13

Introduction to Stationary Distributions

We first briefly review the classification of states in a Markov chain with a quick example and then begin the discussion of the important notion of *stationary distributions*.

First, let's review a little bit with the following

Example: Suppose we have the following transition matrix:

$$
\mathbf{P} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10 \end{array}
\begin{array}{c}
\begin{array}{cccccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \end{array} \\
\left[\begin{array}{cccccccccc}
 & & & & & & 1 & & & \\
 & .3 & .3 & .1 & .3 & & & & & \\
 & & .6 & & & & & .4 & & \\
 & & & 1 & & & & & & \\
 & .4 & & & & .3 & & .3 & & \\
 & & & .9 & & & .1 & & & \\
 & & & & & & & & & 1 \\
 & .8 & & & & & & .2 & & \\
 & & & & & & & & 1 & \\
1 & & & & & & & & &
\end{array}\right]
\end{array}.
$$

Determine the equivalence classes, the period of each equivalence class, and whether each equivalence class it transient or recurrent.

Solution: The state space is small enough (10 elements) that one effective way to determine classes is to just start following possible paths. When you see 1's in the matrix a good place to start is in a state with a 1 in the corresponding row. If we start in state 1, we see that the path $1 \to 7 \to 10 \to 1$ must be followed with probability 1. This immediately tells us that the set $\{1, 7, 10\}$ is a recurrent class with period 3. Next, we see that if we start in state 9, then we just stay there forever. Therefore, $\{9\}$ is a recurrent class with period 1. Similarly, we can see that $\{4\}$ is a recurrent class with period 1. Next suppose we start in state 2. From state 2 we can go directly to states 2, 3, 4 or 5. We also see that from state 3, we can get to state 2 (by the path $3 \to 8 \to 2$) and from state 5 we can get to state 2 (directly). Therefore, state 2 communicates with states 3 and 5. We don't need to check if state 2 communicates with states 1, 4, 7, 9, or 10 (why?). From state 2 we can get to state 6 (by the path $2 \to 5 \to 6$) but from state 6 we must go to either state 4 or state 7, therefore from state 6 we cannot get to state 2. Therefore, state 2 and 6 do not communicate. Finally, we can see that states 2 and 8 do communicate. Therefore, $\{2, 3, 5, 8\}$ is an equivalence class. It is transient because from this class we can get to state 4 (and never come back). Finally, it's period is 1 because the period of state 2 is clearly 1 (we can start in state 2 and come back to state 2 in 1 step). The only state left that is still unclassified is state 6, which is in a class by itself $\{6\}$ and is clearly transient. Note that $p_{66}(n) = 0$ for all $n > 0$ so the set of times at which we could possibly return to state 6 is the empty set. By convention, we will say that the greatest common divisor of the empty set is infinity, so the period of state 6 is infinity. □

Sometimes a useful technique for determining the equivalence classes in a Markov chain is to draw what is called a *state transition diagram*, which is a graph with one node for each state and with a (directed) edge between nodes i and j if $p_{ij} > 0$. We also usually write the transition probability p_{ij} beside the directed edge between nodes i and j if $p_{ij} > 0$. For example, here is the state transition diagram for the previous example.

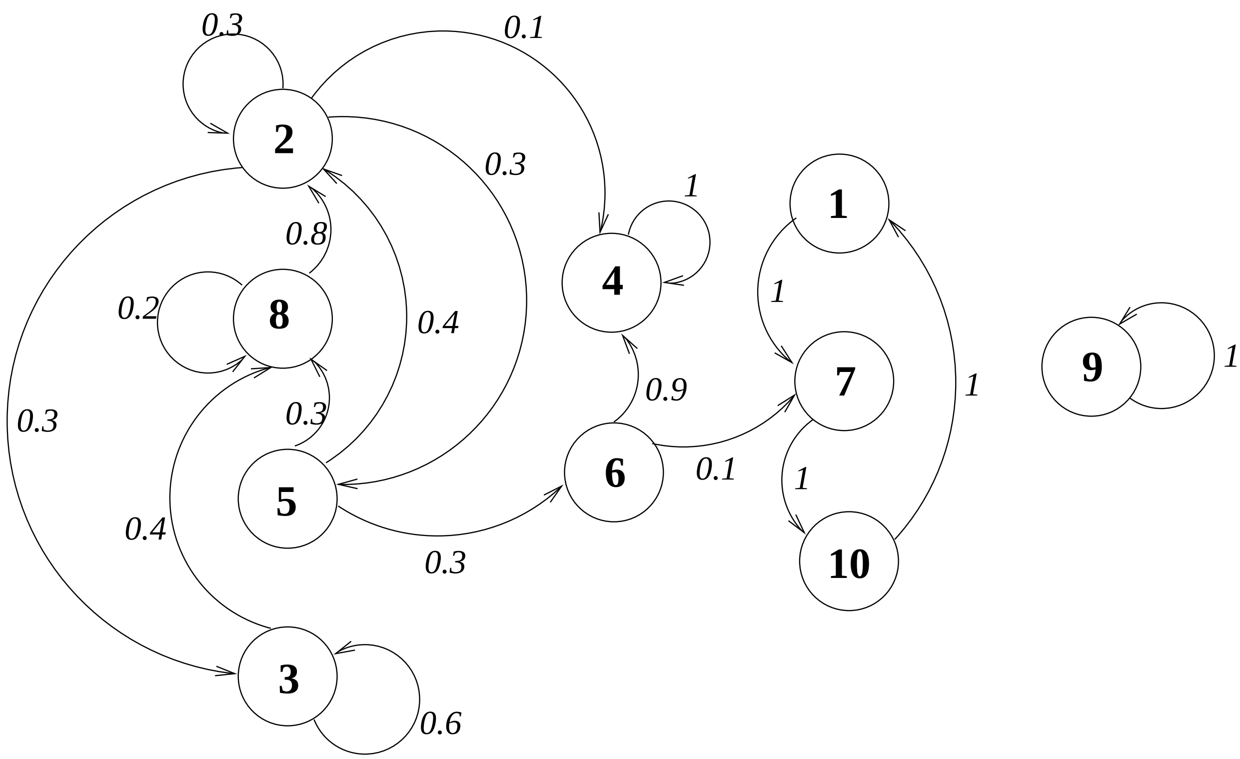

Figure 13.1: State Transition Diagram for Preceding Example

Since the diagram displays all one-step transitions pictorially, it is usually easier to see the equivalence classes with the diagram than just by looking at the transition matrix. It helps if the diagram can be drawn neatly, with, for example, no edges crossing each other.

Usually when we construct a Markov model for some system the equivalence classes, if there are more than one, are apparent or obvious because we *designed* the model so that certain states go together and we designed them to be transient or recurrent.

Other times we may be trying to verify, modify, improve, or just understand someone else's (complicated) model and one of the first things we may want to know is how to classify the states, and it may not be obvious or even easy to determine the equivalence classes if the state space is large and there are many transitions that don't follow a regular pattern. For S finite, the following algorithm determines $T(i)$, the set of states accessible from i, $F(i)$, the set of states from which i is accessible, and $C(i) = F(i) \bigcap T(i)$, the equivalence class of state i, for each state i:

1. For each state $i \in S$, let $T(i) = \{i\}$ and $F(i) = \{\phi\}$, the empty set.
2. For each state $i \in S$, do the following: For each state $k \in T(i)$, add to $T(i)$ all states j such that $p_{kj} > 0$ (if k is not already in $T(i)$. Repeat this step until no further addition is possible.
3. For each state $i \in S$, do the following: For each state $j \in S$, add state j to $F(i)$ if state i is in $T(j)$.
4. For each state $i \in S$, let $C(i) = F(i) \bigcap T(i)$.

Note that if $C(i) = T(i)$ (the equivalence class containing i equals the set of states that are accessible from i), then $C(i)$ is *closed* (hence recurrent since we are assuming S is finite for this algorithm). This algorithm is taken from *An Introduction to Stochastic Processes*, by Edward P. C. Kao, Duxbury Press, 1997. Also in this reference is the listing of a `MATLAB` implementation of this algorithm.

Stationary Markov Chains

Now that we know the general architecture of a Markov chain, it's time to look at how we might analyse a Markov chain to make predictions about system behaviour. For this we'll first consider the concept of a *stationary distribution*. This is distinct from the notion of limiting probabilities, which we'll consider a bit later. First, let's define what we mean when we say that a process is *stationary*.

Definition: A (discrete-time) stochastic process $\{X_n : n \geq 0\}$ is *stationary* if for any time points $i_1, \ldots, i_n$ and any $m \geq 0$, the joint distribution of $(X_{i_1}, \ldots, X_{i_n})$ is the same as the joint distribution of $(X_{i_1+m}, \ldots, X_{i_n+m})$.

So "stationary" refers to "stationary in time". In particular, for a stationary process, the distribution of X_n is the same for all n.

So why do we care if our Markov chain is stationary? Well, if it were stationary *and* we knew what the distribution of each X_n was then we would know a lot because we would know the long run proportion of time that the Markov chain was in any state. For example, suppose that the process was stationary and we knew that $P(X_n = 2) = 1/10$ for every n. Then over 1000 time periods we should expect that roughly 100 of those time periods was spent in state 2, and over N time periods roughly $N/10$ of those time periods was spent in state 2. As N went to infinity, the *proportion* of time spent in state 2 will converge to $1/10$ (this can be proved rigorously by some form of the Strong Law of Large Numbers). One of the attractive features of Markov chains is that we can often make them stationary *and* there is a nice and neat characterization of the distribution of X_n when it is stationary. We discuss this next.

Stationary Distributions

So how do we make a Markov chain stationary? If it can be made stationary (and not all of them can; for example, the simple random walk cannot be made stationary and, more generally, a Markov chain where all states were transient or null recurrent cannot be made stationary), then making it stationary is simply a matter of choosing the right initial distribution for X_0. If the Markov chain is stationary, then we call the common distribution of all the X_n the *stationary distribution* of the Markov chain.

Here's how we find a stationary distribution for a Markov chain.

Proposition: Suppose $\boldsymbol{X}$ is a Markov chain with state space S and transition probability matrix $\mathbf{P}$. If $\boldsymbol{\pi} = (\pi_j, j \in S)$ is a distribution over S (that is, $\boldsymbol{\pi}$ is a (row) vector with $|S|$ components such that $\sum_j \pi_j = 1$ and $\pi_j \geq 0$ for all $j \in S$), then setting the initial distribution of X_0 equal to $\boldsymbol{\pi}$ will make the Markov chain stationary with stationary distribution $\boldsymbol{\pi}$ if

$$\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$$

That is,

$$\pi_j = \sum_{i \in S} \pi_i p_{ij} \quad \text{for all } j \in S.$$

In words, π_j is the dot product between $\boldsymbol{\pi}$ and the jth *column* of $\mathbf{P}$.

Proof: Suppose $\boldsymbol{\pi}$ satisfies the above equations and we set the distribution of X_0 to be $\boldsymbol{\pi}$. Let's set $\boldsymbol{\mu}(n)$ to be the distribution of X_n (that is, $\boldsymbol{\mu}_j(n) = P(X_n = j)$). Then

$$\begin{aligned} \boldsymbol{\mu}_j(n) &= P(X_n = j) = \sum_{i \in S} P(X_n = j | X_0 = i) P(X_0 = i) \\ &= \sum_{i \in S} p_{ij}(n) \pi_i, \end{aligned}$$

or, in matrix notation,

$$\boldsymbol{\mu}(n) = \boldsymbol{\pi} \mathbf{P}(n).$$

But, by the Chapman-Kolmogorov equations, we get

$$\begin{aligned} \boldsymbol{\mu}(n) &= \boldsymbol{\pi} \mathbf{P}^n \\ &= (\boldsymbol{\pi} \mathbf{P}) \mathbf{P}^{n-1} \\ &= \boldsymbol{\pi} \mathbf{P}^{n-1} \\ &\vdots \\ &= \boldsymbol{\pi} \mathbf{P} \\ &= \boldsymbol{\pi} \end{aligned}$$

We'll stop the proof here. □

Note we haven't fully shown that the Markov chain $\boldsymbol{X}$ is stationary with this choice of initial distribution $\boldsymbol{\pi}$ (though it is and not too difficult to show). But we have shown that by setting the distribution of X_0 to be $\boldsymbol{\pi}$, the distribution of X_n is also $\boldsymbol{\pi}$ for all $n \geq 0$, and this is enough to say that π_j can be interpreted as the long run proportion of time the Markov chain spends in state j (if such a $\boldsymbol{\pi}$ exists). We also haven't answered any questions about the existence or uniqueness of a stationary distribution. But let's finish off today with some examples.

