2

Finite Markov Chains

Author: Eric Rieders, Department of the Mathematical Sciences, DePaul University.

Prerequisites: The prerequisites for this chapter are finite probability, matrix algebra, mathematical induction, and sequences. See Sections 2.4, 2.7, 3.1, and 6.1 of *Discrete Mathematics and Its Applications*.

Introduction

Probability theory is the mathematical basis of the study of phenomena that involve an element of chance. There are many situations in which it is not possible to predict exactly what is going to happen. A gambler at the roulette table cannot say with certainty how much money he or she will win (if any!) after one round of play. The exact number of users logged on to a time-share computer at 2 p.m. tomorrow can only be guessed. The eye or hair color of a child is a subject of speculation prior to its birth. Yet, in all three of these cases, we feel that it is reasonable to give the probability of observing a particular outcome.

These probabilities may be based on past experience, such as the number of users "typically" logged on around 2 p.m. They might also derive from a *mathematical model* of the phenomenon of interest. Combinatorial arguments lead one to the computation of probabilities associated with betting at roulette, and models of heredity may be used to predict the likelihood of a child inheriting certain genetic characteristics (such as eye color) from its parents.

The preceding examples may be extended to *sequences* of consecutive observations, where the sequence of outcomes cannot be predicted in advance. The rising and falling fortunes of the gambler; the number of users requiring cpu time at, say, one minute intervals; or the eye color of succeeding generations of children are all examples of random processes. A **random process** is any phenomenon which evolves in time and whose evolution is governed by some random (i.e. chance) mechanism. Questions relating to the likelihood that the gambler reaches a preset goal (before going broke), the response time of the computer, or the chance that one of your grandchildren will have green eyes all may be addressed by analyzing the underlying random process.

In this chapter, we will examine an important kind of random process known as a *Markov chain*. All of the foregoing examples of random processes may in fact be modeled as Markov chains. The basic property of a Markov chain is that in making the best possible predictions about the future behavior of the random process, given the information yielded by a sequence of observations up to the present, only the most recently observed outcome need be considered. Prior observations yield no additional information useful for the purposes of prediction.

Markov chains were first studied systematically by the Russian mathematician Andrei Markov*. In the course of his investigations in probability theory, Markov wished to extend the investigation of the properties of sequences of *independent* experiments (i.e. those for which the outcome of one experiment does not influence, nor is influenced by, any of the other experiments) to an investigation of sequences of experiments for which the present outcome does in fact affect future outcomes. That is, the outcomes of the experiments are "chained" together by the influence each outcome exerts on the probability of observing a particular outcome as the result of the next experiment.

The widespread applicability of the Markov chain model to such diverse fields as population genetics, decision sciences, physics and other fields has gone hand in hand with the large amount of mathematical research concerning the properties of this kind of random process. Markov chain models are simple enough to be analyzed, yet realistic enough to be of genuine use in understanding real random processes. We will touch only on the more basic definitions and properties of Markov chains. For further study, see the references at the end of this chapter.

---

* Andrei Andreevich Markov, born in 1856 in Ryazan, Russia, showed an early mathematical talent. He studied at St. Petersburg University, where he was heavily influenced by the father of the Russian school of probability, P. Chebyshev. Markov remained as a professor at St. Petersburg, where he distinguished himself in a number of mathematical disciplines. Markov died in 1922, his health having suffered from a winter spent teaching high school mathematics in the interior of Russia.

The State Space

We now develop a precise definition of a Markov chain. Suppose that a sequence of "experiments" is to be performed and that each experiment will result in the observation of exactly one outcome from a finite set S of possible outcomes. The set S is called the **state space** associated with the experiments, and the elements of S are called **states**. Typically these experiments are carried out to investigate how some phenomenon changes as time passes by classifying the possible types of behavior into states.

We denote the sequence of observed outcomes by $X_1, X_2, \ldots$ (which presumably we do not know in advance). We denote by X_k the outcome of experiment k. Note that $X_k \in S$. In our description of a random process, we will usually include a state X_0, which is the state in which we initially find the phenomenon. We can think of the sequence $X_0, X_1, X_2, \ldots$ as being the observations of some phenomenon (such as the gambler's net winnings) as it evolves from state to state as time goes by. For our purposes, the random process is simply the sequence of observed states $X_0, X_1, X_2, \ldots$.

Example 1 Set up a Markov chain to model the following gambling situation. A gambler starts with \$2. A coin is flipped; if it comes up heads, the gambler wins \$1, and if it comes up tails the gambler loses \$1. The gambler will play until having gone broke or having reached a goal of \$4. After each play, the observation to be made is how much money the gambler has.

Solution: The possible amounts are $0, 1, 2, 3$ and 4 dollars. Thus, the state space is $S = \{0, 1, 2, 3, 4\}$; the five elements of S describe the status of the gambler's fortune as the game progresses.

Initially (before the first "experiment" begins), the gambler has \$2, and we write $X_0 = 2$ to describe this "initial state". X_1 is the amount of money the gambler has after the first coin toss. Note that the only information required to give predictions about the quantity X_2 (the amount of money the gambler has after the second toss) is in fact the value of X_1. If the coin is fair, we easily see that the probability of winning \$1 on the second toss is 0.5, that is,

$$p(X_2 = X_1 + 1) = 0.5.$$

Similarly,

$$p(X_2 = X_1 - 1) = 0.5.$$

If the gambler wins the first round and then loses the next 3 rounds, we will observe $X_0 = 2$, $X_1 = 3$, $X_2 = 2$, $X_3 = 1$, $X_4 = 0$. For $k \geq 5$, we will have $X_k = 0$, since once the gambler is broke, nothing more happens. In general, X_k is the amount that the gambler has after k coin tosses, where the state space S was described above. We can make probabilistic predictions about the value of X_{k+1} as soon as the value of X_k is known, *without* knowing anything about $X_0, X_1, \ldots X_{k-1}$. In fact, we can state explicitly

$$p(X_k = j | X_{k-1} = i) = \begin{cases} 1 & \text{if } i = j = 0 \text{ or } i = j = 4 \\ \frac{1}{2} & \text{if } 0 \le j = i - 1 \le 3 \text{ or } 1 \le j = i + 1 \le 4 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $p(X_{k+1} = j | X_k = i)$ is a conditional probability (see Section 6.2 of *Discrete Mathematics and Its Applications*), namely, the probability that the gambler will have $\$j$ after the $(k+1)$st coin toss, given that the gambler had $\$i$ after k tosses. □

The probabilities expressed in (1) are called **transition probabilities**. A transition probability is associated with each pair of states in the state space; such a quantity represents the probability of the phenomenon moving from the first state of the pair to the second state.

