CHAPTER 3

Recurrence and Ergodicity

1 Potential Matrix Criterion

1.1 Recurrent and Transient States

Consider a Markov chain taking its values in $E = \mathbb{N}$. There is a possibility that for any initial state $i \in \mathbb{N}$ the chain will never visit i after some finite random time. This is often an undesirable feature. For example, if the chain counts the number of customers waiting in line at a service counter (we shall see Markovian models of waiting lines, or *queues*, at different places in this book), such a behavior implies that the waiting line will eventually go beyond the limits of the waiting facility. In a sense, the corresponding system is unstable.

The good notion of stability for an irreducible HMC is that of *positive recurrence*, when any given state is visited infinitely often and when, moreover, the average time between two successive visits to this state is finite. The principal problem is to find sufficient, and maybe necessary, conditions guaranteeing stability. We begin with the *potential matrix* criterion (necessary and sufficient condition), which is of mainly theoretical interest, and the *stationary distribution criterion*. Further conditions, such as *Foster's theorem*, will be given in Chapter 5.

For the time being, we introduce the relevant definitions. First recall that T_i denotes the *return* time to state i.

Definition 1.1. *Recurrence and Transience*
State $i \in E$ is called *recurrent* if

$$P_i(T_i < \infty) = 1, \tag{1.1}$$

and otherwise it is called *transient*. A recurrent state $i \in E$ is called *positive* recurrent if

$$E_i[T_i] < \infty, \tag{1.2}$$

and otherwise it is called *null* recurrent.

Example 1.1. *Success Runs*
The rule of the game is the following: A coin is tossed repeatedly, and whenever the result is tails (probability $q = 1 - p$), you go one step up the ladder, but if the result is heads, you fall all the way down. If X_n is your position at time n, $\{X_n\}_{n\geq 0}$ forms a homogeneous Markov chain with state space $E = \mathbb{N}$ and is a special case of the chain with the transition graph in Figure 3.1.1.

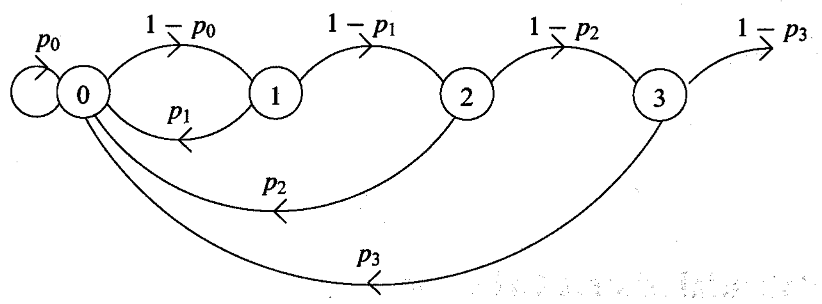

Figure 3.1.1. Transition graph of the success-runs chain

The state space of the chain of Figure 3.1.1 is $E = \mathbb{N}$, and we impose the condition

$$p_i \in (0, 1) \text{ for all } i \in E,$$

which guaranteees irreducibility. We shall compute the probability of returning to state 0 and the mean return time to state 0, and from the expressions obtained, we shall deduce the nature of state 0.

There is just one way of going from state 0 to state 0 in exactly n steps. The corresponding path is $0, 1, 2, \ldots, n-1, 0$, and therefore $P_0(T_0 = 1) = p_0$, and for $n \geq 1$,

$$P_0(T_0 = n) = (1 - p_0) \cdots (1 - p_{n-2}) p_{n-1}.$$

Defining $u_0 = 1$, and for $n \geq 1$,

$$u_n = (1 - p_0) \cdots (1 - p_{n-1}),$$

we see that

$$P_0(T_0 = n) = u_{n-1} - u_n.$$

Since

$$P_0(T_0 < \infty) = \sum_{n=1}^{\infty} P_0(T_0 = n) = \lim_{m\uparrow\infty} \sum_{n=1}^{m} P_0(T_0 = n) = \lim_{m\uparrow\infty} (1 - u_m),$$

we have

$$P_0(T_0 < \infty) = 1 - \lim_{m\uparrow\infty} \prod_{i=0}^{m-1} (1 - p_i).$$

Therefore, in view of a classical result on infinite products (see Theorem 1.9 of the Appendix)

$$P_0(T_0 < \infty) = 1 \Leftrightarrow \prod_{i=0}^{\infty} (1 - p_i) = 0 \Leftrightarrow \sum_{i=0}^{\infty} p_i = \infty. \qquad \diamond$$

In general, it is not easy to check whether a given state is transient or recurrent. One of the goals of the theory of Markov chains is to provide criteria of recurrence. Sometimes, one is happy with just a sufficient condition, or a necessary condition.

The problem of finding useful (easy to check) conditions of recurrence is an active area of research. However, the theory has a few conditions that qualify as useful and are applicable to many practical situations. Although the next criterion is of theoretical rather than practical interest, it can be helpful in a few situations, for instance in the study of recurrence of random walks (see Examples 1.2 and 1.3 below).

1.2 Potential Matrix

The *potential matrix* **G** associated with the transition matrix **P** is defined by

$$\mathbf{G} = \sum_{n\geq 0} \mathbf{P}^n.$$

Its general term

$$g_{ij} = \sum_{n=0}^{\infty} p_{ij}(n) = \sum_{n=0}^{\infty} P_i(X_n = j) = \sum_{n=0}^{\infty} E_i[1_{\{X_n=j\}}] = E_i\left[\sum_{n=0}^{\infty} 1_{\{X_n=j\}}\right]$$

is the average number of visits to state j, given that the chain starts from state i.

Theorem 1.1. *Potential Matrix Criterion*
State $i \in E$ is recurrent if and only if

$$\sum_{n=0}^{\infty} p_{ii}(n) = \infty. \tag{1.3}$$

Proof. Theorem 1.1 merely rephrases Theorem 7.3 of Chapter 2. □

Example 1.2. *1-D Random Walk*
The corresponding Markov chain was described in Example 2.1 of Chapter 2. The nonzero terms of its transition matrix are

$$p_{i,i+1} = p\ ,\ p_{i,i-1} = 1 - p,$$

where $p \in (0, 1)$. We shall study the nature (recurrent or transient) of any one of its states, say, 0. We have $p_{00}(2n + 1) = 0$ and

$$p_{00}(2n) = \frac{(2n)!}{n!n!} p^n (1-p)^n.$$

By Stirling's equivalence formula $n! \sim (n/e)^n \sqrt{2\pi n}$, the above quantity is equivalent to

$$\frac{[4p(1-p)]^n}{\sqrt{\pi n}}, \tag{1.4}$$

and the nature of the series $\sum_{n=0}^{\infty} p_{00}(n)$ (convergent or divergent) is that of the series with general term (1.4). If $p \neq \frac{1}{2}$, in which case $4p(1-p) < 1$, the latter series converges, and if $p = \frac{1}{2}$, in which case $4p(1-p) = 1$, it diverges. In summary, the states of the 1-D random walk are transient if $p \neq \frac{1}{2}$, recurrent if $p = \frac{1}{2}$.

Example 7.6 of Chapter 2 shows that for the *symmetric* ($p = \frac{1}{2}$) 1-D random walk, the states are in fact *null* recurrent. ◇

Example 1.3. *3-D Symmetric Random Walk*
The state space of this HMC is $E = \mathbb{Z}^3$. Denoting by e_1, e_2, and e_3 the canonical basis vectors of $\mathbb{R}^3$ (respectively $(1, 0, 0), (0, 1, 0)$, and $(0, 0, 1)$), the nonnull terms of the transition matrix of the 3-D symmetric random walk are given by

$$p_{x, x\pm e_i} = \frac{1}{6}.$$

We elucidate the nature of state, say, $0 = (0, 0, 0)$. Clearly, $p_{00}(2n + 1) = 0$ for all $n \geq 0$, and (exercise)

$$p_{00}(2n) = \sum_{0 \leq i+j \leq n} \frac{(2n)!}{(i!j!(n-i-j)!)^2} \left(\frac{1}{6}\right)^{2n}.$$

This can be rewritten as

$$p_{00}(2n) = \sum_{0 \leq i+j \leq n} \frac{1}{2^{2n}} \binom{2n}{n} \left(\frac{n!}{i!j!(n-i-j)!}\right)^2 \left(\frac{1}{3}\right)^{2n}.$$

Using the *trinomial formula*

$$\sum_{0 \leq i+j \leq n} \frac{n!}{i!j!(n-i-j)!} \left(\frac{1}{3}\right)^n = 1,$$

we obtain the bound

$$p_{00}(2n) \leq K_n \frac{1}{2^{2n}} \binom{2n}{n} \left(\frac{1}{3}\right)^n,$$

where

$$K_n = \max_{0 \leq i+j \leq n} \frac{n!}{i!j!(n-i-j)!}.$$

For large values of n, K_n is bounded as follows. Let i_0 and j_0 be the values of i, j that maximize $n!/(i!j!(n+j)!)$ in the domain of interest $0 \leq i + j \leq n$. From the definition of i_0 and j_0, the quantities

$$\frac{n!}{(i_0 - 1)!j_0!(n - i_0 - j_0 + 1)!},$$

$$\frac{n!}{(i_0 + 1)!j_0!(n - i_0 - j_0 - 1)!},$$

$$\frac{n!}{i_0!(j_0 - 1)!(n - i_0 - j_0 + 1)!},$$

$$\frac{n!}{i_0!(j_0 + 1)!(n - i_0 - j_0 - 1)!},$$

are bounded by

$$\frac{n!}{i_0!j_0!(n - i_0 - j_0)!}.$$

The corresponding inequalities reduce to

$$n - i_0 - 1 \leq 2j_0 \leq n - i_0 + 1 \text{ and } n - j_0 - 1 \leq 2i_0 \leq n - j_0 + 1,$$

and this shows that for large n, $i_0 \sim n/3$ and $j_0 \sim n/3$. Therefore, for large n,

$$p_{00}(2n) \sim \frac{n!}{(n/3)!(n/3)!2^{2n}e^n} \binom{2n}{n}.$$

By Stirling's equivalence formula, the right-hand side of the latter equivalence is in turn equivalent to

$$\frac{3\sqrt{3}}{2(\pi n)^{3/2}},$$

the general term of a divergent series. State 0 is therefore transient. ◇

Suppose that state $i \in E$ is recurrent, and accessible from state $j \in E$. That is, starting from j, the probability of visiting i at least once is positive (accessibility of i from j), and starting from i, the average number of visits to i is infinite (recurrence of i). Therefore, starting from j the average number of visits to i is infinite:

$$E_j[N_i] = \sum_{n \geq 1} p_{ji}(n) = \infty.$$

Similarly, if i is transient, then for any state $j \in E$,

$$E_j[N_i] = \sum_{n \geq 1} p_{ji}(n) < \infty.$$

1.3 Structure of the Transition Matrix

A theoretical application of the potential matrix criterion is to the proof that recurrence is a (communication) class property.

