_2 + 3f_3 + 4f_4 \dots \\ &= r_0 + r_1 + f_2 + 2f_3 + 3f_4 + \dots \\ &= r_0 + r_1 + r_2 + f_3 + 2f_4 + \dots \\ &= r_0 + r_1 + r_2 + r_3 + f_4 + \dots, \end{aligned}$$

and so on.²

We conclude that if this expected value does not exist, then no stationary distribution exists. Otherwise, the stationary distribution is given by

$$u_i = r_i/\mu \quad \text{for } i = 0, 1, \dots$$

Note that $\mu = 1/u_0$ is the expected number of periods between visits to state 0, because μ is the expected value of d . We can also show that $1/u_k = \mu/r_k$ is the expected number of periods μ_k between visits to state k , for any $k \geq 0$. Indeed, the fact that $u_k = 1/\mu_k$, where u_k is the probability of being in state k in the long run and μ_k is the expected number of periods between visits to state k , is a general feature of Markov Processes with stationary distributions. It is called the *renewal equation*. Because of the renewal theorem, recurrent states in finite Markov processes must be positive recurrent.

Let us prove that $\mu_k = \mu/r_k$ for $k = 2$ in the preceding model, leaving the general case to the reader. From state 2 the Markov process moves to state 0 in two periods, then requires some number j of periods before it moves to some state $k \geq 2$, and then in $k - 2$ transitions moves to state 2. Thus, if we let v be the expected value of j and we let w represent the expected value of k , we have $\mu_k = 2 + v + w - 2 = v + w$. Now v satisfies the recursion equation

$$v = f_0(1 + v) + f_1(2 + v) + r_2(1),$$

²More generally, noting that $r_k = \mathbf{P}[d \geq k]$, suppose x is a random variable on $[0, \infty)$ with density $f(x)$ and distribution $F(x)$. If x has finite expected value, then using integration by parts, we have $\int_0^{\infty} [1 - F(x)] dx = \int_0^{\infty} \int_x^{\infty} f(y) dy dx = \int_0^{\infty} x f(x) dx = \mathbf{E}[x]$.

because after a single move the system remains in state 0 with probability f_0 and the expected number of periods before hitting $k > 1$ is $1 + v$ (the first term), or it moves to state 1 with probability f_1 and the expected number of periods before hitting $k > 1$ is $2 + v$ (the second term), or hits $k > 1$ immediately with probability r_2 (the final term). Solving, we find that $v = (1 + f_1)/r_2$. To find w , note that the probability of being in state k conditional on $k \geq 2$ is f_k/r_2 . Thus $v + w = \mu/r_2$ follows from

$$\begin{aligned} w &= \frac{1}{r_2}(2f_2 + 3f_3 + \dots) \\ &= \frac{1}{r_2}(\mu - 1 - f_1). \end{aligned}$$

13.8 The Infinite Random Walk

The random walk on the line starts at zero and then, with equal probability in each succeeding period, does not move, or moves up or down one unit. It is intuitively clear that in the long run, when the system has “forgotten” its starting point, is equally likely to be in any state. Because there are an infinite number of states, the probability of being in any particular state in the long run is thus zero. Clearly this Markov process is irreducible and aperiodic. It can be shown to be recurrent, so by the ergodic theorem, it must be null-recurrent. This means that even though the Markov random walk returns to any state with probability one, its mean recurrence time is infinite.

Perhaps the fact that the recurrence time for the random walk is infinite explains why individuals tend to see statistical patterns in random data that are not really there. Figure 13.2 plots the random walk for 100 million periods. The result looks biased in favor of forward from about period 20 million to 50 million, backward 75 million, forward 90 million, and forward thereafter. Of course the maximum deviation from the mean (zero) is less than 2% of the total number of periods.

13.9 The Sisyphean Markov Process

As an exercise, consider the following *Sisyphean Markov process*, in which Albert has a piano on his back and must climb up an infinite number of steps $k = 1, 2, \dots$. At step k , with probability b_k , he stumbles and falls all




Figure 13.2. The random walk on the line

the way back to the first step, and with probability $1 - b_k$ he proceeds to the next step. This gives the probability transition matrix

$$P = \begin{bmatrix} b_1 & 1 - b_1 & 0 & 0 & 0 & \dots \\ b_2 & 0 & 1 - b_2 & 0 & 0 & \dots \\ b_3 & 0 & 0 & 1 - b_3 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