Example 2 Suppose that the gambler of Example 1 is given a choice of five envelopes containing $0, 1, \ldots, 4$ dollars, respectively. Find the probability distribution of X_0.

Solution: Assuming the gambler is equally likely to choose any of the proffered envelopes, we have that $p(X_0 = i) = 0.2,\ 0 \le i \le 4$. □

More generally, we might have

$$p(X_0 = i - 1) = q_i \quad 1 \le i \le 5$$

where q_i is the probability that the gambler has $i-1$ dollars to start with, and $\sum_{i=1}^{5} q_i = 1$. The 1×5 matrix

$$\mathbf{Q} = (q_1\ q_2\ q_3\ q_4\ q_5)$$

is called the **initial probability distribution**. The entries of $\mathbf{Q}$ are the probabilities of the various possible initial states of the gambler's wallet. For ease of notation, it is often useful to list the elements of the state space sequentially. In this gambling example, we would write $S = \{s_1, s_2, s_3, s_4, s_5\}$, where s_1 is the state where the gambler is broke (i.e., has 0 dollars), s_2 is the state where the gambler has 1 dollar, and so on. Thus, with probability q_i the gambler is initially in state s_i, for $1 \le i \le 5$.

Example 3 Suppose that we have a package of 1000 seeds of a plant that has the following life cycle. The plant lives exactly one year, bearing a single flower that is either red, yellow or orange. Before dying, the plant produces a single seed which, when it blooms, produces a flower the color of which depends (randomly) on the color of the flower of its "parent". Use Table 1 to model the color of the flower during succeeding years as a Markov chain.

			Offspring		
			Red	Yellow	Orange
			s_1	s_2	s_3
	Red	s_1	0.3	0.5	0.2
Parent	Yellow	s_2	0.4	0.4	0.2
	Orange	s_3	0.5	0.5	0

Table 1. Transition Probabilities

Solution: Here, the state space S is the set of colors red (denoted s_1), yellow (s_2) and orange (s_3), and the random process $X_0, X_1, X_2 \ldots$ is the yearly succession of flower colors produced by succeeding generations of plants beginning with the one that grows from the first seed you have chosen. The rule which describes how the color of the flower of the offspring depends on that of its parent is summarized in Table 1.

This table gives the probability that the offspring will bear a flower of a particular color (corresponding to a column of the table), given the color of its parent's flower (corresponding to a row of the table). For example, the probability that a plant having a red flower produces a plant with a yellow flower is 0.5. This is the probability of the random process moving from state s_1 to state s_2; we denote this as p_{12}. (Note that we are assuming the probability of observing a particular flower color in any given year depends only on the color of the flower observed during the *preceding* year.) The quantity p_{12} is the probability of a transition from state s_1 to state s_2. □

We can use basic facts about conditional probability together with the Markov chain model to make predictions about the behavior of the Markov chain.

Example 4 In the preceding example, suppose you choose a seed at random from the package which has 300 seeds that produce red flowers, 400 producing yellow flowers and the remaining 300 produce orange flowers. Find the probability that the colors observed in the first two years are red and then orange.

Solution: Denote by X_0 the color of the flower produced (the "initial state"); the initial probability distribution is

$$\mathbf{Q} = (0.3 \;\; 0.4 \;\; 0.3).$$

To make predictions about the sequence of flower colors observed we need the information in Table 1 as well as knowledge of the initial probability distribution. With probability 0.3, we have $X_0 = s_1$, i.e., the first flower is red. Given that $X_0 = s_1$, we have $X_1 = s_3$ (the next plant has an orange flower) with probability 0.2. That is, $p(X_1 = s_3 | X_0 = s_1) = 0.2$; this comes from Table 1. Recall that

$$p(X_1 = s_3 | X_0 = s_1) = \frac{p(X_0 = s_1, X_1 = s_3)}{p(X_0 = s_1)}$$

where $p(X_0 = s_1, X_1 = s_3)$ is the probability that the sequence of colors red, orange is observed in the first two years. It is the probability of the *intersection* of the two events $X_0 = s_1$ and $X_1 = s_3$. Rewriting this expression, we then have

$$\begin{aligned} p(X_0 = s_1, X_1 = s_3) &= p(X_0 = s_1) \cdot p(X_1 = s_3 | X_0 = s_1) \\ &= 0.3 \cdot 0.2 \\ &= 0.06 \end{aligned}$$

□

Can we deduce the probability of observing a particular sequence of flower colors during, say, the first 4 years using an argument similar to the one just carried out? We in fact can, using Theorem 1 (later in this chapter), and the fact that the random process in Example 4 is a Markov chain.

Now suppose that $X_0, X_1, X_2, \ldots$ are observations made of a random process (including the initial state) with state space $S = \{s_1, s_2, \ldots, s_N\}$. We denote the probability of observing a particular sequence of states $s_{i_0}, s_{i_1}, \ldots, s_{i_k}$, beginning with the initial one, by

$$p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}).$$

The (conditional) probability of observing the state $s_{i_{k+1}}$ as the $(k+1)$st outcome, given that the first k outcomes observed are $s_{i_0}, s_{i_1}, \ldots, s_{i_k}$, is given by

$$p(X_{k+1} = s_{i_{k+1}} | X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}) = \frac{p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}, X_{k+1} = s_{i_{k+1}})}{p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k})}.$$

Example 5 With reference to Example 1, find an expression which gives the probability that the gambler has \$1 after 2 rounds, given that the gambler started with \$1 and then won the first round.

