25 Ergodicity for Finite-State Discrete-Time Markov Chains

At this point in our discussion of discrete-time Markov chains (DTMCs) with M states, we have defined the notion of a *limiting probability of being in state* j:

$$\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij},$$

where the limiting distribution is

$$\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots, \pi_{M-1}), \quad \text{where} \sum_{i=0}^{M-1} \pi_i = 1.$$

We have also defined the notion of a *stationary distribution*, $\vec{\pi}$, as a distribution that satisfies

$$\vec{\pi} \cdot \mathbf{P} = \vec{\pi} \quad \text{and} \quad \sum_{i=0}^{M-1} \pi_i = 1,$$

or, equivalently,

$$\pi_j = \sum_{i=0}^{M-1} \pi_i P_{ij} \quad \text{and} \quad \sum_{i=0}^{M-1} \pi_i = 1.$$

We also proved Theorem 24.9 for finite-state chains that says that, assuming the limiting distribution exists, the limiting distribution is a stationary distribution and no other stationary distribution exists. This theorem is important because it allows us to simply solve the stationary equations to get the limiting distribution.

In Chapter 24, we did *not* spend time on questions like the following:

1. Under what conditions does the limiting distribution exist?
2. How does π_j, the limiting probability of being in state j, compare with p_j, the long-run time-average fraction of time spent in state j?
3. What can we say about m_{jj}, the mean time between visits to state j, and how is this related to π_j?

This entire chapter is devoted to these and other theoretical questions, all related to the notion of *ergodicity*, to be defined soon. This chapter will only address

ergodicity questions for finite-state chains. Infinite-state chains are deferred to Chapter 26.

25.1 Some Examples on Whether the Limiting Distribution Exists

We dive right into the question of existence of the limiting distribution, with a few examples.

Question: What is an example of a valid two-state transition matrix for which π_j does not exist?

Answer: Figure 25.1 shows an example of a chain with transition matrix

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

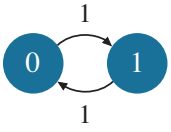

Figure 25.1 *Limiting distribution does not exist.*

The problem is that the chain $\mathbf{P}$ is *periodic*; specifically, a given state is only visited every *other* time step (we will formally define the term "periodic" soon). Observe that $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{jj}$ does not exist, although $\lim_{n\to\infty} \left(\mathbf{P}^{2n}\right)_{jj}$ does exist.

Question: Does this chain have a stationary distribution?

Answer: Yes, the stationary distribution *does* exist. To see this, let's set up the stationary equations $\vec{\pi} \cdot \mathbf{P} = \vec{\pi}$:

$$\begin{aligned} \pi_0 &= \pi_1 \\ \pi_1 &= \pi_0 \\ \pi_0 + \pi_1 &= 1. \end{aligned}$$

Solving these, we get $\vec{\pi} = (\frac{1}{2}, \frac{1}{2})$.

Question: If you walk along the Markov chain for a long time, what fraction of time, p_j, do you spend in state j?

Answer: $p_0 = p_1 = \frac{1}{2}$. These match the stationary probabilities. This is no coincidence. We will see that for any Markov chain, the p_j's satisfy the stationary equations and thus form a stationary distribution.

Question: Is there another two-state example for which the limiting distribution does not exist?

Answer: Consider the transition matrix $\mathbf{Q}$:

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

The corresponding chain is shown in Figure 25.2.

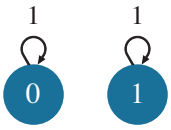

Figure 25.2 *Limiting distribution does not exist.*

The chain $\mathbf{Q}$ has the problem that the limiting state depends on where you start. Recall that the limiting probability of being in state j is supposed to be independent of the start state, i, that is, for transition matrix $\mathbf{Q}$ we want

$$\pi_j = \lim_{n\to\infty} (\mathbf{Q}^n)_{ij}$$

to be independent of i.

However, in our example, if you start in state 1, then you stay there forever, and if you start in state 0, then you stay there forever. Similarly, p_1, the long-run time-average fraction of time spent in state 1, isn't well defined, since it depends on the start state.

Question: What is the stationary distribution of chain $\mathbf{Q}$?

Answer: Chain $\mathbf{Q}$ has an infinite number of stationary distributions!

Examples like these illustrate why we need to differentiate between the stationary probability of being in state j, the limiting probability of being in state j, and the long-run fraction of time spent in state j.

Question: As a final example, does chain $\mathbf{R}$ have limiting probabilities?

$$\mathbf{R} = \begin{bmatrix} 0 & 0 & 1/2 & 1/2 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

Answer: No, chain **R** is also periodic – it is just a little harder to see.

25.2 Aperiodicity

Definition 25.1 *The* **period** *of state j is the greatest common divisor (gcd) of the set of integers n, such that $(\mathbf{P}^n)_{jj} > 0$. A state is* **aperiodic** *if its period is 1. A chain is said to be aperiodic if all of its states are aperiodic.*

To understand the reasoning behind the definition of aperiodic, we recall the Chicken McNugget theorem. Once upon a time, a mathematician walked into McDonald's, hoping to buy food for all his n friends. He wanted to feed them each one chicken nugget (now you know why mathematicians are so skinny). Unfortunately the chicken nugget boxes only came in sizes of 4 nuggets/box or 9 nuggets/box. The mathematician (who was not just skinny but also thrifty) started to wonder if he could express n as a linear combination of 4 and 9, so that no nuggets would go to waste. As often happens, all this thinking led to a theorem, which is called the Chicken McNugget Theorem.

Theorem 25.2 (Chicken McNugget Theorem) *There exists a positive integer n_0, such that, for all integers n, where $n \geq n_0$, we can express n as a non-negative linear combination of 4 and 9. Specifically, we can write:*

$$n = a \cdot 4 + b \cdot 9,$$

where a and b are non-negative integer coefficients.

The Euclidean Number Property extends the Chicken McNugget Theorem to other-sized nugget boxes.

Theorem 25.3 (Euclidean Number Property) *Suppose we're given k positive integers, $i_1, i_2, \ldots, i_k$, where $\gcd(i_1, i_2, \ldots, i_k) = 1$. Then there exists a positive integer n_0, such that for all integers n, where $n \geq n_0$, we can express n as a non-negative linear combination of $i_1, i_2, \ldots, i_k$. Specifically, we can write:*

$$n = a_1 \cdot i_1 + a_2 \cdot i_2 + \cdots + a_k \cdot i_k$$

where the a_i's are non-negative integer coefficients.

Question: Returning to Markov chains, suppose there's a j to j path of length 4 and also one of length 3, as shown in Figure 25.3. Since $\gcd(3,4) = 1$, state j by definition has period 1. But why *intuitively* does state j have period 1?

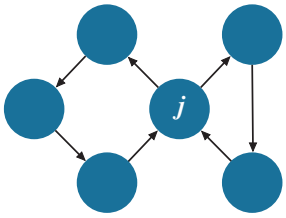

Figure 25.3 *There's a j-to-j path of length 3 and 4.*

Answer: By the Euclidean Number Property we know that for every integer n, greater than some n_0, we can express n as a linear combination of 3 and 4, with non-negative integer coefficients. Thus, there exists a j-to-j path of length n_0, as well as a j-to-j path of length $n_0 + 1$, as well as a j-to-j path of length $n_0 + 2$, and so on. Since there's a j-to-j path of length k for *every* sufficiently large k, we say that the period of j is 1.

Question: Why is it necessary that j be aperiodic for the limiting probability π_j to exist?

Answer: If j has period $d > 1$, then we can't say that there's a j-to-j path of length k for every sufficiently large k (in fact, it turns out we will only end up visiting j once every d steps). But this means that we can't talk about a limiting probability of being in state j independent of the time step n.

25.3 Irreducibility

We've seen that aperiodicity is necessary for the limiting probabilities to exist. Even when a DTMC is aperiodic, there's another problem that could come up: it is possible that the limiting probabilities could depend on the start state, whereas we want

$$\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$$

to be the same for all start states i.

If we also want the limiting probabilities to be independent of the start state, we need one more condition, known as *irreducibility*, which says that from any state one can get to any other state.

Definition 25.4 *State j is* **accessible** *from state i if $(\mathbf{P}^n)_{ij} > 0$ for some $n > 0$. States i and j* **communicate** *if i is accessible from j and vice-versa.*

Definition 25.5 *A Markov chain is* **irreducible** *if all its states communicate with each other.*

Question: Why is irreducibility important for the limiting probabilities to exist?

Answer: The chain might consist of two disconnected components, as in Figure 25.4(a). Here the limiting probability of being in state j depends on the starting state, which is not allowed. Note, however, that irreducibility is not always *necessary* for the existence of the limiting probability. Consider for example Figure 25.4(b), which is also not irreducible, yet the limiting probabilities are all well defined.

(a) No limiting distribution

(b) Limiting distribution exists

Figure 25.4 *Both (a) and (b) show chains which are not irreducible. In (a) the limiting distribution does not exist, because it depends on the start state. In (b) the limiting distribution is* $\vec{\pi} = (0, 1)$.

Question: Do you think that aperiodicity and irreducibility are enough to guarantee the existence of the limiting distribution?

Answer: As we see in Theorem 25.6, for a *finite-state* DTMC, aperiodicity and irreducibility are all that are needed to ensure that the limiting probabilities exist, are positive, sum to 1, and are independent of the starting state. This is convenient, as it is often easy to argue that a DTMC is aperiodic and irreducible.

25.4 Aperiodicity plus Irreducibility Implies Limiting Distribution

Theorem 25.6 (Aperiodicity + irreducibility implies limiting distribution) *Given an aperiodic, irreducible,* finite-state *DTMC with transition matrix* $\mathbf{P}$, *as* $n \to \infty$, $\mathbf{P}^n \to \mathbf{L}$, *where* $\mathbf{L}$ *is a limiting matrix all of whose rows are the same vector,* $\vec{\pi}$. *The vector* $\vec{\pi}$ *has all positive components, summing to* 1.

Question: What does L_{ij} represent?

