Discrete Time Markov Chains

1 Examples

Discrete Time Markov Chain (DTMC) is an extremely pervasive probability model [1]. In this lecture we shall briefly overview the basic theoretical foundation of DTMC. Let us first look at a few examples which can be naturally modelled by a DTMC.

Example 1.1 (Gambler Ruin Problem)**.** A gambler has \$100. He bets \$1 each game, and wins with probability 1/2. He stops playing he gets broke or wins \$1000. Natural questions include: what's the probability that he gets broke? On average how many games are played? This problem is a special case of the so-called *Gambler Ruin* problem, which can be modelled using a *Markov chain* as follows. We will be a hand-wavy before rigorously defining what a DTMC is. Imagine we have 1001 "states" as shown in Figure 1.1, each state i is indexed by the number of dollars the gambler is having. Before playing each game, the gambler has an equal probability of move up to state $i+1$ or down to state i.

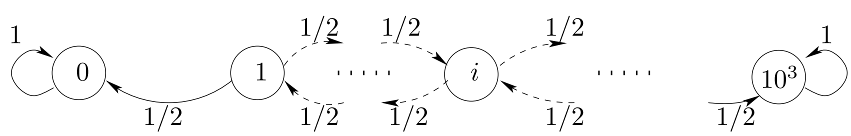

Figure 1: DTMC for the Gambler Ruin Problem

Example 1.2 (Single Server Queue)**.** At each time slot, an Internet router's buffer gets an additional packet with probability p, or releases one packet with probability q, or remains the same with probability r. Starting from an empty buffer, what is the distribution of the number of packets after n slots? As $n \to \infty$, will the buffer be overflowed? As $n \to \infty$, what's the typical buffer size? These are the types of questions that can be answered with DTMC analysis.

2 Basic definitions and properties

A *stochastic process* is a collection of random variables (on some probability space) indexed by some set T: $\{X_t, t \in T\}$. When $T \subseteq \mathbb{R}$, we can think of T as set of points in time, and X_t as the "state" of the process at time t. The *state space*, denoted by I, is the set of all possible values of the X_t. When T is countable we have a *discrete-time stochastic process*. When T is an interval of the real line we have a *continuous-time stochastic process*.

Example 2.1 (Bernoulli process)**.** A sequence $\{X_0, X_1, X_2, \dots\}$ of *independent* Bernoulli random variables with parameter p is called a Bernoulli process. It is not hard to compute the expectations and the variances of the following statistics related to the Bernoulli process.

$$\begin{aligned} S_n &= X_1 + \cdots + X_n \\ T_n &= \text{number of slots from the } (n-1)\text{th 1 to the } n\text{th 1} \\ Y_n &= T_1 + \cdots + T_n \end{aligned}$$

Most useful stochastic processes are not that simple, however. We often see processes whose variables are correlated, such as stock prices, exchange rates, signals (speech, audio and video), daily temperatures, Brownian motion or random walks, etc.

A *discrete time Markov chain* (DTMC) is a discrete-time stochastic process $\{X_n\}_{n\geq 0}$ satisfying the following:

- the state space I is countable (often labeled with a subset of $\mathbb{N}$).
- For all states i, j there is a given probability p_{ij} such that

  $$P\big[X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots X_0 = i_0\big] = p_{ij},$$

  for all $i_0, \dots, i_{n-1} \in I$ and all $n \geq 0$. Implicitly, the p_{ij} satisfy the following

  $$\begin{aligned} p_{ij} &\geq 0, \quad \forall i, j \in I, \\ \sum_{j \in I} p_{ij} &= 1, \quad \forall i \geq 0. \end{aligned}$$

The (could be infinite) matrix $\mathbf{P} = (p_{ij})$ is called the *transition probability matrix* of the chain.

Given a DTMC $\mathbf{P}$ with state space I, let $A \subset I$ be some subset of states. We often want to answer many types of questions about the chain. For example, starting from $X_0 = i \notin A$

- what's the probability A is ever "hit"?
- what's the probability $X_n \in A$ for a given n?
- what's the expected number of steps until A is hit
- what's the probability we'll come back to i?
- what's the expected number of steps until we come back?
- what's the expected number of steps until all states are visited?
- as $n \to \infty$, what's the distribution of where we are? Does the "limit distribution" even exist? If it does, how fast is the convergence rate?

