ECE 581: Finite-State Markov Chains

Henry D. Pfister
Duke University

December 4, 2025

Contents

1 What is a Markov chain?

1.1 Introduction

Definition 1 (Finite-state Markov chain)**.** A *finite-state Markov chain* (FSMC) with n states is a sequence of random variables $X_0, X_1, X_2, \ldots$ where each $X_i \in [n] \coloneqq \{1, 2, \ldots, n\}$ and the following Markov condition holds,

$$\Pr\left(X_{t+1} = j \,|\, X_t = i, X_0, X_1, \ldots, X_{t-1}\right) = \Pr\left(X_{t+1} = j \,|\, X_t = i\right).$$

If $\Pr(X_{t+1} = j \,|\, X_t = i)$ does not depend on t, then the Markov chain is called *time invariant* (or homogeneous).

In the remainder of this note, we assume the FSMC is time invariant and we let $P \in \mathbb{R}^{n\times n}$ denote the *transition-probability matrix* with entries $[P]_{i,j} \coloneqq P_{i,j} = \Pr(X_{t+1} = j \mid X_t = i)$. Since the i-th row of P represents the probability distribution of the next state when the current state is i, we see that $P_{i,j} \geq 0$ and $\sum_{j=1}^{n} P_{i,j} = 1$. The Markov property also implies that

$$\begin{aligned}\Pr\left(X_{t+2} = j \mid X_t = i\right) &= \sum_{k=1}^{n} \Pr\left(X_{t+2} = j \mid X_{t+1} = k\right) \Pr\left(X_{t+1} = k \mid X_t = i\right) \\ &= \sum_{k=1}^{n} P_{k,j} P_{i,k} = \sum_{k=1}^{n} P_{i,k} P_{k,j} = \left[P^2\right]_{i,j}.\end{aligned}$$

Arguing by induction, one also finds that $\Pr(X_{t+m} = j \mid X_t = i) = [P^m]_{i,j}$. Thus, given a fixed starting state, one can calculate the probability of being in state j after m steps by computing the m-th power of a matrix. Using the notation $\underline{\pi}^{(t)} = (\pi_1^{(t)}, \ldots, \pi_n^{(t)})$ with $\pi_i^{(t)} \coloneqq \Pr(X_t = i)$, we also see that

$$\begin{aligned}\pi_j^{(t+1)} &= \sum_{i=1}^{n} \Pr(X_{t+1} = j, X_t = i) \\ &= \sum_{i=1}^{n} \Pr(X_{t+1} = j | X_t = i) \Pr(X_t = i) \\ &= \sum_{i=1}^{n} P_{i,j} \pi_i^{(t)} = \left[\underline{\pi}^{(0)} P^{t+1}\right]_j,\end{aligned}$$

where $\pi_i^{(0)}$ is the probability that the process starts in state i.

Two-state weather model. Let $n = 2$ with states $\{R, S\}$ (rain/sun), and $P = \begin{bmatrix} 0.9 & 0.1 \\ 0.5 & 0.5 \end{bmatrix}$. From any initial distribution $\pi^{(0)}$, we can compute $\pi^{(t)} = \pi^{(0)} P^t$ or simulate by repeated sampling from P's rows.

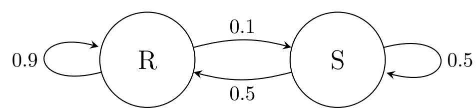

Figure 1: Two-state weather chain.

1.2 Sampling From a Markov Chain

Let the random variable U be uniformly distributed on the interval $[0, 1)$ and let u be a realization of U. For a discrete random variable $X \in [n]$, one can use U to simulate X by assigning subintervals of $[0, 1)$ to each of the n possibilities for X. Let $F_X(x) = \sum_{i=1}^{x} \Pr(X = i)$ for $x \in [n]$ be the cumulative distribution function of X. Then, we can set $X = x$ if $u \in [F_X(x - 1), F_X(x))$ (with $F_X(0) = 0$ by convention). This works because

$$\Pr\left(U \in [F_X(x-1), F_X(x))\right) = F_X(x) - F_X(x-1) = \Pr(X = x).$$

For built-in sampling functions in Matlab and Python, see `randsrc` and `np.random.choice`, respectively.

Similarly, one can simulate a Markov chain by using pseudo-random numbers to generate realizations of the process. Let $u_1, u_2, \ldots$ be a realization of a sequence of independent and identically distributed (i.i.d.) copies of U. Then, a realization $x_1, x_2, \ldots$ of $X_1, X_2, \ldots$ is generated by choosing x_t to be the unique value satisfying

$$\sum_{i=1}^{x_t - 1} \Pr(X_t = i | X_{t-1} = x_{t-1}) \leq u_t < \sum_{i=1}^{x_t} \Pr(X_t = i | X_{t-1} = x_{t-1}).$$

1.3 A Few Simple Examples

1.3.1 What is the chance that a game of Candyland will last m moves?

Candyland is an American boardgame where players draw cards to move and the goal is to reach the candy castle first. It can be played by very young children because it requires neither reading nor counting. Players draw cards randomly and, if a colored card is drawn, they move their piece to the next position of that color. If the card has a picture, they move to the position with that picture. There are also spaces that allow shortcuts or cause delays. A picture of board can be found at:

`https://kim.scarborough.chicago.il.us/images/cl-2010`

1.3.2 What is the chance that a game of Chutes and Ladders lasts m moves?

Chutes and Ladders (aka Snakes and Ladders outside of the US) is a boardgame where a single die is rolled to determine how far you move on a gameboard defined by a grid. Some locations contain ladders that let you skip ahead while others contains chutes (aka snakes) that you move you backwards. Historically, it is based on an ancient game from India that teaches morality by associating ladders with virtues and snakes with vices. For more information, see:

`https://en.wikipedia.org/wiki/Snakes_and_Ladders`

1.3.3 What are the best properties to buy in Monopoly?

Monopoly is an American boardgame where players move around the gameboard buying, selling, and developing properties. Rent is collected from other players who land on your properties. Properties differ both in their expense (e.g., Park Place is valued much more highly than Baltic Avenue) and the chance that players will land on them. Markov chains can be used to estimate how often players will land on each property, which can be used to estimate their value. For way too much information, see:

`http://pfister.ee.duke.edu/courses/ece586/monopoly.pdf`

1.4 Directed Graph Perspective

We can associate a Markov chain defined by P with a directed graph comprised of n vertices labeled by $[n]$ with an edge $i \to j$ whenever $P_{ij} > 0$.

