E-mail: yms@statslab.cam.ac.uk

MARKOV CHAINS

7. Convergence to equilibrium. Long-run proportions

*Convergence to equilibrium for irreducible, positive recurrent, aperiodic chains *and proof by coupling*. Long-run proportion of time spent in a given state.*

Convergence to equilibrium means that, as the time progresses, the Markov chain 'forgets' about its initial distribution λ. In particular, if $\lambda = \delta^{(i)}$, the Dirac delta concentrated at i, the chain 'forgets' about initial state i. Clearly, this is related to properties of the n-step matrix P^n as $n \to \infty$. Consider first the case of a finite chain.

Theorem 7.1. *Suppose that a finite $m \times m$ transition matrix P^n converges, in each entry, to a limiting matrix $\mathbb{P}i = (\pi_{ij})$:*

$$\lim_{n\to\infty} p_{ij}^{(n)} = \pi_{ij}, \ \ \forall \ i, j \in I. \tag{7.1}$$

Then (a) *every row $\pi^{(i)}$ of $\mathbb{P}i$ is an equilibrium distribution*

$$\pi^{(i)} P = \pi^{(i)} \ \text{ or } \ \pi_{ij} = \sum_l \pi_{il} p_{lj}.$$

(b) *If P is irreducible then all rows $\pi^{(i)}$ coincide: $\pi^{(1)} = \ldots = \pi^{(m)} = \pi$. In this case,*

$$\lim_{n\to\infty} \mathbb{P}(X_n = j) = \pi_j \ \ \forall j \in I \text{ and the initial distribution } \lambda.$$

Proof. (a) $\forall$ state j we have

$$\begin{aligned}\left(\pi^{(i)} P\right)_j &= \sum_{l\in I} \pi_{il} p_{lj} \\ &= \sum_l \lim_{n\to\infty} p_{il}^{(n)} p_{lj} = \lim_{n\to\infty} \sum_l p_{il}^{(n)} p_{lj} = \lim_{n\to\infty} p_{ij}^{(n+1)} = \pi_{ij} = \left(\pi^{(i)}\right)_j .\end{aligned} \tag{7.2}$$

(b) If P is irreducible then all rows $\pi^{(i)}$ of $\mathbb{P}i$ coincide as there is a unique equilibrium distribution. Also,

$$\lim_{n\to\infty} \mathbb{P}(X_n = j) = \lim_{n\to\infty} \sum_i \lambda_i p_{ij}^{(n)} = \sum_i \lambda_i \lim_{n\to\infty} p_{ij}^{(n)} = \pi_j. \tag{7.3}$$

□

For a countable chain, our argument in Eqn (7.2) requires a justification of exchanging the order of the limit and summation. I'll omit this argument: the reader can find it in the recommended literature.

We see from Theorem 7.1 that the equilibrium distribution of a chain can be identified from the limit of matrices P^n as $n \to \infty$. More precisely, if we know that P^n converges to a matrix $\mathbb{P}i$ whose rows are equal to each other then these rows give the equilibrium distribution π. We see therefore that convergence $P^n \to \Pi$ where Π has a structure $\begin{matrix}\pi \\ \pi \\ \ldots\end{matrix}\begin{pmatrix} - - - \\ - - - \\ - - - \end{pmatrix}$ is a crucial factor.

So when does $P^n \to \Pi$? A simple counterexample: $P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Here

$$\begin{aligned} P^n &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad n \text{ even}, \\ P^n &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad n \text{ odd}. \end{aligned} \tag{7.4}$$

In this case the equilibrium distribution is unique: $\pi = (1/2, 1/2)$ but there is no convergence $P^n \to \Pi$, as P^n is *periodic* (of period 2).

More generally, consider an $m \times m$ matrix $P = \begin{pmatrix} 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ 1 & 0 & 0 & \ldots & 0 \end{pmatrix}$

corresponding to Fig. 7.1.

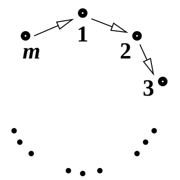

Figure 7.1

Then P^2 will correspond to

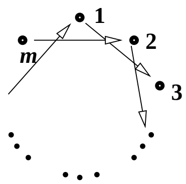

Figure 7.2

Similarly, for higher powers, with the mth power $P^m = \mathbf{I}$:

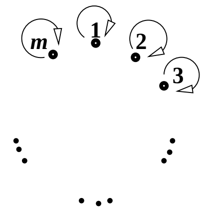

Figure 7.3

The picture will then be repeated mod m. Again, the equilibrium distribution is unique: $\pi = (1/m \ldots, 1/m)$, but the convergence $P^n \to \Pi$ fails.

Definition 7.1. Transistion matrix P is called *aperiodic* if $\forall\ i \in I$

$$p_{ii}^{(n)} > 0 \text{ for all } n \text{ large enough.} \tag{7.5}$$

If in addition, P is irreducible then, $\forall\ i, j \in I$

$$p_{ij}^{(n)} > 0 \text{ for all } n \text{ large enough.} \tag{7.6}$$

Theorem 7.2. *Assume P is irreducible, aperiodic and positive recurrent. Then, as $n \to \infty$,*

$$P^n \to \Pi.$$

The entries of the limiting matrix Π are constant along columns. In other words the rows of Π are repetitions of the same vector π which is the (unique) equilibrium distribution for P. Hence, the irreducible aperiodic and positive recurrent Markov chain forgets its initial distribution: $\forall\ \lambda$ and $j \in I$,

$$\lim_{n\to\infty} \mathbb{P}(X_n = j) = \pi_j.$$

Proof: a sketch. (Non-examinable but useful in many situations.) Consider two Markov chains $(X_n^{(i)})$, which is $\big(\delta^{(i)}, P\big)$, and (X_n^{π}), which is

(π, P). Then

$$p_{ij}^{(n)} = \mathbb{P}_i(X_n^{(i)} = j), \quad \pi_j = \mathbb{P}(X_n^{\pi} = j).$$

To evaluate the difference between these probabilities, we will identify their 'common part', by **coupling** the two Markov chains, i.e. running them together.

