Chapter 2

Markov chains

A *stochastic process* is a mathematical model for a system evolving randomly in time. Depending on the application, time my be modelled as *discrete* (e.g. $0, 1, 2, \ldots$) or *continuous* (e.g. the real interval $[0, \infty)$). Let Formally, a stochastic process is a sequence of random variables X_t, parametrized by t (time), defined on the same probability space and taking values in the same set. In applications, X_t might be the values of an asset (in discrete or continuous time), or the number of customers served by a shop during successive hours (discrete time), or the sequence of house numbers visited by a postman during his round (discrete time).

Later in the course we will look at some continuous-time processes. In this chapter we take time to be discrete (labelled by n) so the process is $X_0, X_1, X_2, \ldots$. The value of the random variable X_n is interpreted as the 'state' of the random system after n time steps.

In this chapter we consider an important class of stochastic processes called *Markov chains*. We shall study various methods to help understand the behaviour of Markov chains, in particular over long times.

2.1 What is a Markov chain?

A Markov chain is a special type of stochastic process with the 'Markov property', which states roughly that *given the present, the future is independent of the past*.

Markov chain theory presents systematic methods to study certain questions, and there are many important models that find widespread applications. We start by discussing an example. Conditional probability plays a central role in the topic.

Example. *Virus mutation.* Suppose that a virus can exist in two different strains α and β and in each generation either stays the same, or with probability $p < 1/2$ mutates to the other strain. Suppose the virus is in strain α initially, what is the probability that it is in the same strain after n generations?

We let X_n be the strain of the virus in the nth generation, which is a random variable with values in $\{\alpha, \beta\}$. We start with $X_0 = \alpha$. A key point here is that the random variables X_n and X_{n+1} are *not independent*; we'll come back to this later.

The probability that we are interested in is

$$p_n = \mathbb{P}(X_n = \alpha \mid X_0 = \alpha).$$

We can try to find this probability by considering what happens at time $n + 1$. We can have $X_{n+1} = \alpha$ in two ways: either (i) $X_n = \alpha$ and then the virus does not mutate, or (ii) $X_n = \beta$ and then the virus mutates. This leads us to the equation

$$p_{n+1} = p_n \cdot (1-p) + (1-p_n) \cdot p. \tag{2.1}$$

More carefully, what we are using here is the law of total probability with the partition $\{X_n = \alpha\}, \{X_n = \beta\}$:

$$\mathbb{P}(X_{n+1} = \alpha) = \mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha)\mathbb{P}(X_n = \alpha) + \mathbb{P}(X_{n+1} = \alpha \mid X_n = \beta)\mathbb{P}(X_n = \beta).$$

But $\mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha) = 1 - p$ (no mutation) and $\mathbb{P}(X_{n+1} = \alpha \mid X_n = \beta) = p$ (mutation), while $\mathbb{P}(X_n = \alpha) = p_n$ (by definition) and $\mathbb{P}(X_n = \beta) = 1 - p_n$ (complementary events). So we verify (2.1).

If we were being more careful, we would explicitly include the condition $X_0 = \alpha$ in all our expressions. This makes the calculation look a bit more cumbersome, but becomes useful for more involved examples. Then our starting point becomes:

$$\begin{aligned}\mathbb{P}(X_{n+1} = \alpha \mid X_0 = \alpha) = {} & \mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha, X_0 = \alpha)\mathbb{P}(X_n = \alpha \mid X_0 = \alpha) \\ & + \mathbb{P}(X_{n+1} = \alpha \mid X_n = \beta, X_0 = \alpha)\mathbb{P}(X_n = \beta \mid X_0 = \alpha).\end{aligned}$$

Now from (2.1) we obtain for the desired quantity p_n the recursion relation

$$p_{n+1} = p + (1-2p) \cdot p_n, \quad (n \geq 0),$$

with the initial condition $p_0 = 1$. This has a unique solution given by

$$p_n = \frac{1}{2} + \frac{1}{2}(1-2p)^n;$$

you can easily check that this works! As $n \to \infty$ this converges to the long term probability that the virus is in strain α, which is $1/2$ and therefore independent of the mutation probability p.

We return briefly to our comment about X_{n+1} and X_n not being independent. We saw that

$$\mathbb{P}(X_{n+1} = \alpha) = p_{n+1} = p + (1-2p)p_n < 1 - p,$$

for $n > 0$, since $p_n < 1$ for $n > 0$. But then, for $n > 0$,

$$\mathbb{P}(X_n = \alpha,\, X_{n+1} = \alpha) = (1-p)\mathbb{P}(X_n = \alpha) > \mathbb{P}(X_{n+1} = \alpha)\mathbb{P}(X_n = \alpha),$$

contradicting the definition of independence. Another way to express the last formula is

$$\mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha) > \mathbb{P}(X_{n+1} = \alpha),$$

for $n > 0$. Indeed, as $n \to \infty$ we know that the right-hand side here tends to $1/2$, while the left-hand side is always $(1-p) > 1/2$.

The theory of Markov chains provides a systematic approach to this and similar questions. The possible values of X_n constitute the *state space* S of the chain: here $S = \{\alpha, \beta\}$. The characteristic feature of a Markov chain is that *the past influences the future only via the present.* For example

$$\mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha, X_{n-1} = \alpha) = \mathbb{P}(X_{n+1} = \alpha \mid X_n = \alpha).$$

Here the state of the virus at time $n+1$ (future) does not depend on the state of the virus at time $n-1$ (past) if the state at time n (present) is already known. ■

2.2 Definition of discrete-time Markov chains

Suppose that S is a *discrete* (i.e. finite or countably infinite) set. A discrete-time Markov chain on S is a sequence of random variables X_n, $n \in \mathbf{Z}^+$, is described by:

(i) An *initial distribution* on the state space S given by $w = (w_i)_{i \in S}$ with $w_i \geq 0$ and $\sum_{i \in S} w_i = 1$.