The recursion equations for this system are

$$\begin{aligned} u_1 &= \sum u_i b_i \\ u_{k+1} &= u_k(1 - b_k) \quad \text{for } k \geq 1, \end{aligned}$$

which are satisfied only if

$$u_1(b_1 + (1 - b_1)b_2 + (1 - b_1)(1 - b_2)b_3 + \dots) = u_1,$$

so either

$$b_1 + (1 - b_1)b_2 + (1 - b_1)(1 - b_2)b_3 + \dots = 1, \quad (13.14)$$

or $u_1 = \infty$ (note that $u_1 \neq 0$). If $b_i = \alpha$ for some $\alpha \in [0, 1]$ and all $i = 1, 2, \dots$, it is easy to see that (13.14) is true (let the left-hand side equal $x < \infty$, subtract b_1 from both sides, and divide by $1 - b_1$; now the left-hand side is just x again; solve for x).

Now, because $\sum_i u_i = 1$, u_1 , which must satisfy

$$u_1[1 + (1 - b_1) + (1 - b_1)(1 - b_2) + \dots] = 1.$$

This implies that the Markov process is ergodic if $b_k = \alpha$ for $\alpha \in (0, 1)$ and indeed $u_i = 1/\alpha$ for $i = 1, \dots$. The Markov process is not ergodic if $b_k = 1/(k + 1)$, however, because the mean time between passages to state 1 is infinite ($b_1 + b_2 + \dots = \infty$).

13.10 Andrei Andreyevich's Two-Urn Problem

After Andrei Andreyevich Markov discovered the processes that bear his name, he proved the ergodic theorem for finite processes. Then he looked around for an interesting problem to solve. Here is what he came up with—this problem had been solved before, but not rigorously.

Suppose there are two urns, one black and one white, each containing m balls. Of the $2m$ balls, r are red and the others are blue. At each time period $t = 1, \dots$ two balls are drawn randomly, one from each urn, and each ball is placed in the other urn. Let state i represent the event that there are $i \in [0, \dots, r]$ red balls in the black urn. What is the probability u_i of state i in the long run?

Let $P = \{p_{ij}\}$ be the $(r + 1) \times (r + 1)$ probability transition matrix. To move from i to $i - 1$, a red ball must be drawn from the black urn, and a blue ball must be drawn from the white urn. This means $p_{i,i-1} = i(m - r + i)/m^2$. To remain in state i , either both balls drawn are red or both are blue, $p_{i,i} = (i(r - i) + (m - i)(m - r + i))/m^2$. To move from i to $i + 1$, a blue ball must be drawn from the black urn, and a red ball must be drawn from the white urn. This means $p_{i,i+1} = (m - i)(r - i)/m^2$. All other transition probabilities are zero.

The recursion equations in this case are given by

$$u_i = u_{i-1}p_{i-1,i} + u_i p_{ii} + u_{i+1}p_{i+1,i} \quad (13.15)$$

for $i = 0, \dots, r + 1$, where we set $u_{-1} = u_{r+2} = 0$. I do not know how Andrei solved these equations, but I suspect he guessed at the answer and then showed that it works. At any rate, that is what I shall do. Our intuition concerning the ergodic theorem suggests that in the long run the probability distribution of red balls in the black urn are the same as if m balls were randomly picked from a pile of $2m$ balls (of which r are red) and put in the black urn. If we write the number of combinations of n things taken r at a time as $\binom{n}{r} = n!/r!(n-r)!$, then u should satisfy

$$u_i = \binom{m}{i} \binom{m}{r-i} / \binom{2m}{r}$$

The denominator in this expression is the number of ways the r red balls can be allocated to the $2m$ possible positions in the two urns, and the numerator is the number of ways this can be done when i red balls are in the black urn. You can check that u now satisfies the recursion equations.

13.11 Good Vibrations

Consider the pure coordination game in the diagram. We can check using the techniques of chapter 6 that there are two pure-strategy equilibria, ll and rr , as well as a mixed strategy equilibrium. If we represent the out-of-equilibrium dynamics of the game using

	l	r
l	5,5	0,0
r	0,0	3,3

a replicator process (see chapter 12), the pure strategy equilibria will be stable and the mixed strategy equilibrium unstable. But the concept of stability that is used, although at first glance compelling and intuitive, may be unrealistic in some cases. The idea is that if we start at the equilibrium ll , and we subject the system to a small disequilibrium shock, the system will move back into equilibrium. But in the real world, dynamical systems may be *constantly* subject to shocks, and if the shocks come frequently enough, the system will not have time to move back close to equilibrium before the next shock comes.