Definition 2. State j is called *reachable* from state i if $[P^m]_{i,j} > 0$ for some $m \in \mathbb{N}$. States i, j *communicate* if they are reachable from the other. This notion of communicating defines an equivalence relation that partitions $[n]$ into *communicating classes*. The chain is *irreducible* if there is only one class and *aperiodic* if $\max_{i \in [N]} \gcd\{m \geq 1 \mid [P^m]_{i,i} > 0\} = 1$.

Definition 3. If the process can become stuck in a single state (e.g., say j is the state at the end of a game), then $P_{j,j} = 1$ and that state is called *absorbing*. A Markov chain is called *absorbing* if an absorbing state is reachable from every state. A set $A \subset [n]$ of states is called *absorbing* if $\sum_{j \in A} P_{i,j} = 1$ for all $i \in A$. In this case, the Markov chain will never leave the set A after entering any state in that set. State i is called *transient* if the expected number of times it will return to itself is finite (e.g., if there is an absorbing set $A \setminus \{i\}$ is reachable from it).

Example 4 (Reducible chain with absorbing set)**.** Consider the $n = 5$ case with transition matrix

$$P = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.4 & 0.5 & 0.1 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The states in this chain can be classified as follows: States 1 and 2 communicate with each other (paths $1 \leftrightarrow 2$), but can enter $\{3, 4\}$ (via $2 \to 3$). They are transient. - States 3 and 4 communicate with each other and form a closed (absorbing) class $\{3, 4\}$ since no probability leaves this set. - State 5 is transient and deterministically moves to 1.

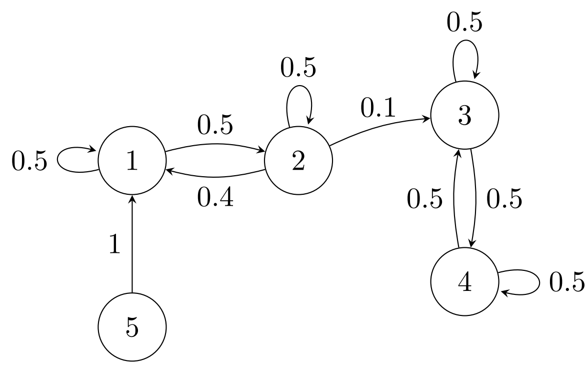

Example 5 (Multiple closed classes, one a singleton absorbing state)**.** Let $n = 5$ with

$$P = \begin{bmatrix} 0.6 & 0.4 & 0 & 0 & 0 \\ 0.3 & 0.6 & 0.1 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The states in this chain can be classified as follows: $\{3, 4\}$ is a closed communicating class (non-trivial absorbing set) and $\{5\}$ is a singleton absorbing state. States $\{1, 2\}$ are transient and eventually reach the set $\{3, 4\}$.

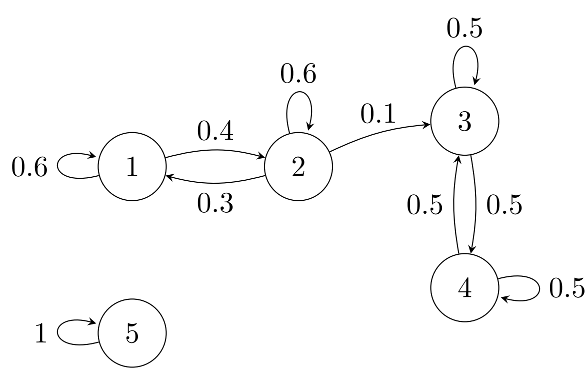

1.5 Hitting Times

Definition 6. For a Markov chain X_t, the *hitting time* (or first hit time) of a set $A \subseteq [n]$ is the random variable

$$H_A(\omega) := \inf\{t \in \mathbb{N}_0 \mid X_t(\omega) \in A\},$$

which equals the first time at which X_t is in the set A. Similarly, the expected time to hit the set A, when starting from $X_0 = i$, is denoted by

$$\eta_{i,A} \coloneqq \mathbb{E}[H_A \mid X_0 = i].$$

For a Markov chain with starting state $X_0 = i$, these notes abuse notation and use $T_{i,j}$ to denote the random variable H_j on the conditional probability space formed by conditioning on $X_0 = i$. Thus, we can write

$$\Pr(T_{i,j} = m) = \Pr(H_j = m \mid X_0 = i) = \Pr\left(X_m = j, X_{m-1} \neq j, X_{m-2} \neq j, \ldots, X_1 \neq j \,|\, X_0 = i\right),$$

where, by convention, $\Pr\big(T_{i,j} = 0\big) = \delta_{i,j}$ with $\delta_{i,j}$ denoting the Kronecker delta function

$$\delta_{i,j} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

If state j is not reachable from state i, then $T_{i,j} = \infty$ (i.e., $\Pr(T_{i,j} = \infty) = 1$) by convention.

Remark 7. By constructing chains where the probability of reaching state j from state i is less than 1, it is also possible to construct examples where $0 < \Pr(T_{i,j} = \infty) < 1$.