(ii) An array of *one-step transition probabilities* p_{ij} for $i, j \in S$ with $p_{ij} \geq 0$ for all i, j and $\sum_{j \in S} p_{ij} = 1$ for all $i \in S$.

Then X_n is a Markov chain if the starting state X_0 is chosen according to w, i.e.,

$$\mathbb{P}(X_0 = i) = w_i, \tag{2.2}$$

for all $i \in S$, and the *Markov property* holds:

$$\mathbb{P}(X_n = i_n \mid X_{n-1} = i_{n-1}, X_{n-2} = i_{n-2}, \ldots, X_0 = i_0) = \mathbb{P}(X_n = i_n \mid X_{n-1} = i_{n-1}) = p_{i_{n-1}, i_n}, \tag{2.3}$$

for any $n \geq 1$ and any $i_n, i_{n-1}, \ldots, i_0 \in S$.

Intuitively, (2.3) says that *given the present, the future is independent of the past*: given the position at time $n-1$, the position at time n depends only on X_{n-1} and does not depend on $X_{n-2}, X_{n-3}, \ldots$.

Also note that $\mathbb{P}(X_n = j \mid X_{n-1} = i) = p_{ij}$ depends *only* on i and j, not on n. So $\mathbb{P}(X_8 = j \mid X_7 = i) = \mathbb{P}(X_1 = j \mid X_0 = i)$. This property is called *time-homogeneity*.

The one-step transition probabilities p_{ij} are conveniently written as a (square) matrix

$$P = (p_{ij})_{i,j \in S}.$$

The properties in (ii) above ensure that the entries in P are all nonnegative and the *rows all sum to* 1 (a matrix with these properties is called a *stochastic matrix*). P is the *one-step transition matrix* (or simply the transition matrix) of the Markov chain.

Definition 2.1. *(X_n) is a Markov chain on state space S with transition matrix $P = (p_{ij})$ and initial distribution $w = (w_i)$ if (2.2) and (2.3) hold.*

The initial distribution vector w gives the probabilities of starting in each state: $\mathbb{P}(X_0 = i) = w_i$. All jumps from i to j in one step occur with probability p_{ij}, so if we fix a state i and observe where we are going in the next time step, then the probability distribution of the next location in S has the probability mass function $(p_{ij})_{j \in S}$.

Example. A simple model of the weather from day to day has two states $S = \{W, D\}$ for Wet and Dry. If it is raining today, the probability that it will rain tomorrow is 0.8. If it is dry today, the probability that it will rain tomorrow is 0.4. Thus our one-step transition probabilities are

$$p_{WW} = 0.8, \quad p_{WD} = 0.2, \quad p_{DW} = 0.4, \quad p_{DD} = 0.6.$$

We can represent this as a transition matrix

$$P = \begin{pmatrix} 0.8 & 0.2 \\ 0.4 & 0.6 \end{pmatrix}.$$

Note that the rows sum to 1. To specify the Markov chain, we would also need to specify an initial distribution, e.g., $\mathbb{P}(X_0 = W) = 1$ — we start on a rainy day. ■

Example. Recall the virus example. The state-space is $S = \{\alpha, \beta\}$. Initially the virus is in strain α, hence the initial distribution is $w = (w_\alpha, w_\beta) = (1, 0)$. The transition matrix is

$$P = \begin{pmatrix} 1-p & p \\ p & 1-p \end{pmatrix}.$$

■

Example. *A maze used for training rats.* There are 5 compartments labelled $1, \ldots, 6$ in the maze.

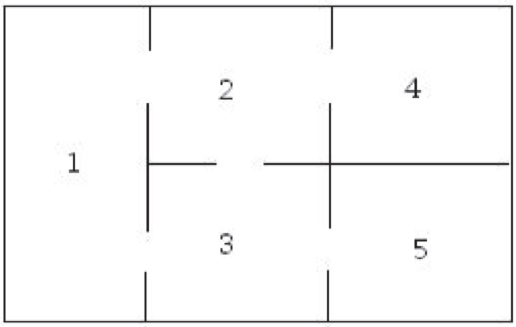

An untrained rat moves in a Markov chain according to the transition matrix

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

■

We will show that (2.2) and (2.3) completely specify the *law of evolution* of the Markov chain, in the following sense.

Theorem 2.2. *Suppose that (X_n) is a Markov chain with transition matrix $P = (p_{ij})$ and initial distribution $w = (w_i)$. Then for any $n \in \mathbf{Z}^+$ and any $i_0, i_1, \ldots, i_n \in S$,*

$$\mathbb{P}(X_0 = i_0, X_1 = i_1, \ldots, X_n = i_n) = w_{i_0} p_{i_0 i_1} \cdots p_{i_{n-1} i_n}.$$

Proof. We use the multiplication rule for conditional probabilities. For clarity we start with $n = 1$.

$$\mathbb{P}(X_1 = i_1, X_0 = i_0) = \mathbb{P}(X_1 = i_1 \mid X_0 = i_0)\mathbb{P}(X_0 = i_0) = p_{i_0 i_1} w_{i_0}.$$