Example: Consider just the recurrent class $\{1, 7, 10\}$ in our first example today. The transition matrix for this class is

$$\mathbf{P} = \begin{array}{c} \\ 1 \\ 7 \\ 10 \end{array} \begin{array}{c} \begin{array}{ccc} 1 & 7 & 10 \end{array} \\ \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \end{array}.$$

Intuitively, the chain spends one third of its time in state 1, one third of its time in state 7, and one third of its time in state 10. One can easily verify that the distribution $\boldsymbol{\pi} = (1/3, 1/3, 1/3)$ satisfies $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, and so $(1/3, 1/3, 1/3)$ is a stationary distribution. □

Remark: Note that in the above example, $p_{ii}(n) = 0$ if n is not a multiple of 3 and $p_{ii} = 1$ if n is a multiple of 3, for all i. Thus, clearly $\lim_{n\to\infty} p_{ii}(n)$ does not exist because these numbers keep jumping back and forth between 0 and 1. This illustrates that limiting probabilities are not exactly the same thing as stationary probabilities. We want them to be! Later we'll give just the right conditions for these two quantities to be equal.

Example: (Ross, p.257 #30). Three out of every four trucks on the road are followed by a car, while only one out of every five cars is followed by a truck. What fraction of vehicles on the road are trucks?

Solution: Imagine sitting on the side of the road watching vehicles go by. If a truck goes by the next vehicle will be a car with probability 3/4 and will be a truck with probability 1/4. If a car goes by the next vehicle will be a car with probability 4/5 and will be a truck with probability 1/5. We may set this up as a Markov chain with two states 0=truck and 1=car, and transition probability matrix

$$\mathbf{P} = \begin{array}{c} \\ 0 \\ 1 \end{array} \begin{array}{c} \begin{array}{cc} 0 & \;\;1 \end{array} \\ \left[\begin{array}{cc} 1/4 & 3/4 \\ 1/5 & 4/5 \end{array}\right] \end{array}.$$

The equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ are

$$\pi_0 = \frac{1}{4}\pi_0 + \frac{1}{5}\pi_1 \quad \text{and} \quad \pi_1 = \frac{3}{4}\pi_0 + \frac{4}{5}\pi_1.$$

Solving, we have from the first equation that $(3/4)\pi_0 = (1/5)\pi_1$, or $\pi_0 = (4/15)\pi_1$. Plugging this into the constraint that $\pi_0 + \pi_1 = 1$ gives us that $(4/15)\pi_1 + \pi_1 = 1$, or $(19/15)\pi_1 = 1$, or $\pi_1 = 15/19$. Therefore, $\pi_0 = 4/19$. That is, as we sit by the side of the road, the long run proportion of vehicles that will be trucks is $4/19$. $\square$

Remark: Note that we need the constraint that $\pi_0 + \pi_1 = 1$ in order to determine a solution. In general, we need the constraint that $\sum_{j \in S} \pi_j = 1$ in order to determine a solution. This is because the system of equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ has just in itself infinitely many solutions (if $\boldsymbol{\pi}$ is a solution then so is $c\boldsymbol{\pi}$ for any constant c). We need the normalization constraint basically to determine c to make $\boldsymbol{\pi}$ a proper distribution over S.

14

Existence and Uniqueness

We now begin to answer some of the main theoretical questions concerning Markov chains. The first, and perhaps most important, question is under what conditions does a stationary distribution exist, and if it exists is it unique? In general a Markov chain can have more than one equivalence class. There are really only 3 combinations of equivalence classes that we need to consider. These are 1) when there is only one equivalence class, 2) when there are two or more classes, all transient, and 3) when there are two or more classes with some transient and some recurrent. As we have mentioned previously when there are two or more classes and they are all recurrent, we can assume that the whole state space is the class that we start the process in, because such classes are closed. We will consider case (3) when we get to Section 4.6 in the text and we will not really consider case (2), as this does not arise very much in practice. Our main focus will be on case (1). When there is only one equivalence class we say the Markov chain is *irreducible*.

We will show that for an irreducible Markov chain, a stationary distribution exists if and only if all states are positive recurrent, and in this case the stationary distribution is unique.

We will start off by showing that if there is at least one recurrent state in our Markov chain, then there exists a solution to the equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, and we will demonstrate that solution by constructing it.

First we'll try to get an intuitive sense of the construction. The basic property of Markov chains can be described as a *starting over* property. If we fix a state k and start out the chain in state k, then every time the chain returns to state k it starts over in a probabilistic sense. We say that the chain *regenerates* itself. Let us call the time that the chain spends moving about the state space from the initial time 0, where it starts in state k, to the time when it first returns to state k, a *sojourn* from state k back to state k. Successive sojourns all "look the same" and so what the chain does during one sojourn should, on average at least, be the same as what it does on every other sojourn. In particular, for any state $i \neq k$, the number of times the chain visits state i during a sojourn should, again on average, be the same as in every other sojourn. If we accept this, then we should accept that the *proportion of time* during a sojourn that the chain spends in state i should be the same, again on average, for all sojourns. But this reasoning then leads us to expect that the proportion of time that the chain spends in state i *over the long run* should be the same as the proportion of time that the chain spends in state i during any sojourn, in particular the first sojourn from state k back to state k. But this is also how we interpret π_i, the stationary probability of state i, as the long run proportion of time the chain spends in state i. So this is how we will construct a vector to satisfy the equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$. We will let the ith component of our solution be the expected number of visits to state i during the first sojourn. This should be proportional to a stationary distribution, if such a distribution exists.

Let us first set our notation. Define

$$\begin{aligned} T_k &= \text{first time the chain visits state } k \text{, starting at time 1,} \\ N_i &= \text{the number of visits to state } i \text{ during the first sojourn,} \\ \rho_i(k) &= E[N_i | X_0 = k]. \end{aligned}$$

Thus, $\rho_i(k)$ is the expected number of visits to state i during the first sojourn from state k back to state k. We define the (row) vector $\boldsymbol{\rho}(k) = (\rho_i(k))_{k \in S}$, whose ith component is $\rho_i(k)$. Based on our previous discussion, our goal now is to show that the vector $\boldsymbol{\rho}(k)$ satisfies $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$. We should mention here that the sojourn from state k back to state k may never even happen if state k is transient because the chain may never return to state k. Therefore, we assume that state k is recurrent, and it is exactly at this point that we need to assume it. Assuming state k is recurrent, then the chain will return to state k with probability 1. Also, the sojourn includes the last step back to state k; that is, during this sojourn, state k is, by definition, visited exactly once. In other words, $\rho_k(k) = 1$ (assuming state k is recurrent).

One other important thing to observe about $\rho_i(k)$ is that if we sum $\rho_i(k)$ over all $i \in S$, then that is the expected length of the whole sojourn. But the expected length of the sojourn is the mean time to return to state k, given that we start in state k. That is, if μ_k denotes the mean recurrence time to state k, then

$$\mu_k = \sum_{i \in S} \rho_i(k).$$

If state k is positive recurrent then this sum will be finite and it will be infinite if state k is null recurrent.

As we have done in previous examples, we will use indicator functions to represent the number of visits to state i during the first sojourn. If we define $I_{\{X_n=i,T_k\geq n\}}$ as the indicator of the event that the chain is in state i at time n and we have not yet revisited state k by time n (i.e. we are still in the first sojourn), then we may represent the total expected number of visits to state i during the first sojourn as

$$\begin{aligned}\rho_i(k) &= \sum_{n=1}^{\infty} E[I_{\{X_n=i,T_k\geq n\}}|X_0=k] \\ &= \sum_{n=1}^{\infty} P(X_n=i,T_k\geq n|X_0=k).\end{aligned}$$

(We are assuming here that $i \neq k$). Purely for the sake of shorter notation we will let $\ell_{ki}(n)$ denote the conditional probability above:

$$\ell_{ki}(n) = P(X_n=i,T_k\geq n|X_0=k)$$

so that now we will write

$$\rho_i(k) = \sum_{n=1}^{\infty} \ell_{ki}(n).$$

We proceed by deriving an equation for $\ell_{ki}(n)$, which will then give an equation for $\rho_i(k)$, and we will see that this equation is exactly the ith equation in $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$. To derive the equation, we intersect the event $\{X_n=i,T_k\geq n\}$ with all possible values of X_{n-1}. Doing this is a special case of the following calculation in basic probability. If $\{B_j\}$ is a partition such that $P(\bigcup_j B_j)=1$ and $B_j\bigcap B_{j'}=\phi$, the empty set, for $j\neq j'$, then for any event A,

$$P(A) = P(A\bigcap(\bigcup_j B_j)) = P(\bigcup_j(A\bigcap B_j)) = \sum_j P(A\bigcap B_j),$$

because the B_j and so the $A\bigcap B_j$ are all disjoint.

For $n = 1$, we have $\ell_{ki}(1) = P(X_1 = i, T_k \geq 1|X_0 = k) = p_{ki}$, the 1-step transition probability from state k to state i. For $n \geq 2$, we let $B_j = \{X_{n-1} = j\}$ and $A = \{X_n = i, T_k \geq n\}$ in the previous paragraph, to get

$$\begin{aligned} \ell_{ki}(n) &= P(X_n = i, T_k \geq n|X_0 = k) \\ &= \sum_{j\in S} P(X_n = i, X_{n-1} = j, T_k \geq n|X_0 = k). \end{aligned}$$

First we note that when $j = k$ the above probability is 0 because the event $\{X_{n-1} = k\}$ implies that the sojourn is over by time $n-1$ while the event $\{T_k \geq n\}$ says that the sojourn is not over at time $n-1$. Therefore, their intersection is the empty set. Thus,

$$\ell_{ki}(n) = \sum_{j\neq k} P(X_n = i, X_{n-1} = j, T_k \geq n|X_0 = k).$$

Next, we note that the event above says that, given we start in state k, we go to state j at time $n-1$ without revisiting state k in the meantime, and then go to state i in the next step. But this is just $\ell_{kj}(n-1)p_{ji}$, and so

$$\ell_{ki}(n) = \sum_{j\neq k} \ell_{kj}(n-1)p_{ji}$$

This is our basic equation for $\ell_{ki}(n)$, for $n \geq 2$. Now, if we sum this over $n \geq 2$ and use the fact that $\ell_{ik}(1) = p_{ki}$ we have

$$\begin{aligned} \rho_i(k) &= \sum_{n=1}^{\infty} \ell_{ki}(n) \\ &= p_{ki} + \sum_{n=2}^{\infty} \sum_{j\neq k} \ell_{kj}(n-1)p_{ji} \\ &= p_{ki} + \sum_{j\neq k} \Big[\sum_{n=2}^{\infty} \ell_{kj}(n-1)\Big] p_{ji}. \end{aligned}$$

But $\sum_{n=2}^{\infty} \ell_{kj}(n-1) = \sum_{n=1}^{\infty} \ell_{kj}(n)$ is equal to $\rho_j(k)$, so we get the equation

$$\rho_i(k) = p_{ki} + \sum_{j \neq k} \rho_j(k) p_{ji}.$$

Now we use the fact that $\rho_k(k) = 1$ to write

$$\begin{aligned} \rho_i(k) &= \rho_k(k) p_{ki} + \sum_{j \neq k} \rho_j(k) p_{ji} \\ &= \sum_{j \in S} \rho_j(k) p_{ji}. \end{aligned}$$

But now we are done, because this is exactly the ith equation in $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$. So we have finished our construction. The vector $\boldsymbol{\rho}(k)$, as we have defined it, has been shown to satisfy the matrix equation $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$. Moreover, as was noted earlier, if state k is a positive recurrent state, then the components of $\boldsymbol{\rho}(k)$ have a finite sum, so that

$$\boldsymbol{\pi} = \boldsymbol{\rho}(k) / \sum_{i \in S} \rho_i(k)$$

is a stationary distribution. We have shown that if our Markov chain has at least one positive recurrent state, then there exists a stationary distribution $\boldsymbol{\pi}$.