Solution: We have $s_{i_0} = 1$, $s_{i_1} = 2$, and $s_{i_2} = 1$. The desired probability is then given by

$$p(X_2 = s_{i_2} | X_0 = s_{i_0}, X_1 = s_{i_1}).$$

□

Markov Chains

The foregoing examples provide motivation for the formal definition of the Markov chain.

Definition 1 A random process with state space $S = \{s_1, s_2, \ldots s_N\}$ and observed outcomes $X_0, X_1, X_2, \ldots$ is called a *Markov chain* with initial probability distribution $\mathbf{Q} = (q_1\ q_2 \ldots\ q_N)$ if

(i) $p(X_{k+1} = s_{i_{k+1}} | X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k})$

$$= p(X_{k+1} = s_{i_{k+1}} | X_k = s_{i_k}) \qquad \text{for } k = 1, 2, \ldots$$

(ii) $p(X_{k+1} = s_j | X_k = s_i) = p_{ij} \qquad$ for $k = 0, 1, 2, \ldots$

(iii) $p(X_0 = i) = q_i \qquad$ for $i = 1, 2, \ldots, N$. □

The numbers p_{ij} are the transition probabilities of the Markov chain. The **Markov property**, (i), says that we need only use the most recent information, namely, $X_k = s_{i_k}$, to determine the probability of observing the state $s_{i_{k+1}}$ as the outcome of the experiment number $k+1$, given that the sequence of states $s_{i_1}, \ldots, s_{i_k}$ were the outcomes of the first k experiments (and the initial state is s_{i_0}).

Property (ii) requires that the underlying random mechanism governing the chance behavior of the random process does not change; the probability of moving from one particular state to another is always the same regardless of when this happens. (Sometimes Markov chains are defined to be random processes satisfying condition (i). Those that also satisfy (ii) are said to have **stationary transition probabilities**.)

It might happen that the initial state of the random process is itself determined by chance, such as in Examples 2 and 3. To compute probabilities associated with a Markov chain whose initial state is unknown, one needs to know the probability of observing a particular state as the initial state. These are the probabilities given in (iii). Note that if the initial state is known, as in Example 1, we still can express this in terms of an initial distribution. For example, if the Markov chain is known to be initially in state s_k, then we have $q_k = 1$ and $q_i = 0$ for $i \neq k$. This says that with probability 1, the Markov chain is initially in state s_k.

It is very useful to arrange the transition probabilities p_{ij} in an $N \times N$ matrix $\mathbf{T}$ (N is the number of elements in the state space), so that the (i, j)th entry will be p_{ij}.

Example 6 Find the matrix of transition probabilities for the Markov chain of Example 3.

Solution: Using Table 1, this matrix is

$$\mathbf{T} = \begin{bmatrix} 0.3 & 0.5 & 0.2 \\ 0.4 & 0.4 & 0.2 \\ 0.5 & 0.5 & 0 \end{bmatrix}.$$

For instance, the entry in the first row and second column, 0.5, is p_{12}, the probability that a red flower will produce an offspring that bears a yellow flower. Since it is impossible for an orange flower to produce an offspring that bears an orange flower, we have $p_{33} = 0$. □

We will show that knowing the transition probabilities p_{ij} and the initial probability distribution $\mathbf{Q}$ of a Markov chain suffices for determining all probabilities of interest in connection with the Markov chain. Indeed, all such probabilities can be computed if we know the probability of observing any specific sequence of outcomes; any other events of interest are made up of such sequences.

Example 7 In Example 2, find the probability of starting out with 3 dollars, and losing the first two rounds, i.e., having 1 dollar left after flipping the coin twice.

Solution: We must compute $p(X_0 = 3, X_1 = 2, X_2 = 1)$. Using the definition of conditional probability, we have

$$\begin{aligned} &p(X_0 = 3, X_1 = 2, X_2 = 1) \\ &\quad = p(X_0 = 3, X_1 = 2)p(X_2 = 1|X_0 = 3, X_1 = 2) \\ &\quad = p(X_0 = 3)p(X_1 = 2|X_0 = 3)p(X_2 = 1|X_0 = 3, X_1 = 2) \\ &\quad = p(X_0 = 3)p(X_1 = 2|X_0 = 3)p(X_2 = 1|X_1 = 2) \quad \text{(Property (i) of Markov chains)} \\ &\quad = q_4 p_{43} p_{32} \\ &\quad = 0.05. \end{aligned}$$

Note that, for example, p_{43} is the probability of going from state s_4, where the gambler has 3 dollars, to state s_3, where the gambler has 2 dollars. □

More generally, we have the following basic result.

Theorem 1 If $X_0, X_1, \ldots, X_k$ denote the first k observed outcomes of the Markov chain with initial probability distribution $\mathbf{Q} = (q_1\ q_2\ \ldots\ q_N)$ then

$$\begin{aligned} &p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}) \\ &\quad = p(X_0 = s_{i_0})p(X_1 = s_{i_1}|X_0 = s_{i_0}) \times \cdots \times p(X_k = s_{i_k}|X_{k-1} = s_{i_{k-1}}) \\ &\quad = q_{i_0}p_{i_0 i_1}p_{i_1 i_2}\ldots p_{i_{k-1}i_k}. \end{aligned}$$

Proof: We will demonstrate this using mathematical induction.

Basis step: The case $k = 1$ is an immediate consequence of the definition of conditional probability.