Answer: The (i, j)th element of $\mathbf{L}$ represents $\lim_{n\to\infty} (\mathbf{P}^n)_{ij}$, namely the limiting probability of being in state j given we started in state i.

Question: What does the ith row of $\mathbf{L}$ represent?

Answer: The ith row of $\mathbf{L}$ is the vector of limiting probabilities $(\pi_0, \pi_1, \ldots, \pi_{M-1})$, where $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$, and M is the number of states in the DTMC.

Question: Why is it important that the rows of $\mathbf{L}$ are the same?

Answer: The fact that row i and row k are the same says that

$$\lim_{n\to\infty} (\mathbf{P}^n)_{ij} = \lim_{n\to\infty} (\mathbf{P}^n)_{kj},$$

which says that the starting state does not affect the limiting probability of being in state j.

As a concrete example of Theorem 25.6, suppose that

$$\mathbf{P} = \begin{bmatrix} 1/2 & 1/3 & 1/6 \\ 1/3 & 1/3 & 1/3 \\ 1/8 & 3/4 & 1/8 \end{bmatrix}.$$

Then Theorem 25.6 is saying that $\mathbf{P}^n$ converges to a matrix $\mathbf{L}$ all of whose rows are the same. That is,

$$\mathbf{P}^n \longrightarrow \begin{bmatrix} 0.34 & 0.43 & 0.23 \\ 0.34 & 0.43 & 0.23 \\ 0.34 & 0.43 & 0.23 \end{bmatrix} = \mathbf{L}.$$

Proof: [Theorem 25.6] The remainder of this section is devoted to the proof of Theorem 25.6. This is a long proof and will require introducing a couple claims along the way. We are trying to show that $\mathbf{P}^n$ converges to a matrix where all rows are the same. Equivalently, we are trying to show that, for any j, the jth column of $\mathbf{P}^n$ converges to a vector whose components are all the same.

Let $\vec{e}$ represent the *column vector* of dimension matching $\mathbf{P}$, whose jth component is 1 and whose remaining components are all 0. That is,

$$\vec{e} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

We are trying to show that

$$\mathbf{P}^n \cdot \vec{e}$$

converges to a vector all of whose components are the same. The idea is to view

$$\mathbf{P}^n \vec{e} = \mathbf{P}(\cdots(\mathbf{P}(\mathbf{P}(\mathbf{P}\vec{e})))).$$

Consider the innermost product $\mathbf{P}\vec{e}$. Because $\mathbf{P}$ is a matrix of probabilities, where each row sums to 1, the effect of multiplying $\vec{e}$ by $\mathbf{P}$ is to replace each component of $\vec{e}$ by a value that is a *weighted average* of all the components. In particular, the effect is to bring all the components of $\vec{e}$ closer together. That is, the difference between the maximum component and the minimum component should decrease.

Here is an example of the effect of successive multiplications by $\mathbf{P}$:

$$\mathbf{P}\vec{e} = \begin{bmatrix} 1/2 & 1/3 & 1/6 \\ 1/3 & 1/3 & 1/3 \\ 1/8 & 3/4 & 1/8 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/3 \\ 1/3 \\ 3/4 \end{bmatrix}.$$

$$\mathbf{P}(\mathbf{P}\vec{e}) = \begin{bmatrix} 1/2 & 1/3 & 1/6 \\ 1/3 & 1/3 & 1/3 \\ 1/8 & 3/4 & 1/8 \end{bmatrix} \cdot \begin{bmatrix} 1/3 \\ 1/3 \\ 3/4 \end{bmatrix} = \begin{bmatrix} 0.40 \\ 0.47 \\ 0.39 \end{bmatrix}.$$

Observe that after just two successive multiplications by $\mathbf{P}$, the components are already quite close!

We now claim that the difference between the maximum and minimum components of $\mathbf{P}^n\vec{e}$ shrinks as we increase n.

Claim 25.7 *Let M_n denote the maximum component of $\mathbf{P}^n\vec{e}$ and let m_n denote the minimum component of $\mathbf{P}^n\vec{e}$. Then*

$$M_n - m_n \leq (1-2s)(M_{n-1} - m_{n-1}), \tag{25.1}$$

where s is the smallest element in $\mathbf{P}$.

Proof: [Claim 25.7] To see intuitively why Claim 25.7 is true, consider the vector $\vec{y} = \mathbf{P}^{n-1}\vec{e}$. By our definition, the maximum component of $\vec{y}$ is M_{n-1} and the minimum is m_{n-1}. Now, if we multiply $\vec{y}$ by $\mathbf{P}$ (obtaining $\mathbf{P}\vec{y} = \mathbf{P}^n\vec{e}$), we are replacing each component of $\vec{y}$ by a weighted average of all the components of $\vec{y}$.

Question: More formally, what is an upper bound on the largest possible component, M_n, in $\mathbf{P} \cdot \vec{y} = \mathbf{P} \cdot \left(\mathbf{P}^{n-1}\vec{e}\right)$?

Answer: The largest possible M_n value is obtained when $\mathbf{P}$ is multiplied by $\vec{y}$ where all but one of the elements of $\vec{y}$ are M_{n-1}, with the remaining one being m_{n-1}, that is, $\vec{y}$ has only one small component.

To maximize M_n, we now want to make sure that the small m_{n-1} component of $\vec{y}$ is multiplied by the *smallest* possible value of $\mathbf{P}$, namely s. To do this, we consider the row, i, of $\mathbf{P}$ that contains s. Suppose s occurs in the jth column of $\mathbf{P}$. Then we make sure that m_{n-1} is likewise in the jth component of $\vec{y}$. This forces m_{n-1} to be multiplied by s. The remaining total weight in row i of $\mathbf{P}$ is $1 - s$, which gets multiplied by only M_{n-1} terms in $\vec{y}$. Thus an upper bound on M_n is given by:

$$M_n \leq s \cdot m_{n-1} + (1 - s) \cdot M_{n-1}. \tag{25.2}$$

Question: What is a lower bound on the smallest possible component, m_n in $\mathbf{P} \cdot \vec{y} = \mathbf{P} \cdot \left(\mathbf{P}^{n-1}\vec{e}\right)$?

Answer: Similarly, the smallest possible m_n value is obtained if all but one of the elements of $\vec{y}$ are m_{n-1}, with the remaining one being M_{n-1}. This time we want to make sure that the M_{n-1} component of $\vec{y}$ is weighted by the smallest possible value of $\mathbf{P}$, namely s. This allows the biggest possible remaining row weight of $1 - s$ to be applied to m_{n-1}. Thus a lower bound on m_n, the smallest component of $\mathbf{P} \cdot \vec{y}$, is:

$$m_n \geq (1 - s) \cdot m_{n-1} + s \cdot M_{n-1}. \tag{25.3}$$

Thus,

$$\begin{aligned} M_n - m_n &\leq (25.2) - (25.3) \\ &= s \cdot m_{n-1} + (1 - s) \cdot M_{n-1} - (1 - s) \cdot m_{n-1} - s \cdot M_{n-1} \\ &= (1 - 2s)(M_{n-1} - m_{n-1}). \end{aligned}$$

■

From Claim 25.7, it seems that the difference between the maximum and minimum elements of $\mathbf{P}^n\vec{e}$ continues to decrease as we continue to multiply by $\mathbf{P}$, until eventually all elements are the same, so we're done with the proof. This is true, except for a small hole ...

Question: Can you see the hole in the argument?

Answer: If $\mathbf{P}$ contains a zero element, then $s = 0$. In this case Claim 25.7 does not result in convergence, because $(1 - 2s) = 1$.

Question: How can this be fixed?

Hint: Even if $\mathbf{P}$ contains some zero elements, what do we know about $\mathbf{P}^n$ for high enough n, given that $\mathbf{P}$ is aperiodic and irreducible?

Answer: When $\mathbf{P}$ is aperiodic and irreducible, we will now show that even if $\mathbf{P}$ contains some zero elements, for all n beyond some point, $\mathbf{P}^n$ has all positive elements.

Claim 25.8 *Given* $\mathbf{P}$ *is aperiodic and irreducible, there exists some* n_0*, such that* $\forall n \geq n_0$*,* $\mathbf{P}^n$ *has all positive elements.*

Proof: [Claim 25.8] The proof is a consequence of the Euclidean Number Property (Theorem 25.3), as follows: Consider an arbitrary (j, j) entry of $\mathbf{P}$.

Question: If $P_{jj} > 0$, can we conclude that $(\mathbf{P}^n)_{jj} > 0$, $\forall n$?

Answer: Yes. The fact that there's a path of length 1 from j to j implies that there's a path of length n from j to j.

So suppose that $P_{jj} = 0$. By irreducibility, there exist paths from j to j. By aperiodicity, the gcd of these j-to-j paths is 1. Suppose, for example, the j-to-j paths have lengths x, y, and z, where $\gcd(x, y, z) = 1$. Hence, by the Euclidean Number Property, $\exists n_0(j, j)$, s.t., $\forall n \geq n_0(j, j)$, n can be expressed as a linear combination of x and y and z with non-negative integer coefficients; hence, $\forall n \geq n_0(j, j)$, there is a path of length n from j to j, and thus the (j, j)th entry of $\mathbf{P}^n$ is positive.

Now repeat this argument for all (i, i) pairs (there are only a finite number).

Next, consider two arbitrary states, i and j, where $i \neq j$. By irreducibility, there is some x s.t. there is a path from i to j of length x. However, since we also know that $\forall n \geq n_0(i, i)$ there is a path of length n from i to i, it follows that $\forall n \geq n_0(i, i)+x$ there's a path of length n from i to j. Define $n_0(i, j) = n_0(i, i)+x$.