Theorem 1.2. *Recurrence Is a Class Property*
If i and j communicate, they are either both recurrent or both transient.

Proof. By definition, i and j communicate if and only if there exist integers M and N such that $p_{ij}(M) > 0$, $p_{ji}(N) > 0$. Going from i to j in M steps, then from j to j in n steps, then from j to i in N steps, is just one way of going from i back to i in $M+n+N$ steps. Therefore, $p_{ii}(M+n+N) \geq p_{ij}(M)p_{jj}(n)p_{ji}(N)$. Similarly, $p_{jj}(N+n+M) \geq p_{ji}(N)p_{ii}(n)p_{ij}(M)$. Therefore, writing $\alpha = p_{ij}(M)p_{ji}(N)$ (a strictly positive quantity), we have $p_{ii}(M+N+n) \geq \alpha p_{jj}(n)$ and $p_{jj}(M+N+n) \geq \alpha p_{ii}(n)$. This implies that the series $\sum_{n=0}^{\infty} p_{ii}(n)$ and $\sum_{n=0}^{\infty} p_{jj}(n)$ either both converge or both diverge. Theorem 1.1 concludes the proof. □

It will be proven later in this chapter that positive recurrence (resp., null recurrence) is also a class property, in the sense that if states i and j communicate and if one of them is positive recurrent (resp., null recurrent), then so is the other.

An irreducible Markov chain has therefore all its states of the same nature: transient, positive recurrent, or null recurrent. We shall therefore call it a transient chain, a positive recurrent chain, or a null recurrent chain, and to determine to which category it belongs, it suffices to study *one* state, selecting the state for which the computations seem easiest (such as state 0 for the chain of Example 1.1).

It follows from the above discussion that there are two types of communication classes: the *transient classes* and the *recurrent classes*. Call T the set of all transient states and R the set of all recurrent states. The set R may be composed of several disjoint communication classes R_1, R_2, etc. Any recurrent communication class, R_1 for instance, is closed. Indeed, if the chain goes from $i \in R_1$ to some $j \in E$, it will have to come back to i, since i is recurrent, and therefore i and j must communicate, so that j must be in R_1. The communication structure of a transition matrix is therefore as shown in Figure 3.1.2.

2 Recurrence and Invariant Measures

The notion of invariant measure plays an important technical role in the recurrence theory of Markov chains. It extends the notion of stationary distribution.

Definition 2.1. *Invariant Measure.*
A nontrivial (that is, nonnull) vector $x = \{x_i\}_{i \in E}$ is called an *invariant measure* of the stochastic matrix $\mathbf{P} = \{p_{ij}\}_{i,j \in E}$ if for all $i \in E$,

$$x_i \in [0, \infty) \tag{2.1}$$

$\mathbf{P} =$

	R_1	R_2	R_3	T	
	$\mathbf{P}_1$	0	0	0	R_1
	0	$\mathbf{P}_2$	0	0	R_2
	0	0	$\mathbf{P}_3$	0	R_3
					T

Figure 3.1.2

and

$$x_i = \sum_{j \in E} x_j p_{ji}. \tag{2.2}$$

(In abbreviated notation, $0 \leq x < \infty$ and $x^T\mathbf{P} = x^T$.)

Theorem 2.1. *Regenerative Form of Invariant Measure*
Let $\mathbf{P}$ be the transition matrix of an irreducible recurrent HMC $\{X_n\}_{n\geq 0}$. Let 0 be an arbitrary state and let T_0 be the return time to 0. Define for all $i \in E$

$$x_i = E_0\left[\sum_{n\geq 1} 1_{\{X_n = i\}} 1_{\{n \leq T_0\}}\right] \tag{2.3}$$

(For $i \neq 0$, x_i is the expected number of visits to state i before returning to 0). Then, for all $i \in E$,

$$x_i \in (0, \infty), \tag{2.4}$$

and x is an invariant measure of $\mathbf{P}$.

Observe that for $n \in [1, T_0]$, $X_n = 0$ if and only if $n = T_0$. Therefore,

$$x_0 = 1. \tag{2.5}$$

Also, $\sum_{i \in E} \sum_{n \geq 1} 1_{\{X_n = i\}} 1_{\{n \leq T_0\}} = \sum_{n \geq 1} \left\{\sum_{i \in E} 1_{\{X_n = i\}}\right\} 1_{\{n \leq T_0\}} = \sum_{n\geq 1} 1_{\{n \leq T_0\}} = T_0$, and therefore

$$\sum_{i \in E} x_i = E_0[T_0]. \tag{2.6}$$

For the proof of Theorem 2.1, we introduce the quantity

$${}_0p_{0i}(n) \stackrel{\text{def}}{=} E_0[1_{\{X_n = i\}} 1_{\{n \leq T_0\}}] = P_0(X_1 \neq 0, \cdots, X_{n-1} \neq 0, X_n = i). \tag{2.7}$$

This is the probability, starting from state 0, of visiting i at time n before returning to 0. From the definition of x,

$$x_i = \sum_{n \geq 1} {}_0p_{0i}(n). \tag{2.8}$$

Proof. (*of Theorem 2.1*) We first prove (2.2). Observe that

$${}_0p_{0i}(1) = p_{0i} \tag{2.9}$$

and, using first-step analysis, for all $n \geq 2$,

$${}_0p_{0i}(n) = \sum_{j \neq 0} {}_0p_{0j}(n-1)p_{ji} \tag{2.10}$$

(see Problem 3.2.1). Summing up all the above equalities, and taking (2.8) into account, we obtain

$$x_i = p_{0i} + \sum_{j \neq 0} x_j p_{ji},$$

that is, (2.2), since $x_0 = 1$ (see (2.5)).

Next we show that $x_i > 0$ for all $i \in E$. Indeed, iterating (2.2), we find $x^T = x^T\mathbf{P}^n$, that is, since $x_0 = 1$,

$$x_i = \sum_{j \in E} x_j p_{ji}(n) = p_{0i}(n) + \sum_{j \neq 0} x_j p_{ji}(n).$$

If x_i were null for some $i \in E$, $i \neq 0$, the latter equality would imply that $p_{0i}(n) = 0$ for all $n \geq 0$, which means that 0 and i do not communicate, in contradiction to the irreducibility assumption.

It remains to show that $x_i < \infty$ for all $i \in E$. As before, we find that

$$1 = x_0 = \sum_{j \in E} x_j p_{j0}(n)$$

for all $n \geq 1$, and therefore if $x_i = \infty$ for some i, necessarily $p_{i0}(n) = 0$ for all $n \geq 1$, and this also contradicts irreducibility. □

Theorem 2.2. *Uniqueness of Invariant Measure*
The invariant measure of an irreducible recurrent stochastic matrix is unique up to a multiplicative factor.

Proof. In the proof of Theorem 2.1 we showed that for an invariant measure y of an irreducible chain, $y_i > 0$ for all $i \in E$, and therefore, one can define, for all $i, j \in E$, the matrix $\mathbf{Q}$ by

$$q_{ji} = \frac{y_i}{y_j} p_{ij}. \tag{2.11}$$

It is a transition matrix, since $\sum_{i \in E} q_{ji} = \frac{1}{y_j} \sum_{i \in E} y_i p_{ij} = \frac{y_j}{y_j} = 1$. The general term of $\mathbf{Q}^n$ is

$$q_{ji}(n) = \frac{y_i}{y_j} p_{ij}(n). \tag{2.12}$$

Indeed, supposing (2.12) true for n,

$$q_{ji}(n+1) = \sum_{k \in E} q_{jk} q_{ki}(n) = \sum_{k \in E} \frac{y_k}{y_j} p_{kj} \frac{y_i}{y_k} p_{ik}(n)$$
$$= \frac{y_i}{y_j} \sum_{k \in E} p_{ik}(n) p_{kj} = \frac{y_i}{y_j} p_{ij}(n+1),$$

and (2.12) follows, by induction, for all $n \geq 1$.