Induction step: Assuming the truth of the result for some particular k, we must then deduce its truth for $k + 1$. We have

$$\begin{aligned} &p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}, X_{k+1} = s_{i_{k+1}}) \\ &\quad = p(X_{k+1} = s_{i_{k+1}}|X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}) \\ &\qquad\qquad \times p(X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}) \\ &\quad = p(X_{k+1} = s_{i_{k+1}}|X_0 = s_{i_0}, X_1 = s_{i_1}, \ldots, X_k = s_{i_k}) \\ &\qquad\quad p(X_0 = i_0)p(X_1 = s_{i_1}|X_0 = s_{i_0}) \times \cdots \times p(X_k = s_{i_k}|X_{k-1} = s_{i_{k-1}}) \\ &\quad = p(X_{k+1} = s_{i_{k+1}}|X_k = s_{i_k})p(X_0 = i_0)p(X_1 = s_{i_1}|X_0 = s_{i_0}) \times \cdots \\ &\qquad\qquad \times p(X_k = s_{i_k}|X_{k-1} = s_{i_{k-1}}), \end{aligned}$$

which is the result. (The definition of conditional probability was used at the first step, the induction assumption at the second step, and property (i) at the third step.) ■

Example 8 Compute the probability of the sequence of outcomes described in Example 1 ($X_0 = 2$, $X_1 = 3$, $X_2 = 2$, $X_3 = 1$, $X_4 = 0$) using Theorem 1.

Solution: Note that $q_3 = 1$ (state s_3 occurs when the gambler has \$2), since we are certain that we begin with \$2. We have

$$\begin{aligned} p(X_0 = 2, X_1 = 3, X_2 = 2, X_3 = 1, X_4 = 0) &= q_3 p_{34} p_{43} p_{32} p_{21} \\ &= 1 \cdot 0.5 \cdot 0.5 \cdot 0.5 \cdot 0.5 \\ &= 0.0625. \end{aligned}$$ □

Example 9 With reference to Example 4, what is the probability that, having planted one of the seeds in the package (chosen at random), we observe the sequence of flower colors "red, red, yellow, red" over the first four years?

Solution: Using the notation for the state space and Theorem 1, we compute

$$\begin{aligned} &p(X_0 = s_1,\ X_1 = s_1,\ X_2 = s_2,\ X_3 = s_1) \\ &\qquad = q_1 p_{11} p_{12} p_{21} \\ &\qquad = 0.3 \cdot 0.3 \cdot 0.5 \cdot 0.4 \\ &\qquad = 0.018. \end{aligned}$$ □

Example 10 Referring again to Example 4, compute the probability that after having planted a randomly chosen seed, we must wait 3 years for the first *red* flower.

Solution: This means that one of the following sequences of outcomes must have occurred:

$$s_2s_2s_1 \quad s_2s_3s_1 \quad s_3s_2s_1 \quad s_3s_3s_1.$$

As in Example 9, we can use Theorem 1 to compute the probabilities of observing these sequences of outcomes; they are (respectively)

$$0.064, 0.04, 0.06, 0.$$

Thus, the "event" that the first red flower is observed during the third year is the sum of these, 0.164. Notice that knowledge of the initial distribution is essential for computing the desired probability. In this computation, we wish to suggest how knowledge of the probability that a particular sequence is observed allows one to compute the probability of a particular event related to the random process, by simply adding up the probabilities of all the sequences of outcomes which give rise to the event. □

Long-term Behavior

A very important aspect of the analysis of random processes involves discerning some regularity in the long-term behavior of the process. For example, can we be assured that in Example 1 the game will necessarily end after a finite number of coin tosses, or is it conceivable that the game could be indefinitely prolonged with the gambler never going broke and never reaching his $4 goal? In Example 4, we might wish to predict the proportion of plants with a particular flower color among the succeeding generations of the 1000 plants, after planting all 1000 seeds. In fact, it turns out that these proportions tend to certain equilibrium values (that after a period of time will remain virtually unchanged from one year to the next), and that these proportions are not affected by the initial distribution.

We will see that the Markov chains of Examples 1 and 3 exhibit dramatic differences in their long-term behavior. The gambler's ultimate chances of reaching a particular state (like winning four dollars) depend very much on the initial distribution; this is in contrast to the situation in Example 3. For Markov chains, the kinds of questions we have just suggested can in fact be answered once the initial probability distribution and the transition probabilities are known.

We first determine the conditional probability that the random process is in state s_j after k experiments, given that it was initially in state s_i, i.e., we compute $p(X_k = s_j | X_0 = s_i)$. Theorem 2 shows that these probabilities may be found by computing the appropriate power of the matrix of transition probabilities.

Theorem 2 Let $\mathbf{T} = [p_{ij}]$ be the $N \times N$ matrix whose (i, j)th entry is the probability of moving from state s_i to state s_j. Then

$$\mathbf{T}^k = [p_{ij}^{(k)}], \tag{5}$$

where $p_{ij}^{(k)} = p(X_k = s_j | X_0 = s_i)$.

Remark: Note that $p_{ij}^{(k)}$ is *not* the kth power of the quantity p_{ij}. It is rather the (i, j)th entry of the matrix $\mathbf{T}^k$.

Proof: We will prove the Theorem for the case $n = 2$. (The case $n = 1$ is immediate from the definition of $\mathbf{T}$.) The general proof may be obtained using mathematical induction (see Exercise 5). First, note that

$$p(X_2 = s_j, X_0 = s_i) = \sum_{n=1}^{N} p(X_2 = s_j, X_1 = s_n, X_0 = s_i)$$

since X_1 must equal exactly one of the elements of the state space. That is, the N events $\{X_2 = s_j, X_1 = s_n, X_0 = s_i\}$, $(1 \le n \le N)$, are disjoint events whose union is $\{X_2 = s_j, X_0 = s_i\}$. Using this fact, we compute

$$\begin{aligned}
p_{ij}^{(2)} &= p(X_2 = s_j | X_0 = s_i) \\
&= \frac{p(X_2 = s_j, X_0 = s_i)}{p(X_0 = s_i)} \\
&= \sum_{n=1}^{N} \frac{p(X_2 = s_j, X_1 = s_n, X_0 = s_i)}{p(X_0 = s_i)} \quad \text{(by addition rule for probabilities)} \\
&= \sum_{n=1}^{N} p(X_2 = s_j | X_1 = s_n) \frac{p(X_1 = s_n, X_0 = s_i)}{p(X_0 = s_i)} \qquad \text{(by property (i))} \\
&= \sum_{n=1}^{N} p(X_2 = s_j | X_1 = s_n) p(X_1 = s_n | X_0 = s_i) \\
&= \sum_{n=1}^{N} p_{in} p_{nj}. \qquad \text{(by property (ii))}
\end{aligned}$$