Finally, define

$$n_0 = \max_{i,j}\{n_0(i, j)\}.$$

Now, for all $n \geq n_0$, $\mathbf{P}^n$ has all positive elements. ■

To complete the proof of Theorem 25.6, we now define $\mathbf{P}' = \mathbf{P}^{n_0}$. Then,

$$\mathbf{P}^n = (\mathbf{P}^{n_0})^{n/n_0} = (\mathbf{P}')^{n/n_0}.$$

Now repeat the argument in Claim 25.7, except that rather than the decrease by a factor of $(1 - 2s) < 1$ occurring with each multiplication of $\mathbf{P}$, this decrease only happens every n_0 multiplications of $\mathbf{P}$. However, because $n/n_0 \to \infty$ as $n \to \infty$,

we still have an infinite number of these decreases, meaning that

$$(\mathbf{P}')^{n/n_0} \to \mathbf{L}, \qquad \text{as } n \to \infty.$$

Note that this argument still works even if n/n_0 is a fraction. In that case we define $n = m \cdot n_0 + r$, where $r < n_0$ and use m in place of n/n_0 in our argument. Here,

$$\mathbf{P}^n = \mathbf{P}^r \cdot (\mathbf{P}^{n_0})^m,$$

where the rightmost term converges to $\mathbf{L}$ as $m \to \infty$, and the $\mathbf{P}^r$ term doesn't affect this limit.

To finish off the proof of Theorem 25.6, we note that by Exercise 24.2, all powers of $\mathbf{P}$ have the property that the components of each row sum to 1. Furthermore, because $\mathbf{P}^{n_0}$ has all positive elements, and because multiplying by $\mathbf{P}$ only creates weighted averages of already positive values, then $\mathbf{P} \cdot \mathbf{P}^{n_0}$ still has all positive elements and so forth as we continue to multiply by $\mathbf{P}$. Hence the limiting matrix $\mathbf{L}$ will still have all positive elements and will have the property that the components of each row sum to 1. ■

Summary: We have proven that for any aperiodic, irreducible, finite-state Markov chain, the limiting probabilities exist and are all positive.

Definition 25.9 *We say that a* finite-state *DTMC is* **ergodic** *if it has both desirable properties: aperiodicity and irreducibility. For the case of an infinite-state DTMC, ergodicity requires one more property (see Chapter 26).*

25.5 Mean Time Between Visits to a State

Consider the mean time between visits to state j, which we'll call m_{jj}. It seems that m_{jj} should be related to π_j, the limiting probability of being in state j. Theorem 25.12 shows that m_{jj} and π_j are in fact reciprocals.

Definition 25.10 *Let m_{ij} denote the expected number of time steps needed to first get to state j, given we are currently at state i. Likewise, let m_{jj} denote the expected number of steps between visits to state j.*

Theorem 25.11 *For an irreducible finite-state DTMC, m_{ij} is finite, for all i, j.*

Proof: See Exercise 25.19. ■

Theorem 25.12 *For an irreducible, aperiodic finite-state Markov chain with transition matrix* $\mathbf{P}$*,*

$$\pi_j = \frac{1}{m_{jj}} > 0,$$

where m_{jj} *is the mean time between visits to state* j *and* $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$.

Proof: We derive m_{ij} by conditioning on the first step, as follows:

$$\begin{aligned} m_{ij} &= P_{ij} \cdot 1 + \sum_{k\neq j} P_{ik}(1 + m_{kj}) \\ &= 1 + \sum_{k\neq j} P_{ik} m_{kj}. \end{aligned} \tag{25.4}$$

Likewise,

$$\begin{aligned} m_{jj} &= P_{jj} \cdot 1 + \sum_{k\neq j} P_{jk}(1 + m_{kj}) \\ &= 1 + \sum_{k\neq j} P_{jk} m_{kj}. \end{aligned} \tag{25.5}$$

We will now express (25.4) and (25.5) using matrix notation. All the matrices in this proof are of the same dimension as $\mathbf{P}$. Let $\mathbf{M}$ be a matrix whose (i, j)th entry is m_{ij}. For purposes of the proof, it will be convenient to express $\mathbf{M}$ as a sum of two matrices,

$$\mathbf{M} = \mathbf{D} + \mathbf{N},$$

where $\mathbf{D}$ is a matrix whose entries are all zero, except for its diagonal entries: $d_{jj} = m_{jj}$, and $\mathbf{N}$ is a matrix whose diagonal entries are all zero, but where $N_{ij} = m_{ij}$, $\forall i \neq j$. Finally, let $\mathbf{E}$ be a matrix with *all* entries 1. Then we can express (25.4) and (25.5) as:

$$\mathbf{M} = \mathbf{E} + \mathbf{PN}. \tag{25.6}$$

Rewriting (25.6), we have

$$\begin{aligned} \mathbf{N} + \mathbf{D} &= \mathbf{E} + \mathbf{PN} \\ (\mathbf{I} - \mathbf{P}) \cdot \mathbf{N} &= \mathbf{E} - \mathbf{D}. \end{aligned}$$

From Theorem 25.6, since we have aperiodicity and irreducibility, we know that the limiting distribution, $\vec{\pi}$, exists. Multiplying both sides by $\vec{\pi}$, we have:

$$\vec{\pi} \cdot (\mathbf{I} - \mathbf{P}) \cdot \mathbf{N} = \vec{\pi} (\mathbf{E} - \mathbf{D}) . \tag{25.7}$$

Question: What do we know about the left-hand side of (25.7)?

Hint: Remember that $\vec{\pi}$ is also a stationary distribution, by Theorem 25.6.

Answer:

$$\begin{aligned} \vec{\pi}\mathbf{P} &= \vec{\pi} \\ \Rightarrow \vec{\pi}(\mathbf{I}-\mathbf{P}) &= \vec{0} \\ \Rightarrow \vec{\pi}(\mathbf{I}-\mathbf{P})\mathbf{N} &= \vec{0}. \end{aligned}$$

Thus, from (25.7) we have:

$$\begin{aligned} \vec{0} &= \vec{\pi}(\mathbf{E}-\mathbf{D}) \\ \vec{\pi}\mathbf{E} &= \vec{\pi}\mathbf{D} \\ (1,1,\ldots,1) &= (\pi_0 m_{00}, \pi_1 m_{11}, \ldots, \pi_{M-1} m_{M-1,M-1}) \\ \pi_i m_{ii} &= 1, \quad \forall i \\ \pi_i &= \frac{1}{m_{ii}} > 0, \quad \forall i, \end{aligned}$$

where the last line follows from the fact that m_{ii} is finite by Theorem 25.11. ■

Corollary 25.13 *For an irreducible,* periodic *finite-state Markov chain,*

$$\pi_j^{stationary} = \frac{1}{m_{jj}} > 0,$$

where m_{jj} is the mean time between visits to state j and $\pi_j^{stationary}$ is the stationary probability of being in state j.

Proof: This is an easy consequence of the proof of Theorem 25.12 and is shown in Exercise 25.5. ■

25.6 Long-Run Time Averages

For the purpose of this section, we imagine that we have an ergodic, finite-state DTMC, such as that shown in Figure 25.5.

A **random walk** is a walk through a Markov chain, where we move *indefinitely* between the states of the Markov chain according to the probabilities of the chain. For example, we might start at some state like *Blah*, and next move to *Tired* and then to *Wasted* and from there maybe back to *Blah*, and so on. Of course, you might take a different random walk through the chain, where you again start at state *Blah*, but this time next move to state *Achy* and from there back to state *Blah* and so on. Each random walk is often referred to as a **sample path** in that

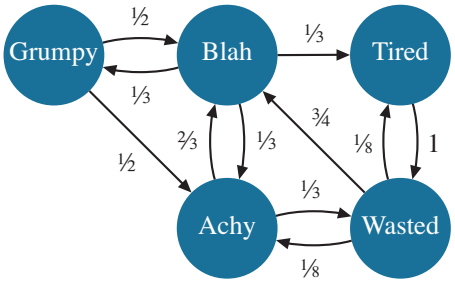

Figure 25.5 *The moods of students in the aftermath of Carnival festivities.*

it depends on the random coin flips. We say sample *path* rather than a sample *point*, because the walk is infinitely long.

Let $N_j(t)$ be the number of visits to state j by time t on our random walk. Our goal is to formally define p_j, the long-run proportion of time that a random walk spends in state j.

Question: How might we define p_j using $N_j(t)$?

Definition 25.14 *Given an irreducible DTMC, we define the* **long-run time-average fraction of time** *that a random walk on the DTMC spends in state j as:*

$$p_j = \lim_{t\to\infty} \frac{N_j(t)}{t},$$

where $N_j(t)$ is the number of times that the random walk enters state j by time t (in the first t time steps).

Question: Why does Definition 25.14 start by specifying that the DTMC is irreducible?

Answer: If the DTMC were not irreducible, then the time-average fraction of time spent in state j might depend on where we start, which would make it undefined.

In this section we ask: *How does p_j compare to π_j?*

Recall the definition of π_j, the **limiting probability** of being in state j:

$$\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}.$$

While p_j is an average over a *single* sample path, π_j is an average over *many*

sample paths. To see this, let's consider the quantity

$$(\mathbf{P}^n)_{ij}.$$

This represents the probability of being in state j after n steps, given that we started in state i. If we consider $n = 1$, $\left(\mathbf{P}^1\right)_{ij} = P_{ij}$, namely the probability that in the first step we move to state j. On the other hand, $\left(\mathbf{P}^2\right)_{ij}$ is the weighted average over M two-step sample paths (the intermediate state could be any of the M states). Similarly, $\left(\mathbf{P}^3\right)_{ij}$ is the weighted average over M^2 three-step sample paths, and so on. We refer to

$$\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$$

as an **ensemble average**, meaning that it is an average over *many* sample paths, in fact an infinite number.

Question: Does $p_j = \pi_j$?

Answer: It is not at all obvious that p_j, the *time-average* fraction of time spent in state j on a *single* sample path, should equal π_j, the *ensemble average* fraction of time spent in state j, averaged over *all* sample paths. The purpose of this section is to prove that, when π_j exists, then, on "almost all" sample paths, $p_j = \pi_j$. We will spend the rest of this section making this claim precise and proving it.

Before we get into it, we note one important way in which π_j and p_j differ.

Question: Recall that aperiodicity was required for π_j to exist. Is aperiodicity required for p_j to exist?

Answer: No. Irreducibility is all that is needed to ensure p_j is well defined.

To prove our claim that $p_j = \pi_j$, we will need to first understand the Strong Law of Large Numbers and then to learn a little renewal theory.

