1 Stationary distributions and the limit theorem

Definition 1.1. The vector $\boldsymbol{\pi}$ is called a *stationary distribution* of a Markov chain with matrix of transition probabilities $\mathbf{P}$ if $\boldsymbol{\pi}$ has entries $(\pi_j : j \in S)$ such that:

(a) $\pi_j \geq 0$ for all j, $\sum_j \pi_j = 1$, and

(b) $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, which is to say that $\pi_j = \sum_i \pi_i p_{ij}$ for all j (the *balance equations*).

Note: This implies that

$$\boldsymbol{\pi}\mathbf{P}^n = \boldsymbol{\pi} \quad \text{for all } n \geq 0,$$

e.g. if X_0 has distribution $\boldsymbol{\pi}$ then X_n has distribution $\boldsymbol{\pi}$ for all n.

Proposition 1.2. *An irreducible chain has a stationary distribution* $\boldsymbol{\pi}$ *if and only if all the states are non-null persistent; in this case,* $\boldsymbol{\pi}$ *is the unique stationary distribution and is given by* $\pi_i = \frac{1}{\mu_i}$ *for each* $i \in S$*, where* μ_i *is the mean recurrence time of* i*.*

We will carry out the proof of this in several steps.

Fix a state k and let $\rho_i(k)$ be the mean number of visits of the chain to the state i between two successive visits to state k:

$$\rho_i(k) = \mathbb{E}\left[N_i | X_0 = k\right],$$

where

$$N_i = \sum_{n=1}^{\infty} \mathbb{1}\{X_n = i\} \cap \{T_k \geq n\}$$

and T_k is the time of the first return to state k. We write $\boldsymbol{\rho}(k)$ for the vector $(\rho_i(k) : i \in S)$. Clearly $T_k = \sum_{i \in S} N_i$, and hence

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

Lemma 1.3. *For any state k of an irreducible persistent chain, the vector $\boldsymbol{\rho}(k)$ satisfies $\rho_i(k) < \infty$ for all i, and furthermore $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$.*

Proof. We show first that $\rho_i(k) < \infty$ when $i \neq k$. Observe that $\rho_k(k) = 1$. We write

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

Clearly $f_{kk}(m+n) \geq l_{ki}(m) f_{ik}(n)$. By irreducibility of the chain, there exists n such that $f_{ik}(n) > 0$. So for such a $n \geq 2$

$$\rho_i(k) = \sum_{m=1}^{\infty} l_{ki}(m) \leq \frac{1}{f_{ik}(n)} \sum_{m=1}^{\infty} f_{kk}(m+n) \leq \frac{1}{f_{ik}(n)} < \infty$$

as required.

Next observe that $l_{ki}(1) = p_{ki}$, and

$$l_{ki}(n) = \sum_{j:j \neq k} \mathbb{P}\left(X_n = i, X_{n-1} = j, T_k \geq n | X_0 = k\right) = \sum_{j:j \neq k} l_{kj}(n-1) p_{ji}$$

Summing over $n \geq 2$, we obtain

$$\rho_i(k) = p_{ki} + \sum_{j:j \neq k} \left(\sum_{n \geq 2} l_{kj}(n-1) p_{ji} \right) = \rho_k(k) p_{ki} + \sum_{j:j \neq k} \rho_j(k) p_{ji},$$

since $\rho_k(k) = 1$. $\square$

We have shown that for any irreducible chain, the vector $\boldsymbol{\rho}(k)$ satisfies $\boldsymbol{\rho}(k) = \boldsymbol{\rho}(k)\mathbf{P}$, and furthermore that the components of $\boldsymbol{\rho}(k)$ are non-negative with sum μ_k. Hence, if $\mu_k < \infty$, the vector $\boldsymbol{\pi}$ with entries $\pi_i = \rho_i(k)/\mu_k$ is a distribution satisfying $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$.

Therefore, every non-null persistent irreducible chain has a stationary distribution.