Clearly, $\mathbf{Q}$ is irreducible, since $\mathbf{P}$ is irreducible (just observe that $q_{ji}(n) > 0$ if and only if $p_{ij}(n) > 0$ in view of (2.12)). Also, $p_{ii}(n) = q_{ii}(n)$, and therefore $\sum_{n \geq 0} q_{ii}(n) = \sum_{n \geq 0} p_{ii}(n)$, and this ensures that $\mathbf{Q}$ is recurrent by the potential matrix criterion. Call $g_{ji}(n)$ the probability, relative to the chain governed by the transition matrix $\mathbf{Q}$, of returning to state i for the first time at step n when starting from j. First-step analysis gives $g_{i0}(n+1) = \sum_{j \neq 0} q_{ij} g_{j0}(n)$ (see Problem 3.2.1), that is, using (2.11),

$$y_i g_{i0}(n+1) = \sum_{j \neq 0} (y_j g_{j0}(n)) p_{ji}.$$

Recall that ${}_0p_{0i}(n+1) = \sum_{j \neq 0} {}_0p_{0j}(n) p_{ji}$, or, equivalently,

$$y_0\, {}_0p_{0i}(n+1) = \sum_{j \neq 0} (y_0\, {}_0p_{0j}(n)) p_{ji}.$$

We therefore see that the sequences $\{y_0\, {}_0p_{0i}(n)\}$ and $\{y_i g_{i0}(n)\}$ satisfy the same recurrence equation. Their first terms ($n = 1$), respectively $y_0\, {}_0p_{0i}(1) = y_0 p_{0i}$ and $y_i g_{i0}(1) = y_i q_{i0}$, are equal in view of (2.11). Therefore, for all $n \geq 1$,

$${}_0p_{0i}(n) = \frac{y_i}{y_0} g_{i0}(n).$$

Summing up with respect to $n \geq 1$ and using $\sum_{n \geq 1} g_{i0}(n) = 1$ ($\mathbf{Q}$ is recurrent), we obtain the announced result $x_i = \frac{y_i}{y_0}$. □

Equality (2.6) and the definition of positive recurrence give the following.

Theorem 2.3. *Positive vs. Null Recurrence.*
An irreducible recurrent HMC is positive recurrent if and only if its invariant measures x satisfy

$$\sum_{i \in E} x_i < \infty. \tag{2.13}$$

Remark 2.1.
An HMC may well be irreducible and possess an invariant measure, and yet not be recurrent. The simplest example is the 1-D nonsymmetric random walk, which was shown to be transient (Example 1.2) and which admits $x_i \equiv 1$ for invariant measure. ◇

3 Positive Recurrence

3.1 Stationary Distribution Criterion

In the previous section, an irreducible Markov chain was assumed recurrent, and it was shown that it has a unique stationary distribution if it is positive recurrent. It was also observed that the existence of an invariant measure is not sufficient for recurrence. It turns out, however, that the existence of a stationary probability distribution is necessary and sufficient for an irreducible chain (not a priori assumed recurrent) to be recurrent positive.

Theorem 3.1. *Stationary Distribution Criterion*
An irreducible homogeneous Markov chain is positive recurrent if and only if there exists a stationary distribution. Moreover, the stationary distribution π is, when it exists, unique, and $\pi > 0$.

Proof. The direct part follows from Theorems 2.1 and 2.3. For the converse part, assume the existence of a stationary distribution π. Iterating $\pi^T = \pi^T \mathbf{P}$, we obtain $\pi^T = \pi^T \mathbf{P}^n$, that is, for all $i \in E$,

$$\pi(i) = \sum_{j \in E} \pi(j) p_{ji}(n).$$

If the chain were transient, then, in view of the potential matrix criterion and the discussion following it, for all states i, j,

$$\lim_{n \uparrow \infty} p_{ji}(n) = 0,$$

and since $p_{ji}(n)$ is bounded by 1 uniformly in j and n, by the dominated convergence theorem for series (see Theorem 1.6 of the Appendix)

$$\pi(i) = \lim_{n \uparrow \infty} \sum_{j \in E} \pi(j) p_{ji}(n) = \sum_{j \in E} \pi(j) \left(\lim_{n \uparrow \infty} p_{ji}(n) \right) = 0.$$

This contradicts the assumption that π is a stationary distribution (in particular, $\sum_{i \in E} \pi(i) = 1$). The chain must therefore be recurrent, and by Theorem 2.3, it is positive recurrent.

The stationary distribution π of an irreducible positive recurrent chain is unique (use Theorem 2.2 and the fact that there is no choice for a multiplicative factor but 1). Also recall that $\pi(i) > 0$ for all $i \in E$ (see Theorem 2.1). □

Theorem 3.2. *Mean Return Time*
Let π be the unique stationary distribution of an irreducible positive recurrent chain, and let T_i be the return time to state i. Then

$$\pi(i) E_i[T_i] = 1. \tag{3.1}$$

Proof. This equality is a direct consequence of expression (2.3) for the invariant measure. Indeed, π is obtained by normalization of x: for all $i \in E$,

$$\pi(i) = \frac{x_i}{\sum_{j \in E} x_j},$$

and in particular, for $i = 0$, using (2.5) and (2.6),

$$\pi(0) = \frac{x_0}{\sum_{j \in E} x_j} = \frac{1}{E_0[T_0]}.$$

Since state 0 does not play a special role in the analysis, (3.1) is true for all $i \in E$. □

The situation is extremely simple when the state space is finite.

Theorem 3.3. *Finite State Space and Positive Recurrence.*
An irreducible HMC with finite state space is positive recurrent.

Proof. We first show recurrence. If the chain were transient, then, from the potential matrix criterion and the observations following it, for all $i, j \in E$,

$$\sum_{n \geq 0} p_{ij}(n) < \infty,$$

and therefore, since the state space is finite

$$\sum_{j \in E} \sum_{n \geq 0} p_{ij}(n) < \infty.$$

But the latter sum is equal to

$$\sum_{n \geq 0} \sum_{j \in E} p_{ij}(n) = \sum_{n \geq 0} 1 = \infty,$$

a contradiction. Therefore, the chain is recurrent. By Theorem 2.1 it has an invariant measure x. Since E is finite, $\sum_{i \in E} x_i < \infty$, and therefore the chain is positive recurrent, by Theorem 3.1. □

A "talk proof" of recurrence in Theorem 3.3 is available: The states cannot be all visited only a finite number of times; otherwise, there would exist a finite random time after which no state is visited!

3.2 Examples

Example 3.1. *Random Walk Reflected at 0*
This chain has the state space $E = \mathbb{N}$ and the transition graph of Figure 3.3.1. It is assumed that p_i (and therefore $q_i = 1 - p_i$) are in the open interval $(0, 1)$ for all $i \in E$, so that the chain is irreducible.
The invariant measure equation $x^T = x^T \mathbf{P}$ takes in this case the form

$$\begin{aligned} x_0 &= x_1 q_1, \\ x_i &= x_{i-1} p_{i-1} + x_{i+1} q_{i+1}, \ i \geq 1, \end{aligned}$$

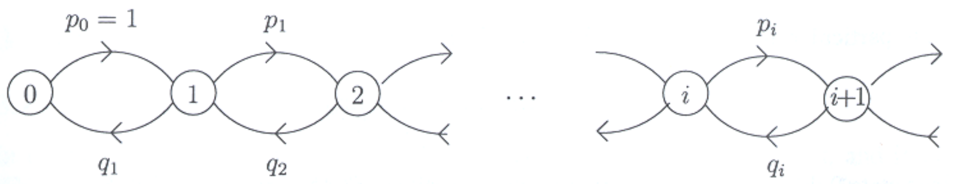

Figure 3.3.1. Reflected random walk

with $p_0 = 1$. The general solution is, for $i \geq 1$,

$$x_i = x_0 \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i}.$$

The positive recurrence condition $\sum_{i \in E} x_i < \infty$ is

$$1 + \sum_{i \geq 1} \frac{p_0 \dots p_{i-1}}{q_1 \dots q_i} < \infty, \tag{3.2}$$

and if it is satisfied, the stationary distribution π is obtained by normalization of the general solution. This gives

$$\pi(0) = \left(1 + \sum_{i \geq 1} \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i}\right)^{-1}, \tag{3.3}$$

and for $i \geq 1$,

$$\pi(i) = \pi(0) \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i}. \tag{3.4}$$

In the special case where $p_i = p$, $q_i = q = 1 - p$, the positive recurrence condition becomes $1 + \frac{1}{q} \sum_{j \geq 0} \left(\frac{p}{q}\right)^j < \infty$, that is to say $p < q$, or equivalently,

$$p < \frac{1}{2}.$$

◇

Example 3.2. *Success Runs and Machine Replacement*
The result of Example 1.1 will be derived once again, this time via the stationary distribution criterion. We shall set $q_i = 1 - p_i$. Equality $x^T = x^T \mathbf{P}$ takes the form

$$x_0 = p_0 x_0 + p_1 x_1 + p_2 x_2 + \cdots,$$

and for $i \geq 1$,

$$x_i = q_{i-1} x_{i-1}.$$

Therefore, leaving aside the first equality, for $i \geq 1$,

$$x_i = (q_0 q_1 \cdots q_{i-1}) x_0.$$

Discarding the possibility $x_0 \leq 0$, which would imply that x is negative or null, the first equation is satisfied if and only if

$$1 = p_0 + q_0 p_1 + q_0 q_1 p_2 + \cdots,$$

that is, since $q_0 q_1 \cdots q_{n-1} p_n = q_0 q_1 \cdots q_{n-1} - q_0 q_1 \cdots q_n$,

$$\prod_{i=0}^{\infty} q_i = 0. \tag{3.5}$$

Since $q_i = 1 - p_i$ and $p_i \in (0, 1)$, the convergence criterion for infinite products (see Theorem 1.9 of the Appendix) tells that this is in turn equivalent to

$$\sum_{i=0}^{\infty} p_i = \infty.$$

The divergence of the series $\sum_{i=0}^{\infty} p_i$ is therefore a necessary and sufficient condition of existence of an invariant measure.

Recall, however, that the existence of an invariant measure does not imply recurrence (see Remark 2.1). But existence of a stationary distribution does imply recurrence (and actually positive recurrence) by the stationary distribution criterion.