This last expression is the (i, j)th entry of the matrix $\mathbf{T}^2$. ■

Example 11 What happens in the long run to the gambler in Example 1? In particular, give the long term probabilities of the gambler either reaching the $4 goal or going broke.

Solution: From (1), the matrix of transition probabilities is

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

According to Theorem 2, computing the powers of T will provide the probabilities of finding the gambler in specific states, given the gambler's initial state.

$$\mathbf{T}^2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.5 & 0.25 & 0 & 0.5 & 0 \\ 0.25 & 0 & 0.5 & 0 & 0.25 \\ 0 & 0.25 & 0 & 0.25 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

$$\mathbf{T}^3 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.625 & 0 & 0.25 & 0 & 0.125 \\ 0.25 & 0.25 & 0 & 0.25 & 0.25 \\ 0.125 & 0 & 0.25 & 0 & 0.625 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

$$\mathbf{T}^{20} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.749 & 0 & 0.001 & 0 & 0.250 \\ 0.499 & 0.001 & 0 & 0.001 & 0.499 \\ 0.250 & 0 & 0.001 & 0 & 0.749 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

$$\mathbf{T}^{50} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.750 & 0 & 0 & 0 & 0.250 \\ 0.500 & 0 & 0 & 0 & 0.500 \\ 0.250 & 0 & 0 & 0 & 0.750 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

(We have rounded the entries to three decimal places.) The matrices $\mathbf{T}^k$ for $k > 50$ will be the same as $\mathbf{T}^{50}$ (at least up to 3 decimal place accuracy). According to Theorem 2, the probability of the gambler going broke after three rounds, given that he starts with $1 is the entry of $\mathbf{T}^3$ in the second row and first column, 0.625. If we want to know the probability that, starting with $2, we will have $3 after 20 rounds (i.e. we want to compute $p_{34}^{(20)}$), we need only look at the entry in the third row and fourth column of the matrix $\mathbf{T}^{20}$, which is 0.001. Notice that the probability of reaching the goal of $4 (after many rounds) depends upon the state in which the gambler is initially observed. Thus, it is with probability 0.75 that the goal is reached if the gambler starts with \$3, 0.5 if he starts with \$2 and 0.25 if he starts with \$1. These are the probabilities that comprise the fifth column of the matrix $\mathbf{T}^k$, for large k ($k = 50$, say). Similar statements concerning the chances of the gambler eventually going broke can easily be made. □

The following result may be deduced as a corollary of Theorem 2. It describes how to compute the probability that the Markov chain will be in any particular state after k experiments, thus allowing one to make predictions about the future behavior of the Markov chain.

Corollary 1 Suppose that a Markov chain has initial probability distribution $\mathbf{Q}$ and $N \times N$ matrix $\mathbf{T}$ of transition probabilities. Then the $1 \times N$ matrix $\mathbf{QT}^k$, denoted $\mathbf{Q}_k$, has as ith entry the probability of observing the Markov chain in state s_i after k experiments. ■

Example 12 Suppose now that the gambler initially gets to pick one of ten envelopes, 6 of them containing \$1, and the other four, \$2. Find the probabilities of the gambler being each possible state after one round of gambling, and estimate the long-term chances of the gambler reaching the \$4 goal.

Solution: The initial probability distribution is

$$(0 \;\; 0.6 \;\; 0.4 \;\; 0 \;\; 0),$$

i.e., the gambler has a probability of 0.6 of starting with \$1 and 0.4 of starting with \$2. By Corollary 1,

$$\mathbf{Q}_1 = \mathbf{QT} = (0.2 \;\; 0.3 \;\; 0.2 \;\; 0.3 \;\; 0).$$

This means that after the first round of gambling, the probabilities that the gambler has 0,1,2,3, or \$4 are, respectively, 0.2, 0.3, 0.2, 0.3, 0. The probability is 0 that the goal of \$4 has been reached.

We have observed that for large k, $\mathbf{T}^{(k)}$ will be almost exactly equal to $\mathbf{T}^{(50)}$. This means that the long-term chances of the gambler reaching the \$4 goal (or going broke) are approximately equal to the probabilities of being in these states after 50 coin tosses. These probabilities will depend on the initial distribution. Using Corollary 1, if

$$\mathbf{Q} = (0 \;\; 0.6 \;\; 0.4 \;\; 0 \;\; 0),$$

we obtain

$$\mathbf{Q}_{50} = (0.6 \;\; 0 \;\; 0 \;\; 0 \;\; 0.4).$$

This means that with this initial distribution, the gambler's long-term chances of going broke are 0.6 and are 0.4 of reaching the \$4 goal. □

In the previous example, if the initial distribution is changed to

$$\mathbf{Q} = (0 \;\; 0 \;\; 0.4 \;\; 0.6 \;\; 0),$$

then we find that

$$\mathbf{Q}_{50} = (0.35 \;\; 0 \;\; 0 \;\; 0 \;\; 0.65);$$

and if the initial distribution were

$$\mathbf{Q} = (0 \;\; 0.4 \;\; 0 \;\; 0.6 \;\; 0),$$

then

$$\mathbf{Q}_{50} = (0.45 \;\; 0 \;\; 0 \;\; 0 \;\; 0.55).$$

Thus, we see that the probability distribution of this Markov chain after many rounds of play depends very much on its initial distribution. This should be contrasted to the behavior of the Markov chain of Example 3 (see Example 13).