Proposition 1.4. *For any irreducible, persistent Markov chain with matrix of transition probabilities* $\mathbf{P}$*, there exists a positive solution* $\mathbf{x}$ *to the equation*

$$\mathbf{x} = \mathbf{x}\mathbf{P},$$

which is unique up to a multiplicative constant. The chain is non-null if $\sum_i x_i < \infty$ *and null if* $\sum_i x_i = \infty$*.*

We've seen all of this except for the uniqueness claim, which we won't establish – although it isn't difficult.

Proof of Proposition (1.2) Suppose the π is a stationary distribution of the chain. If all states are transient then $p_{ij}(n) \to 0$ as $n \to \infty$, for all i and j. So

$$\pi_j = \sum_i \pi_i p_{ij}(n) \to 0 \quad \text{as} \quad n \to \infty \qquad \text{for all } i \text{ and } j,$$

which contradicts $\sum_j \pi_j = 1$. $\square$

We show next that the existence of $\boldsymbol{\pi}$ implies that all states are non-null and that $\pi_i = \frac{1}{\mu_i}$ for each i. Suppose that X_0 has distribution $\boldsymbol{\pi}$, so that $\mathbb{P}\{X_0 = i\} = \pi_i$ for each i. Then

$$\pi_j \mu_j = \mathbb{P}\{X_0 = j\} \sum_{n=1}^{\infty} \mathbb{P}\{T_j \geq n | X_0 = j\} = \sum_{n=1}^{\infty} \mathbb{P}\{T_j \geq n, X_0 = j\}.$$

However, $\mathbb{P}\{T_j \geq 1, X_0 = j\} = \mathbb{P}\{X_0 = j\}$, and for $n \geq 2$,

$$
\begin{aligned}
&\mathbb{P}\{T_j \geq n, X_0 = j\} \\
&= \mathbb{P}\{X_0 = j, X_m \neq j \text{ for } 1 \leq m \leq n-1\} \\
&= \mathbb{P}\{X_m \neq j \text{ for } 1 \leq m \leq n-1\} - \mathbb{P}\{X_m \neq j \text{ for } 0 \leq m \leq n-1\} \\
&= \mathbb{P}\{X_m \neq j \text{ for } 0 \leq m \leq n-2\} - \mathbb{P}\{X_m \neq j \text{ for } 0 \leq m \leq n-1\} \\
&\qquad\qquad \text{by stationarity} \\
&= a_{n-2} - a_{n-1}
\end{aligned}
$$

where $a_n = \mathbb{P}\{X_m \neq j \text{ for } 0 \leq m \leq n\}$.

Sum over n to obtain

$$\pi_j \mu_j = \mathbb{P}\{X_0 = j\} + \mathbb{P}\{X_0 \neq j\} - \lim_{n\to\infty} a_n = 1 - \lim_{n\to\infty} a_n.$$

However, $a_n \to \mathbb{P}\{X_m \neq j \text{ for all } m\} = 0$ as $n \to \infty$, by the persistence of j. We have shown that

$$\pi_j \mu_j = 1,$$

so that $\mu_j = \frac{1}{\pi_j} < \infty$ if $\pi_j > 0$. To see that $\pi_j > 0$ for all j, suppose on the contrary that $\pi_j = 0$ for some j.

Then

$$0 = \pi_j = \sum_i \pi_i p_{ij}(n) \geq \pi_i p_{ij}(n) \qquad \text{for all } i \text{ and } n,$$

yielding that $\pi_i = 0$ whenever $i \to j$. The chain is assumed irreducible, so that $\pi_i = 0$ for all i in contradiction of the fact that π_i's sum to 1. Hence $\mu_j < \infty$ and all states of the chain are non-null. Furthermore, we see that π_j are specified uniquely as $\frac{1}{\mu_j}$.

Thus, if $\boldsymbol{\pi}$ exists then it is unique and all the states of the chain are non-null persistent. Conversely, if the states of the chain are non-null persistent then the chain has a stationary distribution given by Lemma (1.3).