1.5.1 Special Case of Absorbing States

If state j is absorbing, then the distribution of the hitting time $T_{i,j}$ satisfies the following simple formula,

$$\Pr\left(T_{i,j} \leq m\right) = \Pr\left(X_m = j \,|\, X_0 = i\right) = \left[P^m\right]_{i,j}.$$

This follows from initializing $X_0 = i$ and observing the equality of the two events "$T_{i,j} \leq m$" and "$X_m = j$". Since $X_0 = i$, if $X_m = j$, then we clearly have $T_{i,j} \leq m$. On the other hand, if $T_{i,j} \leq m$, then $X_t = j$ for some $t \leq m$ and, thus, $X_m = j$ because state j is absorbing. Hence, for $m \geq 1$, we find that

$$\Pr(T_{i,j} = m) = \Pr\left(T_{i,j} \leq m\right) - \Pr\left(T_{i,j} \leq m-1\right) = \left[P^m - P^{m-1}\right]_{i,j}.$$

1.5.2 General Case

Lemma 8 (Expected hitting-time recursion). *For a finite Markov chain with transition matrix P, let $\eta_{i,j}$ be the expected hitting time (or time to hit) state j when starting in state i. By convention, the $i = j$ case is defined by $\eta_{i,i} = 0$ for all $i \in [n]$. For $i \neq j$, the expected hitting times satisfy the linear equations defined by*

$$\eta_{i,j} = 1 + \sum_{k=1}^{n} P_{i,k}\eta_{k,j}.$$

Proof. This follows from

$$\begin{aligned} \eta_{i,j} &= \mathbb{E}[T_{i,j} \mid X_0 = i] \\ &= \sum_{k=1}^{n} \Pr(X_1 = k \mid X_0 = i)\, \mathbb{E}[T_{i,j} \mid X_0 = i, X_1 = k] \\ &= \sum_{k=1}^{n} P_{i,k}\big(1 + \eta_{k,j}\big) \\ &= 1 + \sum_{k=1}^{n} P_{i,k}\eta_{k,j}. \end{aligned}$$

where the second step uses the total probability law, the Markov property, and time invariance. ☐

Remark 9. Since $\eta_{i,j} = \infty$ whenever j is not reachable from i, one must be somewhat careful when solving these equations but they can be used to solve $\eta_{i,j}$ whenever it is finite. In Section 1.6, this issue is avoided by focusing on the number of visits to transient states before absorption.

Example 10. What is the distribution of the number of fair coin tosses before one observes 3 heads in a row? To solve this, consider a 4-state Markov chain with transition probability matrix

$$P = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 \\ 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

where $X_t = i$ if the last $i-1$ tosses were heads. Then, $X_t = 4$ is an absorbing state that is reached when the last three tosses are heads.

Solution. For the above Markov chain starting in state 1, the hitting time to observe 3 heads is given by H_4. Then, we have $\eta_{i,4} = \mathbb{E}[H_4 \,|\, X_0 = i]$ for $i \in [4]$ and recall that where $\eta_{4,4} = 0$. Using Lemma 8, we can relate these expectations to each other. Dropping the subscript 4, we use the shorthand $\eta_i = \eta_{i,4}$ to write

$$\begin{aligned} \eta_1 &= 1 + 0.5\,\eta_1 + 0.5\,\eta_2, \\ \eta_2 &= 1 + 0.5\,\eta_1 + 0.5\,\eta_3, \\ \eta_3 &= 1 + 0.5\,\eta_1 + 0.5\,\eta_4. \end{aligned}$$

Solving these equations, we find that $\eta_2 = 1 + 0.5\,\eta_1 + 0.5\,\eta_3 = 1.5 + 0.75\,\eta_1$ and substituting into $\eta_1 = 1 + 0.5\,\eta_1 + 0.5\,\eta_2$ gives $\eta_1 = 2 + \eta_2 = 2 + 1.5 + 0.75\,\eta_1$. Hence, $0.25\,\eta_1 = 3.5$ and $\eta_1 = 14$. Thus, the expected number of tosses to see three heads in a row is 14.

Lemma 11 (Hitting-time CDF recursion)**.** *For a finite Markov chain with transition matrix P, let $T_{i,j}$ be the first hitting time of state j starting from i. Define $\phi_{i,j}^{(m)} \coloneqq \Pr(H_j \le m \mid X_0 = i) = \Pr(T_{i,j} \le m)$. Then, for all $m \ge 0$, we have*

$$\phi_{i,j}^{(m+1)} = \begin{cases} 1, & i = j, \\ \sum_{k=1}^{n} P_{i,k}\, \phi_{k,j}^{(m)}, & i \ne j, \end{cases}$$

with initialization $\phi_{i,j}^{(0)} = \delta_{i,j}$. The probability of eventually reaching state j from state i is given by

$$\phi_{i,j} \coloneqq \Pr(H_j < \infty \mid X_0 = i) = \lim_{m\to\infty} \phi_{i,j}^{(m)},$$

which is also a fixed point of the recursion.

Proof. Consider any $i, j \in [n]$. If $i = j$, then $T_{i,j} = 0$ almost surely, so $\phi_{i,j}^{(m)} = 1$ for all $m \ge 0$. If $i \ne j$, then we can condition on the first step to see that

$$\begin{aligned} \phi_{i,j}^{(m+1)} &= \Pr(H_j \le m+1 \mid X_0 = i) \\ &= \sum_{k=1}^{n} \Pr(X_1 = k \mid X_0 = i)\, \Pr(T_{k,j} \le m \mid X_0 = k) \\ &= \sum_{k=1}^{n} P_{i,k}\, \phi_{k,j}^{(m)}, \end{aligned}$$

Figure 2: Miniature chutes and ladders game

where the second step uses the total probability law, the Markov property, and time invariance. Since $\phi_{i,j}^{(m)}$ is non-decreasing in m, its limit as $m \to \infty$ exists and equals $\phi_{i,j} = \Pr(H_j < \infty \mid X_0 = i)$. Moreover, taking the limit of the recursion shows that $\phi_{i,j}$ is a fixed point of the recursion. □

Exercise 1. Write a computer program (e.g., in Python, Matlab, ...) to compute $\Pr(T_{1,4} = m)$ for $m = 1, 2, \ldots, 100$ and use this to compute and print an estimate of the expected number of tosses $\mathbb{E}[T_{1,4}]$. Write a computer program that generates 500 realizations from this Markov chain. Then, use them to plot a histogram of $T_{1,4}$ and compute/print an estimate of the expected number of tosses $\mathbb{E}[T_{1,4}]$.