The entries in the matrices $\mathbf{T}^{20}$ and $\mathbf{T}^{50}$ suggest that in the long run the probability is (essentially) 0 that the gambler will not have either gone broke or reached his goal, *irrespective* of his initial state. The only reasonably large entries in these matrices are in columns 1 and 5, corresponding to the probabilities of entering states s_1 and s_5 respectively. In fact, what our computations indicate is indeed the case. As the number of rounds gets very large, the probability of being in states s_2, s_3, or s_4 (corresponding to a fortune of \$1, \$2 or \$3, respectively) tends to 0. In this setting, this means that the game cannot last indefinitely; eventually, the gambler will reach \$4 or go broke trying. The states s_1 and s_5 thus have a different character than the other 3 states; eventually the gambler enters one of these two states and can never leave it (since we have assumed the game will stop when one of these states is reached).

Example 13 Analyze the long-term behavior of the Markov chain of Example 3.

Solution: We have already used Table 1 to write down the matrix of transition probabilities of the Markov chain of Example 2:

$$\mathbf{T} = \begin{bmatrix} 0.3 & 0.5 & 0.2 \\ 0.4 & 0.4 & 0.2 \\ 0.5 & 0.5 & 0 \end{bmatrix}.$$

We then compute

$$\mathbf{T}^2 = \begin{bmatrix} 0.39 & 0.45 & 0.16 \\ 0.38 & 0.46 & 0.16 \\ 0.35 & 0.45 & 0.2 \end{bmatrix} \quad \mathbf{T}^3 = \begin{bmatrix} 0.377 & 0.455 & 0.168 \\ 0.378 & 0.454 & 0.168 \\ 0.385 & 0.455 & 0.160 \end{bmatrix}$$

$$\mathbf{T}^4 = \begin{bmatrix} 0.379 & 0.455 & 0.166 \\ 0.379 & 0.455 & 0.166 \\ 0.378 & 0.455 & 0.168 \end{bmatrix} \quad \mathbf{T}^{19} = \begin{bmatrix} 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \end{bmatrix}$$

$$\mathbf{T}^{20} = \begin{bmatrix} 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \end{bmatrix}.$$

□

We also observe that all larger powers of **T** are the same (to within three decimal place accuracy), as was true in Example 11. Thus, in the long run, the behavior of the Markov chain "settles down" in the sense that the probability of observing a particular state remains essentially constant after an initial period of time.

The states of a Markov chain are typically classified according to a criterion suggested by the last two examples.

Definition 2 A state s_i of a Markov chain is called *recurrent* if, given that the chain is in this state at some time, it will return to this state infinitely often with probability 1, i.e.

$$p(X_n = s_i \text{ for infinitely many } n > k | X_k = s_i) = 1.$$

States that are not recurrent are called **transient**; these are the states which will not be observed after enough experiments have been performed. □

None of the states in Example 3 are transient, whereas in Example 1, states s_2, s_3, and s_4 are transient. We have not proved this, but the truth of this fact is certainly suggested by the computations, since the probability of being in any of these states after only 20 rounds is essentially zero. Criteria for determining if a state is transient are developed, for example, in [2].

To summarize, in Example 1 there are two states (s_1 and s_5) which are recurrent, and the remaining three are transient. Consequently, the long term behavior of this random process may be succinctly described by saying that eventually the Markov chain will reach one of the two recurrent states and stay there. Furthermore, the probability of reaching a particular recurrent state depends upon the initial state of the random process.

In Example 2, all states are recurrent. The fact that the numbers in the columns of the higher powers of the matrix of transition probabilities are the same means that the probability of observing the particular state corresponding to the given column after enough experiments have been performed is the same, regardless of the initial state. For example, the probability of observing a red flower (i.e. state s_1) after 19 (or 20 or more) years is 0.379, irrespective of the color of the flower produced by the seed that is originally planted! If all 1000 seeds are planted, then after 20 years we expect to see approximately 379 red flowers, 454 yellow and 167 orange flowers. The same is true of all succeeding years, although we don't expect to find the same colors in the same places; it is the overall proportion of these colors that remains constant.

The random process of Example 3 belongs to a special class of Markov chain which we now define.

Definition 3 A Markov chain with matrix of transition probabilities $\mathbf{T}$ is called *regular* if for some k, $\mathbf{T}^k$ contains all positive entries. □

Regular Markov chains exhibit a long-term behavior quite different from that of Markov chains that have transient states (such as that of Example 1). In particular, it turns out that irrespective of the initial distribution $\mathbf{Q}$, the entries of $\mathbf{Q}_k$ tend toward specific values as k gets large. That is, the probabilities of observing the various states in the long term can be accurately predicted, without even knowing the initial probability distribution. To see why this is so, note that in Example 13 we showed that the Markov chain in Example 3 is regular. In contrast to the Markov chain discussed in Example 12, the probability of observing a particular state after a long period of time does *not* depend on the state in which the random process is initially observed. If $\mathbf{Q} = (q_1\ q_2\ q_3)$ is an arbitrary initial probability distribution for this Markov chain, we have $q_1 + q_2 + q_3 = 1$. Now for large k,

$$\mathbf{T}^k = \begin{bmatrix} 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \\ 0.379 & 0.454 & 0.167 \end{bmatrix},$$

so that

$$\begin{aligned}\mathbf{Q}_k &= \mathbf{Q}\mathbf{T}^k \\ &= (\ (q_1+q_2+q_3)0.379 \quad (q_1+q_2+q_3)0.454 \quad (q_1+q_2+q_3)0.167) \\ &= (0.379\ \ 0.454\ \ 0.167).\end{aligned}$$

Thus, we see that it does not matter what the proportions of colors in the original package of seeds were. In the long run, the proportion of colors tends to 0.379 red, 0.454 yellow and 0.166 orange. This distribution of flower colors is called the *equilibrium distribution* for the Markov chain.