Proposition 1.5. *If $i \leftrightarrow j$ then i is null persistent if and only if j is null persistent.*

Proof. Let $C(i)$ be the irreducible closed equivalence class of states which contains the non-null persistent state i. Suppose that $X_0 \in C(i)$. Then $X_n \in C(i)$ for all n, and Lemma (1.3) and Proposition (1.2) combine to tell us that all the states in $C(i)$ are non-null. □

Proposition 1.6. *Let $s \in S$ be any state of an irreducible chain. The chain is transient if and only if there exists a non-zero solution $\{y_i : i \neq s\}$, satisfying $|y_i| \leq 1$ for all i, to the equations*

$$y_i = \sum_{j : j \neq s} p_{ij} y_j, \qquad i \neq s.$$

1.0.1 Example: Random walk with retaining barrier

A particle performs a random walk on the non-negative integers with a retaining barrier at 0. The transition probabilities are

$$p_{0,0} = q, \qquad p_{i,i+1} = p \quad \text{for } i \geq 0 \qquad p_{i,i-1} = q \quad \text{for } i \geq 1,$$

Let $\rho = p/q$.

(a) If $q < p$, take $s = 0$ to see that $y_i = 1 - \frac{1}{\rho^i}$ satisfies the equation in Proposition (1.6), and so the chain is transient.

(b) Solve the equation $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ to find that there exists a stationary distribution, with $\pi_j = \rho^j(1-\rho)$, if and only if $q > p$. Thus the chain is non-null persistent if and only if $q > p$.

(c) If $q = p = \frac{1}{2}$ the chain is persistent since symmetric random walk is persistent (just reflect negative excursions of a symmetric random walk into the positive half-line). Solve the equation $\mathbf{x} = \mathbf{x}\mathbf{P}$ to find that $x_i = 1$ for all i is the solution, unique up to a multiplicative constant. However, $\sum_i x_i = \infty$ so that the chain is null by Proposition (1.4).

Theorem 1.7. *For an irreducible aperiodic chain, we have that*

$$p_{ij}(n) \to \frac{1}{\mu_j} \quad \textit{as} \quad n \to \infty, \qquad \textit{for all } i \textit{ and } j.$$

Proof. If the chain is transient the the result is trivial.

Suppose X is an irreducible, aperiodic, non-null, persistent Markov chain. Construct the "coupled chain" $Z = (X, Y)$, as an ordered pair $X = \{X_n : n \geq 0\}$, $Y = \{Y_n : n \geq 0\}$ of *independent* Markov chains, each having state space S and transition matrix $\mathbf{P}$. Then $Z = \{Z_n = (X_n, Y_n) : n \geq 0\}$ takes values in $S \times S$, and it is easy to check that Z is a Markov chain with transition probabilities

$$\begin{aligned} p_{ij,kl} &= \mathbb{P}\left(Z_{n+1} = (k,l) | Z_n = (i,j)\right) \\ &= \mathbb{P}\left(X_{n+1} = k | X_n = i\right) \mathbb{P}\left(Y_{n+1} = l | Y_n = j\right) \qquad \text{by independence} \\ &= p_{ik} p_{jl} \end{aligned}$$

Since X is irreducible and aperiodic then Z is also irreducible. Since X is non-null persistent it has a unique stationary distribution $\boldsymbol{\pi}$, and it is easy to see that Z has a stationary distribution $\boldsymbol{\nu} = (\nu_{ij} : i, j \in S)$ given by $\nu_{ij} = \pi_i \pi_j$; thus Z is also non-null persistent.