The evolutionary models considered in chapters 10 and 12 are certainly subject to continual random “shocks,” because agents are paired randomly, play mixed strategies with stochastic outcomes, and update their strategies by sampling the population. We avoided considering the stochastic nature of these processes by implicitly assuming that random variables can be replaced by their expected values, and mutations occur infrequently compared with the time to restore equilibrium. But these assumptions need not be appropriate.

We may move to stochastic differential equations, where we add a random error term to the right-hand side of an equation such as (11.1). This approach is very powerful, but uses sophisticated mathematical techniques, including stochastic processes and partial differential equations.³ Moreover, applications have been confined mainly to financial economics. Applying the approach to game theory is very difficult, because stochastic differential equations with more than one independent variable virtually never have a closed-form solution. Consider the following alternative approach, based on the work of H. Peyton Young (1998) and others. We start by modeling adaptive learning with and without errors.

³For relatively accessible expositions, see Dixit 1993 and Karlin and Taylor 1981.

13.12 Adaptive Learning

How does an agent decide what strategy to follow in a game? We have described three distinct methods so far in our study of game theory. The first is to determine the expected behavior of the other players and choose a best response (“rational expectations”). The second is to inherit a strategy (e.g., from one’s parents) and blindly play it. The third is to mimic another player by switching to the other player’s strategy, if it seems to be doing better than one’s own. But there is a fourth, and very commonly followed, *modus operandi*: follow the history of how other players have played against you in the past, and choose a strategy for the future that is a best response to the past play of others. We call this *adaptive learning*, or *adaptive expectations*.

To formalize this, consider an evolutionary game in which each player has limited memory, remembering only $h = \{h_1, h_2, \dots, h_m\}$, the last m moves of the players with whom he has been paired. If the player chooses the next move as a best response to h , we say the player follows adaptive learning.

Suppose, for instance, two agents play the coordination game in section 13.11, but the payoffs to ll and rr are both 5, 5. Let $m = 2$, so the players look at the last two actions chosen by their opponents. The best response to ll is thus l , the best response to rr is r , and the best response to rl or lr is any combination of l and r . We take this combination to be: play l with probability 1/2 and r with probability 1/2. There are 16 distinct “states” of the game, which we label $abcd$, where each of the letters can be l or r , b is the previous move by player 1, a is player 1’s move previous to this, d is the previous move by player 2, and c is player 2’s move previous to this. For instance, $llrl$ means player 1 moved l on the previous two rounds, whereas player 2 moved first r and then l .

We can reduce the number of states to 10 by recognizing that because we do not care about the order in which the players are counted, a state $abcd$ and a state $cdab$ are equivalent. Eliminating redundant states, and ordering the remaining states alphabetically, the states become $llll$, $lllr$, $llrl$, $llrr$, $lrlr$, $lrrl$, $lrrr$, $rlrl$, $rlrr$, and $rrrr$. Given any state, we can now compute the probability of a transition to any other state on the next play of the game. For instance, $llll$ (and similarly $rrrr$) is an *absorbing* state in the sense that, once it is entered, it stays there forever. The state $lllr$ goes to states $llrl$ and $lrrl$, each with probability 1/2. The state $llrl$ goes either to $llll$ where it stays forever, or to $lllr$, each with probability

1/2. The state $lr lr$ goes to $rlrl$ and $rrrr$ each with probability 1/4, and to $rlrr$ with probability 1/2. And so on.

We can summarize the transitions from state to state in a 10×10 matrix $M = (m_{ij})$, where $m_{abcd,efgi}$ is the probability of moving from state $abcd$ to state $efgi$. We call M a *probability transition matrix*, and the dynamic process of moving from state to state is a *Markov process* (§13.1). Because matrices are easier to describe and manipulate if their rows and columns are numbered, we will assign numbers to the various states, as follows: $llll = 1, ll lr = 2, \dots, rrrr = 10$. This gives us the following probability transition matrix:

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0 & 0.5 & 0 & 0 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0.5 & 0.25 \\ 0 & 0 & 0.25 & 0.25 & 0 & 0.25 & 0.25 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0.5 \\ 0.25 & 0.5 & 0 & 0 & 0.25 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.5 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Also, if we represent the 10 states by the 10 10-dimensional row vectors $\{v_1, \dots, v_{10}\}$, where $v_1 = (1, 0, \dots, 0)$, $v_2 = (0, 1, 0, \dots, 0)$, and so on, then it is easy to see that, if we are in state v_i in one period, the probability distribution of states in the next period is just $v_i M$, meaning the product of v_i , which is a 1×10 row vector, and M , which is a 10×10 matrix, so the product is another 1×10 row vector. It is also easy to see that the sum of the entries in $v_i M$ is unity and that each entry represents the probability that the corresponding state will be entered in the next period.