Now that we have shown that a stationary distribution exists if there is at least one positive recurrent state, the next thing we want to show is that if a stationary distribution does exist, then all states must be positive recurrent and the stationary distribution is unique.

First, we can show that if a stationary distribution exists, then the Markov chain cannot be transient. If $\boldsymbol{\pi}$ is a stationary distribution, then $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$. Multiplying both sides by $\mathbf{P}^{n-1}$ we get $\boldsymbol{\pi}\mathbf{P}^{n-1} = \boldsymbol{\pi}\mathbf{P}^n$. But we can reduce the left hand side down to $\boldsymbol{\pi}$ by successively applying the relationship $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$. Therefore, we have the relationship that $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}^n$ for any $n \geq 1$, which in a more detailed form is

$$\pi_j = \sum_{i \in S} \pi_i p_{ij}(n),$$

for any $i, j \in S$ and all $n \geq 1$, where $p_{ij}(n)$ is the n-step transition probability from state i to state j.

Now consider what happens when we take the limit as $n \to \infty$ in the above equality. When we look at

$$\lim_{n \to \infty} \sum_{i \in S} \pi_i p_{ij}(n),$$

if we can take the limit inside the summation, then we could use the fact that $\lim_{n \to \infty} p_{ij}(n) = 0$ for all $i, j \in S$ if all states are transient (recall the Corollary we showed at the end of Lecture 10), to conclude that π_j must equal zero for all $j \in S$. It turns out we *can* take the limit inside the summation, but we should be careful because the summation is in general an infinite sum, and limits cannot be taken inside infinite sums in general (recall the example that $+\infty = \lim_{n \to \infty} \sum_{i=1}^{\infty} 1/n \neq \sum_{i=1}^{\infty} \lim_{n \to \infty} 1/n = 0$). The fact that we can take the limit inside the summation here is a consequence of the fact that we can uniformly bound the vector $(\pi_i p_{ij}(n))_{i \in S}$ by a summable vector (uniformly means we can find a bound that works for all n). In particular, since $p_{ij}(n) \leq 1$ for all n, we have that $\pi_i p_{ij}(n) \leq \pi_i$ for all $i \in S$. The fact that this allows us to take the limit inside the summation is an instance of a more general result known as the *bounded convergence theorem*. This is a well-known and useful result in probability, but we won't invoke its use here, as we can show directly that we can take the limit inside the summation, as follows. Let F be any finite subset of the state space S. Then we can write

$$\begin{aligned}\lim_{n\to\infty}\sum_{i\in S}\pi_i p_{ij}(n) &= \lim_{n\to\infty}\sum_{i\in F}\pi_i p_{ij}(n) + \lim_{n\to\infty}\sum_{i\in F^c}\pi_i p_{ij}(n) \\ &\le \lim_{n\to\infty}\sum_{i\in F}\pi_i p_{ij}(n) + \sum_{i\in F^c}\pi_i,\end{aligned}$$

from the inequality $p_{ij}(n) \le 1$. But for the first finite summation, we can take the limit inside, so we get that the limit of the first sum (over F) is 0. Therefore,

$$\lim_{n\to\infty}\sum_{i\in S}\pi_i p_{ij}(n) \le \sum_{i\in F^c}\pi_i,$$

for any finite subset F of S. But since $\sum_{i\in S}\pi_i = 1$ is a convergent sum, for any $\epsilon > 0$, we can take the set F so large (but still finite) to make $\sum_{i\in F^c}\pi_i < \epsilon$. This implies that

$$\lim_{n\to\infty}\sum_{i\in S}\pi_i p_{ij}(n) \le \epsilon$$

for every $\epsilon > 0$. But the only way this can be true is if the above limit is 0. Therefore, going back to our original argument, we see that if all states are transient, this implies that $\pi_j = 0$ for all $j \in S$. This is clearly impossible since the components of $\boldsymbol{\pi}$ must sum to 1. Therefore, if a stationary distribution exists for an irreducible Markov chain, all states must be recurrent.

We end here with another attempt at some intuitive understanding, this time of why the stationary distribution $\boldsymbol{\pi}$, if it did exist, might be unique. In particular, let us try to see why we might expect that $\pi_i = 1/\mu_i$, where μ_i is the mean recurrence time to state i. Suppose we start the chain in state i and then observe the chain over N time periods, where N is large. Over those N time periods, let n_i be the number of times that the chain revisits state i. If N is large, we expect that n_i/N is approximately equal to π_i, and indeed should converge to π_i as N went to infinity. On the other hand, if the times that the chain returned to state i were *uniformly* spread over the times from 0 to N, then each time state i was visited the chain would return to state i after N/n_i steps. For example, if the chain visited state i 10 times in 100 steps and the times it returned to state i were uniformly spread, then the chain would have returned to state i every 100/10=10 steps. In reality, the return times to state i vary, perhaps a lot, over the different returns to state i. But if we average all these return times (meaning the arithmetic average), then this average behaves very much like the return time when all the return times are the same. So we should expect that the average return time to state i should be close to N/n_i, when N is very large (note that as N grows, so does n_i), and as N went to infinity, the ratio N/n_i should actually converge to μ_i, the mean return time to state i. Given these two things, that π_i should be close to n_i/N and μ_i should be close to N/n_i, we should expect their product to be 1; that is, $\pi_i\mu_i = 1$, or $\pi_i = 1/\mu_i$. Note that if this relationship holds, then this directly relates the stationary distribution to the null or positive recurrence of the chain, through the mean recurrence times μ_i. If π_i is positive, then μ_i must be finite, and hence state i must be positive recurrent. Also, the stationary distribution must be unique, because the mean recurrence times are unique. Next we will prove more rigorously that the relationship $\pi_i \mu_i = 1$ does indeed hold and we will furthermore show that if the stationary distribution exists then *all* states must be positive recurrent.

15

Existence and Uniqueness (cont'd)

Previously we saw how to construct a vector $\boldsymbol{\rho}(k)$ that satisfies the equations $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$, when $\mathbf{P}$ is the transition matrix of an irreducible, recurrent Markov chain. Note that we didn't need the chain to be positive recurrent, just recurrent. As an example, consider the simple random walk with $p = 1/2$. We have seen that this Markov chain is irreducible and null recurrent. The transition matrix is

$$\mathbf{P} = \begin{bmatrix} \ddots & \ddots & \ddots & & & & \\ & \frac{1}{2} & 0 & \frac{1}{2} & & & \\ & & \frac{1}{2} & 0 & \frac{1}{2} & & \\ & & & \frac{1}{2} & 0 & \frac{1}{2} & \\ & & & & \ddots & \ddots & \ddots \end{bmatrix},$$

and one can easily verify that the vector $\boldsymbol{\pi} = (\ldots, 1, 1, 1, \ldots)$ satisfies $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ (any constant multiple of $\boldsymbol{\pi}$ will also work). However, $\boldsymbol{\pi}$ cannot be a stationary distribution because its components sum to infinity. Today we will show that if a stationary distribution exists for an irreducible Markov chain, then it must be a positive recurrent Markov chain. Moreover, the stationary distribution is unique.

Last time we gave a (hopefully) intuitive argument as to why, if a stationary distribution did exist, we might expect that $\pi_i\mu_i = 1$, where μ_i is the mean time to return to state i, given that we start in state i. We'll prove this rigorously now. So assume that a stationary distribution $\boldsymbol{\pi}$ exists, and let the initial distribution of X_0 be $\boldsymbol{\pi}$, so that we make our process stationary. Let T_i be the first time we enter state i, starting from time 1 (this is the same definition of T_i as in the last lecture). So we have that

$$\mu_i = E[T_i | X_0 = i]$$

and also

$$\mu_i \pi_i = E[T_i | X_0 = i] P(X_0 = i).$$

We wish to show that this equals one, and the first thing we do is write out the expectation, but in a somewhat nonstandard form. The random variable T_i is defined on the nonnegative integers, and there is a useful way to represent the mean of such a random variable, as follows:

$$\begin{aligned} E[T_i | X_0 = i] &= \sum_{k=1}^{\infty} k P(T_i = k | X_0 = i) \\ &= \sum_{k=1}^{\infty} \Big(\sum_{n=1}^{k} (1) \Big) P(T_i = k | X_0 = i) \\ &= \sum_{n=1}^{\infty} \sum_{k=n}^{\infty} P(T_i = k | X_0 = i) \\ &= \sum_{n=1}^{\infty} P(T_i \geq n | X_0 = i), \end{aligned}$$

by interchanging the order of summation in the third equality.

So we have that

$$\begin{aligned}\mu_i\pi_i &= \sum_{n=1}^{\infty} P(T_i \geq n | X_0 = i) P(X_0 = i) \\ &= \sum_{n=1}^{\infty} P(T_i \geq n, X_0 = i).\end{aligned}$$

Now for $n = 1$, we have $P(T_i \geq 1, X_0 = i) = P(X_0 = i)$, while for $n \geq 2$, we write

$$P(T_i \geq n, X_0 = i) = P(X_{n-1} \neq i, X_{n-2} \neq i, \ldots, X_1 \neq i, X_0 = i)$$

Now for any events A and B, we have that

$$P(A \bigcap B) = P(A) - P(A \bigcap B^c),$$

which follows directly from $P(A) = P(A \bigcap B) + P(A \bigcap B^c)$. With $A = \{X_{n-1} \neq i, \ldots, X_1 \neq i\}$ and $B = \{X_0 = i\}$ we get

$$\begin{aligned}\mu_i\pi_i &= P(X_0 = i) + \sum_{n=2}^{\infty} \Big(P(X_{n-1} \neq i, \ldots, X_1 \neq i) \\ &\qquad - P(X_{n-1} \neq i, \ldots, X_1 \neq i, X_0 \neq i) \Big) \\ &= P(X_0 = i) + \sum_{n=2}^{\infty} \Big(P(X_{n-2} \neq i, \ldots, X_0 \neq i) \\ &\qquad - P(X_{n-1} \neq i, \ldots, X_1 \neq i, X_0 \neq i) \Big)\end{aligned}$$

where we did a shift in index to get the last expression. This shift is allowed because we are assuming the process is stationary.