One way is to run both chains independently. It means that we consider the Markov chain (Y_n) on $I \times I$, with states (k, l) where $k, l \in I$, with the transition probabilities

$$p^Y_{(k,l)(u,v)} = p_{ku} p_{lv}, \quad k, l, u, v \in I, \tag{7.7}$$

and with the initial distribution

$$\mathbb{P}(Y_0 = (k, l)) = \mathbf{1}(k = i)\pi_l, \quad k, l \in I.$$

However, a better way for us is to run the chain (W_n) where the transition probabilities are

$$p^W_{(k,l)(u,v)} = \begin{cases} p_{ku} p_{lv}, & \text{if } k \neq l, \\ p_{ku}\mathbf{1}(u = v), & \text{if } k = l, \end{cases} \quad k, l, u, v \in I, \tag{7.8}$$

with the same initial distribution

$$\mathbb{P}(W_0 = (k, l)) = \mathbf{1}(k = i)\pi_l, \quad k, l \in I. \tag{7.9}$$

Indeed, Eqn (7.8) determines a transition probability matrix on $I \times I$: all entries $p^W_{(k,l)(u,v)} \geq 0$ and the sum along a row equals one. In fact,

$$\sum_{u,v \in I} p^W_{(k,l)(u,v)} = \begin{cases} \sum_u p_{ku} \sum_v p_{lv}, & \text{if } k \neq l \\ \sum_u p_{ku}, & \text{if } k = l \end{cases} = 1.$$

Further, the partial summation gives the original transitional probabilities P

$$\sum_{v \in I} p^W_{(k,l)(u,v)} = p_{ku}, \quad \sum_{u \in I} p^W_{(k,l)(u,v)} = p_{lv}.$$

Pictorially, the two components of the chain (W_n) behave individually like $(X_n^{(i)})$ and (X_n^{π}); together they evolve independently (i.e. as in (Y_n)) until the (random) time T when they coincide

$$T = \inf \left[n \geq 1 : \; X_n^{(i)} = X_n^{\pi} \right],$$

after which they stay together. Therefore,

$$p_{ij}^{(n)} - \pi_j = \mathbb{P}^W \left(X_n^{(i)} = j \right) - \mathbb{P}^W \left(X_n^{\pi} = j \right).$$

Writing

$$\mathbb{P}^W \left(X_n^{(i)} = j \right) = \mathbb{P}^W \left(X_n^{(i)} = j, T \leq n \right) + \mathbb{P}^W \left(X_n^{(i)} = j, T > n \right) \tag{7.10}$$

and

$$\mathbb{P}^W \left(X_n^{\pi} = j \right) = \mathbb{P}^W \left(X_n^{\pi} = j, T \leq n \right) + \mathbb{P}^W \left(X_n^{\pi} = j, T > n \right), \tag{7.11}$$

we see that the first summands cancel each other:

$$\mathbb{P}^W \left(X_n^{(i)} = j, T \leq n \right) = \mathbb{P}^W \left(X_n^{\pi} = j, T \leq n \right),$$

as the events $\left\{ X_n^{(i)} = j, T \leq n \right\}$ and $\{ X_n^{\pi} = j, T \leq n \}$ coincide. Hence

$$p_{ij}^{(n)} - \pi_j = \mathbb{P}^W \left(X_n^{(i)} = j, T > n \right) - \mathbb{P}^W \left(X_n^{\pi} = j, T > n \right)$$

and

$$\left| p_{ij}^{(n)} - \pi_j \right| \leq \mathbb{P}^W(T > n) = \mathbb{P}^Y(T > n). \tag{7.12}$$

The last bound is called the coupling inequality.

Thus, it suffices to check that $\mathbb{P}^W(T > n) \to 0$, i.e. $\mathbb{P}(T < \infty) = 1$. This is established by using the fact that the original matrix P is irreducible and aperiodic. (I omit the details.) □

In the case of a finite irreducible aperiodic chain it is possible to establish that the rate of convergence of $p_{ij}^{(n)}$ to π_j is geometric. In fact, in this case ∃ $m \geq 1$ and $\rho \in (0, 1)$ such that

$$p_{ij}^{(m)} \geq \rho \;\; \forall \text{ states } i, j. \tag{7.13}$$

Theorem 7.3. *If P is finite irreducible and aperiodic then ∀ states i, j*

$$\left| p_{ij}^{(n)} - \pi_j \right| \leq (1 - \rho)^{n/m - 1}, \tag{7.14}$$

where m and ρ are as in (7.13).