Definition 4 A Markov chain with matrix $\mathbf{T}$ of transition probabilities is said to have an *equilibrium distribution* $\mathbf{Q}_e$ if $\mathbf{Q}_e\mathbf{T} = \mathbf{Q}_e$. □

Note that if $\mathbf{Q}_e\mathbf{T} = \mathbf{Q}_e$, then

$$\mathbf{Q}_e\mathbf{T}^2 = (\mathbf{Q}_e\mathbf{T})\mathbf{T} = \mathbf{Q}_e\mathbf{T} = \mathbf{Q}_e$$

and in general (by the principle of mathematical induction),

$$\mathbf{Q}_e\mathbf{T}^k = \mathbf{Q}_e$$

for all positive integers k. According to the Corollary to Theorem 2, this means that if the initial distribution of the Markov chain is an equilibrium distribution, then the probability of observing the various possible states of the random process does not change with the passage of time.

For the Markov chain of Example 3, we deduced that the equilibrium distribution is

$$(0.379 \quad 0.454 \quad 0.167);$$

furthermore, there is no other equilibrium distribution. To summarize the observations made concerning this Markov chain, we find that all three possible flower colors can be expected to be repeatedly observed, regardless of the seed chosen the first year for planting. If many seeds are planted, the proportion of the flowers bearing red, yellow, and orange flowers in the long run will be, respectively, 0.379, 0.454, and 0.167.

If a random process is known to be a regular Markov chain, then it is of great interest to determine its equilibrium distribution, since this gives a good idea of the long-term behavior of the process. Notice that the entries of the rows of the higher powers of $\mathbf{T}$ are in fact the entries of the equilibrium distribution. This is what in general happens with regular Markov chains: the equilibrium distribution will appear to within any desired accuracy as the entries of the rows of high enough powers of the matrix of transition probabilities. See Exercise 13 for more on this point.

It is no coincidence that the regular Markov chain of Example 3 possesses an equilibrium distribution. We conclude our discussion with the statement of a very important result, the proof of which may be found in any reference on Markov chains, such as [2].

Theorem 3 Every regular Markov chain has a unique equilibrium distribution. ■

Thus, the analysis of the long-term behavior of regular Markov chains can always be carried out by computing the equilibrium distribution.

Suggested Readings

1. W. Feller, *An Introduction to Probability Theory and its Applications*, Vol.1, Third Edition, John Wiley & Sons, Hoboken, N.J., 1968. (Chapter 15 deals with Markov chains. This text contains a wealth of ideas and interesting examples in probability theory and is considered a classic.)

2. J. Kemeny and J. Snell *Finite Markov Chains*, Springer-Verlag, New York, 1983. (This text provides a rigorous and thorough development of the subject, and is directed to the serious undergraduate. Some applications are given.)

3. J. Kemeny, J. Snell and G. Thompson *Introduction to Finite Mathematics*, Third Edition, Prentice Hall, Upper Saddle River, N.J., 1974. (This was probably the first finite mathematics text specifically written to introduce non-science majors to this subject. The basic ideas of a Markov chain are laid out, and many applications can be found throughout the text.)

4. K. Trivedi, *Probability and Statistics with Reliability, Queueing and Computer Science Applications*, Second Edition, John Wiley & Sons, Hoboken, N.J., 2001. (This intermediate-level text provides many applications relevant to computer science.)

Exercises

1. In Example 4, suppose one seed is selected at random from the package.

a) Find the probability of observing an orange flower two years after observing a red flower.

b) Find the probability that in the first four years the flower colors are: yellow, orange, red, red.

c) Find the probability of having to wait exactly three years to observe the first yellow flower.

2. In Example 1, suppose the coin being flipped is not necessarily fair, i.e., the probability of heads is p and the probability of tails is $1 - p$, where $0 < p < 1$. Find the matrix of transition probabilities in this case.

3. Two jars contain 10 marbles each. Every day, there is a simultaneous interchange of two marbles between the two jars, i.e., two marbles are chosen, one from each jar, and placed in the other jar. Suppose that initially, the first jar contained two red marbles and eight white, and that the second jar initially contained all white marbles. Let X_k be the number of red marbles in the first jar after k days (so $X_0 = 2$).

a) Explain why this random process, with observed outcomes X_0, X_1, X_2, ..., is a Markov chain. Find the state space of the process.

⋆b) Find the matrix of transition probabilities of this Markov chain.

c) What is the probability that after three days, there is one red marble in each jar, i.e., $X_3 = 1$?

4. If $\mathbf{T} = [p_{ij}]$ is the matrix of transition probabilities of some Markov chain, explain why the sum of the entries in any of the rows of $\mathbf{T}$ is 1 (i.e. $\sum_{j=1}^{N} p_{ij} = 1$ for $1 \leq i \leq N$).

5. Give the complete proof of Theorem 2.

6. a) Are either of the following matrices the matrix of transition probabilities of a regular Markov chain?

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ 0.5 & 0.3 & 0.2 \\ 0.1 & 0.9 & 0 \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 0 & 0.5 & 0.5 \\ 1 & 0 & 0 \\ 0.3 & 0 & 0.7 \end{bmatrix}.$$

b) If $\mathbf{T}$ is the matrix of transition probabilities of a regular Markov chain, prove that there is an integer k such that the entries of the matrix $\mathbf{T}^m$ are all positive for any $m \geq k$.

7. In a certain country, it has been observed that a girl whose mother is an active voter will with probability 0.5 also vote . A girl whose mother *doesn't* vote is found to become a nonvoter with probability 0.9.

a) Model this phenomenon as a Markov chain with two states. Describe the state space and find the matrix $\mathbf{T}$ of transition probablities.

b) Find the equilibrium distribution $\mathbf{Q}_e = (p\ q)$ of this Markov chain by solving the system of equations $\mathbf{Q}_e\mathbf{T} = \mathbf{Q}_e,\ p + q = 1$. Give an interpretation of the equilibrium distribution in this setting.