Now, instead of starting from a specific state, it is common to encounter situations where the initial distribution X_0 is given and we want to repeat the above questions. There are many more questions we can ask about the chain.

Before we move on, let's fix several other basic notions and terminologies. A *measure* on the state space I is a vector λ where $\lambda_i \geq 0$, for all $i \in I$. A measure is called a *distribution* if $\sum_{i\in I} \lambda_i = 1$. For any event F, define

$$\text{Prob}_i[F] = \text{Prob}[F \mid X_0 = i]$$

For any random variable Z, define

$$\text{E}_i[Z] = \text{E}[Z \mid X_0 = i]$$

If we know λ is the distribution of X_0, then we also write

$$(X_n)_{n\geq 0} = \text{Markov}(\mathbf{P}, \lambda).$$

3 Multistep transition probabilities and matrices

Define

$$p_{ij}^{(n)} = \text{Prob}_i[X_n = j],$$

which is the probability of going from i to j in n steps. Also define the n-step transition probability matrix

$$\mathbf{P}^{(n)} = (p_{ij}^{(n)}).$$

The following lemma is a straightforward application of the law of total probabilities

Lemma 3.1 (Chapman-Kolmogorov Equations)**.**

$$p_{ij}^{(m+n)} = \sum_{k\in I} p_{ik}^{(n)} p_{kj}^{(m)}, \quad \forall n, m \geq 0, i, j \in I.$$

It follows that

$$\mathbf{P}^{(n)} = \mathbf{P}^n$$

Corollary 3.2. *If λ (a row vector) is the distribution of X_0, then $\lambda\mathbf{P}^n$ is the distribution of X_n*

Exercise 1. Prove the above Lemma and Corollary.

4 Classification of States

A state j is said to be *reachable* from state i if $p_{ij}^{(n)} > 0$ for some $n \geq 0$. In that case, we write $i \rightsquigarrow j$. Two states i and j are said to *communicate* if $i \rightsquigarrow j$ and $j \rightsquigarrow i$, in which case we write $i \leftrightarrow j$.

Exercise 2. Show that communication is an *equivalence relation*, namely it satisfies three properties

- (Reflexive) $i \leftrightarrow i$, for all $i \in I$
- (Transitive) $i \leftrightarrow j$ and $j \leftrightarrow k$ imply $i \leftrightarrow k$.
- (Symmetric) $i \leftrightarrow j$ implies $j \leftrightarrow i$

The fact that communication is an equivalence relation means the relation partitions the state space I into equivalent classes called *communication classes*. A chain is said to be *irreducible* if there is only one class

We can visualize a transition graph representing the chain. The graph has a directed edge from i to j iff $p_{ij} > 0$. The communication classes are *strongly connected components* of this graph. See Figure 2 for an illustration of communication classes.

A *closed* class C is a class where $i \in C$ and $i \rightsquigarrow j$ imply $j \in C$ (i.e., there no escape!). A state i is *absorbing* if $\{i\}$ is a closed class.

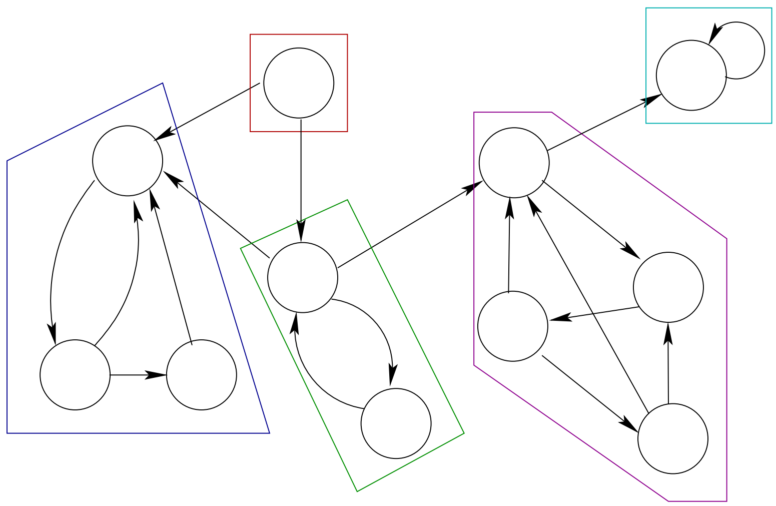

Figure 2: Communication classes

For any state i, define the *first passage time* to be $T_i = \inf\{n \geq 1 \mid X_n = i\}$. (Thus, $X_0 = i$ doesn't count!) Also define the first passage probabilities

$$\begin{aligned} f_{ij}^{(n)} &= \text{Prob}_i[X_n = j \ \wedge \ X_s \neq j, \forall s = 1, .., n-1] = \text{Prob}_i[T_j = n] \\ f_{ij} &= \sum_{n=1}^{\infty} f_{ij}^{(n)} = \text{Prob}_i[T_j < \infty] \end{aligned}$$