Proof. (Non-examinable but useful in many situations.) Repeat the scheme of the proof of Theorem 7.2: we have to assess $\mathbb{P}^Y(T > n)$. But in the finite case, we can write

$$\mathbb{P}^W_{(k,l)}(T \le m) \ge \sum_{u \in I} p^{(m)}_{ku} p^{(m)}_{lu} \ge \rho \sum_{u \in I} p^{(m)}_{lu} = \rho,$$

i.e.

$$\mathbb{P}^W_{(k,l)}(T > m) \le (1-\rho) \ \ \forall\, k, l \in I.$$

Then, by the strong Markov property

$$\mathbb{P}^W(T > n) \le \mathbb{P}^W\left(T > \left[\frac{n}{m}\right] m\right) \le \mathbb{P}^W(T > m)^{[n/m]}$$

and the assertion of Theorem 7.3 follows. □

Examples. 7.1. Consider an $m \times m$ stochastic matrix whose rows are cyclic shifts of a given stochastic vector $(p_1, \ldots, p_m)$ where $p_1, \ldots, p_m > 0$ and $p_1 + \ldots + p_m = 1$:

$$P = \begin{pmatrix} p_1 & p_2 & \cdots & p_{m-1} & p_m \\ p_2 & p_3 & \cdots & p_m & p_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_m & p_1 & \cdots & p_{m-1} & \end{pmatrix}$$

Since all states communicate directly, this matrix is irreducible and aperiodic; moreover, the value $m_0 = \min\left[n : \ p^{(n)}_{ij} > 0 \ \forall\, i, j \in I\right] = 1$. The equilibrium distribution is unique: $\pi = (1/m, \ldots, 1/m)$. By Theorems 7.1 and 7.2, $P^n \to \Pi$ geometrically fast:

$$\left| p^{(n)}_{ij} - \pi_j \right| \le (1-\rho)^n$$

where $\rho = \min\left[p_1, \ldots, p_m\right] \in (0, 1)$.

7.2. (Card shuffling) The problem of shuffling a pack of cards is important not only in gambling but in a number of other application. See Example Sheet 2.

Remark 7.1. For a transient or null recurrent irreducible aperiodic chain, matrix P^n converges to a zero matrix:

$$\lim_{n \to \infty} P^n = \mathbf{O}.$$

We will not give here the formal proof of this assertion. (For a transient case the proof is based on the fact that the series $\sum\limits_{n \ge 1} p^{(n)}_{ii} < \infty$.)

Definition 7.2. Consider the number of visits to state i before time n:

$$V_i(n) = \sum_{k=0}^{n-1} \mathbf{1}\left(X_k = i\right). \tag{7.17}$$

The limit (if it exists)

$$\lim_{n \to \infty} \frac{V_i(n)}{n} \tag{7.18}$$

is called the *long-run proportion* of the time spent in state i.

Theorem 7.4. *$\forall$ state $i \in I$:*

$$\mathbb{P}_i\left(\lim_{n \to \infty} \frac{V_i(n)}{n} = r_i\right) = 1, \tag{7.19}$$

where

$$r_i = \begin{cases} \pi_i, & \text{if } i \text{ is positive recurrent,} \\ 0, & \text{if } i \text{ is null recurrent or transient.} \end{cases} \tag{7.20}$$

Proof. First, suppose that state i is transient. Then, as we know, the total number V_i of visits to i is finite with probability 1. See Eqns (5.8), (5.18). Hence, $V_i/n \to 0$ as $n \to \infty$ with probability 1. As $0 \le V_i(n) \le V_i$, we deduce that $V_i(n)/n \to 0$ as $n \to \infty$ with probability 1.

Now let i be recurrent. Then the times $T_i^{(1)}$, $T_i^{(2)}$, ... between successive returns to state i are finite with $\mathbb{P}_i$-probability 1. By Theorem 6.5, they are IID random variables, with mean value m_i equal to $1/\pi_i$ in the positive recurrent case and to ∞ in the null recurrent case. Obviously,

$$T_i^{(1)} + \cdots + T_i^{(V_i(n))} \ge n,$$

but

$$T_i^{(1)} + \cdots + T_i^{(V_i(n)-1)} \le n-1,$$

see Fig. 7.4. So, we can write:

$$\frac{1}{V_i(n)}\left(T_i^{(1)} + \cdots + T_i^{(V_i(n)-1)}\right) \le \frac{n}{V_i(n)} \le \frac{1}{V_i(n)}\left(T_i^{(1)} + \cdots + T_i^{(V_i(n))}\right). \tag{7.21}$$

By Theorem 6.6, on an event of $\mathbb{P}_i$-probability 1, the limit $\lim\limits_{n\to\infty} \frac{1}{n}\sum\limits_{l=1}^{n} T_i^{(l)} = m_i$ holds:

$$\mathbb{P}_i\left(\frac{1}{n}\sum_{l=1}^{n} T_i^{(l)} \to m_i, \ \text{ as } n \to \infty\right) = 1. \tag{7.22}$$

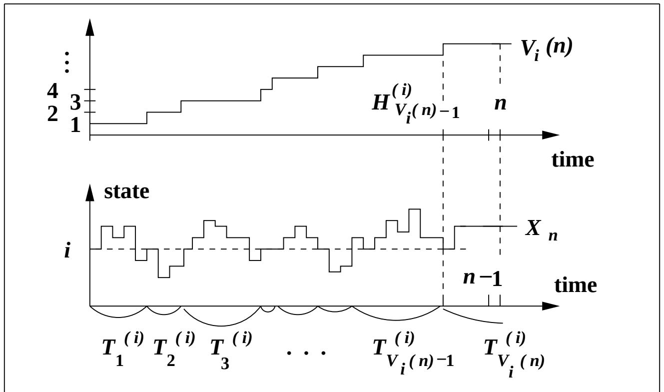

Figure 7.4

Next, as i is recurrent, sequence $(V_i(n))$ increases indefinitely, again on an event of $\mathbb{P}_i$-probability 1:

$$\mathbb{P}_i\Big(V_i(n) \nearrow \infty, \ \text{ as } \ n \to \infty\Big) = 1. \tag{7.23}$$