Now suppose that $X_0 = i$ and $Y_0 = j$, so that $Z_0 = (i, j)$. Choose any state $s \in S$ and let

$$T = \min\{n \geq 1 : Z_n = (s, s)\}$$

denote the time that Z first hits (s, s). Note that $\mathbb{P}\{T < \infty\} = 1$. Starting from $Z_0 = (X_0, Y_0) = (i, j)$

$$\begin{aligned} p_{ik}(n) &= \mathbb{P}\{X_n = k\} \\ &= \mathbb{P}\{X_n = k, T \leq n\} + \mathbb{P}\{X_n = k, T > n\} \\ &= \mathbb{P}\{Y_n = k, T \leq n\} + \mathbb{P}\{X_n = k, T > n\}, \end{aligned}$$

since given $T \leq n$, X_n and Y_n are identically distributed. Also,

$$\begin{aligned} &\leq \mathbb{P}\{Y_n = k\} + \mathbb{P}\{T > n\} \\ &= p_{jk}(n) + \mathbb{P}\{T > n\}. \end{aligned}$$

This, and the related inequality with i and j interchanged, yields

$$|p_{ik}(n) - p_{jk}(n)| \leq \mathbb{P}\{T > n\} \to 0 \quad \text{as} \quad n \to \infty;$$

therefore,

$$p_{ik}(n) - p_{jk}(n) \to 0 \quad \text{as} \quad n \to \infty \qquad \text{for all } i, j \text{ and } k.$$

Thus, if $\lim_{n\to\infty} p_{ik}(n)$ exists, then it does not depend on i. To show that it exists, write

$$\pi_k - p_{jk}(n) = \sum_i \pi_i \left(p_{ik}(n) - p_{jk}(n)\right) \to 0 \quad \text{as } n \to \infty.$$

□

1.1 Examples

1.1.1 Example: The age of a renewal process

Initially an item is put into use, and when it fails it is replaced at the beginning of the next time period by a new item. Suppose that the lives of the items are independent and each will fail in its ith period of use with probability $P_i, i \geq 1$, where the distribution $\{P_i\}$ is aperiodic and $\sum_i iP_i < \infty$. Let X_n denote the age of the item in use at time n — that is, the number of periods (including the nth) it has been in use.

Then if we let

$$\lambda_i = \frac{P_i}{\sum_{j=i}^{\infty} P_j}$$

denote the probability that a unit that has lived for $i-1$ time units fails on the ith time unit, then $\{X_n, n \geq 0\}$ is a Markov chain with transition probabilities given by

$$\mathbb{P}_{i,1} = \lambda_i = 1 - \mathbb{P}_{i,i+1}, \quad i \geq 1.$$

Hence the limiting probabilities are such that

$$\pi_1 = \sum_i \pi_i \lambda(i),$$

$$\pi_{i+1} = \pi_i (1 - \lambda_i), \quad i \geq 1.$$

Since

$$(1-\lambda_i) = \frac{\sum_{j=i+1}^{\infty} P_j}{\sum_{j=i}^{\infty} P_j},$$

iterating yields

$$\begin{aligned}
\pi_{i+1} &= \pi_i(1-\lambda_i) \\
&= \pi_{i-1}(1-\lambda_i)(1-\lambda_{i-1}) \\
&= \pi_1(1-\lambda_1)(1-\lambda_2)\cdots(1-\lambda_i) \\
&= \pi_1 \sum_{j=i+1}^{\infty} P_j \\
&= \pi_1 \mathbb{P}\{X \geq i+1\},
\end{aligned}$$

where X is the life of an item.

Using $\sum_{i=1}^{\infty} \pi_i = 1$ yields

$$1 = \pi_1 \sum_{i=1}^{\infty} \mathbb{P}\{X \geq i\}$$

or

$$\pi_1 = \frac{1}{\mathbb{E}[X]}$$

and hence

$$\pi_i = \frac{\mathbb{P}\{X \geq i\}}{\mathbb{E}[X]}, \quad i \geq 1.$$

This is an example of a *size-biased* distribution.

1.1.2 Example: Poisson births

Suppose that during each time period, each member of a population dies independently of the others with probability p, and that a Poisson(λ) number of new members join the population each time period. If we let X_n denote the number of members of the population at the beginning of period n, then it is easy to see that $\{X_n, n = 1, \ldots\}$ is a Markov chain.