Exercise 2. Consider the miniature chutes and ladders game shown in Figure 2. Assume a player starts on the space labeled 1 and plays by rolling a fair four-sided die and then moves that number of spaces. If a player lands on the bottom of a ladder, then they automatically climb to the top. If a player lands at the top of a slide, then they automatically slide to the bottom. This process can be modeled by a Markov chain with $n = 16$ states where each state is associated with a square where players can start their turn (i.e., players never start at the bottom of a ladder or the top of a slide). To finish the game, players must land exactly on space 20 (i.e., if your roll would take you beyond 20, then no move is made). Compute the transition probability matrix P of the implied Markov chain. For this Markov chain, use the program from Exercise 1 to compute and plot the cumulative distribution of the number turns a player takes to finish (i.e., the probability $\Pr\left(T_{1,20} \leq m\right)$ where $T_{1,20}$ is the hitting time from state 1 to state 20). Compute and print the mean $\mathbb{E}[T_{1,20}]$. Use the program from Exercise 1 to generate 500 realizations from this Markov chain. Then, use them to plot a histogram of $T_{1,20}$ and compute/print an estimate of the expected number of tosses $\mathbb{E}[T_{1,20}]$.

1.6 Canonical Form and the Fundamental Matrix

One property of a transient state is that, starting from it, the expected number of return visits is finite (equivalently, it does not belong to any closed communicating class). In an *absorbing chain* there is at least one absorbing state and every nonabsorbing state is transient. For a target absorbing state j and initial state i, the *absorption probability* is defined to be

$$\Pr(H_j < \infty \mid X_0 = i) = \phi_{i,j},$$

where $\phi_{i,j}$ is defined in Lemma 11.

After relabeling states so that all r transient states come first and the s absorbing states last, the transition matrix has the canonical block form

$$P = \begin{bmatrix} Q & R \\ 0 & I_s \end{bmatrix},$$

where $Q \in \mathbb{R}^{r\times r}$ gives transient-to-transient transitions and $R \in \mathbb{R}^{r\times s}$ gives transient-to-absorbing transitions. This defines a kind of generalized geometric distribution where all transitions from transient states to absorbing states in k steps are defined by $\Pr(X_k = j \mid X_0 = i) = [Q^{k-1}R]_{i,j-r}$.

To understand this better, we note that

$$\Pr(H_A > m \mid X_0 = i) = \Pr(X_m \notin A \mid X_0 = i) = \sum_{j=1}^{r} [Q^m]_{i,j}$$

because $[Q^m]_{i,j}$ is probability of going from i to j in exactly m steps without being absorbed into A. Computing the mean by summing the complementary CDF[1] shows that

$$\begin{aligned} \eta_{i,A} = \mathbb{E}[H_A \mid X_0 = i] &= \sum_{k=1}^{\infty} \Pr(H_A \geq k \mid X_0 = i) \\ &= \sum_{k=1}^{\infty}\sum_{j=1}^{r} [Q^{k-1}]_{i,j} = \sum_{j=1}^{r} \left[\sum_{k=1}^{\infty} Q^{k-1} \right]_{i,j}. \end{aligned}$$

For an absorbing chain, all eigenvalues of Q satisfy $|\lambda| < 1$, so $(I_r - Q)^{-1}$ is well-defined. For this reason, we can define the *fundamental matrix*

$$N \coloneqq (I_r - Q)^{-1} = \sum_{k=0}^{\infty} Q^k$$

and observe that the expected time to hit A starting from i (i.e., the expected time to absorption starting from state $i \in [r]$) is equal to the sum

$$\eta_{i,A} = \sum_{j=1}^{r} N_{ij}.$$

Similarly, the probability of being absorbed by state $j \in \{r+1, \ldots, n\}$ when starting from state $i \in [r]$ equals

$$\phi_{i,j} = \sum_{k=0}^{\infty} [Q^k R]_{i,j-r} = [NR]_{i,j-r}.$$

Example 12 (Gambler's ruin)**.** Consider states $1, 2, \ldots, n$ with absorbing set $A = \{1, n\}$ and, for $2 \leq i \leq n-1$, transitions $P_{i,i+1} = p$ and $P_{i,i-1} = q = 1-p$. For $2 \leq i \leq n-1$, let $h_i \coloneqq \Pr(\text{hit } n \text{ before } 1 \mid X_0 = i)$ and let $d_i \coloneqq \mathbb{E}[\text{time to hit } A \mid X_0 = i]$. A gambler betting 1 dollar per game starts with $i-1$ dollars, and decides to leave if they reach $n-1$ dollars, is ruined if they reach 0 dollars before reaching the $n-1$ dollars.

The following lemma gives closed-form expressions for the probability of ruin using the fundamental matrix method.

[1] For a non-negative integer-valued random variable with cdf $F_X(x)$, we have $\mathbb{E}[X] = \sum_{i=1}^{\infty}(1 - F_X(i))$.

Lemma 13 (Gambler's ruin solution via the fundamental matrix)**.** *For the gambler's ruin chain above, the absorption probabilities and expected durations are*

$$h_i = \begin{cases} \dfrac{i-1}{n-1}, & p = q = \frac{1}{2}, \\ \dfrac{1-(q/p)^{i-1}}{1-(q/p)^{n-1}}, & p \neq q, \end{cases} \qquad d_i = \begin{cases} (i-1)(n-i), & p = q = \frac{1}{2}, \\ \dfrac{i-1}{q-p} - \dfrac{n-1}{q-p} \cdot \dfrac{1-(q/p)^{i-1}}{1-(q/p)^{n-1}}, & p \neq q. \end{cases}$$

Proof. For $i \notin A$, it is easy to verify that $h_i = p\,h_{i+1} + q\,h_{i-1}$ for $2 \leq i \leq n-1$ with boundary conditions $h_1 = 0$, $h_n = 1$. For the symmetric case ($p = q = \frac{1}{2}$), the relation becomes $h_{i+1} - h_i = h_i - h_{i-1}$ and the solutions are affine $h_i = ai + b$. Using $h_1 = 0$ and $h_n = 1$ gives $a = \frac{1}{n-1}$ and $b = -\frac{1}{n-1}$, hence $h_i = \frac{i-1}{n-1}$. For the biased case ($p \neq q$), we can solve the homogeneous second-order linear recurrence

$$h_i = p\,h_{i+1} + q\,h_{i-1}$$

that is satisfied when $i \in \{2, \ldots, n-1\}$. We seek solutions of the form $h_i = z^{\,i}$, which leads to the characteristic equation

$$z^{\,i} = p\,z^{\,i+1} + q\,z^{\,i-1} \quad \Longleftrightarrow \quad pz^2 - z + q = 0.$$