Then we can put in (7.22) a summation up to $V_i(n)$, instead of n and, correspondingly, divide by the factor $V_i(n)$:

$$\lim_{n\to\infty} \frac{1}{V_i(n)} \sum_{l=1}^{V_i(n)} T_i^{(l)} = m_i.$$

This relation holds on the intersection of the two aforementioned events of probability 1, which obviously has again $\mathbb{P}_i$-probability 1. On the same event,

$$\lim_{n\to\infty} \frac{1}{V_i(n)} \sum_{l=1}^{V_i(n)-1} T_i^{(l)} = m_i.$$

In other words, Eqns (7.22) and (7.23) together yield that

$$\mathbb{P}_i\left(\frac{1}{V_i(n)} \sum_{l=1}^{V_i(n)-1} T_i^{(l)} \to m_i \ \text{ and } \ \frac{1}{V_i(n)} \sum_{l=1}^{V_i(n)} T_i^{(l)} \to m_i, \ \text{ as } n \to \infty\right) = 1. \tag{7.24}$$

But then, owing to (7.21), still on the same intersection of two events of $\mathbb{P}_i$-probability 1, the ratio $n/V_i(n)$ tends to m_i, i.e. the inverse ratio $V_i(n)/n$ tends to $r_i = 1/m_i$. This gives (7.19), (7.20) and completes the proof of Theorem 7.4.

Remark 7.3. A careful analysis of the proof of Theorem 7.4 shows that if P is irreducible and positive recurrent, then we can claim that in (7.19) the probability distribution $\mathbb{P}_i$ can be replaced by $\mathbb{P}_j$, or, in fact, by the distribution $\mathbb{P}$ generated by an arbitrary initial distribution λ. This is possible because sums $T_i^{(1)} + \cdots + T_i^{(n)}$ still behave asymptotically as if the RVs $T_i^{(l)}$ were IID. (In reality, the distribution of the first RV, $T_i^{(1)} = T_i = H_1$, will be different and depend on the choice of the initial state.)

Theorem 7.5. *Let P be a finite irreducible transition matrix. Then for any initial distribution λ and a bounded function f on I:*

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{V(f,n)}{n} = \pi(f)\right) = 1, \tag{7.25}$$

where

$$\pi(f) = \sum_{i\in I} \pi_i f(i). \tag{7.26}$$

Proof. The proof of Theorem 7.5 is a re-finement of that of Theorem 7.4. More precisely, (7.25) is equivalent to

$$\mathbb{P}\left(\lim_{n\to\infty} \left|\frac{V(f,n)}{n} - \pi(f)\right| = 0\right) = 1.$$

In other words, we have to check that on an event of $\mathbb{P}$-probability 1,

$$\left|\frac{V(f,n)}{n} - \pi(f)\right| \to 0, \quad \text{as} \quad n\to\infty. \tag{7.27}$$

Writing $V(f,n) = \sum_{i\in I} V_i(n) f(i)$ and $\pi(f) = \sum_{i\in I} \pi_i f(i)$, we can transform

and bound the left-hand side in (7.27) as follows

$$\left|\frac{V(f,n)}{n} - \pi(f)\right| = \left|\sum_{i\in I}\left(\frac{V_i(n)}{n} - \pi_i\right) f(i)\right| \le \sum_{i\in I}\left|\frac{V_i(n)}{n} - \pi_i\right| \big|f(i)\big|.$$

We know that, $\forall\ i \in I$, on an event of $\mathbb{P}_i$-probability 1, $V_i(n)/n \to \pi_i$. Remark 7.1 allows us to claim convergence $V_i(n)/n \to \pi_i$ on an event of $\mathbb{P}_j$-probability 1 (that is, regardless of the choice of the initial state), or, even stronger, on an event of $\mathbb{P}$-probability 1, where P is the distribution of the (λ, P) Markov chain with any initial distribution λ. Then (7.25) follows, which completes the proof of Theorem 7.5.

Example 7.3. (Markov Chains, Part IIA, 2002, A401M) Write an essay on the long-time behaviour of discrete time Markov chains on a finite state space. Your essay should include discussion of the convergence of probabilities as well as almost-sure behaviour. You should also explain what happens when the chain is not irreducible.

Solution. The state space splits into open classes O_1, ..., O_j and closed classes C_{j+1}, ..., C_{j+l}. If $l = 1$ (a unique closed class), it is irreducible. Starting from an open class, say O_i, we end up in closed class C_k with probability h_i^k. These probabilities satisfy

$$h_i^k = \sum_{r=1}^{j+l} \widehat{p}_{ir} h_r^k.$$

Here, $\widehat{p}_{ir}$ is the probability that we exit class O_i to class O_r or C_r, and for $r = j+1$, ..., $j+l$: $h_r^k = \delta_{r,k}$.

The chain has a unique equlibrium distribution $\pi^{(r)}$ concentrated on C_r, $r = j+1$, ..., $j+l$ (hence, a unique equilibrium distribution when $l = 1$). Any equilibrium distribution is a mixture of the equilibrium distributions $\pi^{(r)}$.

Starting in C_r, we have, for any function f on C_r:

$$\frac{1}{n}\sum_{t=0}^{n} f(X_t) \to \sum_{i\in C_r} \pi_i^{(r)} f(i) \ \text{ almost surely.}$$

Moreover, in the aperiodic case (where gcd $\{n : p_{aa}^{(n)} > 0\} = 1$ for some $a \in C_r$), $\forall\ i_0 \in C_r$:

$$\mathbb{P}(X_n = i | X_0 = i_0) \to \pi_i^r,$$

and the convergence is with a geometric speed.

8. Detailed balance and reversibility

Time reversal, detailed balance, reversibility; random walk on a graph.