Definition 4.1 (Recurrence and transience)**.** State i is said to be *recurrent* (also called *persistent*) if $f_{ii} = \text{Prob}_i[T_i < \infty] = 1$. State i is *transient* if $f_{ii} = \text{Prob}_i[T_i < \infty] < 1$.

Let V_i be the *number of visits to* i, namely

$$V_i := \sum_{n=0}^{\infty} 1_{\{X_n = i\}}.$$

The following two theorems characterizes recurrent and transient states.

Theorem 4.2. *Given a DTMC* $\mathbf{P}$ *and a state* i*, the following are equivalent*

1. *i is recurrent*

2. $f_{ii} = \text{Prob}_i[T_i < \infty] = 1$

3. $\mathrm{Prob}_i[X_n = i$ *for infinitely many* $n] = 1$

4. $\mathrm{E}_i[V_i] = \infty$

5. $\sum_{n\geq 0} p_{ii}^{(n)} = \infty$

Theorem 4.3. *Given a DTMC* $\mathbf{P}$ *and a state* i*, the following are equivalent*

1. *i is transient*

2. $f_{ii} = \mathrm{Prob}_i[T_i < \infty] < 1$

3. $\mathrm{Prob}_i[X_n = i$ *for infinitely many* $n] = 0$

4. $\mathrm{E}_i[V_i] < \infty$

5. $\sum_{n\geq 0} p_{ii}^{(n)} < \infty$

To prove the last two theorems, we need something called the *strong Markovian property.* A random variable T taking values in $\mathbb{N}$ is called a *stopping time* of a DTMC $(X_n)_{n\geq 0}$ if the event $\{T = n\}$ can be determined by looking only at $X_0, \cdots, X_n$ (We need measure theory to be more rigorous on this definition.) For example, the *first passage time* $T_i = \inf\{n \geq 1 : X_n = i\}$ is a stopping time, while the *last exit time* $L_A = \sup\{n : X_n \in A\}$ is not a stopping time.

Theorem 4.4 (Strong Markov Property)**.** *Suppose* T *is a stopping time of a DTMC* $(X_n)_{n\geq 0}$. *Then, conditioned on* $T < \infty$ *and* $X_T = i$*, the sequence* $(X_{T+n})_{n\geq 0}$ *behaves exactly like the Markov chain with initial state* i.

Proof of the previous two theorems. By strong Markovian:

$$\mathrm{E}_i[V_i] = \sum_{n=1}^{\infty} n f_{ii}^{n-1}(1 - f_{ii}) = \frac{1}{1 - f_{ii}}.$$

On the other hand,

$$\mathrm{E}_i[V_i] = \mathrm{E}_i\left[\sum_{n=0}^{\infty} 1_{\{X_n = i\}}\right] = \sum_{n=0}^{\infty} \mathrm{E}_i[1_{\{X_n=i\}}] = \sum_{n=0}^{\infty} p_{ii}^{(n)}.$$

□

Theorem 4.5. *Recurrence and transience are class properties, i.e. in a communication class* C *either all states are recurrent or all states are transient.*

Intuitively, suppose i and j belong to the same class. If i is recurrent and j is transient, each time the process returns to i there's a positive chance of going to j. Thus, the process cannot avoid j forever.

Theorem 4.6. *Let* $\mathbf{P}$ *be a DTMC, then*

(i) Every recurrent class is closed

(ii) Every finite, closed class is positive recurrent

Intuitively, if the class is not closed, there's an escape route, and thus the class cannot be recurrent. In a finite and closed class, it cannot be the case that every state is visited a finite number of times. So, the chain is recurrent.