If the system starts in state i at $t = 0$, $v_i M$ is the probability distribution of the state it is in at $t = 1$. The probability distribution of the state the system at $t = 2$ can be written as

$$v_i M = p_1 v_1 + \dots + p_{10} v_{10}.$$

Then, with probability p_j the system has probability distribution $v_j M$ in the second period, so the probability distribution of states in the second period is

$$p_1 v_1 M + \dots + p_{10} v_{10} M = v_i M^2.$$

Similar reasoning shows that the probability distribution of states after k periods is simply $v_i M^k$. Thus, just as M is the probability transition matrix for one period, so is M^k the probability transition matrix for k periods. To find out the long-run behavior of the system, we therefore want to calculate

$$M^* = \lim_{k \rightarrow \infty} M^k.$$

I let Mathematica, the computer algebra software package, calculate M^k for larger and larger k until the entries in the matrix stopped changing or became vanishingly small, and I came up with the following matrix:

$$M^* = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 \\ 5/6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/6 \\ 1/2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/2 \\ 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2/3 \\ 1/2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/2 \\ 1/6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 5/6 \\ 2/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 \\ 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2/3 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

In other words, no matter where you start, you end up in one of the absorbing states, which is a Pareto-optimal Nash equilibrium. We call pure-strategy Nash equilibria in which all players choose the same strategy *conventions* (Young 1998). We conclude that *adaptive learning leads with probability 1 to a convention*.

13.13 Fictitious Play

Let \mathcal{G} be an n -player normal form game with pure strategy set S_i for each player type i , and with payoffs $\pi_i(s_1, \dots, s_n)$ for player type i when each player type j chooses $s_j \in S_j$. We assume the game is played repeatedly in time periods $t = 1, 2, \dots$ by players who are randomly assigned to groups of size n to play \mathcal{G} . Let $p_i^t(s_i)$ be the fraction of i -type players who play $s_i \in S_i$ in period t . We take $p^0 = (p_1^0, \dots, p_n^0)$ to be arbitrarily given, and in each period $t > 1$, each player i chooses $s_i^t \in S_i$ that is a best response to the population averages $p_i^{t-1}(s^{t-1})$. The resulting system is a Markov process known as *fictitious play* (?). The Markov processes analyzed in the

§13.12 is an example of fictitious play in a two-player game with a two-player population.

We say \mathcal{G} has the *fictitious play property* (Young 1998), if every sequence p^1, p^2, \dots generated by fictitious play is a Nash equilibrium of \mathcal{G} . Thus, when a game has the fictitious play property, there is a certain sense in which players' merely being rational (i.e., choosing best responses) leads in the long run to their playing Nash equilibria.

13.14 The Steady State of a Markov Process

There is a simpler way to compute M^* in the previous case. The computation also gives a better intuitive feel for the steady-state solution to the adaptive learning dynamical system generated by a pure coordination game. We know that whatever state we start the system in, we will end up in either state $llll$ or state $rrrr$. For state $abcd$, let $P[abcd]$ be the probability that we end up in $llll$ starting from $abcd$. Clearly, $P[llll] = 1$ and $P[rrrr] = 0$. Moreover, $P[lllr] = P[llrl]/2 + P[lrrl]/2$, because $lllr$ moves to either $llrl$ or to $lrrl$ with equal probability. Generalizing, you can check that, if we define

$$v = (P[llll], P[lllr], \dots, P[rrrr])'$$

the column vector of probabilities of being absorbed in state $llll$, then we have

$$Mv = v.$$

If we solve this equation for v , subject to $v[1] = 1$, we get

$$v = (1, 2/3, 5/6, 1/2, 1/3, 1/2, 1/6, 2/3, 1/3, 0)'$$

which then must be the first column of M^* . The rest of the columns are zero, except for the last, which must have entries so the rows each sum up to unity. By the way, I would not try to solve the equation $Mv = v$ by hand unless you're a masochist. I let Mathematica do it (v is a *left eigenvector* of M , so Mathematica has a special routine for finding v easily).

13.15 Adaptive Learning II

Now consider the pure coordination game illustrated in section 13.11, where the ll convention Pareto-dominates the rr convention. How does

adaptive learning work in such an environment? We again assume each player finds a best response to the history of the other player's previous two moves. The best response to ll and rr are still l and r , respectively, but now the best response to rl or lr is also l . Now, for instance, $lllr$ and $lrlr$ both go to $llll$ with probability 1. The probability transition matrix now becomes as shown.