25.6.1 Strong Law of Large Numbers

The Strong Law of Large Numbers (SLLN) is an extremely important result in probability theory, but it is difficult to prove. We refer the interested reader to [22].

Theorem 25.15 (SLLN) *Let $X_1, X_2, \ldots$ be a sequence of independent, identically distributed (i.i.d.) random variables each with finite mean $\mathbf{E}[X]$. Let $S_n = \sum_{i=1}^n X_i$. Then, with probability 1,*

$$\lim_{n\to\infty} \frac{S_n}{n} = \mathbf{E}[X].$$

While we omit the proof in this book, we will spend time discussing the meaning of the result. Let's consider, for example, that

$$X_i \sim \text{Bernoulli}(0.5),$$

that is X_i represents the ith flip of a fair coin. Here, S_n represents the sum of the first n coinflips, and $\frac{S_n}{n}$ represents the average over the first n coinflips. SLLN says that, when n gets large, this average should converge to 0.5.

At first this sounds entirely obvious. After all, what else could the average be?

Looking a little closer, we note that SLLN says this happens **"with probability 1."** The term "with probability 1" is roughly saying that the statement is true on *almost every* sample path. A sample path here refers to a sequence of instances of $X_1, X_2, X_3, \ldots$ Each sample path is infinitely long, and there are infinitely many sample paths (there are two values possible for each X_i). More precisely, the statement "with probability 1" says that if we consider the number of "bad" sample paths on which the convergence doesn't happen and divide that by the total number of sample paths, then:

$$\frac{\text{Number bad sample paths up to length } n}{\text{Total number sample paths up to length } n} \to 0 \text{ as } n \to \infty.$$

Let's consider whether this makes sense.

Question: What's an example of a "bad" sample path?

Answer: 00000 . . . or 11111 . . .

Question: Are there a finite or infinite number of bad sample paths?

Answer: Infinite.

Question: Is the number of bad sample paths countably infinite or uncountably infinite?

Answer: Uncountably infinite. Here's how to see this. Let's refer to the sequence 110 as a "red car" and to the sequence 101 as a "blue car" (Figure 25.6). Now any sequence made up of red and blue cars is clearly bad, because it has twice as many 1's as 0's. However, there are an uncountable number of possible sequences of red and blue cars (by Cantor's diagonalization argument [11]).

Figure 25.6 *Any sequence of red and blue cars is a bad sample path.*

Given that there are an uncountably infinite number of bad sample paths, it should be a little clearer why it's not so obvious that the fraction of bad sample paths goes to 0. This explains the power of SLLN.

25.6.2 A Bit of Renewal Theory

Definition 25.16 *A* **renewal process** *is any process for which the times between events are i.i.d. random variables, with a non-negative distribution X.*

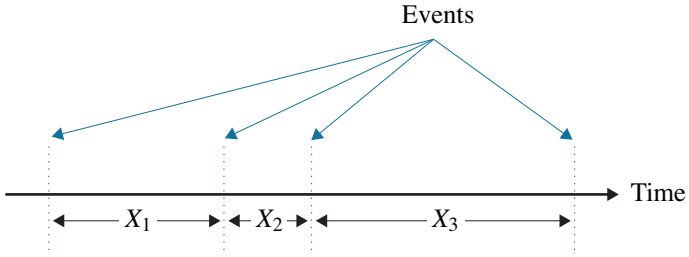

Figure 25.7 *A renewal process. $X_i \sim X$, for all i.*

An example of a renewal process is shown in Figure 25.7. Let $N(t)$ denote the number of renewal events by time t. Then, we have the following theorem:

Theorem 25.17 (Renewal Theorem) *For a renewal process, if $\mathbf{E}[X] > 0$ is the mean time between renewals, where $\mathbf{E}[X]$ is finite, we have*

$$\lim_{t\to\infty} \frac{N(t)}{t} = \frac{1}{\mathbf{E}[X]} \text{ with probability 1.} \tag{25.8}$$

Proof: The basic idea in this proof is to apply SLLN, which gives us the convergence on all sample paths with probability 1 (abbreviated, w.p.1). Let S_n be the time of the nth event. Then we have, $\forall t$,

$$\begin{array}{ccccc} S_{N(t)} & \leq & t & < & S_{N(t)+1} \\ \frac{S_{N(t)}}{N(t)} & \leq & \frac{t}{N(t)} & < & \frac{S_{N(t)+1}}{N(t)}. \end{array}$$

Looking at the leftmost term, we have:

$$\frac{S_{N(t)}}{N(t)} = \frac{\sum_{i=1}^{N(t)} X_i}{N(t)} \longrightarrow \mathbf{E}[X] \text{ as } t \to \infty \quad \text{w.p.1 (SLLN).}$$

Looking at the rightmost term, we have:

$$\frac{S_{N(t)+1}}{N(t)} = \frac{S_{N(t)+1}}{N(t)+1} \cdot \frac{N(t)+1}{N(t)} \longrightarrow \mathbf{E}[X] \cdot 1 \text{ as } t \to \infty \quad \text{w.p.1 (SLLN).}$$

So, by the sandwich theorem, the center term likewise converges to $\mathbf{E}[X]$, namely:

$$\frac{t}{N(t)} \longrightarrow \mathbf{E}[X] \text{ w.p.1,}$$

which implies that

$$\Rightarrow \frac{N(t)}{t} \longrightarrow \frac{1}{\mathbf{E}[X]} \text{ as } t \to \infty \text{ w.p.1.} \qquad \blacksquare$$

25.6.3 Equality of the Time Average and Ensemble Average

We are finally ready to relate p_j, the time-average fraction of time that a DTMC spends in state j, to π_j, the limiting probability of being in state j.

Theorem 25.18 *For a finite-state irreducible DTMC, with probability* 1,

$$p_j = \frac{1}{m_{jj}}.$$

For a finite-state, irreducible, and aperiodic DTMC, with probability 1,

$$p_j = \pi_j.$$

Proof: By Theorem 25.11, we know that m_{jj} is finite. Thus we can apply the Renewal Theorem (Theorem 25.17) to say that

$$p_j = \lim_{t\to\infty} \frac{N_j(t)}{t} = \frac{1}{m_{jj}} \quad \text{w.p.1,}$$

where $N_j(t)$ is the number of visits to state j by time t.

Now, if we have both irreducibility and aperiodicity, we can invoke Theorem 25.12 which says that

$$\pi_j = \frac{1}{m_{jj}}.$$

Thus, $\pi_j = p_j$, w.p.1. $\blacksquare$

25.7 Summary of Results for Ergodic Finite-State DTMCs

So far we've seen that for a finite-state DTMC which is both aperiodic and irreducible, the limiting distribution, $\vec{\pi}$ exists. This $\vec{\pi}$ is also the unique stationary distribution and furthermore represents the time-average probabilities of being in each state. In Theorem 25.19 we summarize all the results we've seen about ergodic finite-state DTMCs.

Theorem 25.19 (Summary theorem for ergodic, finite-state DTMCs) *In a finite-state DTMC, the word* ergodic *refers to two properties: aperiodic and irreducible. Given an ergodic finite-state chain, the following results hold:*

- *(Theorem 25.6) The limiting distribution exists and has all-positive components.*
- *(Theorem 25.12)* $\pi_j^{limiting} = \frac{1}{m_{jj}}$.
- *(Theorem 24.9) The stationary distribution is unique and is equal to the limiting distribution.*
- *(Theorem 25.18) Time-average* $p_j = \frac{1}{m_{jj}}$, *w.p.*1.
- *Putting it all together, we have that:*

$$0 < \frac{1}{m_{jj}} = \pi_j^{limiting} = \pi_j^{stationary} = p_j, \text{ w.p.1.}$$

25.8 What If My DTMC Is Irreducible but Periodic?

So life is great when your DTMC is ergodic. But suppose instead you have a finite-state DTMC that is irreducible but periodic.

For any *periodic* chain, the limiting distribution does not exist (because the probability of being in a state depends on the time step).

However, it turns out that if the finite-state DTMC is irreducible, that alone suffices to ensure that the stationary distribution *exists* and is *unique* [35, section 9.8]. We saw an example of such an irreducible periodic chain in Figure 25.1. For such chains, the stationary distribution represents the long-run time-average proportion of time spent in each state, that is, the p_j's.

Very roughly the proof in [35, section 9.8] starts with the observation that when a chain is irreducible, all states have the *same* period d (see Exercise 25.17). Thus, it turns out that we can divide all the states into d residue classes, where some states are visited at times 0 mod d, some at times 1 mod d, . . ., and some

are visited at times $d - 1 \bmod d$. Thus, while $\lim_{n\to\infty} (\mathbf{P}^n)_{ij}$ does not exist, $\lim_{n\to\infty} (\mathbf{P}^{nd})_{ij}$ does exist, where d is the period of the chain. Thus we can think of the limiting distribution as existing if we only observe the chain every dth time step; and when the limiting distribution exists, we get a unique stationary distribution.

Since the case of irreducible, periodic finite-state DTMCs comes up quite a bit, we provide another summary theorem with everything you need to know about this case.

Theorem 25.20 (Summary for irreducible, periodic, finite-state DTMCs) *For a finite-state DTMC that is irreducible, but periodic:*

- *The limiting distribution does not exist (it depends on the time step).*
- *The stationary distribution exists and is unique [35].*
- *(Theorem 25.11) For every state j, m_{jj} is finite.*
- *(Corollary 25.13) $\pi_j^{stationary} = \frac{1}{m_{jj}}$.*
- *(Theorem 25.18) Time-average $p_j = \frac{1}{m_{jj}}$, w.p.1.*
- *Putting it all together, we have that:*

$$0 < \frac{1}{m_{jj}} = \pi_j^{stationary} = p_j, \text{ w.p.1.}$$

25.9 When the DTMC Is Not Irreducible

In the case of a finite-state DTMC that is *not* irreducible, the limiting distribution may or may not exist.