Infinite closed class could be transient or recurrent. Consider a *random walk on* $\mathbb{Z}$, where $p_{i,i+1} = p$ and $p_{i+1,i} = 1 - p$, for all $i \in \mathbb{Z}$, $0 < p < 1$ The chain is an infinite and closed class. For any state i, we have

$$\begin{aligned} p_{ii}^{(2n+1)} &= 0 \\ p_{ii}^{(2n)} &= \binom{2n}{n} p^n (1-p)^n \end{aligned}$$

Hence,

$$\begin{aligned} \sum_{n=0}^{\infty} p_{ii}^{(2n)} &= \sum_{n=0}^{\infty} \binom{2n}{n} p^n (1-p)^n \\ &\approx \sum_{n \geq n_0} \frac{\sqrt{4\pi n}(2n/e)^{2n}(1+o(1))}{2\pi n (n/e)^{2n}(1+o(1))} p^n (1-p)^n \\ &\approx \sum_{n \geq n_0} \frac{1}{\sqrt{\pi n}} (4p(1-p))^n (1+o(1)). \end{aligned}$$

which is ∞ if $p = 1/2$ and finite if $p \neq 1/2$.

Theorem 4.7. *In an irreducible and recurrent chain, $f_{ij} = 1$ for all i, j*

This is true due to the following reasoning. If $f_{ij} < 1$, there's a non-zero chance of the chain starting from j, getting to i, and never come back to j. However, j is recurrent!

Example 4.8 (Birth-and-Death Chain)**.** Consider a DTMC on state space $\mathbb{N}$ where

- $p_{i,i+1} = a_i$, $p_{i,i-1} = b_i$, $p_{ii} = c_i$
- $a_i + b_i + c_i = 1, \ \forall i \in \mathbb{N}$, and implicitly $b_0 = 0$
- $a_i, b_i > 0$ for all i, except for b_0

The chain is called a birth-and-death chain.

The question is, when is this chain transient/recurrent? To answer this question, we need some results about computing hitting probabilities

Let $\mathbf{P}$ be a DTMC on I. Let $A \subseteq I$. The *hitting time* H^A is defined to be

$$H^A := \inf\{n \geq 0 \ : \ X_n \in A\}.$$

The probability of hitting A starting from i is

$$h_i^A := \mathrm{Prob}_i[H^A < \infty].$$

If A is a closed class, the h_i^A are called the *absorption probabilities*. The *mean hitting time* μ_i^A is defined by

$$\mu_i^A := \mathrm{E}_i[H^A] = \sum_{n<\infty} n\mathrm{Prob}[H^A = n] + \infty\mathrm{Prob}[H^A = \infty].$$

Theorem 4.9. *The vector* $(h_i^A : i \in I)$ *is the minimal non-negative solution to the following system*

$$\begin{cases} h_i^A = 1 & \textit{for } i \in A \\ h_i^A = \sum_{j \in I} p_{ij} h_j^A & \textit{for } i \notin A \end{cases}$$

Here, "minimal" means if x *is another non-negative solution then* $x_i \geq h_i^A$ *for all* i*.*

Theorem 4.10. *The vector* $(\mu_i^A : i \in I)$ *is the minimal non-negative solution to the following system*

$$\begin{cases} \mu_i^A = 0 & \textit{for } i \in A \\ \mu_i^A = 1 + \sum_{j \notin A} p_{ij} \mu_j^A & \textit{for } i \notin A \end{cases}$$

Back to the Birth-and-Death Chain, we can use the above theorem to derive whether the chain is transient or recurrent. Note that

$$f_{00} = c_0 + a_0 h_1^{\{0\}} = 1 - a_0(1 - h_1^{\{0\}}).$$

The system of equations is

$$\begin{cases} h_0^{\{0\}} = 1 \\ h_i^{\{0\}} = a_i h_{i+1}^{\{0\}} + c_i h_i^{\{0\}} + b_i h_{i-1}^{\{0\}} & \text{for } i \geq 1 \end{cases}$$

Define $d_n := \frac{b_1 \ldots b_n}{a_1 \ldots a_n}$, $n \geq 1$, and $d_0 = 1$

When $\sum_{n=0}^{\infty} d_n = \infty$, $h_i^{\{0\}} = 1, \forall i$ is the solution.