Under condition (3.5), there exists an invariant measure, and this measure has finite mass ($\sum_{i=0}^{\infty} x_i < \infty$) if and only if

$$1 + \sum_{n=1}^{\infty} \left(\prod_{i=0}^{n-1} q_i\right) < \infty. \tag{3.6}$$

The stationary distribution is then given by

$$\pi(0) = \left(1 + \sum_{n=1}^{\infty} \left(\prod_{i=0}^{n-1} q_i\right)\right)^{-1} \tag{3.7}$$

and for $i \geq 1$,

$$\pi(i) = \left(\prod_{j=0}^{i-1} q_j\right) \pi_0. \tag{3.8}$$

In Problem 3.1.2, the reader is invited to verify that the success-runs chain of Example 1.1, and the machine-replacement chain of Example 1.1, Chapter 2, are the same if one sets

$$p_i = \frac{P(U = i + 1)}{P(U > i)}.$$

Inequality (3.6) then reads $E[U] < \infty$, and (3.7) and (3.8) give

$$\pi(i) = \frac{P(U > i)}{E[U]}. \tag{3.9}$$

◇

The stationary distribution criterion can also be used to prove instability.

Example 3.3. *Instability of ALOHA*
A typical situation in a multiple-access satellite communications system is the following. Users—each one identified with a message—contend for access to a single-channel satellite communications link for the purpose of transmitting messages. Two or more messages in the air at the same time jam each other, and are not successfully transmitted. The users are somehow able to detect a collision of this sort and will try to retransmit later the message involved in a collision. The difficulty in such communications systems resides mainly in the absence of cooperation among users, who are all unaware of the intention to transmit of competing users.

The *slotted* ALOHA protocol imposes on the users the following rules (see Fig. 3.3.2):

(i) Transmissions and retransmissions of messages can start only at equally spaced moments; the interval between two consecutive (re-)transmission times is called a *slot*; the duration of a slot is always larger than that of any message.

(ii) All *backlogged* messages, i.e., those messages having already tried unsuccessfully—maybe more than once—to get through the link, require retransmission independently of one another with probability $\nu \in (0, 1)$ at each slot. This is the so-called *Bernoulli retransmission policy*.

(iii) The *fresh messages*—those presenting themselves for the first time—immediately attempt to get through.

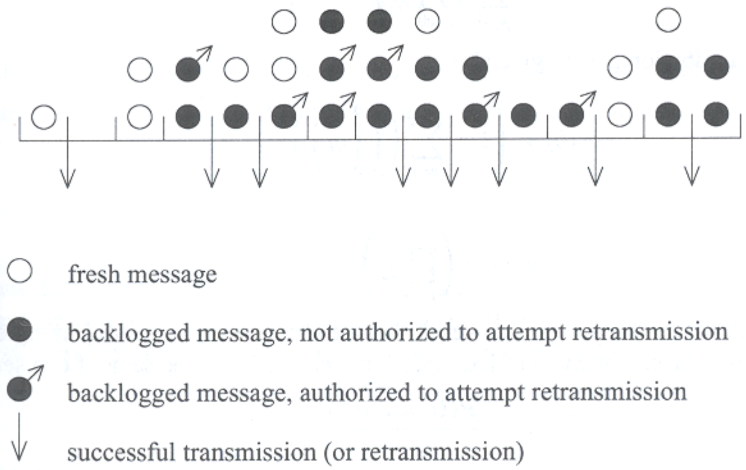

Figure 3.3.2. The ALOHA protocol

Let X_n be the number of backlogged messages at the beginning of slot n. The backlogged messages behave independently, and each one has probability ν of attempting retransmission

in slot n. In particular, if there are $X_n = k$ backlogged messages, the probability that i among them attempt to retransmit in slot n is

$$b_i(k) = \binom{k}{i} \nu^i (1-\nu)^{k-i}. \tag{3.10}$$

Let A_n be the number of fresh requests for transmission in slot n. The sequence $\{A_n\}_{n\geq 0}$ is assumed i.i.d with the distribution

$$P(A_n = j) = a_j. \tag{3.11}$$

The quantity

$$\lambda = E[A_n] = \sum_{i=1}^{\infty} i a_i \tag{3.12}$$

is the *traffic intensity*. We suppose that $a_0 + a_1 \in (0, 1)$, so that $\{X_n\}_{n\geq 0}$ is an irreducible HMC. Its transition matrix is

$$p_{ij} = \begin{cases} b_1(i)a_0 & \text{if } j = i-1, \\ [1 - b_1(i)]a_0 + b_0(i)a_1 & \text{if } j = i, \\ [1 - b_0(i)]a_1 & \text{if } j = i+1, \\ a_{j-i} & \text{if } j \geq i+2. \end{cases} \tag{3.13}$$

The proof of (3.13) is by accounting. For instance, the first line corresponds to one among the i backlogged messages having succeeded to retransmit, and for this there should be no fresh arrival (probability a_0) and only one of the i backlogged messages allowed to retransmit (probability $b_1(i)$). The second line corresponds to one of the two events "no fresh arrival and zero or strictly more than two retransmission requests from the backlog" and "zero retransmission request from the backlog and one fresh arrival."

Our objective in this example is to show that the system using the Bernoulli retransmission policy is *not stable*, in the sense that the chain $\{X_n\}_{n\geq 0}$ is *not positive recurrent*. Later on, in Example 4.2, a remedy to this situation will be proposed. To prove unstability, we must, in view of Theorem 3.1, contradict the existence of a stationary distribution π.

If such a stationary distribution existed, it should satisfy the balance equations

$$\pi(i) = \pi(i)\{[1 - b_1(i)]a_0 + b_0(i)a_1\} + \pi(i-1)[1 - b_0(i-1)]a_1 + \pi(i+1)b_1(i+1)a_0 + \sum_{\ell=2}^{\infty} \pi(i-\ell)a_\ell$$

where $\pi(j) = 0$ if $j < 0$. Writing

$$P_N = \sum_{i=0}^{N} \pi(i)$$

and summing up the balance equations from $i = 0$ to N, we obtain

$$P_N = \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + \sum_{\ell=0}^{N} a_\ell P_{N-\ell}.$$

This in turn gives

$$P_N(1-a_0) = \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + \sum_{\ell=1}^{N} a_\ell P_{N-\ell}.$$

But since P_N increases with N and $\sum_{\ell=1}^{N} a_\ell \le \sum_{\ell=1}^{\infty} a_\ell = 1 - a_0$, we have

$$\sum_{\ell=1}^{N} a_\ell P_{N-\ell} \le P_{N-1}(1-a_0),$$

and therefore

$$P_N(1-a_0) \le \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + P_{N-1}(1-a_0),$$

from which it follows that

$$\frac{\pi(N+1)}{\pi(N)} \ge \frac{1 - a_0 - b_0(N)a_1}{b_1(N+1)a_0}.$$

Using expression (3.10), we obtain

$$\frac{\pi(N+1)}{\pi(N)} \ge \frac{(1-a_0) - (1-\nu)^N a_1}{(N+1)\nu(1-\nu)^N a_0}.$$

For all values of $\nu \in (0, 1)$, the right-hand side of this inequality eventually becomes infinite, and this contradicts the equality $\sum_{N=1}^{\infty} \pi(N) = 1$ and the inequalities $\pi(N) > 0$ that π should satisfy as the stationary distribution of an irreducible Markov chain. ◇

4 Empirical Averages

4.1 Ergodic Theorem

This subsection is devoted to the ergodic theorem for Markov chains. It gives conditions which guarantee that empirical averages of the type

$$\frac{1}{N}\sum_{k=1}^{N} f(X_k, \ldots, X_{k+L})$$

converge to probabilistic averages.

Proposition 4.1.
Let $\{X_n\}_{n\ge 0}$ be an irreducible recurrent HMC, and let x denote the canonical invariant measure associated with state $0 \in E$,

$$x_i = E_0\left[\sum_{n\ge 1} 1_{\{X_n = i\}} 1_{\{n \le T_0\}}\right], \tag{4.1}$$

where T_0 is the return time to 0. Define for $n \ge 1$

$$\nu(n) = \sum_{k=1}^{n} 1_{\{X_k = 0\}}. \tag{4.2}$$

Let $f : E \to \mathbb{R}$ be such that

$$\sum_{i\in E} |f(i)| x_i < \infty. \tag{4.3}$$

Then, for any initial distribution μ, P_μ-a.s.,

$$\lim_{N\uparrow\infty} \frac{1}{\nu(N)} \sum_{k=1}^{N} f(X_k) = \sum_{i\in E} f(i) x_i. \tag{4.4}$$

◇

Before the proof, we shall harvest the most interesting consequences.

Theorem 4.1. *Ergodic Theorem*
Let $\{X_n\}_{n\ge 0}$ be an irreducible positive recurrent Markov chain with the stationary distribution π, and let $f : E \to \mathbb{R}$ be such that

$$\sum_{i\in E} |f(i)| \pi(i) < \infty. \tag{4.5}$$

Then for any initial distribution μ, P_μ-a.s.,

$$\lim_{n\uparrow\infty} \frac{1}{N} \sum_{k=1}^{N} f(X_k) = \sum_{i\in E} f(i)\pi(i). \tag{4.6}$$

◇

Proof. Apply Proposition 4.1 to $f \equiv 1$. Condition (4.3) is satisfied, since in the positive recurrent case, $\sum_{i\in E} x_i < \infty$. Therefore, P_μ-a.s.,

$$\lim_{N\uparrow\infty} \frac{N}{\nu(N)} = \sum_{j\in E} x_j.$$

Now, f satisfying (4.5) also satisfies (4.3), since x and π are proportional, and therefore, P_μ-a.s.,

$$\lim_{N\uparrow\infty} \frac{1}{\nu(N)} \sum_{k=1}^{N} f(X_k) = \sum_{i\in E} f(i)x_i .$$

Combination of the above equalities gives, P_μ-a.s.,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{k=1}^{N} f(X_k) = \lim_{N\to\infty} \frac{\nu(N)}{N} \frac{1}{\nu(N)} \sum_{k=1}^{N} f(X_k) = \frac{\sum_{i\in E} f(i)x_i}{\sum_{j\in E} x_j},$$

from which (4.6) follows, since π is obtained by normalization of x. □

Corollary 4.1.
Let $\{X_n\}_{n\geq 1}$ be an irreducible positive recurrent Markov chain with the stationary distribution π, and let $g : E^{L+1} \to \mathbb{R}$ be such that

$$\sum_{i_0,\ldots,i_L} |g(i_0,\ldots,i_L)|\pi(i_0)p_{i_0i_1}\cdots p_{i_{L-1}i_L} < \infty. \tag{4.7}$$