Let $(X_0, X_1, \ldots)$ be a Markov chain and fix $N \geq 1$. What can we say about the time reversal of (X_n), i.e. the family $(X_{N-n}, n = 0, 1, \ldots, N) = (X_N, X_{N-1}, \ldots, X_0)$?

Theorem 8.1. *Let (X_n) be a (π, P) Markov chain where $\pi = (\pi_i)$ is an equilibrium distribution for P with $\pi_i > 0\ \forall\ i \in I$. Then:* (a) $\forall\ N \geq 1$, *the time reversal $(X_N, X_{N-1}, \ldots, X_0)$ is a $(\pi, \widehat{P})$ Markov chain where $\widehat{P} = (\widehat{p}_{ij})$ has*

$$\widehat{p}_{ij} = \frac{\pi_j}{\pi_i}\, p_{ji}. \tag{8.1}$$

(b) *If P is irreducible then so is $\widehat{P}$.*

Proof. (a) First, observe that $\widehat{P}$ is a stochastic matrix, that is, $\widehat{p}_{ij} \geq 0$ and

$$\sum_j \widehat{p}_{ij} = \frac{1}{\pi_i} \sum_j \pi_j p_{ji} = \frac{1}{\pi_i}\, \pi_i = 1.$$

Next, π is $\widehat{P}$-invariant

$$\sum_i \pi_i \widehat{p}_{ij} = \sum_i \pi_j p_{ji} = \pi_j \sum_i p_{ji} = \pi_j.$$

Now pull the factor $\pi_\bullet$ through the product

$$\begin{aligned}
\mathbb{P}\left(X_N = i_N, \ldots, X_0 = i_0\right) &= \mathbb{P}\left(X_0 = i_0, \ldots, X_N = i_N\right) \\
&= \pi_{i_0} p_{i_0 i_1} \ldots p_{i_{N-1} i_N} \\
&= \widehat{p}_{i_1 i_0} \pi_{i_1} \ldots p_{i_{N-1} i_N} \\
&= \widehat{p}_{i_1 i_0} \widehat{p}_{i_2 i_1} \pi_{i_2} \ldots \\
&= \widehat{p}_{i_1 i_0} \ldots \widehat{p}_{i_N i_{N-1}} \pi_{i_N} \\
&= \pi_{i_N} \widehat{p}_{i_N i_{N-1}} \ldots \widehat{p}_{i_1 i_0}.
\end{aligned}$$

We see that (X_{N-n}) is a $(\pi, \widehat{P})$ Markov chain.

(b) If P is irreducible then any pair of states i, j is connected, that is $\exists$ a path $i = i_0$, i_1, ..., $i_n = j$ with

$$\begin{aligned} 0 < p_{i_0 i_1} \dots p_{i_{n-1} i_n} &= (1/\pi_{i_0}) \pi_{i_0} p_{i_0 i_1} \dots p_{i_{n-1} i_n} \\ &= (1/\pi_{i_0}) \widehat{p}_{i_1 i_0} \pi_{i_1} \dots p_{i_{n-1} i_n} = \dots \\ &= (1/\pi_{i_0}) \widehat{p}_{i_1 i_0} \dots \widehat{p}_{i_n i_{n-1}} \pi_{i_n}. \end{aligned}$$

So, $\widehat{p}_{i_1 i_0} \dots \widehat{p}_{i_n i_{n-1}} > 0$, and j, i are connected in $\widehat{P}$. □

The case where chain (X_{N-n}) has the same distribution as (X_n) is of a particular interest

Theorem 8.2. *Let (X_n) be a Markov chain. The following properties are equivalent:*

(i) $\forall$ $n \geq 1$ *and states* i_0, ..., i_n:

$$\mathbb{P}(X_0 = i_0, \dots, X_n = i_n) = \mathbb{P}(X_0 = i_n, \dots X_n = i_0). \tag{8.2}$$

(ii) (X_n) *is in equilibrium, i.e.* $(X_n) \sim (\pi, P)$ *where* π *is an equilibrium distribution for* P, *and*

$$\pi_i p_{ij} = \pi_j p_{ji} \ \forall \text{ states } i, j \in I. \tag{8.3}$$

Proof. (i) $\Rightarrow$ (ii). Take $n = 1$,

$$\mathbb{P}(X_0 = i, X_1 = j) = \mathbb{P}(X_0 = j, X_1 = i),$$

and sum over j

$$\sum_j \mathbb{P}(X_0 = i, X_1 = j) = \mathbb{P}(X_0 = i) = \lambda_i,$$

$$\sum_j \mathbb{P}(X_0 = j, X_1 = i) = \mathbb{P}(X_1 = i) = (\lambda P)_i \,.$$

So, $\lambda_i = (\lambda P)_i$ $\forall$ i, i.e. $\lambda P = \lambda$. Hence, the chain is in equilibrium $\lambda = \pi$. Next, $\forall$ i, j

$$\mathbb{P}(X_0 = i, X_1 = j) = \pi_i p_{ij} = \mathbb{P}(X_0 = j, X_1 = i) = \pi_j p_{ji}.$$

(ii) $\Rightarrow$ (i). Write

$$\mathbb{P}(X_0 = i_0, \dots, X_n = i_n) = \pi_{i_0} p_{i_0 i_1} \dots p_{i_{n-1} i_n}$$

and use Eqns (8.3) to pull $\pi_\bullet$ through the product

$$\begin{aligned} \pi_{i_0} p_{i_0 i_1} \dots p_{i_{n-1} i_n} &= p_{i_0 i_1} \pi_{i_1} \dots p_{i_{n-1} i_n} = \dots \\ = p_{i_0 i_1} \dots p_{i_n i_{n-1}} \pi_{i_n} &= \pi_{i_n} p_{i_n i_{n-1}} \dots p_{i_0 i_1} \\ &= \mathbb{P}(X_0 = i_n, \dots, X_n = i_0). \end{aligned}$$

□

Definition 8.1 A Markov chain (X_n) satisfying (8.2) is called *reversible*. Eqns (8.3) are called *detailed balance* equations (DBEs). So, the assertion of Theorem 8.2 reads; a Markov chain is reversible if and only if it is in equilibrium, and the DBEs are satisfied.