For examples of chains which are not irreducible and the limiting distribution *does not* exist, see Figure 25.2 and Figure 25.4(a). Generally, a lack of existence happens if the DTMC consists of two completely disconnected components. In such situations, the limiting probability of being in state j is not independent of the starting state i. Note that while the limiting distribution doesn't exist, in Exercise 25.20 we prove that (at least one) stationary always exists for any finite-state chain.

An example of a chain which is not irreducible, yet the limiting distribution nevertheless *exists*, is given in Figure 25.4(b), where the limiting distribution is $\vec{\pi} = (0, 1)$, even though the chain is not irreducible and the period is undefined. In cases when the limiting distribution *does* exist, it is no longer the case that the limiting probability of every state j is positive, as we had in Theorem 25.12,

since some states may not be reachable, or there may be an "absorbing" state (or states), from which one never leaves, as is the case in Figure 25.4(b).

Even if the entire chain is not irreducible, the chain can still be subdivided into irreducible components (sometimes individual states), where an irreducible component may function as its own ergodic chain.

In the next section, we will encounter some examples of chains that are not irreducible and illustrate the above points.

25.10 An Application: PageRank

We now consider an application of finite-state DTMCs and some of the ergodicity concepts that we've been studying.

Question: How many web search engines can you name?

Answer: Here are a few: W3Catalog (1993), WebCrawler (1994), Lycos (1994), AltaVista (1995), Excite (1995), Yahoo! (1995), Google (1998), Bing (2009).

The goal of a web search engine is not just to find a page that contains the item that you're searching for, but to find *the best* page that contains that item. For example, your name might appear on a lot of web pages: chess tournaments, swim competitions, theater productions, etc. Every search engine will show all these different pages. However what makes a search engine good is its ability to **rank** the pages, showing the most important pages first, so that someone searching for you will first see your Homepage or Linked In page, rather than that picture of you as a third grader.

Of course, how can a search engine know exactly which of the thousand pages is the most relevant one?

A common solution is to rank the pages in order of the number of links to that page (often called **backlinks** of the page), starting with the page that has the highest number of pointers into it. We refer to this strategy as *citation counting*.

Citation counting is a very commonly used measure of importance. For example, many tenure decisions are determined not by your number of publications, but by the number of citations to your publications.

Question: Suppose that we could determine the number of backlinks of each page (number of links pointing to the page). Why would that *not* necessarily be a good measure of the importance of the page?

Answer:

(1) Not all links are equal. If a page is pointed to from *cnn.com*, that link should be counted much more than if a page is pointed to from Joe Schmo's page.

(2) The citation counting scheme is easily tricked. Suppose I want my web page to have a high rank. I simply create a thousand pages that each point to my web page. Now my web page has a thousand pointers into it, so it should be ranked highly. (Hmmm ... not a bad way to handle the tenure citation issue too).

Okay, so citation counting is not the best of schemes. While it is insufficient to just count the number of pages pointing into a page p, we might do better by weighting each pointer by the number of pages pointing into it.

Question: Why is this system also easy to fool?

Answer: I can again create a thousand dummy web pages and have them all point to each other, in a clique, as well as pointing to my page. Now my web page has a high number of backlinks, all of which also have a high number of backlinks.

Google's PageRank Solution: Google's solution is to define PageRank recursively: *"A page has high rank if the sum of the ranks of its backlinks is high."* Observe that this covers both the case when a page has many backlinks and when a page has a few highly ranked backlinks.

Question: It is easy to say that "a page has high rank if the sum of the ranks of its backlinks is high," but how does that help us figure out the rank of a page?

Answer: The "aha" that the Google founders made was to realize that the recursive definition is actually saying

$$\pi_j = \sum_{i=1}^{n} \pi_i P_{ij},$$

where n is the number of pages.

That is, the only way for page j to have high limiting probability is if the pages i pointing into j have high limiting probability. Remind you of anything?

The rank of a page is thus just its stationary probability in a Markov chain!

Algorithm 25.21 (Google's PageRank algorithm)

1. *Create a DTMC transition diagram where there is one state for each web page and there is an arrow from state i to state j if and only if page i has a link to page j.*
2. *If page i has $k > 0$ outgoing links, then set the probability on each outgoing arrow from state i to be $1/k$.*
3. *Solve the DTMC to determine stationary probabilities. Pages are then ranked based on their stationary probabilities (higher probability first).*

This simple algorithm was the original basis behind the entire Google company. Today, Google has incorporated additional heuristics.

Example 25.22 (Well-behaved web graph)

Suppose the entire web consists of the three pages shown in Figure 25.8(a). Then the corresponding DTMC transition diagram is shown in Figure 25.8(b).

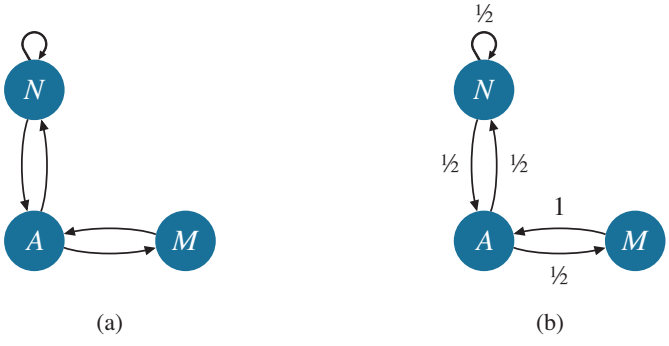

Figure 25.8 *(a) Links between web pages. (b) Corresponding DTMC transition diagram.*

We now solve the stationary equations:

$$\begin{aligned}
\pi_A &= \frac{1}{2}\pi_N + \pi_M \\
\pi_N &= \frac{1}{2}\pi_A + \frac{1}{2}\pi_N \\
\pi_M &= \frac{1}{2}\pi_A \\
1 &= \pi_A + \pi_M + \pi_N.
\end{aligned}$$

This results in: $\pi_A = \pi_N = \frac{2}{5}$; $\pi_M = \frac{1}{5}$.

Intuition behind the PageRank algorithm: Imagine that each page initially has one unit of importance. At each round, each page shares whatever importance it has among its successors. Pages with a lot of incoming links will receive lots of importance (will be visited frequently in the DTMC).

25.10.1 Problems with Real Web Graphs

Unfortunately, PageRank does not work well on all web graphs. Consider the following two examples.

Example 25.23 (Dead end or spider trap)

Consider Figure 25.8(a), where this time there is either no outgoing link from page M (in this case M is called a "dead end") or there is a self-loop at state M (in this case M is called a "spider trap"). In either case, Figure 25.9 shows the corresponding DTMC transition diagram.

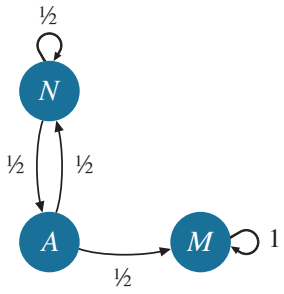

Figure 25.9 *DTMC for a web graph with a dead end or spider trap at M.*

The stationary equations are:

$$\begin{aligned}
\pi_N &= \frac{1}{2}\pi_A + \frac{1}{2}\pi_N \\
\pi_M &= \frac{1}{2}\pi_A + \pi_M \\
\pi_A &= \frac{1}{2}\pi_N \\
\pi_A + \pi_N + \pi_M &= 1.
\end{aligned}$$

The solution to these equations is $\pi_M = 1$, $\pi_N = 0 = \pi_A$. These are also the limiting probabilities (note that the start state does not matter). Somehow this solution is very unsatisfying. Just because person M chooses to be anti-social and not link to anyone else, it should not follow that person M is the only important

person on the web. Our solution does not match our intuitive view of surfing a web graph.

Example 25.24 (Two spider traps)

Now imagine that both M and N are anti-social and link only to themselves. The resulting DTMC transition diagram is shown in Figure 25.10.

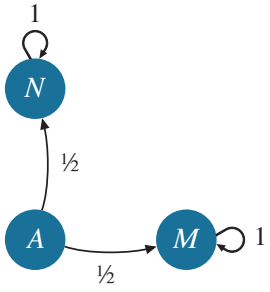

Figure 25.10 *DTMC for a web graph with two spider traps.*

The corresponding stationary equations are:

$$\begin{aligned} \pi_N &= \frac{1}{2} \cdot \pi_A + \pi_N \\ \pi_M &= \frac{1}{2} \cdot \pi_A + \pi_M \\ \pi_A &= 0 \\ \pi_A + \pi_N + \pi_M &= 1. \end{aligned}$$

Again our graph is not irreducible. Observe that there are now an infinite number of possible stationary solutions. This is because the limiting probabilities depend on the start state. Again the solution is very unsatisfying.

25.10.2 Google's Solution to Dead Ends and Spider Traps

Google's initial solution to dead ends and spider traps is to "tax" each page some fraction of its "importance" and then distribute that taxed importance equally among all pages in the web graph. This "tax" keeps the DTMC from getting trapped in a dead end or spider trap.

Figure 25.11 shows the effect of applying a 30% tax on the DTMC of Figure 25.9. First, every original transition is multiplied by 70%. Then, for each state s in an

M-state chain, we add a transition of weight $\frac{30\%}{M}$ from state s to every other state, including itself. Thus in the three-state chain in Figure 25.9, we add a transition of weight 10% from each state to every other state.

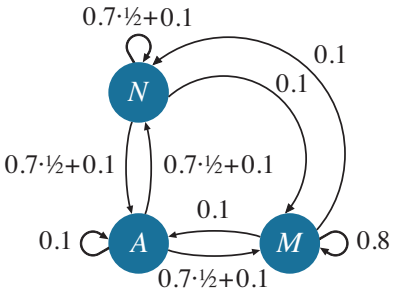

Figure 25.11 *Corresponding DTMC transition diagram.*

Observe that the spider trap is now no longer a problem, and we can easily solve for the limiting probabilities:

$$\pi_A = 0.19 \qquad \pi_M = 0.55 \qquad \pi_N = 0.26.$$

The problem now is that these limiting probabilities are highly dependent on the amount of tax!

