13

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 chains (§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 Chains

A *finite Markov chain* is a dynamical system that in each time period $t = 0, 1, \ldots$ 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 chain, and each p_{ij} is called a *transition probability*. A *denumerable Markov chain* has an infinite number of states $t = 1, 2, \ldots$, and is otherwise the same. If we do not care whether the finite or denumerable case obtains, we speak simply of a *Markov chain*.

Many games can be viewed as Markov chains. 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 chain 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 (§13.8).

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 chain 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 chain, where the transition probabilities are calculated from the birth and death rates of the beetles.

We are interested in the long-run behavior of Markov chains. In particular, we are interested in the behavior of systems that we expect will attain a long-run equilibrium of some type independent from its initial conditions. If such an equilibrium exists, we say the Markov chain is *ergodic*. In an ergodic system, history does not matter: every initial condition leads to the same long-run behavior. Nonergodic systems are history dependent. It is intuitively reasonable that the repeated prisoner's dilemma and the trading model described previously are ergotic. The gambler model is not ergodic, because the system could end up with either player bankrupt.[1] What is your intuition concerning the beetle population, if there is a positive probability that a female has no offspring in a breeding season?

It turns out that there is a very simple and powerful theorem that tells us exactly when a Markov chain is ergodic and provides a simple characterization of the long-run behavior of the system. To develop the machinery needed to express and understand this theorem, we will define a few terms. Let $p_{ij}^{(m)}$ be the probability of being in state j in m periods if the chain 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

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

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 $k \to \infty$. Let us call this limit (supposing it exists) $P^* = \{p_{ij}^*\}$. Now P^* must have two properties. First, because the long-run behavior of the system cannot depend on where it started, $p_{ij}^* = p_{i'j}^*$ for any two states i and i'. This means that all the rows of P^* must be the same. Let us denote the (common value of the) rows by $u = \{u_1, \ldots, u_n\}$, so u_j is the probability that the Markov chain will be in state j in the long run. The second fact is that

$$PP^* = P \lim_{k\to\infty} P^k = \lim_{k\to\infty} P^{k+1} = P^*.$$

This means u must satisfy

$$u_j = \lim_{m\to\infty} p_{ij}^{(m)} \quad \text{for } i = 1, \ldots, n \tag{13.1}$$

[1] Specifically, you can show that the probability that player 1 wins is $k_1/(k_1 + k_2)$, and if player 1 has wealth w at some point in the game, the probability he will win is $w/(k_1 + k_2)$.

$$u_j = \sum_i u_i p_{ij} \tag{13.2}$$

$$\sum_k u_k = 1, \tag{13.3}$$

for $j = 1, \ldots, n$. Note that (13.2) can be written in matrix notation as $u = uP$, so u is a *left eigenvector* of P. The first equation says that u_j is the limit probability of being in state j starting from any state, the second 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, and the final equation says u is a probability distribution over the states of the Markov chain. The *recursion equations* (13.2) and (13.3) are often sufficient to determine u, which we call the *invariant distribution* or *stationary distribution* of the Markov chain.

In the case where the Markov chain is finite, the preceding description of the stationary distribution is a result of the *Frobenius-Perron* theorem (Horn and Johnson 1985), which says that P has a maximum eigenvalue of unity, and the associated left eigenvector, which is the stationary distribution $(u_1, \ldots, u_n)$ for P, exists and has nonnegative entries. Moreover, if P^k is strictly positive for some k (in which case we say P is irreducible), then the stationary distribution has strictly positive entries.

In case a Markov chain is not ergodic, it is informative to know the whole matrix $P^* = (p^*_{ij})$, because p_{ij} tell you the probability of being absorbed by state j, starting from state i. The Frobenius-Perron theorem is useful here also, because it tells us that all the eigenvalues of P are either unity or strictly less than unity in absolute value. Thus, if $D = (d_{ij})$ is the $n \times n$ diagonal matrix with the eigenvalues of P along the diagonal, then $D^* = \lim_{k\to\infty} D^k$ is the diagonal matrix with zeros everywhere except unity where $d_{ii} = 1$. But, if M is the matrix of left eigenvectors of P, then $MPM^{-1} = D$, which follows from the definitions, implies $P^* = M^{-1}D^*M$. This equation allows us to calculate P^* rather easily.