Similarly, for $n = 2$,

$$\begin{aligned} \mathbb{P}(X_0 = i_0, X_1 = i_1, X_2 = i_2) &= \mathbb{P}(X_2 = i_2 \mid X_1 = i_1, X_0 = i_0)\mathbb{P}(X_1 = i_1 \mid X_0 = i_0)\mathbb{P}(X_0 = i_0) \\ &= p_{i_1 i_2} p_{i_0 i_1} w_{i_0}, \end{aligned}$$

as claimed. The same argument works for general n. □

2.3 Stationary distributions

Let π be a probability mass function on S, i.e., $\pi = (\pi_i)_{i\in S}$ with $\pi_i \geq 0$ and $\sum_{i\in S} \pi_i = 1$. The distribution π is called a *stationary distribution* or an *invariant distribution* for a Markov chain with transition matrix P if, thinking of π as a row vector,

$$\pi P = \pi,$$

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

If such a π exists, then we can use it as an initial distribution to the following effect:

$$\mathbb{P}(X_1 = j) = \sum_{i\in S} \mathbb{P}(X_0 = i,\, X_1 = j) = \sum_{i\in S} \pi_i p_{ij} = \pi_j = \mathbb{P}(X_0 = j),$$

and, more generally, if we start from the stationary distribution,

$$\mathbb{P}(X_n = j) = \pi_j,$$

for all n and all j. In other words, if we start with X_0 distributed according to π, then X_n remains distributed according to π for all times n. In this sense the distribution of the Markov chain is in *equilibrium.*

Note that π is a *left eigenvector* of P corresponding to the eigenvalue 1. We return later in this chapter to the question of when such a π can be found.

2.4 Further examples

Example. *Random walk on a finite graph.* A particle is moving on the graph below by starting on the top left vertex and at each time step moving along one of the adjacent edges to a neighbouring vertex, choosing the edge with equal probability and independently of all previous movements.

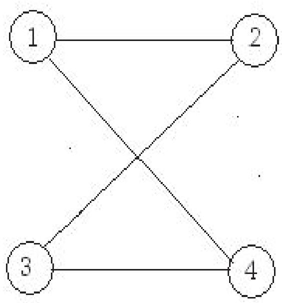

There are four vertices, which we enumerate from left to right by $\{1, \ldots, 4\}$. At time $n = 0$ we are in vertex 1. Hence $w = (1, 0, 0, 0)$ is the initial distribution. Each vertex has exactly two neighbours, so that it jumps to each neighbour with probability 1/2. This gives

$$P = \begin{pmatrix} 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \end{pmatrix}.$$

Is there a stationary distribution $\pi = (\pi_1, \pi_2, \pi_3, \pi_4)$? We attempt to solve $\pi P = \pi$, which in this case gives the system of simultaneous equations

$$\begin{aligned}
\frac{1}{2}\pi_2 + \frac{1}{2}\pi_4 &= \pi_1 \\
\frac{1}{2}\pi_1 + \frac{1}{2}\pi_3 &= \pi_2 \\
\frac{1}{2}\pi_2 + \frac{1}{2}\pi_4 &= \pi_3 \\
\frac{1}{2}\pi_1 + \frac{1}{2}\pi_3 &= \pi_4.
\end{aligned}$$

So $\pi_1 = \pi_3$ and $\pi_2 = \pi_4$. Then the first equation says $\pi_1 = \pi_2$. Hence $\pi_1 = \pi_2 = \pi_3 = \pi_4$. We also want $\pi_1 + \pi_2 + \pi_3 + \pi_4 = 1$. So in this case there is a stationary distribution $\pi = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. ■

Example. *Simple random walk.* A particle jumps about at random on the set $\mathbf{Z}$ of integers. At time 0 the particle is in fixed position $s \in \mathbf{Z}$. At each time $n \in \mathbf{Z}^+$ a coin with probability p of heads and $q = 1-p$ of tails is tossed. If the coin falls heads, then the particle jumps one position to the right, if the coin falls tails the particle jumps one position to the left. For $n \in \mathbf{Z}^+$ the position X_n of the particle is a Markov chain with initial distribution

$$w_i = \mathbb{P}(X_0 = i) = \begin{cases} 1 & \text{if } i = s \\ 0 & \text{if } i \neq s \end{cases},$$

and one-step transition probabilities

$$p_{ij} = \mathbb{P}(X_{n+1} = j \mid X_n = i) = \begin{cases} p & \text{if } j - i = +1 \\ q & \text{if } j - i = -1 \\ 0 & \text{if } |j - i| \neq 1 \end{cases}.$$

This model is related to the *gambler's ruin* problem: A gambler starts with fortune £s, and repeatedly plays a game in which he wins £1 (with probability p) or loses £1 (with probability q). We stop the game when $X_n = 0$ (ruin!) or when $X_n = m$ for some predetermined target wealth $m > s$. What is the probability that the gambler reaches his target before going bankrupt? We'll return to such problems later. ■

The main aim of this part of the course is to study behaviour of a Markov chain X_n for large times n. Fundamental questions include:

- After a long time, what is the probability of finding X_n in a particular state?
- How much time does the Markov chain spend in each of the different states? What is the chain's favourite state? Does the answer depend on the starting position?
- How long does it take, on average, to get from some given state to another?