To find the stationary distribution of this chain, suppose that X_0 has a Poisson distribution with mean α. Since each of these X_0 individuals will independently be alive at the beginning of the next period with probability $1 - p$, by the Poisson marking theorem, the number of them that are still in the population at time 1 is a Poisson random variable with mean $\alpha(1 - p)$.

As the number of new members that join the population at time 1 is an independent Poisson random variable with mean λ, it thus follows that X_1 is a Poisson random variable with mean $\alpha(1-p)+\lambda$. Hence, if

$$\alpha = \alpha(1-p)+\lambda$$

then the chain is stationary. By uniqueness of the stationary distribution, we can conclude that the stationary distribution is Poisson with mean λ/p. That is,

$$\pi_j = e^{-\lambda/p}(\lambda/p)^j/j!, \quad j = 0, 1, \ldots$$

1.1.3 Example: the Gibbs sampler

Let $p(x_1, \ldots, x_n)$ be the joint probability mass function of the random vector $(X_1, \ldots, X_n)$. In some cases, it may be difficult to directly sample from such a distribution, but relatively easy to sample from the conditional distributions of each coordinate X_i given the values of all of the other coordinates $X_j, j \neq i$.

In this case, we can generate a random vector whose probability mass function is approximately $p(x_1, \ldots, x_n)$ by constructing a Markov chain whose stationary distribution is p as follows.

Let $X^0 = (X_1^0, \ldots, X_1^0)$ be any vector for which $p(X_1^0, \ldots, X_n^0) > 0$.

First we generate a random variable whose distribution is the conditional distribution of the first coordinate X_1 given that $X_j = X_j^0, \quad j = 2, \ldots, n$, and call its value X_1^1.

Next, generate a random variable whose distribution is the conditional distribution of X_2 given that $X_1 = X_1^1$, and $X_j = X_j^0, j = 3, \ldots, n$, and call its value X_2^1.

Continue in this fashion until we have a whole new vector $\mathbf{X}^1 = (X_1^1, \ldots, X_n^1)$. Then, repeat the process, this time starting with $\mathbf{X}^1$ in place of $\mathbf{X}^0$, to obtain the new vector $\mathbf{X}^2$, and so on. It is easy to see that the sequence of vectors $\mathbf{X}^j, j \geq 0$ is a Markov chain. We claim that its stationary distribution is $p(x_1, \ldots, x_n)$.

To verify the claim, suppose that X^0 has probability mass function $p(x_1, \ldots, x_n)$. Then it is easy to see that at any point in this algorithm the vector $X_1^j, \ldots, X_{i-1}^j, X_i^{j-1}, \ldots, X_n^{j-1}$ will be the value of a random variable with mass function $p(x_1, \ldots, x_n)$. For instance, letting X_i^j be the random variable that takes on the value denoted by x_i^j then

$$\begin{aligned}
&\mathbb{P}\left\{X_1^1 = x_1, X_j^0 = x_j, j = 2, \ldots, n\right\} \\
&\quad = \mathbb{P}\left\{X_1^1 = x_1 | X_j^0 = x_j, j = 2, \ldots, n\right\} \\
&\qquad \times \quad \mathbb{P}\left\{X_j^0 = x_j, j = 2, \ldots, n\right\} \\
&\quad = \mathbb{P}\{X_1 = x_1 | X_j = x_j, j = 2, \ldots, n\} \\
&\qquad \times \quad \mathbb{P}\{X_j = x_j, j = 2, \ldots, n\} \\
&\quad = p(x_1, \ldots, x_n).
\end{aligned}$$

Therefore, $p(x_1, \ldots, x_n)$ is a stationary probability distribution, so provided that the Markov chain is irreducible and aperiodic, we can conclude that it is the limiting probability vector for the Gibbs sampler.