8. A state s_k of a Markov chain is called *absorbing* if once the random process enters the state, it remains there, i.e.

$$p_{kj} = \begin{cases} 1, & \text{if } k = j; \\ 0, & \text{otherwise} \end{cases}$$

a) Are there any absorbing states for the random process described in Example 1?

b) Are there any absorbing states for the random process described in Example 3?

c) Is it possible for a *regular* Markov chain to have absorbing states? *Hint*: Look at Exercise 6 b).

⋆9. a) Suppose that A and B are events such that exactly one of A or B occurs, i.e., $P(A \cup B) = 1$ and $P(A \cap B) = 0$. If E is any event, prove that

$P(E) = P(A)P(E|A) + P(B)P(E|B)$.

b) Referring to Example 1, let r_k be the probability that the gambler reaches his \$4 goal before going broke, given that he starts out with \$$k$ ($0 \leq k \leq 4$). Note that $r_0 = 0$ and $r_4 = 1$. For $0 < k < 4$, use part a) to determine a recurrence relation satisfied by the r_k. *Hint*: Let E be the event that, starting with \$$k$, the gambler reaches his goal before going broke, A the event that he wins the first round, and B the event that he loses the first round.

c) Extend the result of part b) to the situation where the gambler's goal is to reach \$$N$ before going broke.

d) Verify that $r_k = k/N,\ 0 \leq k \leq N$, satisfies the recurrence relation found in part c).

10. A drunkard decides to take a walk around the perimeter of the Pentagon. As he reaches each corner of the building (there are 5 corners!), he flips a coin to determine whether to go back to the last corner or to go to the next one.

a) Find the appropriate state space for this random process. Model the position of the drunkard as a Markov chain with 5 states. In particular, find the appropriate transition probability matrix.

b) Show that the equilibrium distribution of this Markov chain is given by (0.2 0.2 0.2 0.2 0.2). Give a description of the long-term behavior of the Markov chain.

c) Suppose that the drunkard has found his way to a square building, and uses the same method of walking around it. Find the matrix of transition probabilities in this case, and determine the equilibrium distribution, if it exists.

$\star$**11.** A dice game has the following rules: Two dice are rolled, and the sum is noted. If doubles are rolled at any time (including the first), the player loses. If he hasn't rolled doubles initially, he continues to roll the dice until he either rolls the same sum he started out with (in which case he wins) or rolls doubles and loses. The game ends when the player either wins or loses.

a) Find the appropriate state space for the random process described above, and discuss why it is a Markov chain. *Hint*: You will need to include a state corresponding to winning and one corresponding to losing.

b) Find the initial distribution of the Markov chain.

c) Find the matrix of transition probabilities.

12. Let S be the state space of a Markov chain with transition probabilities p_{ij}. A state s_j is said to be *accessible* from state s_i if $p_{ij}^n > 0$ for some $n \geq 0$, i.e., s_j can be reached (eventually) if the random process starts out from state s_i. (We will agree that every state is accessible from itself). States s_i and s_j are said to *communicate* if s_i is accessible from s_j and s_j

is accessible from s_i. When the states communicate, we write $s_i \leftrightarrow s_j$.

$\star$ a) Show that the relation "$\leftrightarrow$" is an equivalence relation on S.

b) If s_k is an absorbing state, show that the equivalence class $[s_k]$ contains only s_k itself.

c) If s_i is a transient state and s_j is a recurrent state, is it possible that $s_i \in [s_j]$?

d) Find the equivalence classes under $\leftrightarrow$ in Examples 1 and 3.

$\star$ e) If the Markov chain is *regular*, what can you say about the equivalence classes?

$\star\star$**13.** Show that there are many (in fact, infinitely many) equilibrium distributions for the Markov chain of Example 1.

$\star\star$**14.** Let $\mathbf{T}$ be the $N \times N$ matrix of transition probabilities of a regular Markov chain. It is a fact (mentioned in this chapter) that for large k, each row of $\mathbf{T}^k$ will be very close to the equilibrium distribution $\mathbf{Q}_e = (r_1 \ r_2 \ \ldots \ r_N)$. More precisely, if a number e (e stands for "error") is specified, then there is a number k_0 such that $k > k_0$ implies that for each $1 \leq j \leq N$, $|p_{ij}^{(k)} - r_j| \leq e, \quad 1 \leq i \leq N$.

a) Show that if $\mathbf{Q}$ is *any* initial probability distribution, then $\mathbf{Q}_k$ will be close to the equilibrium distribution for large k.

b) Refer to Corollary 1. There, $\mathbf{Q}_k$ is obtained from the equation $\mathbf{Q}_k = \mathbf{Q}\mathbf{T}^k$. If we want to compute $\mathbf{Q}_k$ by first computing $\mathbf{T}^k$ with k matrix multiplications and then computing $\mathbf{Q}_k = \mathbf{Q}\mathbf{T}^k$, how many arithmetic operations (i.e. multiplications and additions) would be required?

c) Prove that $\mathbf{Q}_k$ can be computed from $\mathbf{Q}_{k-1}$ recursively: $\mathbf{Q}_k = \mathbf{Q}_{k-1}\mathbf{T}$. Use this to find a quicker way to compute $\mathbf{Q}_k$ than that suggested in part b). Determine the number of arithmetic operations required with this method.

d) Refer to Exercise 7. Find $\mathbf{Q}_e$ using the method you developed in part c).

Computer Projects

1. Write a computer program to find the equilibrium distribution of any regular Markov chain using the method suggested in Exercise 14c. This is known as the *power method*.

2. Write a computer program that simulates a Markov chain. It should take as input the matrix of transition probabilities and the initial distribution, and give as output the sequence $X_0, X_1, \ldots$. Such a sequence is known as a **sample path** of the Markov chain.