We are almost done now. To make notation a bit less clunky, let's define

$$a_n \equiv P(X_n \neq i, \ldots, X_0 \neq i).$$

Our expression for $\mu_i \pi_i$ can now be written as

$$\begin{aligned}\mu_i \pi_i &= P(X_0 = i) + \sum_{n=2}^{\infty} (a_{n-2} - a_{n-1}) \\ &= P(X_0 = i) + a_0 - a_1 + a_1 - a_2 + a_2 - a_3 + \ldots\end{aligned}$$

The above sum is what is called a *telescoping* sum because of the way the partial sums collapse. Indeed, the nth partial sum is

$$P(X_0 = i) + a_0 - a_n,$$

so that the infinite sum (by definition the limit of the partial sums) is

$$\mu_i \pi_i = P(X_0 = i) + a_0 - \lim_{n \to \infty} a_n.$$

Two facts give us our desired result that $\mu_i \pi_i = 1$. The first is the simple fact that $a_0 = P(X_0 \neq i)$, so that

$$P(X_0 = i) + a_0 = P(X_0 = i) + P(X_0 \neq i) = 1.$$

The second fact is that

$$\lim_{n \to \infty} a_n = 0.$$

This fact is not completely obvious. To see this, note that this limit is the probability that the chain *never* visits state i. Suppose the chain starts in some arbitrary state j. Because j is recurrent by the Markov property it will be revisited infinitely often with probability 1. Since the chain is irreducible there is some n such that $p_{ji}(n) > 0$. Thus on each visit to j there is some positive probability that i will be visited after a finite number of steps. So the situation is like flipping a coin with a positive probability of heads. It is not hard to see that a heads will eventually be flipped with probability one.

Thus, we're done. We've shown that $\mu_i \pi_i = 1$ for any state i. Note that the only thing we've assumed is that the chain is irreducible and that a stationary distribution exists. The fact that $\mu_i \pi_i = 1$ has several important implications. One, obviously, is that

$$\mu_i = \frac{1}{\pi_i}.$$

That is, the mean time to return to state i can be computed by determining the stationary probability π_i, if possible. Another implication is that if a stationary distribution $\boldsymbol{\pi}$ exists, then it must be unique, because the mean recurrence times μ_i are obviously unique. The third important implication is that

$$\pi_i = \frac{1}{\mu_i}.$$

This immediately implies that if state i is positive recurrent (which means by definition that $\mu_i < \infty$), then $\pi_i > 0$. In fact, we're now in a position to prove that positive recurrence is a class property (recall that when we stated this "fact", we delayed the proof of it till later. That later is now). We are still assuming that a stationary distribution exists. As we have seen before, this implies that

$$\pi_j = \sum_{i \in S} \pi_i p_{ij}(n),$$

for every $n \geq 1$ and every $j \in S$. Suppose that $\pi_j = 0$ for some state j. Then, that implies that

$$0 = \sum_{i \in S} \pi_i p_{ij}(n),$$

for that particular j, and for every $n \geq 1$.

But since the state space is irreducible (all states communicate with one another), for every i there is some n such that $p_{ij}(n) > 0$. This implies that π_i must be 0 for every $i \in S$. But this is impossible because the π_i must sum to one. So we have shown that *if a stationary distribution exists, then π_i must be strictly positive for every i.* This implies that all states must be positive recurrent. So, putting this together with our previous result that we can construct a stationary distribution if at least one state is positive recurrent, we see that if one state is positive recurrent, then we can construct a stationary distribution, and then this implies that all states must be positive recurrent. In other words, positive recurrence is a class property. Of course, this then implies that null recurrence is also a class property.

Let's summarize the main results that we've proved over the last two lectures in a theorem:

Theorem. For an irreducible Markov chain, a stationary distribution $\boldsymbol{\pi}$ exists if and only if all states are positive recurrent. In this case, the stationary distribution is unique and $\pi_i = 1/\mu_i$, where μ_i is the mean recurrence time to state i.

So we can't make a transient or a null recurrent Markov chain stationary. Also, if the Markov chain has two or more equivalence classes (we say the Markov chain is *reducible*), then in general there will be many stationary distributions. One of the Stat855 problems is to give an example of this. In these cases, there are different questions to ask about the process, as we shall see. Also note that there are no conditions on the period of the Markov chain for the existence and uniqueness of the stationary distribution. This is not true when we consider limiting probabilities, as we shall also see.

Example: (Ross, p.229 #26, extended). Three out of every four trucks on the road are followed by a car, while only one out of every five cars is followed by a truck. If I see a truck pass me by on the road, on average how many vehicles pass before I see another truck?

Solution: Recall that we set this up as a Markov chain in which we imagine sitting on the side of the road watching vehicles go by. If a truck goes by the next vehicle will be a car with probability 3/4 and will be a truck with probability 1/4. If a car goes by the next vehicle will be a car with probability 4/5 and will be a truck with probability 1/5. If we let X_n denote the type of the nth vehicle that passes by (0 for truck and 1 for car), then $\{X_n : n \geq 1\}$ is a Markov chain with two states (0 and 1) and transition probability matrix

$$\mathbf{P} = \begin{array}{c} \\ 0 \\ 1 \end{array} \begin{array}{c} \begin{array}{cc} 0 & 1 \end{array} \\ \left[\begin{array}{cc} 1/4 & 3/4 \\ 1/5 & 4/5 \end{array}\right] \end{array}.$$

The equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ are

$$\pi_0 = \frac{1}{4}\pi_0 + \frac{1}{5}\pi_1 \quad \text{and} \quad \pi_1 = \frac{3}{4}\pi_0 + \frac{4}{5}\pi_1,$$

which, together with the constraint $\pi_0 + \pi_1 = 1$, we had solved previously to yield $\pi_0 = 4/19$ and $\pi_1 = 15/19$. If I see a truck pass by then the average number of vehicles that pass by before I see another truck corresponds to the mean recurrence time to state 0, given that I am currently in state 0. By our theorem, the mean recurrence time to state 0 is $\mu_0 = 1/\pi_0 = 19/4$, which is roughly 5 vehicles. □

16

Example of PGF for π/Some Number Theory

Today we'll start with another example illustrating the calculation of the mean time to return to a state in a Markov chain by calculating the stationary probability of that state, but this time through the use of the probability generating function (pgf) of the stationary distribution.

Example: I'm taking a lot of courses this term. Every Monday I get 2 new assignments with probability $2/3$ and 3 new assignments with probability $1/3$. Every week, between Monday morning and Friday afternoon I finish 2 assignments (they might be new ones or ones unfinished from previous weeks). If I have any unfinished assignments on Friday afternoon, then I find that over the weekend, independently of anything else, I finish one assignment by Monday morning with probability c and don't finish any of them with probability $1 - c$. If the term goes on forever, how many weeks is it before I can expect a weekend with no homework to do?

Solution: Let X_n be the number of unfinished homeworks at the end of the nth Friday after term starts, where $X_0 = 0$ is the number of unfinished homeworks on the Friday before term starts. Then $\{X_n : n \geq 0\}$ is a Markov chain with state space $S = \{0, 1, 2, \ldots\}$. Some transition probabilities are, for example

$$
\begin{aligned}
0 &\rightarrow 0 \text{ with probability } 2/3 \text{ (2 new ones on Monday)}\\
0 &\rightarrow 1 \text{ with probability } 1/3 \text{ (3 new ones on Monday)}\\
1 &\rightarrow 0 \text{ with probability } 2c/3\\
1 &\rightarrow 1 \text{ with probability } c/3 + 2(1-c)/3 = (2-c)/3\\
1 &\rightarrow 2 \text{ with probability } (1-c)/3,
\end{aligned}
$$

and, in general, if I have i unfinished homeworks on a Friday afternoon, then the transition probabilities are given by

$$
\begin{aligned}
i &\rightarrow i-1 \text{ with probability } 2c/3,\\
i &\rightarrow i \qquad \text{ with probability } c/3 + 2(1-c)/3 = (2-c)/3,\\
i &\rightarrow i+1 \text{ with probability } (1-c)/3
\end{aligned}
$$

The transition probability matrix for this Markov chain is given by

$$
\mathbf{P} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ \vdots \end{array}
\begin{array}{c}
\begin{array}{cccccccc} 0 & 1 & 2 & 3 & 4 & \cdots & & \end{array}\\
\left[\begin{array}{cccccccc}
2/3 & 1/3 & 0 & \cdots & & & & \\
q & r & p & 0 & \cdots & & & \\
0 & q & r & p & 0 & \cdots & & \\
0 & 0 & q & r & p & 0 & \cdots & \\
0 & 0 & 0 & \ddots & \ddots & \ddots & & \\
\vdots & \vdots & \vdots & & & & &
\end{array}\right]
\end{array},
$$

where

$$
\begin{aligned}
q &= 2c/3\\
r &= (2-c)/3\\
p &= (1-c)/3
\end{aligned}
$$

and $q + r + p = 1$. In the parlance of Markov chains, this process is an example of a *random walk with a reflecting barrier at 0*.

We should remark here that it's not at all clear that this Markov chain chain always has a stationary distribution for every $c \in [0, 1]$. On the one hand, if $c = 1$, so that I always do a homework over the weekend if there is one to do, then I will never have more than one unfinished homework on a Friday afternoon. This case corresponds to $p = 0$, and we can see from the transition matrix that states $\{0, 1\}$ will be a closed, positive recurrent class, while the states $\{2, 3, \ldots\}$ will be a transient class of states. On the other extreme, if $c = 0$, so that I never do a homework on the weekend, then every time I get 3 new homeworks on a Monday, my backlog of unfinished homeworks will increase by one permanently. In this case $q = 0$ and one can see from the transition matrix that I never reduce my number of unfinished homeworks, and eventually my backlog of unfinished homeworks will go off to infinity. We call such a system *unstable*. Stability can often be a major design issue for complex systems that service jobs/tasks/processes (generically customers). A stochastic model can be invaluable for providing insight into the parameters affecting the stability of a system. For our example here, there should be some threshold value c_0 such that the system is stable for $c > c_0$ and unstable for $c < c_0$. One valuable use of stationary distributions comes from the mere fact of their existence. If we can find those values of c for which a stationary distribution exists, then it is for those values of c that the system is stable.

So we look for a stationary distribution. Note that if we find one, then the answer to our question of how many weeks do we have to wait on average for a homework-free weekend is $\mu_0 = 1/\pi_0$, the mean recurrence time to state 0, our starting state. A stationary distribution $\boldsymbol{\pi} = (\pi_0, \pi_1, \ldots)$ must satisfy $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, which we write out as

$$\begin{aligned}
\pi_0 &= \frac{2}{3}\pi_0 + q\pi_1 \\
\pi_1 &= \frac{1}{3}\pi_0 + r\pi_1 + q\pi_2 \\
\pi_2 &= p\pi_1 + r\pi_2 + q\pi_3 \\
&\vdots \\
\pi_i &= p\pi_{i-1} + r\pi_i + q\pi_{i+1} \\
&\vdots
\end{aligned}$$

A direct attack on this system of linear equations is possible, by expressing π_i in terms of π_0, and then summing π_i over all i to get π_0 using the constraint that $\sum_{i=0}^{\infty} = 1$. However, this approach is somewhat cumbersome. A more elegant approach is to use the method of generating functions. This method can often be applied to solve a linear system of equations, especially when there are an infinite number of equations, in situations where each equation only involves variables "close to one another" (for example, each of the equations above involves only two or three consecutive variables) and all, or almost all, of the equations have a regular form (as in $\pi_i p\pi_{i-1} + r\pi_i + q\pi_{i+1}$).

By multiplying the ith equation above by s^i and then summing over i we collapse the above infinite set of equations into just a single equation for the generating function.