It also follows from the proceeding that $p(x_1, \ldots, x_n)$ would be the limiting probability vector even if the Gibbs sampler were not systematic in first changing the value of X_1, then X_2, and so on., Indeed, even if the component whose value was to be changed was always randomly determined, then $p(x_1, \ldots, x_n)$ would remain a stationary distribution, and would thus be the limiting probability mass function provided that the resulting chain is aperiodic and irreducible.

2 Exercises

1) Each day one of n possible elements is requested; it is the ith one with probability $P_i, i \geq 1, \sum_{i=1}^{n} P_i = 1$. These elements are at all times arranged in an ordered list that is revised as follows: the element selected is moved to the front of the list, and the relative positions of all other elements remain unchanged. Define the state at any time to be the ordering of the list at that time.

(a) Argue that the above is Markov chain.

(b) For any state $(i_1, \ldots, i_n)$ (which is a permutation of $(1, 2, \ldots, n)$) let $\pi(i_1, \ldots, i_n)$ denote the limiting probability. Argue that

$$\pi(i_1, \ldots, i_n) = P_{i_1} \frac{P_{i_2}}{1 - P_{i_1}} \cdots \frac{P_{i_{n-1}}}{1 - P_{i_1} - \cdots - P_{i_{n-2}}}.$$

2) Let $\{X_n, n \geq 0\}$ be a Markov chain with stationary probabilities π_j, $j \geq 0$. Suppose that $X_0 = 0$ and define

$$T = \min\{n > 0 : X_n = 0\}.$$

Let $Y_j = X_{T-j}, j = 0, 1, \ldots, T$. Show that $\{Y_j, j = 0, \ldots, T\}$ is distributed as the states visited by a Markov chain (the "reversed" Markov chain) with transition probabilities $P^*_{ij} = \pi_j P_{ji} / \pi_i$ started in state 0 and watched until it returns to 0.

3) Consider a finite Markov chain on the state space $\{0, 1, 2, \ldots, N\}$ with transition probability matrix $\mathbf{P} = (P_{ij})_{i,j=0}^{N}$, and divide the state space into the three classes $\{0\}$, $\{1, 2, \ldots, N-1\}$ and $\{N\}$. Let 0 and N be absorbing states, both accessible from all states in $1, \ldots, N-1$, and let $\{1, 2, \ldots, N-1\}$ be a transient class.

Let k be a transient state. Define an auxiliary process (the "return process") with transition matrix $\widetilde{P}$ by altering the first and last row of $\mathbf{P}$ so that $\widetilde{P}_{0k} = \widetilde{P}_{Nk} = 1$ and leave the other rows unchanged.

The return process is clearly irreducible. Prove that the expected time until absorption μ_k with initial state k in the original process equals $1/(\pi_0 + \pi_N) - 1$ where $\pi_0 + \pi_N$ is the stationary probability of being in state 0 or N for the return process.
Hint: use the relation between stationary probabilities and expected recurrence times to states.

3 Reversibility

Suppose that $\{X_n : 0 \leq n \leq N\}$ is an irreducible, non-null, persistent Markov chain, with transition matrix $\mathbf{P}$ and stationary distribution $\boldsymbol{\pi}$. Suppose further that X_n has distribution $\boldsymbol{\pi}$ for every n. Define the 'reversed chain' Y by $Y_n = X_{N-n}$ for $0 \leq n \leq N$.

Proposition 3.1. *The sequence* Y *is a Markov chain with* $\mathbb{P}\{Y_0 = i\} = \pi_i$ *and*

$$\mathbb{P}\left\{Y_{n+1} = j | Y_n = i\right\} = \frac{\pi_j}{\pi_i} p_{ji}.$$

We call the chain Y the *time reversal* of chain X, and we say that X is *reversible* if X and Y have the same transition probabilities.