This quadratic has distinct roots $z_1 = 1$ and $z_2 = \frac{q}{p}$ since $p \neq q$ implies $p \neq q$ and hence $z_1 \neq z_2$. Therefore, the general solution of the recurrence has the form

$$h_i = A + B\left(\frac{q}{p}\right)^{i-1}.$$

Imposing the boundary condition $h_1 = 0$ gives

$$h_1 = A + B = 0 \quad \Longrightarrow \quad A = -B.$$

From $h_n = 1$, we get

$$h_n = A + B\left(\frac{q}{p}\right)^{n-1} = 1.$$

Substituting $A = -B$ shows

$$-B + B\left(\frac{q}{p}\right)^{n-1} = 1 \quad \Longrightarrow \quad B\left(\left(\frac{q}{p}\right)^{n-1} - 1\right) = 1 \quad \Longrightarrow \quad B = \frac{1}{\left(\frac{q}{p}\right)^{n-1} - 1}.$$

The stated general solution follows from simplifying.

The formulas for the expected duration $d_i = \mathbb{E}[\text{time to hit } \{1, n\} \mid X_0 = i]$ follow from an analogous recurrence

$$d_i = 1 + p\,d_{i+1} + q\,d_{i-1}, \quad 2 \leq i \leq n-1,$$

with $d_1 = d_n = 0$ and are obtained in the same way, so we omit the algebraic details. □

Remark 14 (Connection to the fundamental matrix method)*.* The recurrence solution above is the explicit form of the general *fundamental matrix* approach for absorbing Markov chains. If we label the transient states $\{2, \ldots, n-1\}$ consecutively, the transition matrix takes the canonical form $P = \begin{bmatrix} Q & R \\ 0 & I \end{bmatrix}$, where Q records transitions among transient states and the two columns of

$R = [\, r_1 \; r_n \,] \in \mathbb{R}^{(n-2)\times 2}$ give transition probabilities into the absorbing states 1 and n. Then, the probabilities of hitting n before 1 when starting from each transient state are given by

$$\begin{bmatrix} h_2 \\ \vdots \\ h_{n-1} \end{bmatrix} = N r_n, \qquad N = (I - Q)^{-1}$$

using the fundamental matrix formula. This is equivalent to the linear system

$$(I - Q)h = r_n.$$

Written entrywise, this system is exactly the second-order recurrence $h_i = p\, h_{i+1} + q\, h_{i-1}$ with boundary conditions $h_1 = 0$ and $h_n = 1$. Thus, solving the characteristic equation simply gives a closed-form solution of the linear system underlying the fundamental matrix formula.

2 Recurrent Markov Chains

2.1 Recurrence and Stationary Distributions

From Definition 2, we know that communicating states are reachable from each other and that, if all pairs of states are communicating, then the chain is called irreducible.

Definition 15 (Recurrent)**.** A state is *recurrent* if it is expected to return to itself infinitely many times. All states in an irreducible finite-state Markov chain are recurrent. A state is *transient* if this expectation is finite (i.e, $\sum_{t=1}^{\infty} \Pr(X_t = i \mid X_0 = i) < \infty$).

Definition 16 (Stationary distribution)**.** A probability vector $\underline{\pi}$ is stationary if $\underline{\pi}P = \underline{\pi}$. Equivalently, $\underline{\pi}$ is a left eigenvector for eigenvalue 1 normalized by $\sum_{i=1}^{n} \pi_i = 1$.

Theorem 17 (Existence and uniqueness)**.** *For a finite-state irreducible Markov chain, there exists a unique stationary distribution $\underline{\pi}$ with $\pi_i > 0$ for all $i \in [n]$.*

Sketch of Proof. Applying the Perron–Frobenius Theorem to an irreducible nonnegative matrix shows that all eigenvalues have modulus at most 1 and there is a unique simple eigenvalue of 1. This eigenvalue is associated with a unique positive left eigenvector. Normalizing this left eigenvector the unique stationary distribution. □

If the transition-probability matrix is positive (i.e., the process can transition to any state in one step), then the Markov chain is obviously irreducible and it has a unique stationary distribution. If the n-step transition-probability matrix P^n is positive, then the stationary distribution $\underline{\pi}$ also satisfies the steady-state limit

$$\pi_i = \lim_{t\to\infty} \Pr(X_t = i),$$

which equals the expected fraction of time that the process spends in state i.

One can find the stationary distribution by first rewriting $\underline{\pi}P = \underline{\pi}$ as

$$(I - P)^T \underline{\pi}^T = \underline{0}.$$

Then, one can solve for $\underline{\pi}^T$ (up to normalization) by applying row reduction to $(I - P)^T$ and computing the one-dimensional basis of the null space. This can be computed in Matlab and Python using the functions `null` and `scipy.linalg.null_space`. Lastly, one must normalize the resulting basis vector by enforcing the condition $\sum_{j=1}^{n} \pi_j = 1$ (e.g., by dividing by the sum of its entries).

Example 18. In a certain city, it is said that the weather is rainy with a 90% probability if it was rainy the previous day and with a 50% probability if it not rainy the previous day. If we assume that only the previous day's weather matters, then we can model the weather of this city by a Markov chain with $n = 2$ states whose transitions are governed by

$$P = \begin{bmatrix} 0.9 & 0.1 \\ 0.5 & 0.5 \end{bmatrix}.$$

Under this model, what is the steady-state probability of rainy weather?

To find this, we solve for the stationary distribution. As described above, we write

$$(I - P)^T = \begin{bmatrix} 0.1 & -0.5 \\ -0.1 & 0.5 \end{bmatrix} \implies \begin{bmatrix} 1 & -5 \\ 0 & 0 \end{bmatrix},$$

where $\implies$ denotes row reduction. Thus, $\pi_1 - 5\pi_2 = 0$ and $\pi_1 + \pi_2 = 1$ imply that $\pi_1 = 1 - \pi_2 = 5/6$ is the steady state-probability of rainy weather.