25.10.3 Evaluation of the PageRank Algorithm and Practical Considerations

PageRank is intended to give an indication of the popularity of a page. This works well when the graph is irreducible, but it is problematic when there are spider traps or dead ends. The taxation solution for solving the spider trap problem seems ad hoc. If the tax is too small, then we still end up with too high a limiting probability at the spider trap state (as in $\pi_M = 0.55$ in Section 25.10.2). Thus we need to use a high tax. Yet a high tax seems totally unrealistic, because it leads to every state being of equal weight.

There's also the practical consideration: How does Google go about solving the DTMC for the stationary probabilities, given that it is a huge (finite) DTMC? Solving such a large number of simultaneous equations seems difficult.

Question: Is there another approach to obtain the limiting probabilities?

Answer: Yes, we can take powers of $\mathbf{P}$, the transition probability matrix. This turns out to be faster when $\mathbf{P}$ is large and sparse and only an approximate solution is needed. This is the approach employed by Google.

25.11 From Stationary Equations to Time-Reversibility Equations

Thus far, to derive the limiting distribution of a DTMC, we solve the stationary equations. The purpose of this section is to consider a few *alternative systems of equations*. We will introduce "balance equations," which are only a small twist on stationary equations, and then introduce "time-reversibility equations," which are entirely different and sometimes greatly simplify the process.

All this is best illustrated via an example. Consider the DTMC in Figure 25.12 and its corresponding stationary equations.

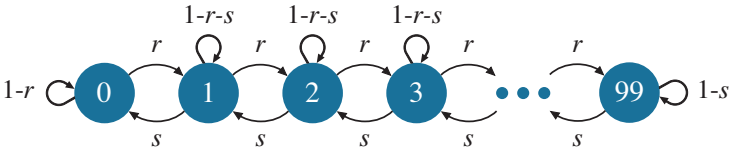

Figure 25.12 *A finite-state DTMC. Assume that* $0 < r, s < 1$.

Stationary equations for DTMC in Figure 25.12:

$$\begin{aligned}
\pi_0 &= \pi_0(1-r) + \pi_1 s \\
\pi_1 &= \pi_0 r + \pi_1(1-r-s) + \pi_2 s \\
\pi_2 &= \pi_1 r + \pi_2(1-r-s) + \pi_3 s \\
&\cdots \\
\pi_i &= \pi_{i-1} r + \pi_i(1-r-s) + \pi_{i+1} s \\
&\cdots \\
\pi_{99} &= \pi_{98} r + \pi_{99}(1-s) \\
\sum_{j=0}^{99} \pi_j &= 1.
\end{aligned}$$

These stationary equations are solvable (see Exercise 25.6), but are cumbersome.

Now consider an alternative to stationary equations, called **balance equations**.

Definition 25.25 *The* **balance equations** *for a Markov chain equate the total rate of leaving each state with the total rate of entering the state.*

Question: For a DTMC with transition probability matrix $\mathbf{P}$, what is the rate of leaving state i?

This may be hard to think about, so let's start with an easier question:

Question: For a DTMC with transition probability matrix $\mathbf{P}$, what is the rate of transitions from state i to state j?

Answer:

$$\pi_i P_{ij} = \text{rate of transitions from state } i \text{ to state } j.$$

To see this, note that the "rate" of transitions from state i to state j is defined as the number of transitions per time step that have their start point in i and end point in j. To understand this quantity, observe that the DTMC is in state i for π_i fraction of all time steps. For P_{ij} fraction of those time steps, the DTMC will next move to state j. Hence, for $\pi_i P_{ij}$ fraction of all time steps, the DTMC is in state i *and* will move to state j in the next transition. Thus, if we look over t time steps (let t be large), then $\pi_i P_{ij} t$ total transitions will have their start point in i and their end point in j. Dividing by t, we see that the *rate* of transitions (number of transitions per time step) that go directly from i to j is $\pi_i P_{ij}$.

Question: So what is the total rate of transitions out of state i?

Answer: The expression $\sum_j \pi_i P_{ij}$ represents the total rate of transitions out of state i, including possibly returning right back to state i (if there are self-loops in the chain). If we want the total rate of transitions out of state i not including returning back to i, then we write: $\sum_{j \neq i} \pi_i P_{ij}$.

Definition 25.26 *The* **balance equations** *for a DTMC with transition matrix* $\mathbf{P}$ *is the set of equations*

$$\sum_{j \neq i} \pi_i P_{ij} = \sum_{j \neq i} \pi_j P_{ji} \qquad \textit{and} \qquad \sum_i \pi_i = 1. \tag{25.9}$$

These hold for every state i. They equate (balance) the rate that we leave state i to go to a state other than i, with the rate that we enter state i from a state other than i.

Balance equations for DTMC in Figure 25.12:

$$\begin{aligned}
\pi_0 r &= \pi_1 s \\
\pi_1 (r+s) &= \pi_0 r + \pi_2 s \\
\pi_2 (r+s) &= \pi_1 r + \pi_3 s \\
&\cdots \\
\pi_i (r+s) &= \pi_{i-1} r + \pi_{i+1} s \\
&\cdots \\
\pi_{99} (s) &= \pi_{98} r \\
\sum_{j=0}^{99} \pi_j &= 1.
\end{aligned}$$

It is easy to see that the balance equations for Figure 25.12 are equivalent to the stationary equations (we've basically just ignored the self-loops in the chain to create simpler equations). Intuitively, the balance equations make sense because every time we leave state i, we cannot again leave state i until we first return to state i.

Theorem 25.27 *Given a DTMC with transition matrix* $\mathbf{P}$*, the balance equations for the DTMC are equivalent to the stationary equations. Thus, satisfying either set of equations is equally good.*

Proof: Recall the stationary equation for state i:

$$\pi_i = \sum_j \pi_j P_{ji}. \tag{25.10}$$

We also know that

$$\pi_i = \pi_i \sum_j P_{ij} = \sum_j \pi_i P_{ij}. \tag{25.11}$$

Combining (25.10) and (25.11), we have:

$$\pi_i = \sum_j \pi_i P_{ij} = \sum_j \pi_j P_{ji}. \tag{25.12}$$

We now subtract $\pi_i P_{ii}$ from both sides of (25.12):

$$\sum_j \pi_i P_{ij} - \pi_i P_{ii} = \sum_j \pi_j P_{ji} - \pi_i P_{ii}$$

$$\sum_{j \neq i} \pi_i P_{ij} = \sum_{j \neq i} \pi_j P_{ji}.$$

Hence we obtain the balance equations. ■

Balance equations can also be applied to a set of states as well as to a single state. For example, if a Markov chain is divided into two sets of states – call these S and S^c (here S^c denotes the complement of S) – then we can write equations equating the rate of transitions (the "flux") from S to S^c with the rate of transitions from S^c to S.

Question: Why does it make sense that the total flux from S to S^c should equal that from S^c to S?

Answer: The argument is identical to what we observed for a single state. Every time a transition takes us from S to S^c, we have left the states in S. We therefore cannot have another transition from S to S^c until we reenter the states in S, but this requires a transition from S^c to S.

We now return to the DTMC in Figure 25.12, and try to write even simpler equations. Such equations are referred to as **time-reversibility equations**.

Definition 25.28 *The* **time-reversibility equations** *for a DTMC with transition matrix* $\mathbf{P}$ *is the set of equations*

$$\pi_i P_{ij} = \pi_j P_{ji}, \quad \forall i, j \qquad \text{and} \qquad \sum_i \pi_i = 1. \tag{25.13}$$

These equations apply to every pair of states, i, j*. Specifically, there is one equation written for each pair of state,* i, j*. They equate the rate of transitions from* i *to* j *with the rate of transitions from* j *to* i*.*

Time-reversibility equations for DTMC in Figure 25.12:

$$\begin{aligned}
\pi_0 r &= \pi_1 s \\
\pi_1 r &= \pi_2 s \\
\pi_2 r &= \pi_3 s \\
&\cdots \\
\pi_i r &= \pi_{i+1} s \\
&\cdots \\
\pi_{98} r &= \pi_{99} s \\
\sum_{j=0}^{99} \pi_j &= 1.
\end{aligned}$$

The time-reversibility equations are *much* simpler than the stationary equations.

Question: Are the time-reversibility equations above equivalent to the stationary equations or balance equations that we've seen?

Answer: No!

While the time-reversibility equations look very different from the stationary and balance equations, it turns out that they *do* yield the correct stationary distribution for the chain in Figure 25.12. This seems impossible, but try it!

Question: Given an aperiodic, irreducible DTMC, are the time-reversibility equations always satisfied?

Answer: No.

Question: What's an example of a chain where the time-reversibility equations are not satisfied?

Answer: Imagine a chain which is irreducible, but where there is an edge from

i to j, but no edge from j to i. Then the rate of transitions from j to i is by definition 0, although the rate of transitions from i to j is non-zero.

In Theorem 25.29, we prove that *if* we can find π_i's that satisfy the time-reversibility equations, then those π_i's are the stationary probabilities. In that case, we say that the chain is called "time-reversible." If we can't find π_i's that satisfy the time-reversibility equations, this does not imply that there's no stationary distribution. It just means that we have to start from scratch with the (more complicated) stationary equations.

Theorem 25.29 (Time-reversibility implies stationarity) *For a DTMC with transition matrix* $\mathbf{P}$*, suppose we can find* $x_0, x_1, x_2, \ldots$ *such that,* $\forall i, j$*:*

$$x_i P_{ij} = x_j P_{ji} \qquad \text{and} \qquad \sum_i x_i = 1. \tag{25.14}$$

Then the vector $\vec{x} = (x_0, x_1, x_2, \ldots)$ *is a stationary distribution, and we say that the DTMC is* **time-reversible**.

Proof:

$$\begin{aligned} & x_i P_{ij} = x_j P_{ji}, \qquad \forall i, j \\ \Rightarrow & \sum_i x_i P_{ij} = \sum_i x_j P_{ji} \\ \Rightarrow & \sum_i x_i P_{ij} = x_j \sum_i P_{ji} \\ \Rightarrow & \sum_i x_i P_{ij} = x_j. \end{aligned}$$

Hence, together with $\sum_i x_i = 1$, the x_j's satisfy the stationary equations. ■

Remark 1: In some books, the definition of "time-reversible" requires additionally that the chain be ergodic, but we won't be making ergodicity a requirement.
Remark 2: Theorem 25.29 does not require that the number of states is finite.