Proof. The crucial step is the stationarity of X:

$$
\begin{aligned}
&\mathbb{P}\{Y_{n+1} = i_{n+1} | Y_n = i_n, Y_{n-1} = i_{n-1}, \ldots, Y_0 = i_0\} \\
&\quad = \frac{\mathbb{P}\{Y_k = i_k, 0 \le k \le n+1\}}{\mathbb{P}\{Y_k = i_k, 0 \le k \le n\}} \\
&\quad = \frac{\mathbb{P}\{X_{N-n-1} = i_{n+1}, X_{N-n} = i_n, \ldots, X_N = i_0\}}{\mathbb{P}\{X_{N-n} = i_n, \ldots, X_N = i_0\}} \\
&\quad = \frac{\pi_{i_{n+1}} p_{i_{n+1}, i_n} p_{i_n, i_{n-1}} \cdots p_{i_1, i_0}}{\pi_{i_n} p_{i_n, i_{n-1}} \cdots p_{i_1, i_0}} \\
&\quad = \frac{\pi_{i_{n+1}} p_{i_{n+1}, i_n}}{\pi_{i_n}}.
\end{aligned}
$$

□

Let $X = \{X_n : 0 \leq n \leq N\}$ be an irreducible Markov chain such that X_n has the stationary distribution $\boldsymbol{\pi}$ for all n. The chain is called *reversible* if the transition matrices of X and its time-reversal Y are the same, which is to say that

$$\pi_i p_{ij} = \pi_j p_{ji} \qquad \text{for all } i,j.$$

These equations are called the *detailed balance* equations.

Proposition 3.2. *Let* $\mathbf{P}$ *be the transition matrix of an irreducible chain* X*, and suppose that there exists a distribution* $\boldsymbol{\pi}$ *such that* $\pi_i p_{ij} = \pi_j p_{ji}$ *for all* $i,j \in S$*. Then* $\boldsymbol{\pi}$ *is a stationary distribution of the chain.*

Proof. Suppose that $\boldsymbol{\pi}$ satisfies the conditions above. Then

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

and so $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, whence π is stationary. □

3.1 Reversible Examples

3.1.1 Example: Ehrenfest model of diffusion

Two containers A and B are placed adjacent to each other and gas is allowed to pass through a small aperture joining them. A total of m gas molecules is distributed between the containers. We assume that at each epoch of time one molecule, picked uniformly at random from the m available, passes through this aperture. Let X_n be the number of molecules in container A after n units of time has passed. Clearly $\{X_n\}$ is a Markov chain with transition matrix

$$p_{i,i+1} = 1 - \frac{i}{m}, \qquad p_{i,i-1} = \frac{i}{m} \quad \text{if} \quad 0 \le i \le m.$$

Rather than solve the equation $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ to find the stationary distribution, we look for solutions of the detailed balance equations

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

.

This is solved by $\pi_i = \binom{m}{i}(\frac{1}{2})^m$, which is therefore the stationary distribution.

3.1.2 Example: the Metropolis algorithm

Let a_j ,$j = 1, \ldots, m$ be positive numbers and let $A = \sum_{j=1}^{m} a_j$. Suppose that m is large and that A is difficult to compute, and suppose we ideally want to simulate the values of a sequence of independent random variables whose probabilities are $p_j = a_j/A$, for $j = 1, \ldots, m$.

Similar to the Gibbs sampler, one way of simulating a sequence of random variables whose distributions converge to $\{p_j, j = 1, \ldots, m\}$ is to find a Markov chain that is both easy to simulate and whose limiting probabilities are the p_j. The *Metropolis algorithm* provides an approach for accomplishing this task.

Let $\mathbf{Q}$ be any irreducible transition probability matrix on the integers $1, 2, \ldots, n$ such that $q_{ij} = q_{ji}$ for all i and j. Now define a Markov chain $\{X_n, n \geq 0\}$ as follows. If $X_n = i$, then generate a random variable that is equal to j with probability $q_{ij}, i, j = 1, \ldots, m$. If this random variable takes on the value j, then set X_{n+1} equal to j with probability $\min\{1, a_j/a_i\}$, and set it equal to i otherwise.