Exercise 3. Write a program to compute the stationary distribution of a Markov chain when it is unique. Consider a game where the gameboard has 8 different spaces arranged in a circle. During each turn, a player rolls two 4-sided dice and moves clockwise by a number of spaces equal to their sum. Define the transition matrix for this 8-state Markov chain and compute its stationary distribution.

Next, suppose that one space is special (e.g., state-1 of the Markov chain) and a player can only leave this space by rolling doubles (i.e., when both dice show the same value). Again, the player moves clockwise by a number of spaces equal to their sum. Define the transition matrix for this 8-state Markov chain and compute its stationary probability distribution.

2.2 Convergence to equilibrium

Theorem 19 (Convergence via spectral decomposition)**.** *If a finite Markov chain is irreducible and aperiodic with stationary distribution $\underline{\pi}$, then for any initial distribution $\underline{\mu}$,*

$$\lim_{t\to\infty} \underline{\mu} P^t = \underline{\pi}$$

componentwise. Moreover, there exist constants $C < \infty$ and $\rho \in (0,1)$ such that $\|\underline{\mu} P^t - \underline{\pi}\|_2 \leq C\,\rho^t \|\underline{\mu} - \underline{\pi}\|_2$ for all $t \geq 0$.

Sketch via Jordan form. Since the chain is aperiodic and irreducible, 1 is a simple eigenvalue of P with right eigenvector $\underline{1}$ and left eigenvector $\underline{\pi}$, and all other eigenvalues satisfy $|\lambda| < 1$. Let $P = SJS^{-1}$ be the Jordan decomposition where $J_{1,1} = 1$, the first column of S is $\underline{1}$ and the first row of S^{-1} is $\underline{\pi}$. Then $P^t = SJ^tS^{-1}$ and J^t converges to the projector $\underline{e}_1^\top \underline{e}_1$ onto the eigenspace of eigenvalue 1 while the other Jordan blocks decay geometrically (i.e., $\|S(J^t - \underline{e}_1^\top \underline{e}_1)S^{-1}\|_2 \leq C\rho^t$). Hence $\underline{\mu} P^t \to \underline{\pi}$ and the bound follows from

$$\begin{aligned}
\|\underline{\mu} P^t - \underline{\pi} P^t\|_2 &= \|(\underline{\mu} - \underline{\pi})(SJ^tS^{-1})\|_2 \\
&= \|(\underline{\mu} - \underline{\pi})S(\underline{e}_1^\top \underline{e}_1 + J^t - \underline{e}_1^\top \underline{e}_1)S^{-1}\|_2 \\
&= \|(\underline{\mu} - \underline{\pi})(\underline{1}^\top \underline{\pi} + S(J^t - \underline{e}_1^\top \underline{e}_1)S^{-1})\|_2 \\
&= \|(\underline{\mu} - \underline{\pi})S(J^t - \underline{e}_1^\top \underline{e}_1)S^{-1}\|_2 \\
&\leq \|\underline{\mu} - \underline{\pi}\|_2 \, \|S(J^t - \underline{e}_1^\top \underline{e}_1)S^{-1}\|_2 \\
&\leq C\,\rho^t \|\underline{\mu} - \underline{\pi}\|_2.
\end{aligned}$$

□

Example 20 (Cycle $\mathbb{Z}/N\mathbb{Z}$ with lazy walk)**.** For $\tilde{P} = \frac{1}{2}I + \frac{1}{4}S^+ + \frac{1}{4}S^-$ (right/left circular shifts), eigenvalues are $\lambda_k = \frac{1}{2} + \frac{1}{2}\cos(2\pi k/N)$. Hence the second-largest eigenvalue satisfies $1 - \lambda_* = \Theta(1/N^2)$ and deviations from stationarity decay like $\lambda_*^t = \exp(-\Theta(t/N^2))$. Equivalently, the time to get within an ε-neighborhood of equilibrium in $\|\cdot\|_2$ scales as $\Theta\big(N^2 \log(1/\varepsilon)\big)$.

2.3 Detailed Balance and Reversibility

Definition 21 (Time reversal)**.** For a Markov chain $X_0, X_1, \ldots$, its *time reversal* $Y_0, Y_1, \ldots$ satisfies $\Pr(X_t = i, X_{t+1} = j) = \Pr(Y_t = j, Y_{t+1} = i)$ for all $i, j \in [n]$ and $t \in \mathbb{N}_0$. If the original chain has transition probability matrix P and positive stationary distribution $\underline{\pi}$, then Bayes' rule implies that the reversed chain has transition probability matrix $Q_{j,i} = \pi_i P_{i,j}/\pi_j$.

Definition 22 (Detailed balance)**.** For a Markov chain with transition probability matrix P, a probability vector $\underline{\pi}$ satisfies the *detailed balance* condition if $\pi_i P_{ij} = \pi_j P_{ji}$ for all $i, j \in [n]$.

Lemma 23. *If a Markov chain satisfies the detailed balance condition with a strictly positive distribution vector $\underline{\pi}$, then it is reversible and $\underline{\pi}$ is stationary. Moreover, if the underlying graph is connected (i.e., the state space cannot be partitioned into two nonempty sets), then the chain is irreducible with $\underline{\pi}$ as the unique stationary distribution.*

Proof. Reversibility follows because the reversed transition matrix satisfies $Q_{i,j} = \pi_i P_{i,j}/\pi_j = P_{i,j}$ by detailed balance. Note that this definition uses Q to define a forward time Markov chain which is equivalent to running the P chain in reverse. Summing $\pi_i P_{i,j} = \pi_j P_{j,i}$ over i yields $\underline{\pi}P = \underline{\pi}$, so $\underline{\pi}$ is stationary. If the chain is not disconnected, then either state i is reachable from j or vice-versa for all $i, j \in [n]$. By reversibility, if one direction holds, then the other direction holds and all states are communicating. Hence, the chain is irreducible. Uniqueness under irreducibility then follows from Theorem 17. This condition is required though because, if the chain is composed of two disconnected sets of states each defining irreducible reversible chains and $\underline{\pi}$ is a convex combination of their stationary distributions, then the detailed balance condition will hold but the chain is not irreducible. □

Example 24 (Random walk on an undirected graph)**.** Let $G = (V, E)$ be a connected simple undirected graph with n vertices and degrees d_i. The Markov chain formed by walking the graph on uniform random edges has $P_{i,j} = 1/d_i$ if $(i, j) \in E$ and 0 otherwise. Then, $\pi_i \propto d_i$ satisfies detailed balance, so normalizing gives a stationary distribution $\pi_i = d_i / \sum_{k=1}^n d_k$.