When $\sum_{n=0}^{\infty} d_n < \infty$, we have the following solution

$$\begin{cases} h_0^{\{0\}} = 1 \\ h_i^{\{0\}} = \frac{\sum_{j=i}^{\infty} d_j}{\sum_{j=0}^{\infty} d_j} < 1 & \text{for } i \geq 1 \end{cases}$$

Thus,

- the birth-and-death chain is recurrent ($f_{00} = 1$) when $\sum_{j=0}^{\infty} d_j = \infty$
- the birth-and-death chain is transient ($f_{00} < 1$) when $\sum_{j=0}^{\infty} d_j < \infty$

To briefly summarize,

- We often only need to look at closed classes (that's where the chain will eventually end up). Thus, we can then consider irreducible chains instead.
- Let **P** be an irreducible chain. Then,
  1. If **P** is finite, then **P** is recurrent.
  2. If **P** is infinite, then **P** could be either transient or recurrent.

5 Stationary Distributions

We will later examine the behavior of DMTCs "in the limit," i.e. after we run it for a long time. A distribution λ is a *stationary* (also *equilibrium* or *invariant*) distribution if $\lambda\mathbf{P} = \lambda$.

Theorem 5.1. *Let* $\mathbf{P}$ *be a DTMC,*

(i) Let $(X_n)_{n\geq 0} = \text{Markov}(\mathbf{P}, \lambda)$*, where* λ *is stationary, then* $(X_{n+m})_{n\geq 0} = \text{Markov}(\mathbf{P}, \lambda)$ *for any fixed* m.

(ii) In a **finite** *DTMC, suppose for some* $i \in I$ *we have*

$$\lim_{n\to\infty} p_{ij}^{(n)} = \pi_j, \forall j \in I,$$

then $\pi = (\pi_j : j \in I)$ *is an invariant distribution.*

In an infinite DTMC, it is possible that $\lim_{n\to\infty} p_{ij}^{(n)}$ exists for all i, j, producing a vector π for each i, yet π is **not** a distribution. For example, consider the DTMC with state space $\mathbb{Z}$ and

$$p_{i,i+1} = p = 1 - q = 1 - p_{i,i-1}, \quad \forall i \in \mathbb{Z}.$$

(This is called a *random walk on* $\mathbb{Z}$.) Then,

$$\lim_{n\to\infty} p_{ij}^{(n)} = 0, \forall i, j.$$

We often want to know when an irreduciable DTMC has a stationary distribution. The following concepts are needed to answer this question. Define

$$\mu_{ij} = \sum_{n=1}^{\infty} n f_{ij}^{(n)} = \mathrm{E}_i[T_j]$$

Definition 5.2 (Positive and null recurrence)**.** A recurrent state i is *positive recurrent* if $\mu_{ii} < \infty$. A recurrent state i is *null recurrent* if $\mu_{ii} = \infty$.

Example 5.3 (Positive recurrence)**.** Consider the following chain

$$\mathbf{P} = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}, \text{for } 0 < p < 1.$$

Let the states be 0 and 1, then

$$\begin{aligned} f_{00}^{(1)} &= f_{11}^{(1)} = p \\ f_{00}^{(n)} &= f_{11}^{(n)} = (1-p)^2 p^{n-2}, \quad n \geq 2 \end{aligned}$$

Both states are recurrent. Moreover,

$$\mu_{00} = \mu_{11} = p + \sum_{n=2}^{\infty} n(1-p)^2 p^{n-2} = 2.$$

Hence, both states are positive recurrent.

Example 5.4 (Null recurrence)**.** Consider another Markov chain with $I = \mathbb{N}$ where $p_{01} = 1$, and

$$\begin{aligned} P_{m,m+1} &= \frac{m}{m+1}, \quad \forall m \geq 1 \\ P_{m,0} &= \frac{1}{m+1}, \quad \forall m \geq 1. \end{aligned}$$

Then,

$$\begin{aligned} f_{00}^{(1)} &= 0 \\ f_{00}^{(n)} &= \frac{1}{n(n-1)} \\ f_{00} = \sum_{n=1}^{\infty} f_{00}^{(n)} &= 1 \\ \mu_{00} = \sum_{n=1}^{\infty} n f_{00}^{(n)} &= \sum_{n=1}^{\infty} \frac{1}{n} = \infty. \end{aligned}$$

Consequently, 0 is a null recurrent state.

Theorem 5.5. *Positive and null recurrence are class properties, i.e. in a recurrent communication class either all states are positive recurrent or all states are null recurrent.*

Why is it true? Suppose i and j belong to the same class. If i is recurrent and j is transient, each time the process returns to i there's a positive chance of going to j. Thus, the process cannot avoid j forever.