A few examples are useful to get a feel for the recursion equations. Consider first the n-state Markov chain called the *random walk on a circle*, in which there are n states, and from any state $t = 2, \ldots, n-1$ 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 the problem is given by

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

The recursion equations for this system are given by

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

Clearly, this set of equations has solution $u_i = 1/n$ for $i = 1, \ldots, n$. Prove that this solution is unique by showing that if some u_i is the largest of the $\{u_k\}$, then its neighbors are equally large.

Consider next a closely related n-state Markov chain called the *random walk on the line with reflecting barriers*, in which from any state $2, \ldots, n-1$ 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

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

The recursion equations for this system are given by

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

It is easy to check directly that $u_i = 1/(n-1)$ for $i = 2, \ldots, n-1$, and $u_1 = u_n = 1/2(n-1)$. In fact, there is a general method for solving difference equations of this type, as described in section 13.6.

We can use the same methods to find other characteristics of a Markov chain. Consider, for instance, 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

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

We can rewrite the recursion equation as

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

We can solve this, using the techniques of section 13.6. 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 $c = 2/3$. 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

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

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 $w(1 - p_i) - wp_i = i$.

For an example of a denumerable Markov chain, 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 $k = 0, 1, \ldots$. This means that the animal surely finds enough food to subsist for at least one day. This is a Markov chain 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

$$u_i = u_\text{o} f_i + u_{i+1}$$

for $i \geq 0$. If we let $r_k = f_k + f_{k+1} + \ldots$ 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 $u_k = r_k u_\text{o}$ satisfies the recursion equations; that is,

$$r_i u_\text{o} = u_\text{o} f_i + r_{i+1} u_\text{o}.$$

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

$$\begin{aligned} \mathbf{E}d &= 1f_\text{o} + 2f_1 + 3f_2 + 4f_3 + 5f_4 + \ldots \\ &= r_\text{o} + f_1 + 2f_2 + 3f_3 + 4f_4 \ldots \\ &= r_\text{o} + r_1 + f_2 + 2f_3 + 3f_4 + \ldots \\ &= r_\text{o} + r_1 + r_2 + f_3 + 2f_4 + \ldots \\ &= r_\text{o} + r_1 + r_2 + r_3 + f_4 + \ldots, \end{aligned}$$

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

$$u_i = r_i/\mu \qquad \text{for } i = 0, 1, \ldots .$$

Note that $\mu = 1/u_o$ 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 chains with stationary distributions. It is called the *renewal equation*.

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 chain 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

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

because after a single move the system remains in state 0 with probability f_o 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

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

[2]More generally, noting that $r_k = \mathrm{P}[d \geq k]$, 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 $\int_0^\infty [1-F(x)]dx = \int_0^\infty \int_x^\infty f(y)dydx = xf(x)|_0^\infty + \int_0^\infty xf(x)dx = \mathbf{E}[x]$.

13.2 The Ergodic Theorem for Markov Chains

When are equations (13.1)-(13.3) true, and what exactly do they say? To answer this, we will need a few more concepts. Throughout, we let M be a finite or denumerable Markov chain with transition probabilities $\{p_{ij}\}$. We say a state j can be *reached* from a state i if $p_{ij}^{(m)} > 0$ for some positive integer m. We say a pair of states i and j *communicates* if each is reached from the other. We say a Markov chain is *irreducible* if every pair of states communicates.