Then for all initial distributions μ, P_μ-a.s.

$$\lim \frac{1}{N} \sum_{k=1}^{N} g(X_k, X_{k+1},\ldots,X_{k+L}) = \sum_{i_0,i_1,\ldots,i_L} g(i_0,i_1,\ldots,i_L)\pi(i_0)p_{i_0i_1}\cdots p_{i_{L-1}i_L}. \tag{4.8}$$

◇

Proof. Apply Theorem 4.1 to the snake chain $\{(X_n, X_{n+1},\ldots,X_{n+L})\}_{n\geq 0}$ (see Problems 2.2.4, 2.4.6, and 2.5.2), which is irreducible recurrent and admits the stationary distribution

$$\pi(i_0)p_{i_0i_1}\cdots p_{i_{L-1}i_L}.$$

□

Note that

$$\sum_{i_0,i_1,\ldots,i_L} g(i_0,i_1,\ldots,i_L)\pi(i_0)p_{i_0i_1}\cdots p_{i_{L-1}i_L} = E_\pi[g(X_0,\ldots,X_L)]$$

Proof. (*of Proposition 4.1.*) Let $T_0 = \tau_1, \tau_2, \tau_3, \ldots$ be the successive return times to state 0, and define

$$U_p = \sum_{n=\tau_p+1}^{\tau_{p+1}} f(X_n).$$

In view of the regenerative cycle theorem (Theorem 7.4 of Chapter 2), $\{U_p\}_{p\geq 1}$ is an i.i.d sequence. Moreover, assuming $f \geq 0$ and using the strong Markov property,

$$\begin{aligned} E[U_1] &= E_0\left[\sum_{n=1}^{T_0} f(X_n)\right] \\ &= E_0\left[\sum_{n=1}^{T_0}\sum_{i\in E} f(i)1_{\{X_n=i\}}\right] = \sum_{i\in E} f(i)E_0\left[\sum_{n=1}^{T_0} 1_{\{X_n=i\}}\right] \\ &= \sum_{i\in E} f(i)x_i . \end{aligned}$$

By hypothesis, this quantity is finite, and threfore the strong law of large numbers applies, to give

$$\lim_{n\uparrow\infty} \frac{1}{n} \sum_{p=1}^{n} U_p = \sum_{i\in E} f(i)x_i ,$$

that is,

$$\lim_{n\uparrow\infty} \frac{1}{n} \sum_{k=T_0+1}^{\tau_{n+1}} f(X_k) = \sum_{i\in E} f(i)x_i . \tag{4.9}$$

Observing that

$$\tau_{\nu(n)} \leq n < \tau_{\nu(n)+1},$$

we have

$$\frac{\sum_{k=1}^{\tau_{\nu(n)}} f(X_k)}{\nu(n)} \leq \frac{\sum_{k=1}^{n} f(X_k)}{\nu(n)} \leq \frac{\sum_{k=1}^{\tau_{\nu(n)+1}} f(X_i)}{\nu(n)}.$$

Since the chain is recurrent, $\lim_{n\uparrow\infty} \nu(n) = \infty$, and therefore, from (4.9), the extreme terms of the above chain of inequality tend to $\sum_{i\in E} f(i)x_i$ as n goes to ∞, and this implies (4.4). The case of a function f of arbitrary sign is obtained by considering (4.4) written separately for $f^+ = \max(0, f)$ and $f^- = \max(0, -f)$, and then taking the difference of the two equalities obtained this way. The difference is not an undetermined form $\infty - \infty$ due to hypothesis (4.3). □

The version of the ergodic theorem for Markov chains featured in Theorem 4.1 is a kind of strong law of large numbers, and it can be used in simulations to compute, when π is unknown, quantities of the type $E_\pi[f(X_0)]$.

4.2 Examples

Example 4.1. *Fixed-Age Retirement*
We adopt the machine replacement interpretation of the success-runs Markov chain (Example 1.1). Assume positive recurrence. A visit of the chain to state 0 corresponds to a breakdown of a machine, and therefore, in view of the ergodic theorem,

$$\pi(0) = \lim_{N\uparrow\infty} \frac{1}{N} \sum_{k=1}^{N} 1_{\{X_k=0\}}$$

is the empirical frequency of breakdowns. Recall that

$$\pi(0) = E_0[T_0]^{-1},$$

where T_0 is the return time to 0. Here,

$$E_0[T_0] = E[U],$$

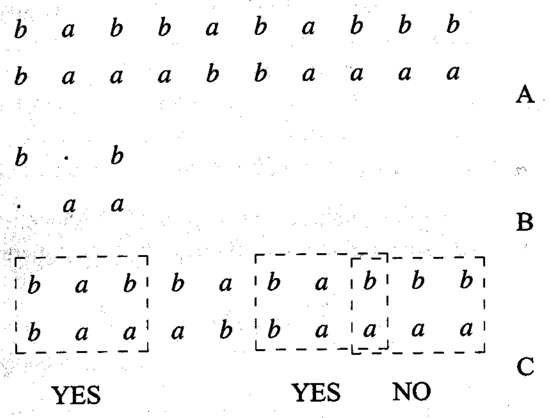

CHAPTER 4

Long Run Behavior

1 Coupling

1.1 Convergence in Variation

Consider an HMC that is irreducible and positive recurrent. In particular, if its initial distribution is the stationary distribution, it keeps the same distribution at all times. The chain is then said to be in the *stationary regime*, or in *equilibrium*, or in *steady state*.

A question arises naturally: What is the long-run behavior of the chain when the initial distribution μ is *arbitrary*? For instance, will it *converge to equilibrium*, and in which sense?

When the HMC is reducible, another type of problem is of interest. Suppose, for instance, that the set of transient states is not empty and that each remaining state is absorbing. One may want to compute the probability of reaching a given absorbing state when the initial state is transient, or the probability of remaining forever in the transient set. In this special case, where all recurrent states are absorbing, the probability of leaving the transient set is exactly the probability of converging. We are dealing here with almost-sure convergence.

For an ergodic HMC, the type of convergence of interest is not almost-sure convergence but convergence in variation of the distribution at time n to the stationary distribution. This type of convergence is relative to a metric structure that we proceed to define.

Definition 1.1. *Distance in Variation*
Let E be a countable space and let α and β be probability distributions on E. The *distance in variation* $d_V(\alpha, \beta)$ between α and β is defined by

$$d_V(\alpha, \beta) = \frac{1}{2}|\alpha - \beta| = \frac{1}{2}\sum_{i \in E}|\alpha(i) - \beta(i)|. \tag{1.1}$$

The distance in variation between two random variables X and Y with values in E and respective distributions $\mathcal{L}(X)$ and $\mathcal{L}(Y)$ is $d_V(\mathcal{L}(X), \mathcal{L}(Y))$, and it is denoted with a slight abuse of notation by $d_V(X, Y)$.

That d_V is indeed a distance is clear.

Lemma 1.1.
Let X and Y be two random variables with values in the same countable space E. Then

$$\sup_{A \subset E} |P(X \in A) - P(Y \in A)| = \sup_{A \subset E} \{P(X \in A) - P(Y \in A)\} = d_V(X, Y). \tag{1.2}$$

Proof. For the first equality observe that for each A there is a B such that $|P(X \in A) - P(Y \in A)| = P(X \in B) - P(Y \in B)$ (take $B = A$ or $\bar{A}$). For the second equality, write

$$P(X \in A) - P(Y \in A) = \sum_{i \in E} 1_A(i)\{P(X = i) - P(Y = i)\}$$

and observe that the right-hand side is maximal for

$$A = \{i \in E;\ P(X = i) > P(Y = i)\}$$

Also, for any $A \subset E$,

$$\sum_{i \in E} 1_A(i)\{P(X = i) - P(Y = i)\} + \sum_{i \in E} 1_{\bar{A}}(i)\{P(X = i) - P(Y = i)\} = 0$$

because $\sum_{i \in E}\{P(X = i) - P(Y = i)\} = 0$. For the specific set A above, $P(X = i) - P(Y = i)$ equals $|P(X = i) - P(Y = i)|$ on A, and equals $-|P(X = i) - P(Y = i)|$ on $\bar{A}$. Therefore, for this particular set A,

$$\begin{aligned}\sum_{i \in E} 1_A(i)\{P(X = i) - P(Y = i)\} &= \sum_{i \in E} 1_A(i)|P(X = i) - P(Y = i)| \\ &= \sum_{i \in E} 1_{\bar{A}}(i)|P(X = i) - P(Y = i)| \\ &= \frac{1}{2}\sum_{i \in E} |P(X = i) - P(Y = i)|. \end{aligned}$$

□

For two probability distributions α and β on the countable set E, let $\mathcal{D}(\alpha, \beta)$ be the collection of random vectors (X, Y) taking their values in $E \times E$, and with marginal distributions α and β, that is,

$$\alpha = \mathcal{L}(X)\ ,\ \beta = \mathcal{L}(Y). \tag{1.3}$$

Theorem 1.1. *Maximal Coincidence*
For any $(X, Y) \in \mathcal{D}(\alpha, \beta)$,

$$P(X = Y) \le 1 - d_V(\alpha, \beta), \tag{1.4}$$

and equality in (1.4) is attained by some pair $(X, Y) \in \mathcal{D}(\alpha, \beta)$, which is then said to *realize maximal coincidence*.