Let $G(s) = \sum_{i=0}^{\infty} s^i \pi_i$ denote the generating function of the stationary distribution $\boldsymbol{\pi}$. If we multiply the ith equation in $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ by s^i and sum over i, we obtain

$$\sum_{i=0}^{\infty} s^i \pi_i \;=\; \frac{2}{3}\pi_0 + \frac{1}{3}\pi_0 s + p\sum_{i=2}^{\infty} s^i \pi_{i-1} + r\sum_{i=1}^{\infty} s^i \pi_i + q\sum_{i=0}^{\infty} s^i \pi_{i+1}$$

The left hand side is just $G(s)$ while the sums on the right hand side are not difficult to express in terms of $G(s)$ with a little bit of manipulation. In particular,

$$\begin{aligned} p\sum_{i=2}^{\infty} s^i \pi_{i-1} &= ps\sum_{i=2}^{\infty} s^{i-1}\pi_{i-1} = ps\sum_{i=1}^{\infty} s^i \pi_i \\ &= ps\sum_{i=0}^{\infty} s^i \pi_i - ps\pi_0 = psG(s) - ps\pi_0 \end{aligned}$$

Similarly,

$$r\sum_{i=1}^{\infty} s^i \pi_i = r\sum_{i=0}^{\infty} s^i \pi_i - r\pi_0 = rG(s) - r\pi_0$$

and

$$\begin{aligned} q\sum_{i=0}^{\infty} s^i \pi_{i+1} &= \frac{q}{s}\sum_{i=0}^{\infty} s^{i+1}\pi_{i+1} = \frac{q}{s}\sum_{i=1}^{\infty} s^i \pi_i \\ &= \frac{q}{s}\sum_{i=0}^{\infty} s^i \pi_i - \frac{q}{s}\pi_0 = \frac{q}{s}G(s) - \frac{q}{s}\pi_0. \end{aligned}$$

Therefore, the equation we obtain for $G(s)$ is

$$G(s) = \frac{2}{3}\pi_0 + \frac{s}{3}\pi_0 + psG(s) - ps\pi_0 + rG(s) - r\pi_0 + \frac{q}{s}G(s) - \frac{q}{s}\pi_0.$$

Collecting like terms, we have

$$G(s)\left[1-ps-r-\frac{q}{s}\right]=\pi_0\left[\frac{2}{3}+\frac{s}{3}-ps-r-\frac{q}{s}\right].$$

To get rid of the fractions, we'll multiply both sides by $3s$, giving

$$\begin{aligned} G(s)[3s-3ps^2-3rs-3q] &= \pi_0[2s+s^2-3ps^2-3rs-3q] \\ \Rightarrow \quad G(s) &= \frac{\pi_0(2s+s^2-3ps^2-3rs-3q)}{3s-3ps^2-3rs-3q}. \end{aligned}$$

In order to determine the unknown π_0 we use the boundary condition $G(1)=1$, which must be satisfied if $\boldsymbol{\pi}$ is to be a stationary distribution. This boundary condition also gives us a way to check for the values of c for which the stationary distribution exists. If a stationary distribution does not exist, then we will not be able to satisfy the condition $G(1)=1$. Plugging in $s=1$, we obtain

$$G(1)=\frac{\pi_0(2+1-3p-3r-3q)}{3-3p-3r-3q}.$$

However, we run into a problem here due to the fact that $p+r+q=1$, which means that $G(1)$ is an indeterminate form

$$G(1)=\left[\frac{0}{0}\right].$$

Therefore, we use L'Hospital's rule to determine the limiting value of $G(s)$ as $s\to 1$. This gives

$$\begin{aligned} \lim_{s\to 1} G(s) &= \pi_0\frac{\lim_{s\to 1}(2+2s-6ps-3r)}{\lim_{s\to 1}(3-6ps-3r)} \\ &= \pi_0\frac{4-6p-3r}{3-6p-3r}. \end{aligned}$$

We had previously defined our quantities p, r and q in terms of c to make it easier to write down the transition matrix $\mathbf{P}$, but now we would like to re-express these back in terms of c to make it simpler to see when $\lim_{s\to 1} G(s) = 1$ is possible. Recall that $p = (1-c)/3$, $r = (2-c)/3$ and $q = 2c/3$, so that $4 - 6p - 3r = 4 - 2(1-c) - (2-c) = 3c$ and $3 - 6p - 3r = 3c - 1$. So in terms of c, we have

$$\lim_{s\to 1} G(s) = \pi_0 \frac{3c}{3c-1}.$$

In order to have a proper stationary distribution, we must have the left hand side equal to 1 and we must have $0 < \pi_0 < 1$. Together these imply that we must have $3c/(3c-1) > 1$, which will only be true if $3c - 1 > 0$, or $c > 1/3$. Thus, we have found our threshold value of $c_0 = 1/3$ such that the system is stable (since it has a stationary distribution) for $c > c_0$ and is unstable for $c \le c_0$. Assuming $c > 1/3$ so that the system is stable, we may now solve for π_0 through the relationship

$$\begin{aligned} 1 &= \pi_0 \frac{3c}{3c-1} \\ \Rightarrow \quad \pi_0 &= \frac{3c-1}{3c}. \end{aligned}$$

The answer to our original question of what is the mean number of weeks until a return to state 0 is

$$\mu_0 = \frac{1}{\pi_0} = \frac{3c}{3c-1}.$$

Observe that we have found a mean return time of interest, μ_0, in terms of a system parameter, c. More generally, a typical thing we try to do in stochastic modeling is find out how some performance measure of interest depends, explicitly or even just qualitatively, on one or more system parameters. In particular, if we have some control over one or more of those system parameters, then we have a useful tool to help us design our system. For example, if I wanted to design my homework habits so that I could expect to have a homework-free weekend in six weeks, I can solve for c to make $\mu_0 \leq 6$. This gives $\mu_0 = 3c/(3c-1) \leq 6 \Rightarrow 3c \leq 18c - 6$ or $c \geq 2/5$. □

Let us now return to some general theory. We've already proved one of the main general theorems concerning Markov chains, that we emphasized by writing it in a framed box near the end of the previous lecture. This was the theorem concerning the conditions for the existence and uniqueness of a stationary distribution for a Markov chain. We reiterate here that there were no conditions on the period of the Markov chain for that result. The other main theoretical result concerning Markov chains has to do with the limiting probabilities $\lim_{n\to\infty} p_{ij}(n)$. For this result the period does matter. Let's state what that result is now: when the stationary distribution exists *and* the chain is aperiodic (so the chain is irreducible, positive recurrent, and aperiodic), $p_{ij}(n)$ converges to the stationary probability π_j as $n \to \infty$. Note that the limit does not depend on the starting state i. This is quite important. In words, for an irreducible, positive recurrent, aperiodic Markov chain, no matter where we start from and *no matter what our initial distribution is*, if we let the chain run for a long time then the distribution of X_n will be very much like the stationary distribution $\boldsymbol{\pi}$.

An important first step in proving the above limiting result is to show that for an irreducible, positive recurrent, aperiodic Markov chain the n-step transition probability $p_{ij}(n)$ is strictly positive for *all* n "big enough". That is, there exists some integer M such that $p_{ij}(n) > 0$ for all $n \geq M$. To show this we will need some results from basic number theory. We'll state and prove these results now.

Some Number Theory:

If we have an irreducible, positive recurrent, aperiodic Markov chain then we know that for any state j, the greatest common divisor (gcd) of the set of times n for which $p_{jj}(n) > 0$ is 1. If $A_j \equiv \{n_1, n_2, \ldots\}$ is this set of times, then this is an infinite set because, for example, there must be some finite n_0 such that $p_{jj}(n_0) > 0$. But that implies $p_{jj}(2n_0) > 0$ and in general $p_{jj}(kn_0) > 0$ for any positive integer k. For reasons which will become clearer in the next lecture, what we would like to be able to do is take some finite subset of A_j that also has gcd 1 and then show that every n large enough can be written as a linear combination of the elements of this finite subset, where the coefficients of the linear combination are all nonnegative integers. This is what we will show now, through a series of three results.

Result 1: Let $n_1, n_2, \ldots$ be a sequence of positive integers with gcd 1. Then there exists a finite subset $b_1, \ldots, b_r$ that has gcd 1.

Proof: Let $b_1 = n_1$ and $b_2 = n_2$ and let $g =$gcd(b_1, b_2). If $g = 1$ then we are done. If $g > 1$ let $p_1, \ldots, p_d$ be the distinct prime factors of g that are larger than 1 (if $g > 1$ it must have at least one prime factor larger than 1). For each p_k, $k = 1, \ldots, d$, there must be at least one integer from $\{n_3, n_4, \ldots\}$ that p_k does not divide because if p_k divided every integer in this set then, since it also divides both n_1 and n_2, it is a common divisor of all the n's. But this contradicts our assumption that the gcd is 1. Therefore,

choose b'_3 from $\{n_3, n_4, \ldots\}$ such that p_1 does not divide b'_3
choose b'_4 from $\{n_3, n_4, \ldots\}$ such that p_2 does not divide b'_4
$\vdots$
choose b'_{d+2} from $\{n_3, n_4, \ldots\}$ such that p_d does not divide b'_{d+2}.

Note that $b'_3, \ldots, b'_{d+2}$ do not need to be distinct. Let $b_3, \ldots, b_r$ be the distinct integers among $b'_3, \ldots, b'_{d+2}$. Then $b_1, b_2, \ldots, b_r$ have gcd 1 because each p_k does not divide at least one of $\{b_3, \ldots, b_r\}$, so that none of the p_k is a common divisor. On the other hand, the p_k's are the *only* integers greater than 1 that divide both b_1 and b_2. Therefore, there are no integers greater than 1 that divide all of $b_1, \ldots, b_r$. So the gcd of $b_1, \ldots, b_r$ is 1. □

Result 2: Let $b_1, \ldots, b_r$ be a finite set of positive integers with gcd 1. Then there exist integers $a_1, \ldots, a_r$ (not necessarily nonnegative) such that $a_1 b_1 + \ldots + a_r b_r = 1$.

Proof: Consider the set of all integers of the form $c_1 b_1 + \ldots + c_r b_r$ as the c_i range over the integers. This set of integers has some *least positive* element ℓ. Let $a_1, \ldots, a_r$ be such that $\ell = a_1 b_1 + \ldots + a_r b_r$. We are done if we show that $\ell = 1$. To do this we will show that ℓ is a common divisor of $b_1, \ldots, b_r$. Since $b_1, \ldots, b_r$ has gcd 1 by assumption, this shows that $\ell = 1$. We will show that ℓ divides b_i by contradiction. Suppose that ℓ did not divide b_i. Then we can write $b_i = q\ell + R$, where $q \geq 0$ is an integer and the remainder R satisfies $0 < R < \ell$. But then

$$R = b_i - q\ell = b_i - q \sum_{k=1}^{r} a_k b_k = (1 - q a_i) b_i + \sum_{k \neq i} (-q a_k) b_k$$

is also of the form $c_1 b_1 + \ldots + c_r b_r$. But $R < \ell$ contradicts the minimality of ℓ. Therefore, ℓ must divide b_i. □

Our final result for today, the one we are really after, uses Result 2 to show that every integer large enough can be written as a linear combination of $b_1, \ldots, b_r$ with nonnegative integer coefficients.

Result 3: Let $b_1, \ldots, b_r$ be a finite set of positive integers with gcd 1. Then there exists a positive integer M such that for every $n > M$ there exist nonnegative integers $d_1, \ldots, d_r$ such that $n = d_1 b_1 + \ldots + d_r b_r$.