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

To calculate

$$M^* = \lim_{k \rightarrow \infty} M^k$$

is relatively simple, because in this case $M^k = M^4$ for $k \geq 4$. Thus, we have

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

In other words, if you start in state $rrrr$, you stay there; otherwise, after four steps you arrive at $llll$ and remain there forever. We conclude that *with adaptive learning, if the system starts in a nonconventional state, it always ends up in the Pareto-efficient conventional state.*

13.16 Adaptive Learning with Errors

We now investigate the effect on a dynamic adaptive learning system when players are subject to error. Consider the pure coordination game illustrated in section 13.11, but where the payoffs to ll and rr are equal. Suppose each player finds a best response to the history of the other player’s previous two moves with probability $1-\epsilon$, but chooses incorrectly with probability $\epsilon > 0$. The probability transition matrix now becomes

$$M = \begin{bmatrix} a & 2b & 0 & 0 & e & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & c & d & 0 & c & d & 0 & 0 & 0 \\ c & 1/2 & 0 & 0 & d & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & b & e & 0 & a & b & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/4 & 1/2 & 1/4 \\ 0 & 0 & 1/4 & 1/4 & 0 & 1/4 & 1/4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & d & 1/2 & c \\ 1/4 & 1/2 & 0 & 0 & 1/4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & d & d & 0 & c & c & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & e & 2b & a \end{bmatrix},$$

where $a = (1-\epsilon)^2$, $b = \epsilon(1-\epsilon)$, $c = (1-\epsilon)/2$, $d = \epsilon/2$, and $e = \epsilon^2$. Note that now *there are no absorbing states*. To see what happens in the long run, suppose $\epsilon = 0.01$, so errors occur 1% of the time. Using Mathematica to calculate M^* , we find *all the rows are the same*, and each row has the entries

$$(0.442 \ 0.018 \ 0.018 \ 0.001 \ 0.0002 \ 0.035 \ 0.018 \ 0.0002 \ 0.018 \ 0.442)$$

In other words, you spend about 88.4% of the time in one of the conventional states, and about 11.6% of the time in the other states.

It should be intuitively obvious how the system behaves. If the system is in a conventional state, say $llll$, it remains there in the next period with probability $(1-\epsilon)^2 = 98\%$. If one player makes an error, the state moves to $lllr$. If there are no more errors for a while, we know it will return to $llll$ eventually. Thus, it requires multiple errors to “kick” the system to a new convention. For instance, $llll \rightarrow lllr \rightarrow lrrr \rightarrow rrrr$ can occur with just two errors: $llll \rightarrow lllr$ with one error, $lllr \rightarrow lrrr$ with one error, and $lrrr \rightarrow rrrr$ with no errors, but probability 1/2. We thus expect convention flips about every 200 plays of the game.

To test our “informed intuition,” I ran 1000 repetitions of this stochastic dynamical system using Mathematica. Figure 13.3 reports on the result.




Figure 13.3. An agent-based model adaptive learning with errors.

13.17 Stochastic Stability

We define a state in a stochastic dynamical system to be *stochastically stable* if the long-run probability of being in that state does not become zero or vanishingly small as the rate of error ϵ goes to zero. Clearly, in the previous example *llll* and *rrrr* are both stochastically stable and no other state is. Consider the game in section 13.11. It would be nice if the Pareto-dominant equilibrium *ll* were stochastically stable, and no other state were stochastically stable. We shall see that is the case. Now the probability transition matrix becomes

$$M = \begin{bmatrix} a & 2b & 0 & 0 & e & 0 & 0 & 0 & 0 & 0 \\ 0 & 2b & a & 0 & e & 0 & 0 & 0 & 0 & 0 \\ a & 2b & 0 & 0 & e & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & b & e & 0 & a & b & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & b & a & b \\ 0 & 0 & a & b & 0 & b & e & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & b & a & b \\ a & 2b & 0 & 0 & e & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & b & e & 0 & a & b & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & e & 2b & a \end{bmatrix},$$

where $a = (1-\epsilon)^2$, $b = \epsilon(1-\epsilon)$, and $e = \epsilon^2$. Again there are no absorbing states. If $\epsilon = 0.01$, we calculate M^* , again we find *all the rows are the same*, and each row has the entries

$$[0.9605 \quad 0.0198 \quad 0.0198 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0].$$

In other words, the system spends 96% of the time in the Pareto-dominant conventional states and virtually all of the remaining time in “nearby states.” It is clear (though it should be formally proved) that ll is the only stochastically stable state.