Theorem 5.6 (Existence of a Stationary Distribution)**.** *An irreducible DTMC* $\mathbf{P}$ *has a stationary distribution if and only if one of its states is positive recurrent. Moreover, if* $\mathbf{P}$ *has a stationary distribution* π*, then*

$$\pi_i = 1/\mu_{ii}$$

.

The *line of proof* is as follows. Every irreducible and recurrent $\mathbf{P}$ basically has a unique invariant measure (unique up to rescaling). Due to positive recurrence, the measure can be normalized to be come an invariant distribution. This strategy is manifested in the next lemma.

Define the *expected time spent in* i *between visits to* k

$$\gamma_i^{(k)} = \mathrm{E}_k \left[\sum_{n=0}^{T_k - 1} 1_{X_n = i} \right]$$

Lemma 5.7. *If* $\mathbf{P}$ *is irreducible and recurrent, then*

(i) $\gamma_k^{(k)} = 1$

(ii) the vector $\gamma^{(k)} = (\gamma_i^{(k)} \mid i \in I)$ *is an invariant measure, namely*

$$\gamma^{(k)} \mathbf{P} = \gamma^{(k)}$$

(iii) $0 < \gamma_i^{(k)} < \infty$ *for all* $i \in I$

Conversely, if $\mathbf{P}$ *is irreducible and* λ *is an invariant measure with* $\lambda_k = 1$*, then* $\lambda \geq \gamma^{(k)}$*. Moreover, if* $\mathbf{P}$ *is also recurrent, then* $\lambda = \gamma^{(k)}$

6 Limit Behavior

For a state $i \in I$, let $\text{period}(i) = \gcd\{n \geq 1 : p_{ii}^{(n)} > 0\}$. When $\text{period}(i) \geq 2$, state i is said to be *periodic* with period $\text{period}(i)$. When $\text{period}(i) = 1$, state i is *aperiodic*. A DTMC is *periodic* if it has a periodic state. Otherwise, the chain is *aperiodic*. i

Exercise 3. If $i \leftrightarrow j$, then $\text{period}(i) = \text{period}(j)$.

Exercise 4. If i is aperiodic, then $\exists n_0 : p_{ii}^{(n)} > 0$, $\forall n \geq n_0$.

Exercise 5. If $\mathbf{P}$ is irreducible and has an aperiodic state i, then $\mathbf{P}^n$ has all strictly positive entries for sufficiently large n.

An *ergodic* state is an aperiodic and positive recurrent state. Roughly (not technically correct but close), positive recurrence implies the existence of a unique stationary distribution, and aperiodicity ensures that the chain tends to the stationary distribution in the limit. An *ergodic Markov chain* is a Markov chain in which all states are ergodic. (Basically, a "well-behaved" chain.)

Theorem 6.1 (Convergence to equilibrium)**.** *Suppose* $\mathbf{P}$ *is irreducible and ergodic. Then, it has an invariant distribution* π*. Moreover,*

$$\frac{1}{\mu_{jj}} = \pi_j = \lim_{n\to\infty} p_{ij}^{(n)}, \quad \forall j \in I.$$

Thus, π *is the unique invariant distribution of* $\mathbf{P}$*.*

Remark 6.2. There is a generalized version of this theorem for irreducible chains with period $d \geq 2$. (And the chain is not even required to be positive recurrent.)

Exercise 6. Prove that if a DTMC is irreducible then it has at most one stationary distribution.

Let

$$V_i(n) = \sum_{k=0}^{n-1} 1_{\{X_k = i\}}.$$

Theorem 6.3 (Ergodic Theorem)**.** *Let* $\mathbf{P}$ *be an irreducible DTMC. Then*

$$\text{Prob}\left[\lim_{n\to\infty} \frac{V_i(n)}{n} = \frac{1}{\mu_{ii}}\right] = 1$$

Moreover, if $\mathbf{P}$ *is positive recurrent with (unique) invariant distribution* π*, then for any bounded function* $f : I \to \mathbb{R}$

$$\text{Prob}\left[\lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} f(X_k) = \sum_{i\in I} \pi_i f_i\right] = 1,$$

References

[1] J. R. Norris, *Markov chains*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 1998. Reprint of 1997 original.