Proof: From Result 2, there exist integers $a_1, \ldots, a_r$ (which may be positive or negative) such that $a_1 b_1 + \ldots + a_r b_r = 1$. Now choose $M = (|a_1| b_1 + \ldots + |a_r| b_r) b_1$, where $|\cdot|$ denotes absolute value. If $n > M$ then we can write n as $n = M + q b_1 + R$, where $q \geq 0$ is an integer and the remainder R satisfies $0 \leq R < b_1$. If $R = 0$ then we are done as we can choose $d_k = |a_k|$ for $k \neq 1$ and $d_1 = |a_1| + q$. If $0 < R < b_1$, then

$$
\begin{aligned}
n &= M + q b_1 + R(1) \\
&= M + q b_1 + R(a_1 b_1 + \ldots + a_r b_r) \\
&= (|a_1| b_1 + q + R a_1) b_1 + \sum_{k=2}^{r} (|a_k| b_1 + R a_k) b_k \\
&= d_1 b_1 + \ldots + d_r b_r,
\end{aligned}
$$

where $d_1 = q + b_1 |a_1| + R a_1 \geq q + (b_1 - R)|a_1| \geq 0$ since $R < b_1$, and $d_k = b_1 |a_k| + R a_k \geq (b_1 - R)|a_k| \geq 0$ also. □

Result 3 is what we need to show that $p_{jj}(n) > 0$ for all n big enough in an irreducible, positive recurrent, aperiodic Markov chain. We will show this next and continue on to prove our main limit result $p_{ij}(n) \to \pi_j$ as $n \to \infty$.

17

Limiting Probabilities

Last time we ended with some results from basic number theory that will allow us to show that for an irreducible, positive recurrent, aperiodic Markov chain, the n-step transition probability $p_{ij}(n) > 0$ for all n large enough. First, fix any state j. Next, choose a finite set of times $b_1, \ldots, b_r$ such that the gcd of $b_1, \ldots, b_r$ is 1 and $p_{jj}(b_k) > 0$ for all $k = 1, \ldots, r$ (we showed we can do this from our Result 1 from last time). Next, Result 2 tells us we can find integers $a_1, \ldots, a_r$ such that $a_1 b_1 + \ldots + a_r b_r = 1$. Now let n be any integer larger than $M = (|a_1|b_1 + \ldots + |a_r|b_r)b_1$. Then Result 3 tells us there are nonnegative integers $d_1, \ldots, d_r$ such that $n = d_1 b_1 + \ldots + d_r b_r$. But now we have that

$$
\begin{aligned}
p_{jj}(n) &\geq \underbrace{p_{jj}(b_1)\ldots p_{jj}(b_1)}_{d_1 \text{ times}}\underbrace{p_{jj}(b_2)\ldots p_{jj}(b_2)}_{d_2 \text{ times}}\ldots\underbrace{p_{jj}(b_r)\ldots p_{jj}(b_r)}_{d_r \text{ times}} \\
&= p_{jj}(b_1)^{d_1} p_{jj}(b_2)^{d_2} \ldots p_{jj}(b_r)^{d_r} \\
&> 0,
\end{aligned}
$$

where the first inequality above follows because the right hand side is the probability of just a subset of the possible paths that go from state j to state j in n steps, and this probability is positive because $b_1, \ldots, b_r$ were chosen to have $p_{jj}(b_k) > 0$ for $k = 1, \ldots, r$.

More generally, fix any two states i and j with $i \neq j$. Since the chain is irreducible, there exists some m such that $p_{ij}(m) > 0$. But then, by the same bounding argument we may write

$$p_{ij}(m+n) \geq p_{ij}(m)p_{jj}(n) > 0$$

for all n large enough.

Let me remind you again that if the period of the Markov chain is d, where d is larger than 1, then we cannot have $p_{jj}(n) > 0$ for all n big enough because $p_{jj}(n) = 0$ for all n that is not a multiple of d. This is why the limiting probability will not exist. We *can* define a different limiting probability in this case, which we'll discuss later, but for now we are assuming that the Markov chain has period 1 (as well as being irreducible and positive recurrent).

Now we are ready to start thinking about the limit of $p_{ij}(n)$ as $n \to \infty$. We stated in the previous lecture that this limit should be π_j, the stationary probability of state j (where we know that the stationary distribution $\boldsymbol{\pi}$ exists and is unique because we are working now under the assumption that the Markov chain is irreducible and positive recurrent). Equivalently, we may show that the difference $\pi_j - p_{ij}(n)$ converges to 0. We can start off our calculations using the fact that π_j satisfies $\pi_j = \sum_{k\in S} \pi_k p_{kj}(n)$ for every $n \geq 1$ and that $\sum_{k\in S} \pi_k = 1$, to write

$$\begin{aligned}
\pi_j - p_{ij}(n) &= \sum_{k\in S} \pi_k p_{kj}(n) - p_{ij}(n) \\
&= \sum_{k\in S} \pi_k p_{kj}(n) - \sum_{k\in S} \pi_k p_{ij}(n) \\
&= \sum_{k\in S} \pi_k (p_{kj}(n) - p_{ij}(n)).
\end{aligned}$$

So now

$$\begin{aligned}
\lim_{n\to\infty} (\pi_j - p_{ij}(n)) &= \lim_{n\to\infty} \Big[\sum_{k\in S} \pi_k (p_{kj}(n) - p_{ij}(n)) \Big] \\
&= \sum_{k\in S} \pi_k \lim_{n\to\infty} (p_{kj}(n) - p_{ij}(n)),
\end{aligned}$$

where taking the limit inside the (in general, infinite) sum above is justified because the vector $(\pi_k |p_{kj}(n) - p_{ij}(n)|)_{k\in S}$ is uniformly bounded (meaning for every n) by the summable vector $(\pi_k)_{k\in S}$.

Coupling: Our goal now is to show that for any $i, j, k \in S$, we have

$$\lim_{n\to\infty} (p_{kj}(n) - p_{ij}(n)) = 0.$$

This is probably the deepest theoretical result we will prove in this course. The proof uses a technique in probability called *coupling*. This technique has proven useful in a wide variety of probability problems in recent years, and can legitimately be called a "modern" technique. The exact definition of coupling is not important to us right now, but let's see how a coupling argument works for us in our present problem. Suppose that $\boldsymbol{X} = \{X_n : n \geq 0\}$ denotes our irreducible, positive recurrent, aperiodic Markov chain. Let $\boldsymbol{Y} = \{Y_n : n \geq 0\}$ be another Markov chain that is independent of $\boldsymbol{X}$ but with the same transition matrix and the same state space as the $\boldsymbol{X}$ chain. We say that $\boldsymbol{Y}$ is an independent copy of $\boldsymbol{X}$. We will start off our $\boldsymbol{X}$ chain in state i and start off our $\boldsymbol{Y}$ chain in state k. Then, as the argument goes, with probability 1 the $\boldsymbol{X}$ chain and the $\boldsymbol{Y}$ chain will come to a time when they are in the same state, say s. When this happens, we say that the two chains have "coupled" because, due to the Markov property, for any time n that is *after* this coupling time, the distribution of X_n and Y_n will be the same. In particular, their limiting distributions will be the same. This is a real and nontrivial result we are trying to prove here. It is not obvious that the limiting distributions of X_n and Y_n should be the same when the two chains started out in different states, and you should be skeptical of its validity without a proof.

We now give a more rigorous version of the above coupling argument to show that

$$\lim_{n\to\infty} (p_{kj}(n) - p_{ij}(n)) = 0.$$

We start out by defining the "bivariate" process $\boldsymbol{Z} = \{Z_n = (X_n, Y_n) : n \geq 0\}$ (bivariate in the sense that the dimension of Z_n is twice that of X_n), where the processes $\boldsymbol{X}$ and $\boldsymbol{Y}$ are independent (irreducible, positive recurrent, and aperiodic) Markov chains with the same transition matrix $\mathbf{P}$ and the same state space S as described on the previous page. Fix any state $s \in S$. According to the coupling argument, if the process $\boldsymbol{Z}$ starts in state (i, k), it should eventually reach the state (s, s) with probability 1. The first thing we need to do is prove that this is true. We do so by showing that $\boldsymbol{Z}$ is an irreducible, recurrent Markov chain. First we show that $\boldsymbol{Z}$ is a Markov chain. This should actually be intuitively clear, since the chains $\boldsymbol{X}$ and $\boldsymbol{Y}$ are independent. If (i_k, j_k), $k = 0, \ldots, n$, are any $n+1$ states in the state space $S \times S$ of $\boldsymbol{Z}$, then we can work out in detail

$$
\begin{aligned}
& P(Z_n = (i_n, j_n) \mid Z_{n-1} = (i_{n-1}, j_{n-1}), \ldots, Z_0 = (i_0, j_0)) \\
= \; & P(X_n = i_n, Y_n = j_n | X_{n-1} = i_{n-1}, Y_{n-1} = j_{n-1}, \ldots, X_0 = i_0, Y_0 = j_0) \\
= \; & P(X_n = i_n | X_{n-1} = i_{n-1}, Y_{n-1} = j_{n-1}, \ldots, X_0 = i_0, Y_0 = j_0) \\
& \quad \times P(Y_n = j_n | X_{n-1} = i_{n-1}, Y_{n-1} = j_{n-1}, \ldots, X_0 = i_0, Y_0 = j_0) \\
& \hfill \text{(by independence)} \\
= \; & P(X_n = i_n | X_{n-1} = i_{n-1}, \ldots, X_0 = i_0) \\
& \quad \times P(Y_n = j_n | Y_{n-1} = j_{n-1}, \ldots, Y_0 = j_0) \quad \text{(by independence)} \\
= \; & P(X_n = i_n \mid X_{n-1} = i_{n-1}) P(Y_n = j_n \mid Y_{n-1} = j_{n-1}) \\
& \hfill \text{(by the Markov property for } \boldsymbol{X} \text{ and } \boldsymbol{Y}) \\
= \; & P(X_n = i_n \mid X_{n-1} = i_{n-1}, Y_{n-1} = j_{n-1}) \\
& \quad \times P(Y_n = j_n \mid X_{n-1} = i_{n-1} Y_{n-1} = j_{n-1}) \quad \text{(by independence)} \\
= \; & P(X_n = i_n, Y_n = j_n \mid X_{n-1} = i_{n-1}, Y_{n-1} = j_{n-1}) \\
& \hfill \text{(by independence)} \\
= \; & P(Z_n = (i_n, j_n) \mid Z_{n-1} = (i_{n-1}, j_{n-1})).
\end{aligned}
$$

Thus, $\boldsymbol{Z}$ has the Markov property. Next, we show that the $\boldsymbol{Z}$ chain is irreducible. Let (i,k) and (j,ℓ) be any two states in the state space of $\boldsymbol{Z}$. Then the n-step transition probability from state (i,k) to state (j,ℓ) is given by

$$\begin{aligned}
&P(Z_n = (j,\ell) \mid Z_0 = (i,k)) \\
&\quad = P(X_n = j, Y_n = \ell \mid X_0 = i, Y_0 = k) \\
&\quad = P(X_n = j \mid X_0 = i, Y_0 = k) P(Y_n = \ell \mid X_0 = i, Y_0 = k) \\
&\qquad\qquad\qquad\qquad\qquad\qquad \text{(by independence)} \\
&\quad = P(X_n = j \mid X_0 = i) P(Y_n = \ell \mid Y_0 = k) \quad \text{(by independence)} \\
&\quad = p_{ij}(n) p_{k\ell}(n).
\end{aligned}$$

Now we may use our result that there exists some integer M_1 such that $p_{ij}(n) > 0$ for every $n > M_1$ and there exists some integer M_2 such that $p_{k\ell}(n) > 0$ for every $n > M_2$. Letting $M = \max(M_1, M_2)$, we see that $p_{ij}(n)p_{k\ell}(n) > 0$ for every $n > M$. Thus the n-step transition probability in the $\boldsymbol{Z}$ chain, $p_{(i,k),(j,\ell)}(n)$ is positive for every $n > M$. Thus, state (j,ℓ) is accessible from state (i,k) in the $\boldsymbol{Z}$ chain. But since states (i,k) and (j,ℓ) were arbitrary, we see that all states must actually communicate with one another, so that the $\boldsymbol{Z}$ chain is irreducible, as desired.

It is worth remarking at this point that this is the *only* place in our proof that we require the $\boldsymbol{X}$ chain to be aperiodic. It is also worth mentioning that if the $\boldsymbol{X}$ chain were not aperiodic, then the $\boldsymbol{Z}$ chain would in general not be irreducible. Consider, for example, the following.