2.5 n-step transition probabilities

The initial distribution tells us the distribution of X_0. What is the distribution of X_1? We can use the law of total probability, with the partition given by the possible values of X_0, to see that

$$\mathbb{P}(X_1 = j) = \sum_{i \in S} \mathbb{P}(X_1 = j \mid X_0 = i)\mathbb{P}(X_0 = i) = \sum_{i \in S} p_{ij} w_i,$$

by definition of the Markov chain. So the one-step transition probabilities combined with the initial distribution tell us the distribution of X_1. A similar argument shows that

$$\mathbb{P}(X_n = j) = \sum_{i \in S} \mathbb{P}(X_n = j \mid X_0 = i)\mathbb{P}(X_0 = i),$$

so we can work out the distribution of X_n *if* we know $\mathbb{P}(X_n = j \mid X_0 = i)$. Instead of a *one-step* probability, $\mathbb{P}(X_n = j \mid X_0 = i)$ is an *n-step transition probability*. We write

$$p_{ij}^{(n)} = \mathbb{P}(X_n = j \mid X_0 = i).$$

So $p_{ij}^{(1)} = p_{ij}$ that we had before. How can we work out $p_{ij}^{(n)}$? Consider $n = 2$. We can use a partition argument again, summing over all the possible intermediate positions X_1:

$$\begin{aligned}
p_{ij}^{(2)} = \mathbb{P}(X_2 = j \mid X_0 = i) &= \sum_{k \in S} \mathbb{P}(X_2 = j,\, X_1 = k \mid X_0 = i) \\
&= \sum_{k \in S} \mathbb{P}(X_2 = j \mid X_1 = k,\, X_0 = i)\mathbb{P}(X_1 = k \mid X_0 = i) \quad \text{[conditional probability]} \\
&= \sum_{k \in S} \mathbb{P}(X_2 = j \mid X_1 = k)\mathbb{P}(X_1 = k \mid X_0 = i) \quad \text{[Markov property]} \\
&= \sum_{k \in S} p_{kj} p_{ik},
\end{aligned}$$

since these are both one-step transitions. But this is just matrix multiplication!

$$p_{ij}^{(2)} = \sum_{k \in S} p_{kj} p_{ik} = (P^2)_{ij},$$

so $p_{ij}^{(2)}$ is the i, j entry in the matrix $P^2 = P \times P$.

Example. Recall the first example in Section 2.4. In this case we get

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

Intuitively, this makes sense: consider for example the first row of P^2. If one is in state 1, then in 2 steps one can either get back to state 1 (either via state 2 or state 4) or get to state 3 (again via 2 or 4). Hence $p_{11}^{(2)} = p_{13}^{(2)} = 1/2$. It is impossible in two steps to get from 1 to 2 or from 1 to 4. Hence $p_{12}^{(2)} = p_{14}^{(2)} = 0$. ■

We return to the general setting. By a similar argument to that above, we have that more generally,

$$p_{ij}^{(n)} = \mathbb{P}(X_{m+n} = j \mid X_m = i) = (P^n)_{ij}.$$

Moreover, if the vector $w = (w_i)_{i \in S}$ is the initial distribution, we get

$$\begin{aligned}
\mathbb{P}(X_n = j) &= \sum_{k \in S} \mathbb{P}(X_0 = k)\mathbb{P}(X_n = j \mid X_0 = k) \\
&= \sum_{k \in S} w_k (P^n)_{kj}.
\end{aligned}$$

Hence we get

$$\mathbb{P}(X_n = j) = (wP^n)_j.$$

Hence, if we can calculate the matrix power P^n, we can find the distribution of X_n and the n-step transition probabilities.

If n is large (recall that we are particularly interested in the long term behaviour of the process!) it is not advisable to calculate P^n directly, but there are some more efficient methods, which we shall discuss shortly.

First we make an important comment. Suppose that π is a stationary distribution for the Markov chain. Then $\pi P = \pi$, and, iterating,

$$\pi P^n = \pi.$$

We write $\mathbb{P}_\pi$ to mean that we start the Markov chain with initial distribution π, i.e., $\mathbb{P}_\pi(X_0 = i) = \pi_i$. Then $\mathbb{P}_\pi(X_n = i) = \pi_i$ for all i, which means that

$$\lim_{n\to\infty} \mathbb{P}_\pi(X_n = i) = \pi_i. \tag{2.4}$$

We will see later that (2.4) actually holds for *any* initial distribution for a large class of Markov chains.

It is not hard to see that *if* P^n tends to a limit as $n \to \infty$, then the limit is closely connected to any stationary distribution π. Indeed, consider $P^{n+1} = P^n \cdot P$. If $P^n \to \Pi$ for some limiting matrix Π, then $\Pi = \Pi P$, and any row $\mathbf{y}$ of the matrix Π must satisfy $\mathbf{y} = \mathbf{y}P$, i.e., the rows of the limit matrix Π are left eigenvectors of P. But when does P^n tend to a limit? We return to this discussion later.

2.6 Hitting probabilities

As always, X_n is a Markov chain with discrete state space S and one-step transition matrix P. We now introduce a notation which allows us to look at the Markov chain and vary the initial distributions; this point of view turns out to be very useful for certain calculations.

We use $\mathbb{P}_i$ to indicate that the chain starts with X_0 in the fixed state $i \in S$. We talk about the 'law' of the Markov chain started from i. In other words, for all $x_1, \dots, x_n \in S$,

$$\mathbb{P}_i(X_1 = x_1, \dots, X_n = x_n) = p_{ix_1} p_{x_1x_2} \cdots p_{x_{n-1}x_n}.$$

We can also think of $\mathbb{P}_i$ as the law of the Markov chain conditioned on $X_0 = i$, i.e.

$$\mathbb{P}_i(A) = \mathbb{P}(A \mid X_0 = i).$$

The law $\mathbb{P}_w$ of X_n with initial distribution w is then given as the *mixture* or weighted average of the laws $\mathbb{P}_i$, i.e., for the initial distribution $w = (w_i)_{i\in S}$ of the Markov chain and any event A,

$$\mathbb{P}_w(A) = \sum_{i\in S} \mathbb{P}_w(A \cap \{X_0 = i\}) = \sum_{i\in S} \mathbb{P}_w(X_0 = i)\mathbb{P}_i(A) = \sum_{i\in S} w_i \mathbb{P}_i(A).$$

We also use this notation for expected values, i.e., $\mathbb{E}_i$ refers to the expectation with respect to $\mathbb{P}_i$ and $\mathbb{E}_w$ refers to expectation with respect to $\mathbb{P}_w$.