Question: The time-reversibility equations are much simpler than the stationary or balance equations, but they aren't always solvable. For the chain in Figure 25.12, the time-reversibility equations had a solution. What was special about this chain?

Answer: The chain in Figure 25.12 has the property that the rate of transitions from state i to state j is always equal to the rate of transitions from state j to state i. To see this, notice first that if j is anything other than $i+1$ or $i-1$, then the rate of transitions from i to j is zero, and, likewise, the rate of transitions from j to i is zero. Now suppose $j = i+1$. The *number* of transitions from i to

$i+1$ during time t is the same (within 1) of the *number* of transitions from $i+1$ to i. This is because every time we go from state i to $i+1$, we can't repeat that transition until we first go from $i+1$ to i. This translates to the *rates* being the same when we divide by time.

As we'll see in the exercises (see, for example, Exercises 25.9 and 25.18) there are plenty of Markov chains that are time-reversible, but it is not always easy to guess in advance which chains will have this beautiful property. When trying to determine the stationary solution, you *first try to solve the time-reversibility equations*. If those yield a solution, then you're done (your solution also satisfies the stationary equations). If the time-reversibility equations are not solvable, then you'll need to try solving the stationary or balance equations.

Question: A final reminder: Solving the stationary equations, or balance equations, or time-reversibility equations, yields a stationary distribution. What does that tell us about the limiting distribution?

Answer: The fact that we have a stationary distribution, $\vec{\pi}$, does not tell us anything about whether a limiting distribution exists. However, if we have a finite-state, irreducible, aperiodic DTMC, then, by Theorem 25.19, $\vec{\pi}$ is also the limiting distribution.

25.12 Exercises

25.1 **Two finite-state chains**

Figure 25.13 depicts two finite-state chains. For each chain, answer the questions below.

Chain 1 Chain 2

Figure 25.13 *Two finite-state chains for Exercise 25.1.*

(a) Is the DTMC irreducible?
(b) Is the DTMC aperiodic?

(c) Does the DTMC have a limiting distribution? If so, what is it? If not, why not?
(d) Does the DTMC have a stationary distribution? If so, what is it? If not, why not?

Explain each answer by citing the appropriate theorems.

25.2 **Passing around a ball**

[Proposed by Sam Yeom] In answering these questions, cite the theorems that you use in making your claims.

(a) Five people stand in a circle, passing a ball around. Suppose that each person either passes the ball right or left with 50% probability each. What is the stationary distribution? Is this also the limiting distribution?
(b) Five people stand in a circle, passing a ball around. Suppose that each person passes the ball to their right with probability 1. What is the stationary distribution? Is this also the limiting distribution?
(c) Now suppose that the five people are standing in a line. Each person passes the ball to their right or left with 50% probability each, except for the two people at the ends who always pass it to their one neighbor. What is the stationary distribution? Is this also the limiting distribution?
(d) Again the five people are standing in a line. Again each person passes the ball to their right or left with 50% probability each, except for the two people at the ends who always hold on to the ball instead of passing it. What is the stationary distribution? Is this also the limiting distribution?

25.3 **Multiple stationary distributions**

Ishani's finite-state DTMC has multiple stationary distributions. We do not know whether the chain is aperiodic or irreducible. What can we conclude?

(a) Ishani's DTMC has multiple limiting distributions.
(b) Ishani's DTMC has no limiting distribution.
(c) Ishani's DTMC has exactly one limiting distribution.
(d) We can't conclude any of these for sure.

Provide full justification for your answer by citing the appropriate theorems.

25.4 **Practice with the definitions**

Consider each of the two simple DTMCs shown in Figure 25.14. For each chain, please answer the following questions. Justify your answers by citing theorems.

(a) Is the chain aperiodic?
(b) Is the chain irreducible?

(c) Is the chain ergodic?
(d) Does the limiting distribution exist? If so, what is it? If not, why not?
(e) Does one or more stationary distributions exist? If so, what are the stationary distribution(s)?
(f) Is p_c, the time-average fraction of time spent in state c, well defined? If so, what is it?
(g) Consider m_{cc}, the mean time until we again visit state c, given we are in state c. Is m_{cc} well-defined? If so, what is it?

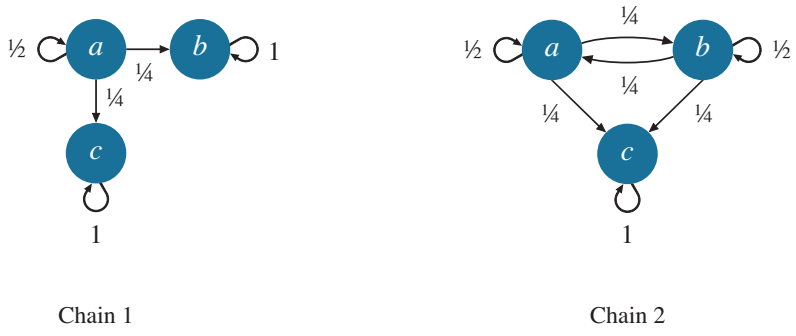

Figure 25.14 *Chains for Exercise 25.4.*

25.5 **Proof of Corollary 25.13**
Prove Corollary 25.13.

25.6 **A simple finite-state chain**
For the DTMC shown in Figure 25.15, explain how we know that the limiting distribution exists by citing theorems from the chapter. Then

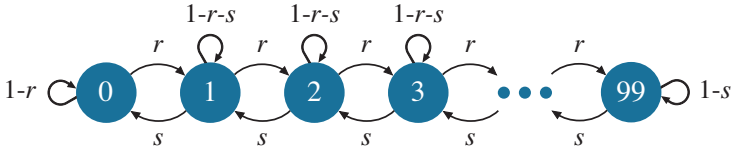

Figure 25.15 *Chain for Exercise 25.6.*

solve for the limiting distribution via these steps:
(a) From the stationary equations, express π_1 in terms of π_0. Then express π_2 in terms of π_0.
(b) You will notice a pattern that will help you make a guess for how to express π_i in terms of π_0 for any i.
(c) Determine π_0 by using $\sum_i \pi_i = 1$ and verify the correctness of your guess.

25.7 **Some example DTMCs**

For each chain shown in Figure 25.16, answer the following questions:

(a) Is the chain irreducible?
(b) Is the chain aperiodic?
(c) Does a stationary distribution exist?
(d) Does the limiting distribution exist?

Provide a one-line explanation for your answer, citing theorems.

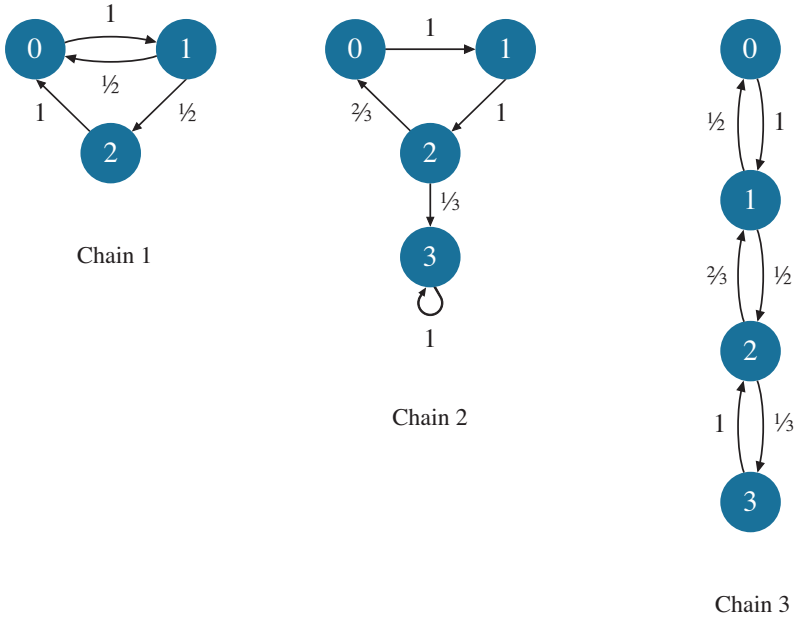

Figure 25.16 *Markov chains for Exercise 25.7.*

25.8 **Caching**

If you think about it, web browsing is basically a Markov chain – the page you will go to next depends on the page you are currently at. Suppose our web server has three pages, and we have the following transition probabilities:

$$\begin{array}{lll} P_{1,1} = 0 & P_{1,2} = x & P_{1,3} = 1 - x \\ P_{2,1} = y & P_{2,2} = 0 & P_{2,3} = 1 - y \\ P_{3,1} = 0 & P_{3,2} = 1 & P_{3,3} = 0, \end{array}$$

where $P_{i,j}$ represents the probability that I will next request page j, given that I last requested page i. Assume that $0 < x < y < \frac{1}{2}$.

Recall that web browsers cache pages so that they can be quickly retrieved later. We will assume that the cache has enough memory to store two pages. Whenever a request comes in for a page that is not cached, the browser will store that page in the cache, replacing the page *least likely* to

be referenced next based on the current request. For example, if my cache contained pages {2,3} and I requested page 1, the cache would now store {1,3} (because $x < 1 - x$).

(a) Find the proportion of time that the cache contains the following pages: (i) {1,2} (ii) {2,3} (iii) {1,3}. [Hint 1: You will need to think carefully about what information you need in your states to create the appropriate DTMC.] [Hint 2: When solving your DTMC, you will find that two of the states are only visited a finite number of times, with probability 1, so the long-run fraction of time spent there is 0. You can thus ignore these states and just solve for the stationary probabilities of the remaining states.]

(b) Find the proportion of requests that are for cached pages.