Example: As a simple example, suppose that the $\boldsymbol{X}$ chain has state space $S = \{0, 1\}$ and transition probability matrix

$$\mathbf{P}_X = \begin{matrix} & \begin{matrix} 0 & 1 \end{matrix} \\ \begin{matrix} 0 \\ 1 \end{matrix} & \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \end{matrix},$$

so that the chain just moves back and forth between states 0 and 1 with probability 1 and so has period 2. Then the chain $\boldsymbol{Z}$ will have state space $S \times S = \{(0,0), (1,1), (0,1), (1,0)\}$ and transition matrix

$$\mathbf{P}_Z = \begin{matrix} & \begin{matrix} (0,0) & (1,1) & (0,1) & (1,0) \end{matrix} \\ \begin{matrix} (0,0) \\ (1,1) \\ (0,1) \\ (1,0) \end{matrix} & \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \end{matrix},$$

From the above matrix it should be clear that the states $\{(0,0), (1,1)\}$ form an equivalence class and the states $\{(0,1), (1,0)\}$ form another equivalence class, so the chain has two equivalence classes and is not irreducible. $\square$

Finally, we show that the $\boldsymbol{Z}$ chain must be recurrent. We do so by demonstrating a stationary distribution for the $\boldsymbol{Z}$ chain. In fact, since $\boldsymbol{Z}$ is irreducible, demonstrating a stationary distribution leads to the stronger conclusion that $\boldsymbol{Z}$ is positive recurrent even though we will only need that $\boldsymbol{Z}$ is recurrent. Let $\boldsymbol{\pi}$ be the stationary distribution of the $\boldsymbol{X}$ (and $\boldsymbol{Y}$) chain. Then we will show that $\pi_{(i,k)} = \pi_i \pi_k$ is the stationary probability of state (i,k) in the $\boldsymbol{Z}$ chain. First, summing over all states $(i,k) \in S \times S$ in the state space of $\boldsymbol{Z}$, we obtain

$$\sum_{(i,k)\in S\times S} \pi_{(i,k)} = \sum_{i\in S}\sum_{k\in S} \pi_i\pi_k = \sum_{i\in S}\pi_i \sum_{k\in S}\pi_k = (1)(1) = 1.$$

Next, we verify that the equations

$$\pi_{(j,\ell)} = \sum_{(i,k)\in S\times S} \pi_{(i,k)} p_{(i,k),(j,\ell)}$$

are satisfied for every $(j,\ell) \in S \times S$. We have

$$\begin{aligned}
\pi_{(j,\ell)} = \pi_j\pi_\ell &= \sum_{i\in S}\pi_i p_{ij} \sum_{k\in S}\pi_k p_{k\ell} \\
&= \sum_{i\in S}\sum_{k\in S} \pi_i\pi_k p_{ij} p_{k\ell} \\
&= \sum_{(i,k)\in S\times S} \pi_{(i,k)} p_{(i,k),(j,\ell)},
\end{aligned}$$

as required. Thus the irreducible chain $\boldsymbol{Z}$ has a stationary distribution, which implies that it is positive recurrent. Recall that our goal was to show that if the $\boldsymbol{Z}$ chain starts out in state (i,k), where (i,k) is any arbitrary state, then it will eventually reach state (s,s) with probability 1. Now that we have shown that $\boldsymbol{Z}$ is irreducible and recurrent, this statement is immediately true by the argument on the bottom of p.134 of these notes.

Thus, if we let T denote the time that the $\boldsymbol{Z}$ chain first reaches state (s, s), then $P(T < \infty | Z_0 = (i, k)) = 1$. Now we are ready to finish off our proof that $p_{ij}(n) - p_{kj}(n) \to 0$ as $n \to \infty$. The following calculations use the following basic properties of events: 1) for any events A, B and C with $P(C) > 0$, we have $P(A \bigcap B | C) \leq P(A|C)$, and 2) for any events A and C with $P(C) > 0$ and any partition $B_1, \ldots, B_n$, we have $P(A|C) = \sum_{m=1}^{n} P(A \bigcap B_m | C)$. For the partition $B_1, \ldots, B_n$ we will use $B_m = \{T = m\}$ for $m = 1, \ldots, n-1$ and $B_n = \{T \geq n\}$. Here's our main calculation:

$$
\begin{aligned}
p_{ij}(n) &= P(X_n = j | X_0 = i) \\
&= P(X_n = j | X_0 = i, Y_0 = k) \quad \text{(by independence)} \\
&= \sum_{m=1}^{n-1} P(X_n = j, T = m | X_0 = i, Y_0 = k) \\
&\qquad\qquad + P(X_n = j, T \geq n | X_0 = i, Y_0 = k) \\
&= \sum_{m=1}^{n-1} P(Y_n = j, T = m | X_0 = i, Y_0 = k) \\
&\qquad\qquad + P(X_n = j, T \geq n | X_0 = i, Y_0 = k) \\
&= P(Y_n = j, T < n | X_0 = i, Y_0 = k) \\
&\qquad\qquad + P(X_n = j, T \geq n | X_0 = i, Y_0 = k) \\
&\leq P(Y_n = j | X_0 = i, Y_0 = k) + P(T \geq n | X_0 = i, Y_0 = k) \\
&= P(Y_n = j | Y_0 = k) + P(T \geq n | X_0 = i, Y_0 = k) \\
&= p_{kj}(n) + P(T \geq n | X_0 = i, Y_0 = k).
\end{aligned}
$$

I hope the only potentially slippery move we made in the above calculation is where we replaced X_n with Y_n in the 4th equality. If you see how that is done, that's good. I'll come back to that later in any case. For now, let's accept it and carry on because we're almost done.

Reiterating the result of that last set of calculations we have

$$p_{ij}(n) \leq p_{kj}(n) + P(T \geq n | X_0 = i, Y_0 = k)$$

which we will write as

$$p_{ij}(n) - p_{kj}(n) \leq P(T \geq n | X_0 = i, Y_0 = k).$$

Now if we interchange the roles of i and k and interchange the roles of $\boldsymbol{X}$ and $\boldsymbol{Y}$ in the previous calculations, then we get

$$p_{kj}(n) - p_{ij}(n) \leq P(T \geq n | X_0 = i, Y_0 = k).$$

Taken together, the last two inequalities imply that

$$|p_{ij}(n) - p_{kj}(n)| \leq P(T \geq n | X_0 = i, Y_0 = k).$$

Now we are basically done because $P(T < \infty | X_0 = i, Y_0 = k) = 1$ implies that

$$\lim_{n \to \infty} P(T \geq n | X_0 = i, Y_0 = k) = 0,$$

and we have our desired result that $p_{ij}(n) - p_{kj}(n) \to 0$ as $n \to \infty$, and then going *way* back to near the beginning of the argument we see that this gives us that $p_{ij}(n) \to \pi_j$ as $n \to \infty$.

Note that the limit result $\lim_{n \to \infty} p_{ij}(n) = \pi_j$ is mostly a theoretical result rather than a computational result. But it's a very important theoretical result. It gives a rigorous justification to using the stationary distribution to analyse the performance of a real system. In practice systems do not start out stationary. What we can say, based on the limit result, is that we can analyse the system based on the stationary distribution when the system has been running for a while. We say that such systems have reached *steady state* or *equilibrium*.

Ok, let's go back now and take a more detailed look at that 4th equality in our calculations a couple of pages back. If you were comfortable with that when you read it, then you may skip over this page of notes. The equality in question was the following:

$$\begin{aligned} &P(X_n = j, T = m | X_0 = i, Y_0 = k) \\ &\qquad = P(Y_n = j, T = m | X_0 = i, Y_0 = k). \end{aligned}$$

So why can we replace X_n with Y_n? The answer in words is that at time m, where $m < n$, both the $\boldsymbol{X}$ and $\boldsymbol{Y}$ processes are in state s. Once we know that, the probability that $Y_n = j$ is the same as the probability that $X_n = j$ because of the Markov property and because $\boldsymbol{X}$ and $\boldsymbol{Y}$ have the same state space and the same transition matrix. We'll do some calculations in more detail now, and we'll use the fact that since the event $\{T = m\}$ implies (i.e. is a subset of) all three events $\{X_m = s\}$, $\{Y_m = s\}$, and $\{X_m = s\} \bigcap \{Y_m = s\}$, we have that $\{T = m\} = \{T = m\} \bigcap \{X_m = s\} = \{T = m\} \bigcap \{Y_m = s\} = \{T = m\} \bigcap \{X_m = s\} \bigcap \{Y_m = s\}$. We may write

$$\begin{aligned} &P\,(X_n = j, T = m | X_0 = i, Y_0 = k) \\ &= P(X_n = j, T = m, X_m = s, Y_m = s | X_0 = i, Y_0 = k) \\ &= P(X_n = j | T = m, X_m = s, Y_m = s, X_0 = i, Y_0 = k) \\ &\qquad \times P(T = m, X_m = s, Y_m = s | X_0 = i, Y_0 = k) \\ &= P(X_n = j | X_m = s, T = m) P(T = m, Y_m = s | X_0 = i, Y_0 = k) \\ &= P(Y_n = j | Y_m = s, T = m) P(T = m, Y_m = s | X_0 = i, Y_0 = k) \\ &= P(Y_n = j | Y_m = s, T = m, X_0 = i, Y_0 = k) \\ &\qquad \times P(T = m, Y_m = s | X_0 = i, Y_0 = k) \\ &= P(Y_n = j, Y_m = s, T = m | X_0 = i, Y_0 = k) \\ &= P(Y_n = j, T = m | X_0 = i, Y_0 = k). \end{aligned}$$

We did an interchange of X_n and Y_n, again in the 4th equality, where we wrote

$$P(X_n = j | X_m = s, T = m) = P(Y_n = j | Y_m = s, T = m),$$

but hopefully in this form the validity of the interchange is more obvious. It should be crystal clear that

$$P(X_n = j | X_m = s) = P(Y_n = j | Y_m = s)$$

holds, since the $\boldsymbol{X}$ and $\boldsymbol{Y}$ chains have the same transition matrix. The extra conditioning on the event $\{T = m\}$ doesn't change either of the above conditional probabilities. It is not dropped from the conditioning only because we want to bring it in front of the conditioning bar later on.

18

Balance and Reversibility

We have said that the stationary probability π_i, if it exists, gives the long run proportion of time in state i. Since every time period spent in state i corresponds to a transition into (or out of) state i, we can also interpret π_i as the long run proportion of transitions that go into (or out of) state i. Also, since p_{ij} is the probability of going to state j given that we are in state i, the product $\pi_i p_{ij}$ is the long run proportion of transitions that go from state i to state j. If we think of a transition from state i to state j as a unit of *flow* from state i to state j, then $\pi_i p_{ij}$ would be the *rate of flow* from state i to state j. Similarly, with this flow interpretation, we have

$$\pi_j = \text{“rate of flow out of state } j\text{”}$$

and

$$\sum_{i \in S} \pi_i p_{ij} = \text{“rate of flow into state } j\text{”}.$$

So the equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ have the interpretation

$$\text{“rate of flow into state } j\text{”} = \text{“rate of flow out of state } j\text{”}$$

for every $j \in S$. That is, the stationary distribution is that vector $\boldsymbol{\pi}$ which achieves balance of flow. For this reason the equations $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ are called the *Balance Equations* or the *Global Balance Equations*.