Remember that in expressions such as $\mathbb{P}(X_{n+1} = x_{n+1} \mid X_n = x_n) = p_{x_nx_{n+1}}$ we do not have to write the index w because if we are told the state at time n we can forget about the initial distribution.

We define the *first hitting time* T_j of a state $j \in S$ by

$$T_j := \min\{n \geq 0 \,:\, X_n = j\},$$

and understand that $T_j := \infty$ if the set is empty. Note that we *do count* X_0, so if we start from j, $T_j = 0$. Also observe the following equality of events,

$$\{T_j = \infty\} = \{X_n \text{ never hits } j\}.$$

Often we are interested in the probability, started from some state $i \in S$, that we reach state j before state k, i.e.,

$$\mathbb{P}_i(T_j < T_k).$$

We can compute such probabilities using the powerful *first step method*. We illustrate with some examples.

Example. Theseus is lost in a (small) Labyrinth consisting of 4 rooms $\{1, 2, 3, 4\}$. He moves randomly according to a Markov chain X_n with $X_0 = 1$ and transition matrix

$$P = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{4} & \frac{1}{2} & 0 & \frac{1}{4} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

States 3 and 4 are called *absorbing states* since once entered, the process never leaves. State 3 is the way out and state 4 is the Minotaur's lair. What is the probability that Theseus escapes?

We are interested in $\mathbb{P}_1(T_3 < T_4)$. In order to solve this problem, we will actually find $y_i = \mathbb{P}_i(T_3 < T_4)$ for all possible starting states i. We set up a system of simultaneous equations for y_i by conditioning on the first step of the Markov chain. Using the law of total probability with the partition $\{X_1 = j\}$, $j \in S$,

$$\begin{aligned} \mathbb{P}_i(T_3 < T_4) &= \sum_{j \in S} \mathbb{P}_i(X_1 = j, T_3 < T_4) \\ &= \sum_{j \in S} \mathbb{P}_i(X_1 = j)\mathbb{P}_i(T_3 < T_4 \mid X_1 = j) \\ &= \sum_{j \in S} p_{ij}\mathbb{P}_j(T_3 < T_4), \end{aligned}$$

using the Markov property. In other words,

$$y_i = \sum_{j \in S} p_{ij} y_j.$$

In our example, we have the *boundary conditions* $y_3 = 1$ and $y_4 = 0$, so our equations become

$$y_1 = \frac{1}{2}y_2 + \frac{1}{2} \cdot 1, \quad y_2 = \frac{1}{4}y_1 + \frac{1}{2}y_2 + \frac{1}{4} \cdot 0.$$

Simplifying the second of these we get $y_1 = 2y_2$, and substituting into the first gives $y_1 = \frac{2}{3}$, $y_2 = \frac{1}{3}$. So Theseus escapes with probability 2/3. ■

The kind of calculation used in the previous example is known as the *first step method*.

Example. *Gambler's ruin.* A gambler bets repeatedly on the toss of a biased coin. On each toss, the gambler wins unit wealth if the coin comes down a head (an event of probability $p \in (0,1)$) and loses unit wealth if the coin is a tail (probability $q = 1-p$). The game stops when the gambler becomes *ruined*, i.e., when his wealth reaches 0, or when his wealth reaches a particular target level m. Suppose the gambler has initial fortune $k \in \{1, 2, \ldots, m-1\}$. What is the probability that the gambler is ruined? Or makes his fortune?

We can formulate this problem as a random walk on $\{0, 1, \ldots, m\}$ with two absorbing states (0 and m). The transition probabilities are $p_{00} = 1$, $p_{mm} = 1$ and for $0 < i < m$, $p_{i\,i+1} = p$, $p_{i\,i-1} = q$. We start at $X_0 = k$. We want to compute

$$y_k = \mathbb{P}(\text{hit } m \text{ before } 0 \mid X_0 = k) = \mathbb{P}_k(T_m < T_0),$$

for $0 \le k \le m$. Clearly $y_0 = 0$ and $y_m = 1$. By conditioning on the first step, we get for $1 \le k \le m-1$:

$$y_k = p y_{k+1} + q y_{k-1}.$$

One way to solve this system of equations is to set $\Delta_k = y_k - y_{k-1}$. Then $q\Delta_k = p\Delta_{k+1}$, or

$$\Delta_{k+1} = \frac{q}{p}\Delta_k,$$

for $1 \le k \le m-1$. So $\Delta_1 = y_1 - y_0 = y_1$ and $\Delta_k = (q/p)^{k-1}\Delta_1 = (q/p)^{k-1}y_1$ for $1 \le k \le m$.

There are two cases that need to be handled separately. First suppose that $p \ne 1/2$, so $(p/q) \ne 1$. Then we have

$$y_k = y_k - y_0 = \sum_{i=1}^{k} \Delta_i = \sum_{i=1}^{k} (q/p)^{i-1} y_1 = y_1 \frac{1-(q/p)^k}{1-(q/p)}.$$

Using the boundary condition $y_m = 1$ gives

$$y_1 = \frac{1-(q/p)}{1-(q/p)^m}.$$

Thus

$$y_k = \frac{1-(q/p)^k}{1-(q/p)^m},$$

when $p \ne 1/2$.