The DBEs are a powerful tool for identification of an ED.

Theorem 8.3. *If λ and P satisfy the DBEs*

$$\lambda_i p_{ij} = \lambda_j p_{ji}, \ \ i, j \in I,$$

then λ is an ED for P, that is $\lambda P = \lambda$.

Proof. Sum over j:

$$\lambda_i \sum_j p_{ij} = \lambda_i,$$

$$\sum_j \lambda_j p_{ji} = (\lambda P)_i.$$

The two expressions are equal $\forall$ i, hence the result. □

So, for a given matrix P, if the DBEs can be solved (that is, a probability distribution that satisfies them can be found), the solution will give an ED. Furthermore, the corresponding Markov chain will be reversible.

An interesting and important class of Markov chains is formed by *random walks on graphs.* We have seen examples of such chains: a birth-death process (a RW on $\mathbb{Z}^1$ or its subset), a RW on a plane square lattice $\mathbb{Z}^2$ and, more generally, a RW on a d-dimensional cubic lattice $\mathbb{Z}^d$. A feature of these examples is that a wandering particle can jump to any of its neighbouring sites; in a symmetric case, the probability of each jump is the same. This idea can be extended to a general *graph*, with directed or non-directed links (edges). Here, we focus on non-directed graphs; a graph is understood as a collection G of *vertices* some of which are joined by non-directed *edges*, or *links*, possibly several. Non-directed means here that the edges can be traversed in both directions; sometimes it's convenient to think that each edge is formed by a pair of opposite arrows.

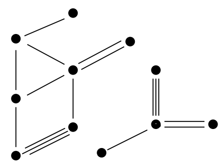

Figure 8.1

A graph is called *connected* if any two distinct vertices are connected with a *path* formed by edges. The *valency* v_i of a vertex i is defined as the number of edges at i. The *connectedness* v_{ij} is the number of edges joining vertices i and j.

The RW on the graph has the following transition matrix $P = (p_{ij})$

$$p_{ij} = \begin{cases} v_{ij}/v_i, & \text{if } i \text{ and } j \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases} \tag{8.4}$$

The matrix P is irreducible if and only if the graph is connected. The vector

$v = (v_i)$ satisfies the DBEs. That is, $\forall$ vertices i, j

$$v_i p_{ij} = v_{ij} = v_j p_{ji}, \tag{8.5}$$

and hence is P-invariant. We obtain the following straightforward result.

Theorem 8.4. *The RW on a graph, with transition matrix P of the form (8.4), could be of all three types: transient (viz., a symmetric nearest-neighbour RW on $\mathbb{Z}^d$ with $d \geq 3$), null recurrent (a symmetric nearest-neighbour RW on $\mathbb{Z}^2$ or $\mathbb{Z}^1$) or positive recurrent. It is positive recurrent if and only if the total valence $\sum_i v_i < \infty$, in which case $\pi_j = v_j \Big/ \sum_i v_i$ is an equilibrium distribution. Furthermore, the chain with equilibrium distribution π is reversible.*

A simple but popular example of a graph is an ℓ-site segment of a one-dimensional lattice: here the valency of every vertex equals 2, except for the endpoints where the valence is 1. See Fig. 8.2 a).

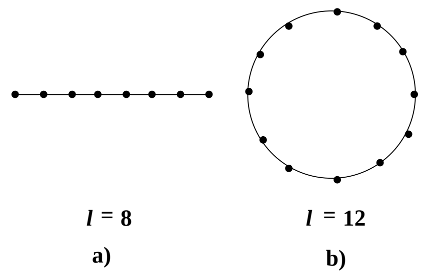

Figure 8.2

An interesting class is formed by graphs with a constant valency: $v_i \equiv v$; again the simplest case is $v = 2$, where ℓ vertices are placed on a circle (or on a perfect polygon or any closed path). See Fig. 8.2 b). A popular example of a graph with a constant valency is a fully connected graph with a given number of vertices, say $\{1, \dots, m\}$: here the valency equals $m-1$, and the graph has $m(m-1)/2$ (non-directed) edges in total. See Fig. 8.3.

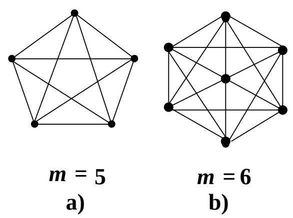

Figure 8.3

Another important example is a regular cube in d dimensions, with 2^d vertices. Here the valency equals d, and the graph has $d2^{d-1}$ (still non-directed) edges joining neigbouring vertices. See Fig. 8.4.

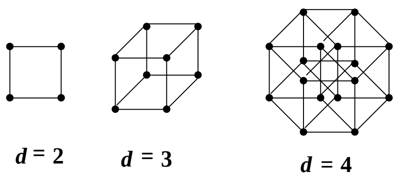

Figure 8.4

Popular examples of infinite graphs of constant valency are lattices and trees.

In the case of a general finite graph of constant valency $v_i = v \ \forall$ vertex i, the sum $\sum\limits_i v_i$ equals $v \times |G|$ where $|G|$ is the number of vertices. Then

probabilities $p_{ij} = p_{ji} = v_{ij}/v, \forall$ neighbouring pair i, j. That is, the transition matrix $P = (p_{ij})$ is Hermitian: $P = P^{\mathrm{T}}$. Furthermore, the equilibrium distribution $\pi = (\pi_i)$ is uniform: $\pi_i = 1/|G|$.