Example 25 (Birth–death chains)**.** Consider a birth–death chain on $\{1, 2, \ldots, n\}$ with transition probabilities $P_{i,i+1} = \lambda_i$, $P_{i,i-1} = \mu_i$, and $P_{ii} = 1 - \lambda_i - \mu_i$ for $1 \le i \le n$, with boundary modifications $\mu_1 = 0$ and $\lambda_n = 0$ (so $P_{11} = 1 - \lambda_1$, $P_{n,n} = 1 - \mu_n$). Assume $\lambda_i, \mu_i \ge 0$ and $\lambda_i + \mu_i \le 1$.

Lemma 26 (Stationary distribution of a birth–death chain)**.** *If the chain is irreducible (i.e., $\lambda_i > 0$ and $\mu_{i+1} > 0$ for all $1 \le i \le n-1$), then it has a unique stationary distribution $\underline{\pi}$ given by*

$$\pi_i = \pi_1 \prod_{k=1}^{i-1} \frac{\lambda_k}{\mu_{k+1}}, \qquad \pi_1^{-1} = 1 + \sum_{i=2}^{n} \prod_{k=1}^{i-1} \frac{\lambda_k}{\mu_{k+1}}.$$

Moreover, $\underline{\pi}$ satisfies detailed balance $\pi_i P_{i,i+1} = \pi_{i+1} P_{i+1,i}$.

Proof. Detailed balance implies $\pi_{i+1} = \pi_i \, \lambda_i/\mu_{i+1}$ for $1 \le i \le n-1$. Iterating yields the product form. Normalizing gives $\sum_{i=1}^n \pi_i = 1$ and the stated expression for π_1. Irreducibility ensures positivity and, by Theorem 17, uniqueness. □

Reversibility simplifies spectral analysis because P is self-adjoint with respect to the weighted inner product $\langle f, g\rangle_\pi = \sum_{i=1}^n \pi_i f(i)g(i)$. Consequently, there exists an orthonormal eigenbasis $\{\psi_k\}_{k=1}^n$ in $\ell_2(\pi)$ with real eigenvalues $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n \geq -1$ and $\psi_1(i) \equiv 1$.

Theorem 27 (L2 contraction for reversible chains)**.** *Let P be reversible with stationary distribution π and let $\lambda_* = \max\{|\lambda_2|, |\lambda_n|\} < 1$. Then, for any f with mean $\langle f, 1\rangle_\pi = 0$,*

$$\|P^t f\|_{2,\pi} \leq \lambda_*^t \|f\|_{2,\pi} \qquad (t \geq 0).$$

Proof sketch. Expand $f = \sum_{k=2}^n a_k \psi_k$. Then, orthonormality gives

$$\|P^t f\|_{2,\pi}^2 = \sum_{k=2}^n a_k^2 \lambda_k^{2t} \leq \lambda_*^{2t} \sum_{k=2}^n a_k^2 = \lambda_*^{2t} \|f\|_{2,\pi}^2.$$

□

Test functions. A convenient way to detect convergence is via test functions: for any $f : [n] \to \mathbb{R}$ and initial state i, track $\mathbb{E}[f(X_t) \mid X_0 = i] = e_i^\top P^t f$ and compare it to the stationary value $\mathbb{E}_\pi[f] = \underline{\pi}^\top f$. If $\mathbb{E}[f(X_t) \mid X_0 = i] \to \mathbb{E}_\pi[f]$ for a family of test functions that separates points (e.g., the indicators $\{\mathbb{1}_{\{j\}}\}_{j=1}^n$), then the entire distribution $\Pr(X_t \in \cdot \mid X_0 = i)$ converges to π.

2.4 Ergodic theorem (LLN for Markov chains)

Averaging and covariance. Given $f : [n] \to \mathbb{R}$, the sample average $S_T = \frac{1}{T}\sum_{t=0}^{T-1} f(X_t)$ estimates the stationary mean $\mathbb{E}_\pi[f] = \sum_i \pi_i f(i)$ when the chain is ergodic. If the chain is started in stationarity (i.e, $X_0 \sim \pi$), then for $\tau \geq 0$, we have

$$\mathbb{E}_\pi[f(X_0)f(X_\tau)] = \sum_{i,j\in[n]} f(i)f(j)\,\pi_i[P^\tau]_{i,j} = \langle f, P^\tau f\rangle_\pi$$

and the covariance obeys

$$\operatorname{Cov}_\pi\left(f(X_0), f(X_\tau)\right) = \langle f, P^\tau f\rangle_\pi - \langle f, 1\rangle_\pi^2.$$

For the rainy–sunny example with $P = \left[\begin{smallmatrix} 0.9 & 0.1 \\ 0.5 & 0.5 \end{smallmatrix}\right]$, let $f = \mathbb{1}_{\{\text{rain}\}}$. Since $\pi = (5/6,\ 1/6)$ and $\lambda_2 = 0.4$, one finds $\operatorname{Cov}_\pi(f(X_0), f(X_\tau)) = \pi_1(1-\pi_1)\cdot 0.4^\tau = \frac{5}{36}\cdot 0.4^\tau$.