If $p = 1/2$, we see $\Delta_k = \Delta_1 = y_1$ for $k \ge 1$, so

$$y_k = k y_1$$

for $1 \le k \le N$. This time $y_m = 1$ implies $y_1 = 1/m$ so that

$$y_k = \frac{k}{m},$$

for $0 \le k \le N$, when $p = 1/2$. ■

Example. *Refinery.* A stream of gas consists of equal numbers of two types of particle: A and B. It is desired to separate the gas into its constituent particles. There is a refining column consisting of chambers labelled $0, 1, 2, \ldots, N$ for some even integer $N \ge 2$, where 0 is at the base of the column and N at the top. Gas is removed from chambers 0 and N. The input gas enters at chamber k. In each time step, particles of type A drift up to the chamber above with probability p_A, and down to the chamber below with probability $1 - p_A$. Similarly, particles of type B drift up to the chamber above with probability p_B, and down to the chamber below with probability $1 - p_B$. Suppose $p_A > 1/2$ but $p_B < 1/2$. Then one would hope to collect all particles of type A at the top of the column and all particles of type B at the base.

Given $p_A = 2/3$, $p_B = 1/3$, what is the optimum choice of k (in terms of N) such that the total proportion of correctly collected particles is maximized? For this optimal k, how many chambers should be used to ensure that at least 99% of particles are correctly separated?

By the gambler's ruin problem, we have

$$\mathbb{P}(\text{particle } A \text{ sorted correctly}) = \frac{1 - 2^{-k}}{1 - 2^{-N}},$$

and

$$\mathbb{P}(\text{particle } B \text{ sorted correctly}) = \frac{1 - 2^{-(N-k)}}{1 - 2^{-N}}.$$

Since 50% of particles are of type A and 50% of type B, in the long run (by the *law of large numbers*), the total proportion of particles sorted correctly is then

$$\frac{1}{2} \cdot \frac{1 - 2^{-k}}{1 - 2^{-N}} + \frac{1}{2} \cdot \frac{1 - 2^{-(N-k)}}{1 - 2^{-N}}.$$

This is maximized when $k = N/2$, assuming N is an even integer. Then, for any given particle,

$$\mathbb{P}(\text{sorted correctly}) = \frac{1 - 2^{-N/2}}{1 - 2^{-N}},$$

and this is greater than 0.99 provided $N \geq 14$. ■

2.7 Expected hitting times

In this section we ask: what is the expected time until we reach a certain state, or set of states? For example, in the gambler's ruin problem, how long do we expect the game to last? Such questions can be answered in a similar manner to the study of hitting probabilities, again using the first step method. Recall $T_j = \min\{n \geq 0 : X_n = j\}$.

Example. In the 3-pub example, what is the expected time to reach the third pub? Remember $S = \{1, 2, 3\}$, $X_0 = 1$, and

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

We want to compute $\mathbb{E}(T_3 \mid X_0 = 1)$. We will actually solve a system of equations for

$$z_i = \mathbb{E}_i(T_3) = \mathbb{E}(T_3 \mid X_0 = i).$$

Condition on the first step to get, for $i \neq 3$,

$$\mathbb{E}_i(T_3) = \sum_{j \in S} \mathbb{P}_i(X_1 = j) \mathbb{E}(T_3 \mid X_1 = j).$$

What is $\mathbb{E}(T_3 \mid X_1 = j)$? Well, this is exactly the same as $\mathbb{E}_j(T_3)$ but starting 1 time step later. So

$$\mathbb{E}_i(T_3) = \sum_{j \in S} \mathbb{P}_i(X_1 = j)\,(1 + \mathbb{E}_j(T_3))\,.$$

So we get

$$z_i = 1 + \sum_{j \in S} p_{ij} z_j$$

imposing the boundary condition $z_3 = 0$. This gives the equations

$$z_1 = 1 + \frac{1}{2} z_1 + \frac{1}{2} z_2, \quad z_2 = 1 + \frac{1}{2} z_1 + \frac{1}{2} \cdot 0.$$

Substituting for z_2 in the first equation gives $\frac{1}{2} z_1 = 1 + \frac{1}{2} + \frac{1}{4} z_1$, so $z_1 = 6$ and $z_2 = 4$. ■

Example. Consider a Markov chain with state space $S = \{A, B, C, D\}$ and jump probabilities given by the diagram.

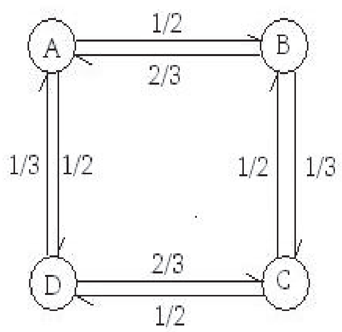

1. Find the expected time until the chain started in C reaches A.
2. What is the probability that the chain started in A reaches the state C before B?

First we solve part 1. Let $z_i = \mathbb{E}_i(T_A)$. By considering the first step and using the Markov property we get

$$\begin{aligned} z_C &= 1 + \frac{1}{2} z_B + \frac{1}{2} z_D \\ z_B &= 1 + \frac{2}{3} 0 + \frac{1}{3} z_C \\ z_D &= 1 + \frac{1}{3} 0 + \frac{2}{3} z_C. \end{aligned}$$

Hence,

$$z_C = 1 + \tfrac{1}{2}(1 + \tfrac{1}{3} z_C) + \tfrac{1}{2}(1 + \tfrac{2}{3} z_C) = 2 + \tfrac{1}{2} z_C,$$

which implies that $\mathbb{E}_C(T_A) = z_C = 4$. The expected time until the chain started in C reaches A is 4.