In Linear Algebra courses, it is asserted that a (complex) Hermitian matrix has an orthonormal basis of eigen-vectors, and its eigen-values are all real. This handy property is nice to retain whenever possible. For a Markov chain, even when P is originally non-Hermitian, it can be 'converted' into a Hermitian matrix, by changing the scalar product. We will explore further this avenue in Sections 12–14.

Example 8.1 (Markov Chains, Part IIA, 2002, A101M and Part IIA, 2002, B101M)

(i) We are given a finite set of airports. Assume that between any two airports, i and j, there are $a_{ij} = a_{ji}$ flights in each direction on every day. A confused traveller takes one flight per day, choosing at random from all available flights. Starting from i, how many days on average will pass until the traveller returns again to i? Be careful to allow for the case where may be no flights at all between two given airports.

(ii) Consider the infinite tree T with root R, where for all $m \geq 0$, all vertices at distance 2^m from R have degree 3, and where all other vertices (except R) have degree 2. Show that the random walk on T is recurrent.

Solution. **(i)** Let $X_0 = i$ be the starting airport, X_n the destination of the nth flight and I denote the set of airports reachable from i. Then (X_n) is an irreducible Markov chain on I, so the expected return time to i, is given by $(1/\pi_i)$, where π is the unique equilibrium distribution. We will show that $1/\pi_i = \sum\limits_{j,k\in I} a_{jk} \Big/ \sum\limits_{k\in I} a_{ik}$

In fact,

$$p_{jk} = \frac{a_{jk}}{\sum\limits_{l\in I} a_{jl}} \quad \text{and} \quad \Big(\sum_{l\in I} a_{jl}\Big) p_{jk} = \Big(\sum_{l\in I} a_{kl}\Big) p_{kj}.$$

So the vector $v=(v_j)$ with $v_j=\sum\limits_{l\in I}a_{jl}$ is in detailed balance with P. Hence

$$\pi_j=\sum_{k\in I}a_{jk}\Big/\sum_{k,l\in I}a_{kl}.$$

(ii) Consider the distance X_n from the root R at time n. Then $(X_n)_{n\geq 0}$ is a birth-death Markov chain with transition

$$q_i=p_i=1/2,\quad \text{if}\quad i\neq 2^m,$$

$$q_i=1/3,\ p_i=2/3,\quad \text{if}\quad i=2^m.$$

By a standard argument for $h_i=\mathbb{P}_i(\text{hit } 0)$

$$h_0=1,\ h_i=p_ih_{i+1}+q_ih_{i-1},\ i\geq 1,$$

$$p_iu_{i+1}=q_iu_i,\ u_i=h_{i-1}-h_i,$$
$$u_{i+1}=\frac{q_i}{p_i}\,u_i=\gamma_iu_1,\ \gamma_i=\frac{q_i\dots q_1}{p_i\dots p_1},$$

and

$$u_1+\dots+u_i=h_0-h_i, h_i=1-A(\gamma_0+\dots+\gamma_{i-1}).$$

The condition $\sum\limits_i\gamma_i=\infty$ forces $A=0$ and hence $h_i=1$ for all i. Here,

$$\gamma_{2^m-1}=2^{-m},$$

so $\sum\limits_i\gamma_i=\infty$ and the walk is recurrent.

The DBEs are a convenient tool to find an equilibrium distribution: if a measure $\lambda\geq\underline{0}$ is in detailed balance with P and has $\sum\limits_i\lambda_i<\infty$, then $\pi_j=\lambda_j/\sum\limits_i\lambda_i$ is an equilibrium distribution.

Example 8.2 Suppose $\pi=(\pi_i)$ forms an ED for transition matrix $P=(p_{ij})$, with $\pi P=\pi$, but the DBE's (8.3) are not satisfied. What is the time reversal of chain (X_n) in equilibrium?

Assume, for definiteness, that P is irreducible, and $\pi_i>0\ \forall\ i\in I$. The answer comes out after we define the transition matrix $P^{\mathrm{RV}}=(p_{ij}^{\mathrm{RV}})$ by

$$\pi_ip_{ij}^{\mathrm{RV}}=\pi_jp_{ij},\quad i,j\in I. \tag{8.6}$$

or

$$p_{ij}^{\mathrm{RV}}=\frac{\pi_j}{\pi_i}p_{ji},\quad i,j\in I. \tag{8.7}$$

Eqns (8.6), (8.7) indeed determine a transition matrix, as $\forall\ i,j\in I$,

$$p_{ij}^{\mathrm{RV}}\geq 0,\quad \text{and}\quad \sum_{j\in I}p_{ij}^{\mathrm{RV}}=\frac{1}{\pi_i}\sum_{j\in I}\pi_jp_{ji}=\frac{1}{\pi_i}\pi_i=1.$$

Next, π gives an ED for P^{RV}: $\forall\ j\in I$,

$$\sum_{i\in I}\pi_ip_{ij}^{\mathrm{RV}}=\sum_{i\in I}\pi_jp_{ji}=\pi_j.$$

Then, repeating the argument from the proof of Theorem 8.1, we obtain that $\forall\ N\geq 1$, the time reversal $(X_{N-n},\ 0\leq n\leq N)$ is a Markov chain in equilibrium, with transition matrix P^{RV} and the same ED π. Symbolically,

$$(X_n^{\mathrm{RV}})\ \sim\ \left(\pi,P^{\mathrm{RV}}\right)-\text{Markovchain}, \tag{8.8}$$

where $(X_n^{\mathrm{RV}})=(X_{N-n})$ stands for the time reversal of T.