Local Balance:

All stationary distributions $\boldsymbol{\pi}$ must create global balance, in the sense just described. If the stationary probabilities $\boldsymbol{\pi}$ also satisfy

$$\pi_i p_{ij} = \pi_j p_{ji},$$

for every $i, j \in S$, then we say that $\boldsymbol{\pi}$ also creates *local balance*. The above equations are called the *Local Balance Equations* (sometimes called the *Detailed Balance Equations*) because they specify balance of flow between every pair of states:

"rate of flow from i to j" $=$ "rate of flow from j to i",

for every $i, j \in S$. If one can find a vector $\boldsymbol{\pi}$ that satisfies local balance, then $\boldsymbol{\pi}$ also satisfies the global balance equations, for

$$\begin{aligned} \pi_i p_{ij} &= \pi_j p_{ji} \\ \Rightarrow \sum_{i \in S} \pi_i p_{ij} &= \sum_{i \in S} \pi_j p_{ji} \\ \Rightarrow \sum_{i \in S} \pi_i p_{ij} &= \pi_j \sum_{i \in S} p_{ji} \\ \Rightarrow \sum_{i \in S} \pi_i p_{ij} &= \pi_j, \end{aligned}$$

for every $j \in S$.

Processes that achieve local balance when they are made (or become) stationary are typically easier to deal with computationally than those that don't. This is because the local balance equations are typically much simpler to solve than global balance equations, because each local balance equation always involves just two unknowns.

Example: In the example from p.139 of the notes in which we used the method of generating functions to obtain information about a stationary distribution, the transition matrix was given by

$$\mathbf{P} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ \vdots \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & 4 & \cdots \end{array} \\ \left[\begin{array}{ccccccc} 2/3 & 1/3 & 0 & \cdots & & & \\ q & r & p & 0 & \cdots & & \\ 0 & q & r & p & 0 & \cdots & \\ 0 & 0 & q & r & p & 0 & \cdots \\ 0 & 0 & 0 & \ddots & \ddots & \ddots & \\ \vdots & \vdots & \vdots & & & & \end{array}\right] \end{array},$$

where

$$\begin{aligned} q &= 2c/3 \\ r &= (2-c)/3 \\ p &= (1-c)/3, \end{aligned}$$

and c is the probability that I do a homework over the weekend if there is at least one to be done. From the transition matrix $\mathbf{P}$ we can write down the local balance equations as

$$\begin{aligned} \pi_0 \frac{1}{3} &= \pi_1 q \\ \pi_1 p &= \pi_2 q \\ &\vdots \\ \pi_i p &= \pi_{i+1} q \end{aligned}$$

Notice that each equation involves only adjacent pairs of states because the process only ever increases or decreases by one in any one step, and the diagonal elements of $\mathbf{P}$ do not enter into the equations because those give the transition probablities from i back to i.

Directly obtaining a recursion from these equations is now simple. We have

$$\begin{aligned} \pi_1 &= \frac{1}{3q}\pi_0, \\ \pi_2 &= \frac{p}{q}\pi_1 = \frac{p}{q}\frac{1}{3q}\pi_0, \end{aligned}$$

and, in general

$$\begin{aligned} \pi_{i+1} &= \frac{p}{q}\pi_i \\ &= \left(\frac{p}{q}\right)^2 \pi_{i-1} \\ &\vdots \\ &= \left(\frac{p}{q}\right)^i \pi_1 \\ &= \left(\frac{p}{q}\right)^i \frac{1}{3q}\pi_0. \end{aligned}$$

To obtain π_0, we can now use the constraint $\sum_{i=0}^{\infty} \pi_i = 1$ to write

$$\begin{aligned} \pi_0 \left[1 + \frac{1}{3q} + \frac{p}{q}\frac{1}{3q} + \left(\frac{p}{q}\right)^2 \frac{1}{3q} + \dots\right] &= 1 \\ \Rightarrow \pi_0 \left[1 + \frac{1}{3q}\sum_{i=0}^{\infty}\left(\frac{p}{q}\right)^i\right] &= 1. \end{aligned}$$

At this point we can see that for a stationary distribution to exist, the infinite sum above must converge, and this is true if and only if $p/q < 1$. In terms of c, this condition is

$$\frac{(1-c)/3}{2c/3} < 1 \Leftrightarrow 1 - c < 2c \Leftrightarrow c > \frac{1}{3},$$

verifying our condition for stability.

Assuming now that $c > 1/3$, we can evaluate the infinite sum as

$$\sum_{i=0}^{\infty} \left(\frac{p}{q}\right)^i = \frac{1}{1 - p/q},$$

which gives

$$\begin{aligned} \pi_0 \left[1 + \frac{1}{3q}\frac{1}{(1 - p/q)}\right] &= 1 \\ \Rightarrow \pi_0 \left[1 + \frac{1}{3(q-p)}\right] &= 1 \\ \Rightarrow \pi_0 \left[\frac{1 + 3(q-p)}{3(q-p)}\right] &= 1, \end{aligned}$$

or

$$\pi_0 = \frac{3(q-p)}{1 + 3(q-p)}.$$

Since $3(q-p) = 3(\frac{2c}{3} - \frac{1-c}{3}) = 3c - 1$, we have

$$\pi_0 = \frac{3c-1}{1+3c-1} = \frac{3c-1}{3c}.$$

Moreover, we also have π_i as

$$\begin{aligned} \pi_i &= \left(\frac{p}{q}\right)^{i-1} \frac{1}{3q}\pi_0 \\ &= \left(\frac{1-c}{2c}\right)^{i-1} \frac{1}{2c}\frac{3c-1}{3c} \\ &= \left(\frac{1-c}{2c}\right)^{i-1} \frac{3c-1}{6c^2}, \end{aligned}$$

a result we didn't obtain explicitly using generating functions. □

As this last example shows, it can be very useful to recognize when local balance might hold. In the example we didn't actually try to guess that it might hold, we just blindly tried to solve the local balance equations and got lucky. But there are a couple of things we can do to see if a Markov chain will satisfy the local balance equations without actually writing down the equations and trying to solve them:

- If there are two state i and j such that $p_{ij} > 0$ but $p_{ji} = 0$, then we can right away conclude that the stationary distribution $\boldsymbol{\pi}$ will not satisfy the local balance equations. This is because the equation

$$\pi_i p_{ij} = \pi_j p_{ji}$$

  will have 0 on the right hand side and, since $p_{ij} > 0$, will only be satisfied if $\pi_i = 0$. But, as we have seen, no stationary distribution can have this.

- If the process $\boldsymbol{X}$ only ever increases or decreases by one (or stays where it is) at each step, then the local balance equations will be satisfied. We have seen this in today's example. To see this more generally, we may refer to the flow interpretation of the local balance equations. Consider any state i. During any fixed interval of time, the number of transitions from i to $i+1$ must be within one of the number of transitions from $i+1$ to i because for each transition from i to $i+1$, in order to get back to state i we must make the transtion from $i+1$ to i. Therefore, in the long run, the proportion of transitions from i to $i+1$ must equal the proportion of transitions from $i+1$ to i. In other words,

$$\pi_i p_{i,i+1} = \pi_{i+1} \pi_i$$

  should be satisfied. But these are exactly the local balance equations in this case.

Reversibility: (Section 4.8)

There is deep connection between local balance and a property of Markov chains (and stochastic processes in general) called *reversibility*, or *time reversibility*. Just as not all Markov chains satisfy local balance, not all Markov chains are reversible.

Keep in mind that we are only talking about stationary Markov chains. Local balance and reversibility (and global balance as well) are properties of only stationary Markov chains. To imagine the notion of reversibility, we start out with a stationary Markov chain and then extend the time index back to $-\infty$, so that now our Markov chain is

$$\boldsymbol{X} = \{X_n : n \in \{\ldots, -2, -1, 0, 1, 2, \ldots\}\}$$

Imagine running the chain backwards in time to obtain a new process

$$\boldsymbol{Y} = \{Y_n = X_{-n} : n \in \{\ldots, -1, 0, 1, \ldots\}\}.$$

The process $\boldsymbol{Y}$ is called the *reversed chain*. Indeed, $\boldsymbol{Y}$ is also a Markov chain. To see this, note that the Markov property for the $\boldsymbol{X}$ chain can be stated in the following way: given the current state of the process, all future states are independent of the entire past up to just before the current time. That is, given X_n, if $k > n$, then X_k is independent of X_m for every $m < n$. But this goes both ways since independence is a symmetric property: if W is independent of Z then Z is independent of W for any random variables W and Z. So we can say: given X_n, if $m < n$, then X_m is independent of X_k for every $k > n$.

Therefore, we can see the Markov property of $\boldsymbol{Y}$, as

$$\begin{aligned} & P\left(Y_{n+1}=j \mid Y_n=i, Y_k=i_k \text{ for } k<n\right) \\ & \quad= P(X_{-(n+1)}=j \mid X_{-n}=i, X_{-k}=i_k \text{ for } k<n) \\ & \quad= P(X_{-(n+1)}=j \mid X_{-n}=i) \\ & \quad= P(Y_{n+1}=j \mid Y_n=i). \end{aligned}$$

So the reversed process $\boldsymbol{Y}$ is a Markov chain. Indeed, it is also stationary *and* has the same stationary distribution, say $\boldsymbol{\pi}$, as the $\boldsymbol{X}$ chain (since, for example, the long run proportion of time the $\boldsymbol{Y}$ chain spends in state i is obviously the same as the long run proportion of time that the $\boldsymbol{X}$ chain spends in state i, for any state i). **However**, the reversed chain $\boldsymbol{Y}$ does *not* in general have the same transition matrix as $\boldsymbol{X}$. In fact, we can explicity compute the transition matrix of the $\boldsymbol{Y}$ chain, using the fact that both the $\boldsymbol{X}$ chain and the $\boldsymbol{Y}$ chain are stationary with common stationary distribution $\boldsymbol{\pi}$. If we let $\mathbf{Q}$ denote the transition matrix of the $\boldsymbol{Y}$ chain (with entries q_{ij}), we have

$$\begin{aligned} q_{ij} &= P(Y_n=j \mid Y_{n-1}=i) \\ &= P(X_{-n}=j \mid X_{-(n-1)}=i) \\ &= \frac{P(X_{-n}=j, X_{-(n-1)}=i)}{P(X_{-(n-1)}=i)} \\ &= \frac{P(X_{-(n-1)}=i \mid X_{-n}=j)P(X_{-n}=j)}{P(X_{-(n-1)}=i)} \\ &= \frac{p_{ji}\pi_j}{\pi_i}, \end{aligned}$$

where p_{ij} is the one-step transition probability from state j to state i in the $\boldsymbol{X}$ chain.

We say of a stationary Markov chain $\boldsymbol{X}$ that it is *reversible*, or *time-reversible*, if the transition matrix of the reversed chain $\boldsymbol{Y}$ is the same as the transition matrix of $\boldsymbol{X}$; that is, $\mathbf{Q} = \mathbf{P}$. Note that the terminology is a little confusing. The reversed chain $\boldsymbol{Y}$ always exists but not every Markov chain $\boldsymbol{X}$ is reversible. Since we have computed q_{ij}, we can see exactly the conditions that will make $\boldsymbol{X}$ reversible:

$$\begin{aligned} \boldsymbol{X} \text{ is reversible} \quad & \text{if and only if} \quad q_{ij} = p_{ij} \\ & \text{if and only if} \quad p_{ij} = \frac{p_{ji}\pi_j}{\pi_i} \\ & \text{if and only if} \quad \pi_i p_{ij} = \pi_j p_{ji}. \end{aligned}$$

So here we see the connection between reversibility and local balance. A Markov chain $\boldsymbol{X}$ is reversible if and only if local balance is satisfied in equilibrium.

So, for example, to prove that a Markov chain $\boldsymbol{X}$ is reversible one can check whether a stationary distribution can be found that satisfies the local balance equations. You are asked to do this on one of the homework problems. In our example today, in finding the stationary distribution through the local balance equations, we have also shown that the process there is reversible.