Proof. For arbitrary $A \subset E$,

$$\begin{aligned} P(X \ne Y) &\ge P(X \in A, Y \in \bar{A}) = P(X \in A) - P(X \in A, Y \in A) \\ &\ge P(X \in A) - P(Y \in A), \end{aligned}$$

and therefore

$$P(X \ne Y) \ge \sup_{A \subset E}\{P(X \in A) - P(Y \in A)\} = d_V(\alpha, \beta).$$

To finish the proof, it suffices to construct $(X, Y) \in \mathcal{D}(\alpha, \beta)$ realizing equality. We shall need the following observations (Problem 4.1.1):

$$\frac{1}{2}|\alpha - \beta| = \sum_{i \in E}(\alpha(i) - \beta(i))^+ = \sum_{i \in E}(\beta(i) - \alpha(i))^+ = 1 - \sum_{i \in E}\min(\alpha(i), \beta(i)). \tag{1.5}$$

Let U, Z, V, and W be independent random variables; U takes its values in $\{0, 1\}$, and Z, V, W take their values in E. The distributions of these random variables is given by

$$\begin{aligned} P(U = 1) &= 1 - d_V(\alpha, \beta), \\ P(Z = i) &= \min(\alpha(i), \beta(i))/(1 - d_V(\alpha, \beta)), \\ P(V = i) &= (\alpha(i) - \beta(i))^+/d_V(\alpha, \beta), \\ P(W = i) &= (\beta(i) - \alpha(i))^+/d_V(\alpha, \beta). \end{aligned}$$

Defining

$$X = UZ + (1 - U)V,\ Y = UZ + (1 - U)W,$$

we have

$$\begin{aligned} P(X = i) &= P(U = 1, Z = i) + P(U = 0, V = i) \\ &= P(U = 1)P(Z = i) + P(U = 0)P(V = i) \\ &= \min(\alpha(i), \beta(i)) + (\alpha(i) - \beta(i))^+ = \alpha(i), \end{aligned}$$

and similarly, $P(Y = i) = \beta(i)$. Therefore, $(X, Y) \in \mathcal{D}(\alpha, \beta)$. Also, $P(X = Y) = P(U = 1) = 1 - d_V(\alpha, \beta)$. □

Example 1.1.
One seeks a pair of $\{0, 1\}$-valued random variables with prescribed marginals

$$P(X = 1) = a\ ,\ P(Y = 1) = b,$$

where $a, b \in (0, 1)$, and such that $P(X = Y)$ is maximal. In the notation of the above theory,

$$\alpha = (1 - a, a)\ ,\ \beta = (1 - b, b),$$

and therefore

$$d_V(\alpha, \beta) = |a - b|.$$

Suppose for definiteness that $a \geq b$. The random U, Z, V, W of the construction of Theorem 1.1 have the following distributions.

$$\begin{aligned} P(U=1) &= 1-a+b, \\ P(Z=1) &= \frac{b}{1-a+b}, \quad P(Z=0) = \frac{1-a}{1-a+b}, \\ V &= 1, \\ W &= 0. \end{aligned}$$

Here $X = UZ + 1 - U$, $Y = UZ$. ◇

Definition 1.2. *Convergence in Variation*
Let $\{\alpha_n\}_{n\geq 0}$ and β be probability distributions on a countable space E. If $\lim_{n\uparrow\infty} d_V(\alpha_n, \beta) = 0$, the sequence $\{\alpha_n\}_{n\geq 0}$ is said to *converge in variation* to the probability distribution β.

Let $\{X_n\}_{n\geq 0}$ be an E-valued stochastic process. If for some probability distribution π on E, the distribution $\mathcal{L}(X_n)$ of the random variable X_n converges in variation to π, i.e., if

$$\lim_{n\uparrow\infty} \sum_{i\in E} |P(X_n = i) - \pi(i)| = 0, \tag{1.6}$$

then $\{X_n\}_{n\geq 0}$ is said to *converge in variation* to π.

There is some abuse of terminology in the above definition (it is the state random variable, not the process, that converges in variation). However, in this book, such abuse turns out to be harmless and very convenient.

If the process $\{X_n\}_{n\geq 0}$ converges in variation to π, then

$$\lim_{n\uparrow\infty} E[f(X_n)] = \pi(f) \tag{1.7}$$

for all bounded functions $f : E \to R$, where

$$\pi(f) = \sum_{i\in E} \pi(i) f(i). \tag{1.8}$$

Indeed, if M is an upper bound of $|f|$, then

$$|E[f(X_n)] - \pi(f)| = \left|\sum_{i\in E} f(i)(P(X_n = i) - \pi(i))\right| \leq M \sum_{i\in E} |P(X_n = i) - \pi(i)|.$$

1.2 The Coupling Method

Coupling is an old idea of Doeblin (1938), revived in Markov-chain theory by the influential work of Griffeath (1975) and Pitman (1974), and brought to fame by Lindvall (1977) who gave a purely probabilistic proof of the renewal theorem (section 3 of the present chapter is devoted to discrete-time renewal theory). The coupling method has a wide range of applications and the reader is directed to the book (Lindvall, 1992) for additional information and historical comments.

Observe that Definition 1.2 concerns only the marginal distributions of the process, not the process itself. Therefore, if there exists another process $\{X'_n\}_{n\geq 0}$ with $\mathcal{L}(X_n) = \mathcal{L}(X'_n)$ for all $n \geq 0$, and if there exists a third process $\{X''_n\}_{n\geq 0}$ such that $\mathcal{L}(X''_n) = \pi$ for all $n \geq 0$, then (1.6) follows from

$$\lim_{n\uparrow\infty} d_V(X'_n, X''_n) = 0. \tag{1.9}$$

This trivial observation is useful because of the resulting freedom in the choice of $\{X'_n\}$ and $\{X''_n\}$. In particular, one can use dependent versions, and the most interesting case occurs when there exists a finite random time τ such that $X'_n = X''_n$ for all $n \geq \tau$. It follows then, as will be proven in Theorem 1.2, that

$$d_V(X'_n, X''_n) \leq P(\tau > n). \tag{1.10}$$

Finiteness of τ is equivalent to $\lim_{n\uparrow\infty} P(\tau > n) = 0$, and therefore (1.9) is a consequence of (1.10).

The above program can be carried out in a variety of situations, and most notably for ergodic Markov chains. For the time being, the preliminary definitions and results relative to the *coupling method* for proving convergence in variation of ergodic chains will be collected.

Definition 1.3. *Coupling*
Two stochastic processes $\{X'_n\}_{n\geq 0}$ and $\{X''_n\}_{n\geq 0}$ taking their values in the same countable state space E are said to *couple* if there exists an almost surely *finite* random time τ such that

$$n \geq \tau \Rightarrow X'_n = X''_n. \tag{1.11}$$

The random variable τ is called a *coupling time* of the two processes.

Theorem 1.2. *Coupling Inequality*
Inequality (1.10) holds for any coupling time τ of $\{X'_n\}_{n\geq 0}$ and $\{X''_n\}_{n\geq 0}$.

Proof. For all $A \subset E$,

$$\begin{aligned} P(X'_n \in A) - P(X''_n \in A) &= P(X'_n \in A,\ \tau \leq n) + P(X'_n \in A,\ \tau > n) \\ &\quad - P(X''_n \in A,\ \tau \leq n) - P(X''_n \in A, \tau > n) \\ &= P(X'_n \in A,\ \tau > n) - P(X''_n \in A,\ \tau > n) \\ &\leq P(X'_n \in A,\ \tau > n) \\ &\leq P(\tau > n). \end{aligned}$$

Inequality (1.10) then follows from Lemma 1.1. □

The framework of the coupling method is now in place. It remains to construct $\{X'_n\}_{n\geq 0}$ and $\{X''_n\}_{n\geq 0}$ that couple and that mimic $\{X_n\}_{n\geq 0}$ and π in the sense that $\mathcal{L}(X'_n) = \mathcal{L}(X_n)$ and $\mathcal{L}(X'_n) = \pi$ for all $n \geq 0$.

2 Convergence to Steady State

2.1 Positive Recurrent Case

The main result concerns ergodic (i.e., irreducible positive recurrent, and *aperiodic*) HMCs.

Theorem 2.1. *Convergence to Steady State*
Let $\mathbf{P}$ be an ergodic transition matrix on the countable state space E. For all probability distributions μ and ν on E,

$$\lim_{n\uparrow\infty} d_V(\mu^T\mathbf{P}^n, \nu^T\mathbf{P}^n) = 0. \tag{2.1}$$

In particular, if ν is the stationary distribution π,

$$\lim_{n\uparrow\infty} |\mu^T\mathbf{P}^n - \pi^T| = 0,$$

and with $\mu = \delta_j$, the probability distribution putting all its mass on j,

$$\lim_{n\uparrow\infty} \sum_{i\in E} |p_{ji}(n) - \pi(i)| = 0.$$

From the discussion preceding Definition 1.3, it suffices to construct two coupling chains with initial distributions μ and ν, respectively.

Theorem 2.2. *Independent Coupling*
Let $\{X_n^{(1)}\}_{n\geq 0}$ and $\{X_n^{(2)}\}_{n\geq 0}$ be two *independent* ergodic HMCs with the same transition matrix $\mathbf{P}$ and initial distributions μ and ν, respectively. Let $\tau = \inf\{n \geq 0;\ X_n^{(1)} = X_n^{(2)}\}$, with $\tau = \infty$ if the chains never intersect. Then τ is, in fact, almost surely finite. Moreover, the process $\{X'_n\}_{n\geq 0}$ defined by

$$X'_n = \begin{cases} X_n^{(1)} & \text{if } n \leq \tau, \\ X_n^{(2)} & \text{if } n \geq \tau \end{cases} \tag{2.2}$$

is an HMC with transition matrix $\mathbf{P}$. (See Fig. 4.2.1.)