25.9 **Practice with balance equations and time-reversibility equations**

Consider the following Markov chains:

$$\mathbf{P}^{(1)} = \begin{pmatrix} 0 & 2/3 & 0 & 1/3 \\ 1/3 & 0 & 2/3 & 0 \\ 0 & 1/3 & 0 & 2/3 \\ 2/3 & 0 & 1/3 & 0 \end{pmatrix}$$

$$\mathbf{P}^{(2)} = \begin{pmatrix} 1/3 & 2/3 & 0 & 0 \\ 1/3 & 0 & 2/3 & 0 \\ 0 & 1/3 & 0 & 2/3 \\ 0 & 0 & 1/3 & 2/3 \end{pmatrix}.$$

(a) Draw the corresponding Markov chains for $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$.

(b) Solve for the time-average fraction of time spent in each state for both $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$. First try to use the time-reversibility equations, and if they do not work, then use the balance equations.

(c) Was $\mathbf{P}^{(1)}$ time-reversible? Was $\mathbf{P}^{(2)}$ time-reversible?

(d) For those chain(s) that were time-reversible, explain why it makes sense that for all states i, j in the chain, the rate of transitions from i to j should equal the rate of transitions from j to i.

25.10 **Data centers, backhoes, and bugs**

Our data center alternates between "working" and "down." There are two reasons why our data center can be down: (1) a backhoe accidentally dug up some cable, or (2) a software bug crashed the machines. Suppose that if the data center is working today, it will be down tomorrow due to backhoe reasons with probability $\frac{1}{6}$ or will be down tomorrow due to a software bug with probability $\frac{1}{4}$. A data center that is down today due to backhoe reasons will be up tomorrow with probability 1. A data center that is down today due to a software bug will be up tomorrow with probability $\frac{3}{4}$.

(a) Draw a DTMC for this problem.

(b) Is your DTMC ergodic? Why or why not?
(c) Is your DTMC time-reversible? Why or why not?
(d) What fraction of time is the data center working?
(e) What is the expected number of days between backhoe failures?

25.11 **CLT versus SLLN**

Consider a sequence of i.i.d. random variables $X_1, X_2, \ldots$ with finite mean $\mathbf{E}[X]$ and finite variance σ. Let $S_n = \sum_{i=1}^{n} X_i$. Now consider the quantity:

$$\frac{S_n - n\mathbf{E}[X]}{n}.$$

What does the Strong Law of Large Numbers (SLLN) say about this quantity as $n \to \infty$? What does the Central Limit Theorem (CLT) say about this quantity as $n \to \infty$? Are they in contradiction?

25.12 **Walks on undirected weighted graphs**

This problem comes up in many areas. Consider any undirected connected graph with weights: $w_{ij} = w_{ji}$ is the weight on edge (i, j) where $w_{ij} \geq 0$, $\forall i, j$. See for example Figure 25.17. A particle moves between nodes in a weighted graph as follows: A particle residing at node i will next move to node j with probability P_{ij}, where

$$P_{ij} = \frac{w_{ij}}{\Sigma_j w_{ij}}.$$

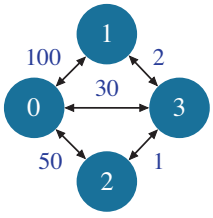

Figure 25.17 *A weighted graph with $M = 4$ nodes describing a particle's motion.*

Your goal is to determine the long-run proportion of time that the particle is in state i.

(a) Play around with the example in Figure 25.17. Which node do you think is visited most often?
(b) You'll now need to guess a solution for a *general weighted graph* and show that your solution satisfies the stationary equations. It will help a lot, both in making your guess and in verifying your guess, if you write out the time-reversibility equations rather than the stationary equations.

25.13 **Finite-state chain with equal weights**
Consider the finite-state chain in Figure 25.18.

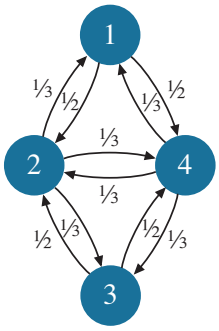

Figure 25.18 *Markov chain for Exercise 25.13.*

(a) Is the chain ergodic?
(b) Is the chain time-reversible?
(c) What is the limiting probability of being in each state?
(d) The finite-state chain in Figure 25.18 has two properties:
  (i) **Balanced weights**: This is the property that the probabilities on each of the arrows leaving a state are equal.
  (ii) **Bidirectional edges**: This is the property that if there's an edge from i to j, then there's also an edge from j to i.

  Look at the structure of the limiting probabilities that you obtained for Figure 25.18. To see the structure, it will help to write these over the same common denominator. Now imagine an *arbitrary* ergodic finite-state chain with n states that has both the "balanced weights" property and the "bidirectional edges" property. What can you say about π_j, the limiting probability of being in state j? Make a guess and verify it.

25.14 **Gas migration**
You have a box with n gas molecules, with a divider in the middle that the molecules can pass through. As shown in Figure 25.19, there is an A side and a B side to the box. Assume that n is even.
All the molecules start out on the A side. Every second, we pick a random gas molecule out of the n molecules and transfer it to the other side.
(a) Determine the proportion of time that the box has the same number of molecules on the A side and the B side. Start by drawing a Markov chain!
(b) Let N_B denote the number of molecules in the B side of the box.
  (i) What is $\mathbf{E}[N_B]$?
  (ii) What is $\mathbf{Var}(N_B)$?

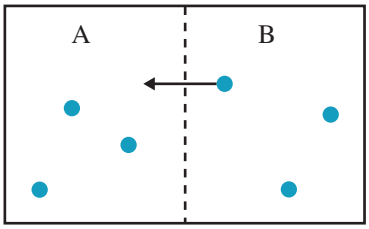

Figure 25.19 *Box of gas molecules for Exercise 25.14.*

[Hint: What do you know about N_B's distribution over a long time?]

25.15 **Randomized chess**

This problem concerns the behavior of various chess pieces as they move randomly around the board. Chess is played on an 8×8 board divided into 64 squares that alternate from white to black. The *king* can move one square in any direction (including the diagonal). The *bishop* can move any number of squares, but only in the diagonal directions. The *knight* moves in an L-shape. That is, the knight moves two squares to either side (left or right) and one square up or down. Or, the knight can move two squares up or down and one square to the side (left or right).

(a) You are given an empty chessboard with a lone king placed in one corner. At each time step, the king will make a uniformly random legal move. Is the corresponding Markov chain for this process irreducible? Is it aperiodic?
(b) What if a bishop is used instead?
(c) What if a knight is used instead?
(d) Now take advantage of Exercise 25.12 on undirected weighted graphs and time-reversibility to calculate the expected time for the king to return to the corner. Think about how hard this would be without time-reversibility. [Hint: The calculation should be very simple.]
(e) Do the same for the bishop.
(f) Do the same for the knight.

25.16 **Interpreting the stationary probabilities as fractions of time**

Assume that you have an irreducible, finite-state DTMC with M states (numbered $0, 1, \ldots, M-1$) and transition matrix $\mathbf{P}$.

Define

$$\phi_j = \lim_{n\to\infty} \frac{\sum_{i=1}^{n} p_j(i)}{n},$$

where $p_j(i)$ is the probability that the chain is in state j at time step i.

(a) What is the meaning of ϕ_j? Please follow these steps:
  (i) What does $\sum_{i=1}^{n} p_j(i)$ mean?

(ii) What does $\frac{\sum_{i=1}^{n} p_j(i)}{n}$ mean?

(iii) What does $\phi_j = \lim_{n\to\infty} \frac{\sum_{i=1}^{n} p_j(i)}{n}$ mean?

(b) Prove that the distribution $\vec{\phi} = (\phi_0, \phi_1, \ldots, \phi_{M-1})$ is a stationary distribution. Please follow these steps:

(i) Express $p_j(i)$ in terms of a sum involving $p_k(i-1)$.

(ii) Show that ϕ_j satisfies the stationary equations.

(iii) Don't forget to prove the needed condition on $\sum_{j=0}^{M-1} \phi_j$.

25.17 **In an irreducible DTMC, do all states have the same period?**

Given an irreducible DTMC, either prove that all states have the same period, or find a counter-example.

25.18 **How rare are time-reversible DTMCs?**

Edward feels that time-reversible chains are very rare. Erica disagrees. Erica claims that it's easy to create time-reversible chains, via the idea of Exercise 25.12.

(a) Consider the DTMC in Figure 25.20 whose transitions are unlabeled. Use what you've learned in Exercise 25.12 to label each edge (i, j) of the DTMC with a transition probability p_{ij} such that $0 < p_{ij} < 1$ and such that the DTMC is time-reversible. Then write the *limiting distribution* of your chain.

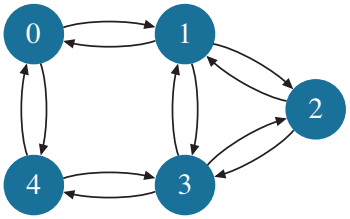

Figure 25.20 *Markov chain for Exercise 25.18.*

(b) How many possible answers are there to question (a)? That is, how many choices of transition probabilities are there that create a time-reversible DTMC? Pick the correct answer and give a one-line explanation:

(i) exactly one

(ii) a finite number

(iii) countably infinite

(iv) uncountably infinite

25.19 **Irreducible finite-state chains have finite mean time to return**

Prove Theorem 25.11: For a finite-state, irreducible DTMC, m_{ij} is finite, for every i, j.

25.20 **Every finite DTMC has at least one stationary distribution**

[Proposed by Misha Ivkov] In this problem we will prove that every *finite-state* DTMC has at least one stationary distribution. Note, we are not making any assumptions about the DTMC.

(a) First, prove that a finite DTMC must have at least one recurrent state.
   (i) Let i be a state in the Markov chain. Argue that there exists some state j such that $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ij} = \infty$.
   (ii) Now argue that $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{jj} = \infty$.

(b) Let j be the recurrent state identified above, and let S be the set of states that are accessible from j. Show that S is an irreducible DTMC.

(c) As explained in Theorem 25.20, since S is irreducible, we know that it has a stationary distribution; let's call that $\vec{\pi}'$. We now define $\vec{\pi}$ as

$$\pi_i = \begin{cases} \pi'_i & \text{if } i \in S \\ 0 & \text{otherwise} \end{cases}.$$

Prove that $\vec{\pi}$ is a stationary distribution for the original DTMC.