For part 2, let $y_i = \mathbb{P}_i(T_C < T_B)$. By the first-step method we get

$$\begin{aligned} y_A &= \frac{1}{2}y_D + \frac{1}{2}0, \\ y_D &= \frac{1}{3}y_A + \frac{2}{3}1. \end{aligned}$$

Hence $2y_A = (1/3)y_A + (2/3)$, which gives $y_A = 2/5$. Hence the probability of hitting C before B when we start in A is $2/5$. ■

2.8 Recap

We are working with a Markov chain X_n on a state space S, with one-step transition probabilities $p_{ij} = \mathbb{P}(X_{n+1} = j \mid X_n = i)$, $i, j \in S$. The transition matrix is $P = (p_{ij})_{i,j\in S}$. A probability distribution $\pi = (\pi_i)_{i\in S}$ is stationary if it solves $\pi P = \pi$ (with π as a row vector). The n-step transition probabilities of the Markov chain are $p_{ij}^{(n)} = \mathbb{P}(X_{m+n} = j \mid X_m = i)$, $i, j \in S$; $p_{ij}^{(n)}$ is in fact the (i, j) entry of the matrix power $P^n = P \times \cdots \times P$.

We want to investigate $p_{ij}^{(n)}$ as $n \to \infty$, i.e., the long-time behaviour.

In several examples we saw (e.g. by diagonalization) that P^n had a limit as $n \to \infty$, and moreover that the limit was the matrix whose rows were each the stationary distribution π, i.e.,

$$p_{ij}^{(n)} \to \pi_j, \tag{2.5}$$

as $n \to \infty$; note this did not depend on i (i.e., where the chain started). When is (2.5) true?

There are (at least) 3 obstructions, and we have encountered some of these in passing already.

(i) There might be no stationary distribution, i.e., $\pi P = \pi$ may not have a solution for any probability distribution π. Then (2.5) could not hold, as there would be no right-hand side!

(ii) Absorbing states cause problems. For example, if a Markov chain on $\{1, 2, 3\}$ has transition matrix

$$P = \begin{pmatrix} 1 & 0 & 0 \\ 1/2 & 0 & 1/2 \\ 0 & 0 & 1, \end{pmatrix}$$

then 1 and 3 are absorbing states, and, for example, $p_{11}^{(n)} = 1$ for all n but $p_{31}^{(n)} = 0$ for all $n \geq 1$. So in this case the starting position is important, whereas in (2.5) we want the limit of $p_{ij}^{(n)}$ to depend only on j and not on i.

(iii) The phenomenon of *periodicity*, which we describe precisely below, also causes a problem. We saw a previous example in which

$$P = \begin{pmatrix} 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \end{pmatrix}.$$

In this example,

$$P^2 = \begin{pmatrix} 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \end{pmatrix},$$

and in fact $P^n = P$ for all odd n but $P^n = P^2$ for all even n. Thus P^n alternates, and so does not have a limit. Note that in this example we already saw that $\pi = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ *was* a stationary distribution for this Markov chain.

2.9 Periodicity and irreducibility

Definition 2.3. *The* period *of a state $i \in S$ is the greatest common divisor of the set of integers n for which it is possible for the Markov chain to go from i back to i in n steps, i.e.,*

$$\textit{period}(i) = \textit{gcd}\{n : p_{ii}^{(n)} > 0\}.$$

Example. Consider the gambler's ruin Markov chain for $p \in (0, 1)$. The states 0 and m are absorbing, while the states $\{1, 2, \ldots, m-1\}$ constitute a communicating class, and all have period 2. Starting from any state $1 \leq i \leq m-1$, it necessarily takes an even number of steps to get back to i. ■

Periodicity complicates the limiting behaviour of a Markov chain: see the random walk on a graph example from Section 2.4.

We say that the state i *communicates* with the state j, and write $i \leftrightarrow j$ if there exist $n, m \geq 0$ with $p_{ij}^{(n)} > 0$ and $p_{ij}^{(m)} > 0$. The relation $\leftrightarrow$ is an equivalence relation on the state space S, meaning that

- $i \leftrightarrow i$,
- $i \leftrightarrow j$ implies $j \leftrightarrow i$,
- $i \leftrightarrow j$ and $j \leftrightarrow k$ implies $i \leftrightarrow k$.

Only the last statement is nontrivial. To prove it assume $i \leftrightarrow j$ and $j \leftrightarrow k$. Then there exist $n_1, n_2 \geq 0$ with $p_{ij}^{(n_1)} > 0$ and $p_{jk}^{(n_2)} > 0$. Then,

$$\begin{aligned} p_{ik}^{(n_1+n_2)} &= \mathbb{P}_i(X_{n_1+n_2} = k) \\ &\geq \mathbb{P}_i(X_{n_1} = j,\, X_{n_1+n_2} = k) \\ &\geq \mathbb{P}_i(X_{n_1} = j)\mathbb{P}_i(X_{n_1+n_2} = k \mid X_{n_1} = j) \\ &= p_{ij}^{(n_1)} p_{jk}^{(n_2)} > 0. \end{aligned}$$

Similarly, there exist m_1, m_2 with $p_{ki}^{(m_1+m_2)} > 0$ and hence $i \leftrightarrow k$.

Since $\leftrightarrow$ is an equivalence relation, we can define the corresponding equivalence classes, which are called *communicating classes*. The class of i consists of all $j \in S$ with $i \leftrightarrow j$.