Proof. Consider the product HMC $\{Z_n\}_{n\geq 0}$ defined by $Z_n = (X_n^{(1)}, X_n^{(2)})$. It takes values in $E \times E$, and the probability of transition from (i,k) to (j,ℓ) in n steps is $p_{ij}(n)p_{k\ell}(n)$. This chain is irreducible. Indeed, since $\mathbf{P}$ is irreducible and *aperiodic*, by Theorem 4.3 of Chapter 2, there exists m such that for all pairs (i,j) and (k,ℓ), $n \geq m$ implies $p_{ij}(n)p_{k\ell}(n) > 0$. This implies that the period of the product chain is 1, again by Theorem 4.3 of Chapter 2.

Clearly, $\{\pi(i)\pi(j)\}_{(i,j)\in E^2}$ is a stationary distribution for the product chain, where π is the stationary distribution of $\mathbf{P}$. Therefore, by the stationary distribution criterion, the product chain is positive recurrent. In particular, it reaches the diagonal of E^2 in finite time, and consequently, $P(\tau < \infty) = 1$.

It remains to show that $\{X'_n\}_{n\geq 0}$ given by (2.2) is an HMC with transition matrix $\mathbf{P}$. This is a consequence of the strong Markov property applied to the product chain. The details are left for the reader (Problem 4.2.1). □

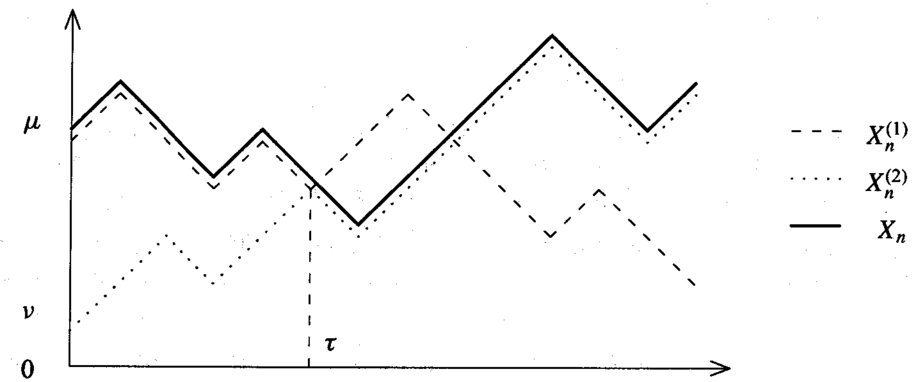

Figure 4.2.1. Independent coupling

Theorem 2.1 concerns ergodic chains, and aperiodicity is needed there to guarantee that the product chain is irreducible (see Problem 4.2.2). For periodic chains, the situation is different, but the result follows directly from the ergodic case.

Theorem 2.3. *Periodic Case*
Let $\mathbf{P}$ be an irreducible positive recurrent transition matrix on the countable space E, with period d. Let π be its stationary distribution. If μ is a probability distribution such that $\mu(C_0) = 1$ for some cyclic class C_0, then

$$\lim_{n\uparrow\infty} \sum_{i\in C_0} |(\mu^T\mathbf{P}^{nd})_i - d\pi(i)| = 0. \tag{2.3}$$

Proof. Consider the restriction of $\mathbf{P}^d$ to C_0, which is irreducible and aperiodic (see Problem 2.4.7). It is positive recurrent, since it has an invariant measure with finite mass, namely π restricted to C_0. It remains to show that $d\pi$ restricted to C_0 is a probability distribution, that is, $\pi(C_0) = 1/d$. By the ergodic theorem,

$$\lim_{N\uparrow\infty} \frac{1}{N} \sum_{n=1}^{N} 1_{\{X_n\in C_0\}} = \pi(C_0),$$

and since $X_n \in C_0$ once every d steps, the left-hand side equals $1/d$. □

2.2 Null Recurrent Case

The last two theorems concern the positive recurrent case. The null recurrent case is a little more difficult.

Theorem 2.4. *Orey's Theorem*
Let $\mathbf{P}$ be an irreducible *null* recurrent transition matrix on E. Then for all $i, j \in E$,

$$\lim_{n\uparrow\infty} p_{ij}(n) = 0. \tag{2.4}$$

Proof. The periodic case follows from the aperiodic case by considering the restriction of $\mathbf{P}^d$ to C_0, an arbitrary cyclic class, and observing that this restriction is also null recurrent. Therefore, $\mathbf{P}$ will be assumed aperiodic.

In this case the product HMC $\{Z_n\}_{n\geq 0} = \{X_n^{(1)}, X_n^{(2)}\}_{n\geq 0}$ defined in the proof of Theorem 2.2 is irreducible and aperiodic. However, it cannot be argued that it is recurrent, although its components are recurrent. One must therefore separate the two possible cases.

If the product chain is transient, its n-step transition probability from (i, i) to (j, j) is $[p_{ij}(n)]^2$, and in view of the discussion of Section 1.2 of Chapter 3, it tends to 0 as $n \to \infty$, and the result is proven.

If the product chain is recurrent, the coupling argument used in the proof of Theorem 2.1 applies and yields

$$\lim_{n\uparrow\infty} |\mu^T \mathbf{P}^n - \nu^T \mathbf{P}^n| = 0 \tag{2.5}$$

for arbitrary initial distributions μ and ν. Suppose now that for some $i, j \in E$, (2.4) is not true. One can then find a sequence $\{n_k\}_{k\geq 0}$ of integers strictly increasing to ∞ such that

$$\lim_{k\uparrow\infty} p_{ij}(n_k) = \alpha > 0.$$

For fixed $i \in E$ chosen as above, the sequence $(\{p_{is}(n_k), s \in E\})_{k\geq 0}$ of vectors of $[0, 1]^E$ is compact in the topology of pointwise convergence. Therefore (see, however, Theorem 1.10 of the Appendix for an elementary proof), there exists a subsequence $\{m_\ell\}_{\ell\geq 0}$ of integers strictly increasing to ∞ and a vector $\{x_s, s \in E\} \in [0, 1]^E$ such that for all $s \in E$,

$$\lim_{\ell\uparrow\infty} p_{is}(m_\ell) = x_s.$$

Now, $x_j = \alpha > 0$, and therefore $\{x_s, s \in E\}$ is nontrivial. Since $\sum_{s\in E} p_{is}(m_\ell) = 1$, it follows from Fatou's lemma (see Theorem 1.8 of the Appendix), that

$$\sum_{s\in E} x_s \leq 1.$$

In (2.5), take $\mu = \delta_i$ (all the mass is on i) and $\nu^T = \delta_i{}^T \mathbf{P}$ to obtain, for all $s \in E$,

$$\lim_{\ell\uparrow\infty} |p_{is}(m_\ell) - p_{is}(m_\ell + 1)| = 0. \tag{2.6}$$

But $p_{is}(m_\ell + 1) = \sum_{k\in E} p_{ik}(m_\ell) p_{ks}$. Therefore, by dominated convergence in (2.6), we obtain for all $s \in E$,

$$x_s = \sum_{k\in E} x_k p_{ks}.$$

In other words, $\{x_s, s \in E\}$ is an invariant measure of $\mathbf{P}$ with finite mass, which implies that $\mathbf{P}$ is positive recurrent, a contradiction. Therefore, (2.4) cannot be contradicted. □

2.3 Thermodynamic Irreversibility

Zermelo's Refutation

We shall now take a small pause, and go back to the Ehrenfest diffusion model. Its fame is due to the insight it gives to the once controversial issue of thermodynamic irreversibility. It is the right time to try to understand this, because we have the notion of recurrence and the theorem of convergence to steady state of ergodic chains.

According to the macroscopic theory of thermodynamics, systems progress in an orderly and irreversible manner towards equilibrium. Consider, for instance, a system of N particles in a box divided in two similar compartments A and B by a fictive membrane. If at the origin of time, all particles are in A, they will rather quickly reorganize themselves, and they will settle to equilibrium, a macroscopic state in which the contents of A and B are thermodynamically equivalent.

Boltzmann claimed that there was an arrow of time in the direction of increasing entropy, and indeed, in the diffusion experiment, equality between the thermodynamic quantities in both compartments corresponds to maximal entropy.

Zermelo, who obviously was not sleeping in the back of the classroom, argued that in view of the time reversibility of the laws of physics, the Boltzmann theory should at least be discussed. Zermelo held a strong position in this controversy. Indeed, there is a famous result of mechanics, Poincaré's recurrence theorem, which implies that in the situation where at time 0 all molecules are in A, then whatever the time T, there will be a subsequent time $t > T$ at which all the molecules will again gather in A. This phenomenon predicted by irrefutable mathematics is, of course, never observed in daily life, where it would imply that the chunk of sugar that one patiently dissolves in one's cup of coffee could escape ingestion by reforming itself at the bottom of the cup.

Convergence vs. Recurrence

Boltzmann's theory was really hurt by this striking and seemingly inescapable argument. Things had to be clarified. Fortunately, Tatiana and Paul Ehrenfest came up with their Markov chain model, and in a sense saved the edifice that Boltzmann had constructed.

Without going into the details, let us just say that the Ehrenfest model is an approximation of the real diffusion phenomenon that is *philosophically correct* from the point of view of statistical mechanics. Also, at first sight, it is subject to Zermelo's attack, presenting both features that the latter found incompatible: an irreversible tendency towards equilibrium, and recurrence. Here the role of Poincaré's recurrence theorem is played by the Markov chain recurrence theorem, stating that an irreducible chain with a stationary distribution visits any fixed state, say 0, infinitely often. As for the irreversible tendency towards equilibrium, one has the theorem of convergence to steady state, according to which the distribution at time n converges to the stationary distribution whatever the initial distribution as n tends to infinity (*stricto sensu* this statement is not true, due to the periodicity of the chain. However, such periodicity is an artifact created by the discretization of time, and it would disappear in the continuous-time model, or in a slight modification of the discrete-time model). Thus, according to Markov-chain theory, convergence to statistical equilibrium and recurrence are not antagonistic, and we are here at the epicenter of Zermelo's refutation.