That is, the transition probabilities of $\{X_n, n \geq 0\}$ are

$$P_{ij} = \begin{cases} q_{ij} \min(1, a_j/a_i) & \text{if } j \neq i, \\ q_{ii} + \sum_{j \neq i} q_{ij}\{1 - \min(1, a_j/a_i)\} & \text{if } j = i. \end{cases}$$

We will now show that the stationary distribution of this Markov chain is given by the p_j.

We will first show that the chain is reversible with stationary probabilities $p_j, j = 1, \ldots, m$ by showing that

$$p_i P_{ij} = p_j P_{ji}.$$

To show this, we must show that

$$p_i q_{ij} \min(1, a_i/a_j) = p_j q_{ji} \min(1, a_j/a_i).$$

Now, $q_{ij} = q_{ji}$ and $a_j/a_i = p_j/p_i$ and so we must verify that

$$p_i \min(1, p_j/p_i) = p_j \min(1, p_i/p_j).$$

This is immediate since both sides of the equation are equal to $\min(p_i, p_j)$.

That these stationary probabilities are also limiting probabilities follows from the fact that since $\mathbf{Q}$ is an irreducible transition probability matrix, $\{X_n\}$ will also be irreducible, and as $P_{ii} > 0$ for some i (except in the trivial case where $p_i \equiv 1/n$), it is also aperiodic.

By choosing a transition probability matrix $\mathbf{Q}$ that is easy to simulate – that is, for each i it is easy to generate the value of a random variable that is equal to j with probability q_{ij} – we can use the preceding to generate a Markov chain whose limiting probabilities are a_j/A, without computing A.

3.1.3 Example: Random walk on a graph

Consider a graph having a positive number w_{ij} associated with each edge (i, j), and suppose that a particle moves from vertex to vertex in the following manner:

If the particle is at vertex i then it will move to vertex j with probability proportional to the outgoing edge weights:

$$P_{ij} = w_{ij} / \sum_j w_{ij}$$

where w_{ij} is 0 if (i, j) is not an edge of the graph. The Markov chain describing the sequence of vertices visited by the particle is called a random walk on an edge weighted graph.

Proposition 3.3. *Consider a random walk on an edge weighted graph with a finite number of vertices. If this Markov chain is irreducible, then in steady state it is time reversible with stationary probabilities given by*

$$\pi_i = \frac{\sum_i w_{ij}}{\sum_j \sum_i w_{ij}}.$$

Proof. The time reversible equations

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

reduce to

$$\frac{\pi_i w_{ij}}{\sum_k w_{ik}} = \frac{\pi_j w_{ji}}{\sum_k w_{jk}}$$

or, equivalently, since $w_{ij} = w_{ji}$

$$\frac{\pi_i}{\sum_k w_{ik}} = \frac{\pi_j}{\sum_k w_{jk}}$$

implying that

$$\pi_i = c \sum_k w_{ik}.$$

Since $\sum \pi_i = 1$, we are done. $\square$

4 Exercises

1) Consider a time-reversible Markov chain on the state space $\{0, 1, 2, \ldots\}$ with transition probabilities P_{ij} and limiting probabilities π_i. Suppose we truncate the chain to the states $\{0, 1, \ldots, M\}$ by defining the transition probabilities

$$\overline{P}_{ij} = \begin{cases} P_{ij} + \sum_{k>M} P_{ik}, & 0 \leq i \leq M, j = i \\ P_{ij}, & 0 \leq i \neq j \leq M \\ 0, & \text{otherwise.} \end{cases}$$

Show that the truncated chain is also time reversible and has limiting probabilities given by

$$\overline{\pi}_i = \frac{\pi_i}{\sum_{i=0}^{M} \pi_i}.$$

2) Suppose M balls are initially distributed among m urns. At each stage one of the balls is selected at random, taken from whichever urn it is in, and placed, again at random, in one of the other $m-1$ urns. Consider the Markov chain whose state at any time is the vector $(n_1, \ldots, n_m)$, where n_i denotes the number of balls in urn i. Guess at the limiting probabilities for this Markov chain and verify your guess, showing at the same time that the Markov chain is time reversible.