We have already encountered the concept of an *absorbing state*, which is a state $i \in S$ with $p_{ii} = 1$. In such a case, $\{i\}$ is a communicating class of size 1.

Definition 2.4. *A Markov chain is called* irreducible *if all states communicate with each other, i.e., if S is a single communicating class.*

A property of a state is a *class property* if whenever it holds for one state, it holds for all states in the same class. A fact that we will not prove is that *periodicity is a class property*. In particular, for an irreducible Markov chain, all states have the same period. So if an irreducible chain has a state with period d, we say that the Markov chain has period d.

Definition 2.5. *If all states of a Markov chain have period* 1*, we say that the Markov chain is* aperiodic*.*

In many cases our Makrov chain will be irreducible and aperiodic, which simplifies the theory.

2.10 Limit theory

As promised, we are now in a position to describe the long-time behaviour of Markov chains under certain conditions. We will see an important connection between stationary distributions and expected hitting times.

We now assume that the Markov chain X_n is irreducible. Then all states have the same period d. Recall that the chain is called *aperiodic* if $d = 1$. We have previously worked with the hitting times $T_i = \min\{n \geq 0 : X_n = i\}$. In the next result we need a slight variation on this, given by $\tau_i = \min\{n \geq 1 : X_n = i\}$. This time we *do not* count X_0, so starting from i, τ_i is the time of the first *return* to i. Quantities $\mathbb{E}_j(\tau_i)$ can be computed by the first step method exactly as for $\mathbb{E}_j(T_i)$, without the boundary conditions.

Theorem 2.6 (Limiting behaviour)**.** *Let X_n be an irreducible Markov chain on state space S. Then the following statements are equivalent:*

- *There is a unique stationary distribution $\pi = (\pi_i)$ satisfying $\pi P = \pi$.*
- *For some $i \in S$, $\mathbb{E}_i(\tau_i) < \infty$.*
- *For all $i \in S$, $\mathbb{E}_i(\tau_i) < \infty$.*

If these conditions hold, the Markov chain is called positive-recurrent*. Moreover, for a positive-recurrent Markov chain,*

(i) For any $i \in S$, the stationary distribution π is given by

$$\pi_i = \frac{1}{\mathbb{E}_i(\tau_i)} > 0.$$

(ii) For all initial distributions w, with probability 1*,*

$$\frac{1}{n}\Big(\#\textit{visits to state } i \textit{ by time } n\Big) \longrightarrow \pi_i.$$

(iii) In the aperiodic case $d = 1$, for all initial distributions w,

$$\lim_{n\to\infty} \mathbb{P}_w(X_n = j) = \pi_j, \ \textit{for all } j \in S.$$

Part (iii) gives the full generalization of (2.4) that we saw earlier. Note also how the statement in part (ii) entails the fact that

$$\frac{\#\text{visits to state } i \text{ by time } n}{\#\text{visits to state } j \text{ by time } n} \longrightarrow \frac{\pi_i}{\pi_j}.$$

So π_i can be interpreted as the long-run proportion of time spent in state i.

We do not prove Theorem 2.6, but instead show how it works in some examples. Also, you can now revisit many of our earlier examples armed with this result. Note that the theorem tells us when π will exist, and usefully relates the two quantities π_i and $\mathbb{E}_i(\tau_i)$ that we can calculate in very different ways! A useful fact is the following:

Theorem 2.7. *If X_n is an irreducible Markov chain on a* finite *state-space S, X_n is positive-recurrent, so has a unique stationary distribution π. If, furthermore, X_n is aperiodic, then each row of P^n tends to π.*

Example. In the weather example, $P = \begin{pmatrix} 0.8 & 0.2 \\ 0.4 & 0.6 \end{pmatrix}$ specifies an irreduicible aperiodic Markov chain with unique stationary distribution $\pi = (\frac{2}{3}, \frac{1}{3})$. (Check: $\pi P = \pi$.) So over a long time period, 2/3 of the days are wet and 1/3 dry (part (ii)). Also, by part (iii), after a long time the probability of a day being wet is 2/3. According to part (i), if $x_i = \mathbb{E}_i(\tau_1)$ then $x_1 = 1/\pi_1 = 3/2$. The linear equations that we get for the x_i using the one step method are

$$x_1 = 1 + 0.8 \cdot 0 + 0.2 \cdot x_2, \quad \text{and} \quad x_2 = 1 + 0.4 \cdot 0 + 0.6 x_2,$$

which give $x_1 = 3/2$ and $x_2 = 5/2$ so are consistent! Note that if we had been calculating $z_i = \mathbb{E}_i(T_1)$ instead, we would have got the same second equation but we would have imposed the boundary condition $z_1 = 0$. ■

Example. We give an example showing that the *aperiodicity* condition in part (iii) of Theorem 2.6 is necessary. Condsider an irreducible Markov chain on state space $S = \{0, 1, 2, 3, 4, 5\}$ with transition matrix

$$P = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ \frac{1}{5} & 0 & \frac{4}{5} & 0 & 0 & 0 \\ 0 & \frac{2}{5} & 0 & \frac{3}{5} & 0 & 0 \\ 0 & 0 & \frac{3}{5} & 0 & \frac{2}{5} & 0 \\ 0 & 0 & 0 & \frac{1}{5} & 0 & \frac{4}{5} \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

This Markov chain has period 2, since from any state it takes an even number of steps to get back to that state. So there cannot be any limit for P^n, since, for example, $p_{ii}^{(n)}$ is zero if n is odd but positive if n is even. In fact there are *two* limit matrices, one for even powers and the other for odd powers. ■