One can show that recurrence is *not observable* for states far from $L = \frac{N}{2}$, assuming that N is even. For instance, the average time to reach 0 from state L is

$$\frac{1}{2L}2^{2L}(1 + O(L)) \tag{2.7}$$

whereas the average time to reach state L from state 0 is less than

$$L + L\log L + O(1). \tag{2.8}$$

(See Chapter III, Section 5 of (Bhattacharya and Waymire, 1990) for a derivation of the above estimates.) With $L = 10^6$ and one unit of mathematical time equal to 10^{-5} second, the return time to equilibrium when compartment A is initially empty is on the order of a second, whereas it would take on the order of

$$\frac{1}{2 \cdot 10^{11}} \times 2^{2^{10^6}} \text{ seconds}$$

to go from L to empty, which is an astronomical time. These numbers teach us not to spend too much time stirring the coffee, or hurry to swallow it for fear of recrystallization of the chunk of sugar. From a mathematical point of view, being in the steady state at a given time does not prevent the chain from being in a rare state, only it is there rarely. The rarity of the state is equivalent to long recurrence times, so long that when there are more than a few particles in the boxes, it would take an astronomical time to witness the effects of Poincaré's recurrence theorem. Note that Boltzmann rightly argued that the recurrence times in Poincaré's theorem are extremely long, but his heuristic arguments failed to convince.

Here is another manifestation of thermodynamic irreversibility in the Ehrenfest model.

Newton's Law of Cooling

Let $g(z, n) = E[z^{X_n}]$ be the generating function of X_n. Using the basic rules of conditional expectation (see Section 7 of Chapter 1), we have

$$\begin{aligned} E[z^{X_{n+1}} \mid X_n] &= E[z^{X_{n+1}} 1_{\{X_{n+1}=X_n-1\}} \mid X_n] + E[z^{X_{n+1}} 1_{\{X_{n+1}=X_n+1\}} \mid X_n] \\ &= E[z^{X_n-1} 1_{\{X_{n+1}=X_n-1\}} \mid X_n] + E[z^{X_n+1} 1_{\{X_{n+1}=X_n+1\}} \mid X_n] \\ &= z^{X_n-1}P(X_{n+1} = X_n - 1 \mid X_n) + z^{X_n+1}P(X_{n+1} = X_n + 1 \mid X_n), \end{aligned}$$

and therefore, taking into account the dynamics of the Ehrenfest model,

$$\begin{aligned} E[z^{X_{n+1}} \mid X_n] &= z^{X_n-1}\frac{X_n}{N} + z^{X_n+1}(1 - \frac{X_n}{N}) \\ &= z \cdot z^{X_n} + \frac{1}{N}(1 - z^2)X_n z^{X_n-1}. \end{aligned}$$

Taking expectations gives

$$g(z, n+1) = zg(z, n) + \frac{1}{N}(1 - z^2)E[X_n z^{X_n-1}],$$

that is,

$$g(z, n+1) = zg(z, n) + \frac{1}{N}(1 - z^2)g'(z, n).$$

Differentiation of the above identity yields

$$g'(z, n+1) = g(z, n) + zg'(z, n) - \frac{2z}{N}g'(z, n) + \frac{1}{N}(1 - z^2)g''(z, n).$$

Letting $z = 1$, we obtain (see Section 5 of Chapter 1)

$$E[X_{n+1}] = 1 + E[X_n] - \frac{2}{N}E[X_n].$$

Supposing N even ($N = 2L$) and rearranging terms, we have

$$E[X_{n+1} - L] = (1 - \frac{1}{L})E[X_n - L],$$

and therefore

$$E[X_n - L] = E[X_0 - L]\left(1 - \frac{1}{L}\right)^n.$$

Supposing $P(X_0 = i) = 1$, we then have

$$E\left[\frac{X_n - L}{L}\right] = \left(\frac{i - L}{L}\right)\left(1 - \frac{1}{L}\right)^n.$$

In the kinetic theory of heat, $\frac{X_n - L}{L}$ is interpreted as the temperature difference between state X_n and the "equilibrum" state L. To account for a large number of particles, we let L tend to infinity and make i depend on L in such a way that

$$\lim_{L\to\infty} \frac{i(L) - L}{L} = \theta(0).$$

The number $\theta(0)$ is interpreted as the initial deviation from the equilibrium temperature in compartment A. Also, the discrete time unit 1 is now interpreted as Δ units of real time. If X_n is the state at real time t, we must then have

$$t = n\Delta.$$

Interpreting $\frac{1}{2L\Delta}$ as the average proportion of the total energy that passes through the membrane per unit of real time, we require that as L tends to infinity, this quantity remains constant, that is,

$$\frac{1}{L\Delta} = \gamma.$$

In particular, Δ tends to zero as L increases to infinity. Therefore, observing that $n = L\gamma t$,

$$\lim_{L\to\infty}\left(1-\frac{1}{L}\right)^n = \lim_{L\to\infty}\left[\left(1-\frac{1}{L}\right)^L\right]^{\gamma t} = e^{-\gamma t},$$

and finally, with $\theta(t) = \lim_{L\to\infty} E[\frac{X_n - L}{L}]$,

$$\theta(t) = \theta(0)e^{-\gamma t}.$$

We conclude this section with a classic reference on thermodynamic irreversibility as probabilists understand it, the article of M. Kac (1947), "Random Walk and the Theory of Brownian Motion", *American Mathematical Monthly*, *54*, 369-391.

2.4 Convergence Rates via Coupling

Knowing that an ergodic Markov chain converges to equilibrium, the next question is, How fast? The first result below is not explicit, but it can be used *in principle* for chains with an infinite number of states.

We recall at this point the meaning of the o (*small o*) symbol. It represents a function defined in a neighborhood of zero and such that $\lim_{t\to 0}\frac{|o(t)|}{t} = 0$.

Theorem 2.5. *Rate of Convergence*
Suppose that the coupling time τ in Theorem 1.2 satisfies

$$E[\psi(\tau)] < \infty \tag{2.9}$$

for some nondecreasing function $\psi : \mathbb{N} \to \mathbb{R}_+$ such that $\lim_{n\uparrow\infty}\psi(n) = \infty$. Then for any initial distributions μ and ν

$$|\mu^T\mathbf{P}^n - \nu^T\mathbf{P}^n| = o\left(\frac{1}{\psi(n)}\right). \tag{2.10}$$

Proof. Since ψ is nondecreasing, $\psi(\tau)1_{\{\tau>n\}} \geq \psi(n)1_{\{\tau>n\}}$, and therefore

$$\psi(n)P(\tau > n) \leq E[\psi(\tau)1_{\{\tau>n\}}].$$

Now,

$$\lim_{n\uparrow\infty} E[\psi(\tau)1_{\{\tau>n\}}] = 0$$

by dominated convergence (see Theorem 3.2 of the Appendix), since $\lim_{n\uparrow\infty}\psi(\tau)1_{\{\tau>n\}} = 0$, by the finiteness of τ, and $\psi(\tau)1_{\{\tau>n\}}$ is bounded by the integrable random variable $\psi(\tau)$. □

Time τ in Theorem 1.2 is the entrance time of the product chain in the diagonal set of $E \times E$. In principle, the distribution of τ can be explicitly computed. However the actual computations are usually difficult when the state space is infinite. The reader will find in (Lindvall, 1992) examples of application of Theorem 2.5 to infinite state space HMCs. When the state space is finite, convergence is exponential (geometric).

Theorem 2.6. *Exponential Convergence Rate of Finite HMCs*
Let $\mathbf{P}$ be an ergodic transition matrix on the *finite* state space E. Then for any initial distributions μ and ν, one can construct two HMCs $\{X_n\}_{n\geq 0}$ and $\{Y_n\}_{n\geq 0}$ on E with the same transition matrix $\mathbf{P}$, and the respective initial distributions μ and ν, in such a way that they couple at a finite time τ such that $E[e^{\alpha\tau}] < \infty$ for some positive α.

Proof. See Problem 4.1.4. □

In Chapter 6, we shall say more on convergence rates in the finite state space case. The reader will find in (Lindvall, 1992) examples of application of Theorem 2.5 for infinite state space HMCs.

3 Discrete-Time Renewal Theory

3.1 Renewal Equation

In the analytic approach to Markov chains, the proof of convergence to steady state of an ergodic HMC is a consequence of a result on power series called the *renewal theorem* by the probabilists. This result forms the matter of the present section. However, the renewal theorem will not be used as the essential step towards the convergence theorem, but on the contrary, it will be obtained as a corollary of the latter.

We start with the basic definitions. Let $\{S_n\}_{n\geq 1}$ be an i.i.d sequence of random variables with values in $\bar{\mathbb{N}} = \{1, 2, \ldots, +\infty\}$ and with the probability distribution

$$P(S_1 = k) = f_k. \tag{3.1}$$

Define for $n \geq 0$,

$$R_{n+1} = R_n + S_{n+1}, \tag{3.2}$$

where R_0 is an arbitrary random variable with values in $\mathbb{N}$ (in particular, $R_0 < \infty$).

Definition 3.1. *Renewal Sequence*
The sequence $\{R_n\}_{n\geq 0}$ is called a *delayed* (by R_0) *renewal sequence* with the *renewal distribution* $\{f_k\}_{k\geq 1}$. If $R_0 \equiv 0$, one speaks of an *undelayed* renewal sequence, or, more simply, of a renewal sequence. If $P(S_1 = \infty) = 0$, the renewal sequence (delayed or not) is called a *proper renewal sequence*, and $\{f_k\}_{k\geq 1}$ is called a *proper renewal distribution*. Otherwise, one speaks of a *defective renewal sequence* and of a *defective renewal distribution*.

The quantity

$$\alpha = P(S_1 = \infty)$$

is the *defect* of the renewal distribution. The random time R_k is the kth *renewal time*, and the sequence $\{S_n\}_{n\geq 1}$ is the *interrenewal sequence*.