It is instructive to remember that P^{RV} was proven to be a stochastic matrix because π is an ED for P while the proof that π is an ED for P^{RV} used only the fact that P is stochastic.

Example 8.3 The detailed balance equations have a useful geometric meaning. Suppose that the state space $I=\{1,\dots,s\}$. Matrix P generates a linear transformation $\mathbb{R}^s\to\mathbb{R}^s$, where vector $\mathbf{x}=\begin{pmatrix}x_1\\ \vdots\\ x_s\end{pmatrix}$ is taken to $P\mathbf{x}$. Assuming P irreducible, let π be the ED, with $\pi_i>0$, $i=1,\dots,s$. Consider a 'tilted' scalar product $\langle\,\cdot\,,\,\cdot\,\rangle_\pi$ in $\mathbb{R}^s$, where

$$\langle\mathbf{x},\mathbf{y}\rangle_\pi=\sum_{i=1}^s x_iy_i\pi_i. \tag{8.9}$$

Then detailed balance equations (8.3) mean that P is self-adjoint (or Hermitian) relative to scalar product $\langle \cdot\,,\,\cdot \rangle_\pi$. that is,

$$\langle \mathbf{x}, P\mathbf{y}\rangle_\pi = \langle P\mathbf{x}, \mathbf{y}\rangle_\pi, \quad \mathbf{x}, \mathbf{y} \in \mathbb{R}^s. \tag{8.10}$$

In fact,

$$\langle \mathbf{x}, P\mathbf{y}\rangle_\pi = \sum_{i,j} x_i p_{ij} y_j \pi_i = \sum_{i,j} x_i p_{ji} y_j \pi_j = \langle P\mathbf{x}, \mathbf{y}\rangle_\pi.$$

The converse is also true: Eqn (8.10) implies (8.3), as we can take as $\mathbf{x}$ and $\mathbf{y}$ the vectors $\boldsymbol{\delta}_i$ and $\boldsymbol{\delta}_j$ with the only non-zero entries 1 at positions i and j, respectively, $\forall\ i, j = 1, \dots, s$.

This observation yields a benefit, as Hermitian matrices have all eigenvalues real, and their eigen-vectors are mutually orthogonal (relative to the scalar product in question, in this instance, $\langle \cdot\,,\,\cdot \rangle_\pi$). We will use this in Section 12.

Remark 8.1. The concept of reversibility and time reversal will be particularly helpful in a continuous-time setting of Part II Applied Probability.

It is now time to give a brief summary of essential results established about various equations emerging in the analysis of Markov chains. We have seen two sets of equations: **(I)** for hitting probabilities h_i^A and mean hitting times k_i^A and **(II)** for equilibrium distributions $\pi = (\pi_i)$ and expected times γ_i^k spent in state i before returning to k. Although they are in a sense similar, there are also differences between them which it is important to remember.

(I.1) For $h_i^j = \mathbb{P}_i(\text{hit } j)$ the equations are

$$h_j^j = 1, \;\; h_i^j = \sum_{l\in I} p_{il} h_l^j = (h^j P^{\mathrm{T}})_i, \;\; i \neq j,$$

where

$$h^j = (h_i^j, i \in I), \quad \text{with} \;\; h_j^j = 1.$$

Here, $h_i^j \equiv 1$ is always a solution

$$\underline{1} P^{\mathrm{T}} = \underline{1}, \quad \text{as} \;\; (\underline{1} P^{\mathrm{T}})_i = \sum_l p_{il} = 1 \;\forall\; i \in I.$$

(I.2) For $k_i^j = \mathbb{E}_i(\text{time to hit } j)$ the equations are

$$k_j^j = 0, \;\; k_i^j = 1 + \sum_{l\in I,\, l\neq j} p_{il} k_l^j = 1 + (k^j P^{\mathrm{T}})_i, \;\; i \neq j.$$

where

$$k^j = (k_i^j, i \in I), \quad \text{with} \;\; k_j^j = 0.$$

Here, taking that $0 \cdot \infty = 0$, $k_i^j = (1 - \delta_{ij})\infty$ is always a solution when the chain is irreducible.

These equations are produced by conditioning on the *first* jump. The vectors h^j and k^j are labelled by the terminal states while their entries h_i^j and k_i^j indicate the initial states. The solution we look for is identified as a minimal non-negative solution satisfying the normalisation constraints $h_j^j = 1$ and $k_j^j = 0$.

(II.1) For

$$\gamma_i^k = \mathbb{E}_k(\text{time spent in } i \text{ before returning to } k)$$

the equations are

$$\gamma_k^k = 1, \;\; \gamma_i^k = \sum_l \gamma_l^k p_{li}, \;\; i \neq k,$$

or

$$\gamma^k = \gamma^k P, \quad \text{when } k \text{ is recurrent.}$$

Here, the conditioning is on the *last* jump, and vectors γ^k are labelled by starting states. The indentification of the solution is by the conditions $\gamma_i^k \geq 0$ and $\gamma_k^k = 1$.

(II.2) Similarly, for an equilibrium distribution (or more generally, an invariant measure)

$$\pi = \pi P.$$

The identification here is through the condition $\pi_i \geq 0$ and $\sum_i \pi_i = 1$.

(II.3) A solution to the detailed balance equations

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

always produces an invariant measure. If in addition, $\sum_i \pi_i = 1$, it gives an equilibrium distribution. As the detailed balance equations are usually easy to solve (when they have a solution), it is a powerful tool which is always worth trying when you need to find an equilibrium distribution.