If M is irreducible, and if a stationary distribution u exists, then all the u_i in (13.1) are *strictly positive*. To see this, suppose some $u_j = 0$. Then by (13.2), if $p_{ij} > 0$, then $p_i = 0$. Thus, any state that reaches j in one period must also have weight zero in u. But a state i' that reaches j in two periods must pass through a state i that reaches j in one period, and because $u_i = 0$, we also must have $u_{i'} = 0$. Extending this argument, we say that any state i that reaches j must have $u_i = 0$, and because M is irreducible, all the $u_i = 0$, which violates (13.3).

Let q_i be the probability that, starting from state i, the system returns to state i in some future period. If $q_i < 1$, then it is clear that with probability one, state i can only occur a finite number of times. Thus, in the long run we must have $u_i = 0$, which is impossible for a stationary distribution. Thus in order for a stationary distribution to exist, we must have $q_i = 1$. We say a state i is *persistent* or *recurrent* if $q_i = 1$. Otherwise, we say state i is *transient*. If all the states of M are recurrent, we say that M is recurrent.

Let μ_i the be the expected number of states before the Markov chain returns to state i. Clearly, if i is transient, then $\mu_i = \infty$, but even if i is persistent, there is no guarantee that $\mu_i < \infty$. We call μ_i the *mean recurrence time* of state i. If the mean recurrence time of state i is μ_i, M is in state i on average one period out of every μ_i, so we should have $u_i = 1/\mu_i$. In fact, this can be shown to be true whenever the Markov chain has a stationary distribution. This is called the *renewal theorem* for Markov chains. We treat the renewal theorem as part of the ergodic theorem. Thus, if M is irreducible, it can have a stationary distribution only if μ_i is finite, so $u_i = 1/\mu_i > 0$. We say a recurrent state i in a Markov chain is *null* if $\mu_i = \infty$, and otherwise we call the state *non-null*. An irreducible Markov chain cannot have a stationary distribution unless all its recurrent states are non-null.

We say state i in a Markov chain is *periodic* if there is some integer $k > 1$ such that $p_{ii}^{(k)} > 0$ and $p_{ii}^{(m)} > 0$ implies m is a multiple of k. Otherwise, we say M is *aperiodic*. It is clear that if M has a non-null, recurrent, periodic state i, then M cannot have a stationary distribution, because we must have $u_i = \lim_{k\to\infty} p_{ii}^{(k)} > 0$, which is impossible unless $p_{ii}^{(k)}$ is bounded away from zero for sufficiently large k.

An irreducible, non-null recurrent, aperiodic Markov chain is called *ergodic*. We have shown that if an irreducible Markov chain is not ergodic, it cannot have a stationary distribution. Conversely, we have the following *ergodic theorem* for Markov chains, the proof of which can be found in Feller (1950).

THEOREM 13.1 *An ergodic Markov chain M has a unique stationary distribution, and the recursion equations (13.1)-(13.3) hold with all $u_i > 0$. Moreover $u_j = 1/\mu_j$ for each state j, where μ_j is the mean recurrence time for state j.*

We say a subset M' of states of M is *isolated* if no state in M' reaches a state not in M'. Clearly an isolated set of states is a Markov chain. We say M' is an *irreducible set* if M' is isolated and all pairs of states in M' communicate. Clearly, an irreducible set is an irreducible Markov chain. Suppose a Markov chain M consists of an irreducible set M' plus a set A of states, each of which reaches M'. Then, if u' is a stationary distribution of M', there is a stationary distribution u for M such that $u_i = u'_i$ for $i \in M'$ and $u_i = 0$ for $i \in A$. We can summarize this by saying that a Markov chain that consists of an irreducible set of states plus a set of transient states has a unique stationary distribution in which the frequency of the transient states is zero and the frequency of recurrent states is strictly positive. We call such a Markov chain *nearly irreducible*, with transient states A and an absorbing set of states M'.

More generally, the states of a Markov chain M can be uniquely partitioned into subsets $A, M_1, M_2 \ldots$ such that for each i, M_i is nearly irreducible and each state in A reaches M_i for some i. The states in A are thus transient, and if each M_i is non-null and aperiodic, it has a unique stationary distribution. However, M does not have a stationary distribution unless it is nearly irreducible.