Theorem 28 (Ergodic theorem)**.** *Let $f : [n] \to \mathbb{R}$ and X_t be irreducible and aperiodic with stationary π. Then, for any X_0, we have*

$$S_T = \frac{1}{T}\sum_{t=0}^{T-1} f(X_t) T \xrightarrow{L^2} \infty \sum_{i=1}^n \pi_i f(i) = \mathbb{E}_\pi[f(X)]. \tag{1}$$

Proof for Stationary Initialization. Assume stationarity (the nonstationary case contributes a negligible term by convergence by exponential convergence to equilibrium). Then, we have

$$\operatorname{Var}\left(\frac{1}{T}\sum_{t=0}^{T-1} f(X_t)\right) = \frac{1}{T^2}\sum_{s,t=0}^{T-1} \operatorname{Cov}_\pi\left(f(X_s), f(X_t)\right) = \frac{1}{T^2}\sum_{\tau=-(T-1)}^{T-1} (T - |\tau|)\,\gamma(\tau),$$

where $\gamma(\tau) = \operatorname{Cov}_\pi(f(X_0), f(X_{|\tau|}))$. Since the chain is aperiodic, $\gamma(\tau)$ decays geometrically (e.g., by the spectral bounds above), hence $\sum_{\tau\geq 0} |\gamma(\tau)| < \infty$ and $\operatorname{Var}(S_T) = O(1/T) \to 0$. Chebyshev's inequality yields $S_T \to \mathbb{E}_\pi[f]$ in L^2 and in probability. □

3 Practical Questions

In practice, the most common questions for Markov chains are:

- How does one draw sample paths from the chain?
- How long will it take for the chain to approach equilibrium?
- What is the correlation between a function of the chain at two points separated by time τ?
- How can one estimate the parameters of a Markov chain from data?

Drawing sample paths. Given a current state x_t, sample x_{t+1} from the categorical distribution given by the x_t-th row of P (e.g., via cumulative sums or np.random.choice). Repeating generates a trajectory; to compute functionals, accumulate $f(x_t)$ along the path.

Time to approach equilibrium. For irreducible aperiodic chains, $\mu P^t \to \pi$ for any initial distribution μ. Quantitatively, the deviation typically decays geometrically at a rate governed by the subdominant eigenvalue modulus $\lambda_* < 1$ (or spectral gap $1-\lambda_*$). Power iteration $\mu^{(t+1)} = \mu^{(t)} P$ provides a practical procedure; stop when $\|\mu^{(t+1)} - \mu^{(t)}\|$ is small.

Autocorrelation of $f(X_t)$. If the chain is started in stationarity,

$$\mathrm{Cov}_\pi(f(X_0), f(X_\tau)) = \underline{\pi}^\top \mathrm{diag}(f)\, P^\tau f - (\underline{\pi}^\top f)^2.$$

The normalized autocorrelation $\rho_f(\tau) = \mathrm{Cov}_\pi(f(X_0), f(X_\tau))/\,\mathrm{Var}_\pi(f)$ often decays like λ_*^τ, and can be estimated empirically from a long run.

Estimating parameters from data. From a trajectory $(X_t)_{t=0}^T$, the maximum-likelihood estimator of P is $\hat{P}_{i,j} = N_{i,j}/N_i$, where $N_{i,j} = \#\{t \in \mathbb{N}_0 \mid X_t = i, X_{t+1} = j\}$ and $N_i = \sum_j N_{i,j}$. For robustness, one can apply Laplace smoothing: $\hat{P}_{i,j} = (N_{i,j} + \alpha)/(\,N_i + \alpha n\,)$. The stationary distribution can be estimated by $\hat{\pi}$ solving $\hat{\pi}\hat{P} = \hat{\pi}$ (or directly via empirical frequencies if the chain is ergodic).

4 Worked Examples

Example 29 (Absorption by linear equations)**.** States $\{1, 2, 3\}$ with target $A = \{3\}$. Let $P = \begin{bmatrix} 0.4 & 0.6 & 0 \\ 0.2 & 0.6 & 0.2 \\ 0 & 0 & 1 \end{bmatrix}$. For hitting probabilities, solve $\phi_{3,3} = 1$ and $\phi_{1,3} = 0.4\,\phi_{1,3} + 0.6\,\phi_{2,3}$, $\phi_{2,3} = 0.2\,\phi_{1,3} + 0.6\,\phi_{2,3} + 0.2\,\phi_{3,3}$, to get $\phi_{1,3} = 3/4$, $\phi_{2,3} = 5/6$. For expected time, use the recursion of Lemma 8: $\eta_{3,3} = 0$, $\eta_{1,3} = 1 + 0.4\,\eta_{1,3} + 0.6\,\eta_{2,3}$, $\eta_{2,3} = 1 + 0.2\,\eta_{1,3} + 0.6\,\eta_{2,3} + 0.2\,\eta_{3,3}$, yielding $\eta_{1,3} = 15/4$, $\eta_{2,3} = 9/2$.

5 Modeling tips and pitfalls

- *State design matters.* The Markov property must hold at the chosen granularity. If it does not, enlarge the state space (e.g., include memory of the last outcome).

- *Absorption vs. reflection.* At boundaries, be explicit: absorb (set $P_{aa} = 1$) or reflect (add mass to P_{ii}).
- *Periodicity.* Periodic chains may fail to converge from certain starts. Use a lazy step ($\tilde{P} = \frac{1}{2}(I + P)$) if needed.
- *Numerics.* For large n, avoid forming powers P^t. Use iterative methods and exploit sparsity.

Exercises

1. (**Classification**) Consider the 6-state chain with transition matrix

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 0.2 & 0.8 & 0 \\ 0.25 & 0 & 0 & 0.35 & 0.4 & 0 \end{bmatrix}.$$

   Partition the states into communicating classes and compute the period of each closed class. Identify any transient states and justify your answer.

2. (**Absorption by linear equations**) Consider a 6-state Markov chain with absorbing states 5 and 6 and transition matrix

$$P = \begin{bmatrix} 0.4 & 0.4 & 0 & 0 & 0.2 & 0 \\ 0.2 & 0 & 0.5 & 0 & 0 & 0.3 \\ 0 & 0.3 & 0 & 0.5 & 0.2 & 0 \\ 0 & 0 & 0.6 & 0 & 0 & 0.4 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

   Let $A = \{5, 6\}$. (a) Form the canonical decomposition with transient block Q (states 1–4) and absorbing block R (transitions from $\{1, 2, 3, 4\}$ into $\{5, 6\}$). Compute the fundamental matrix $N = (I - Q)^{-1}$ and then the absorption-probability matrix $B = NR$. (b) Compute the expected time to absorption $\underline{t} = N\mathbb{1}$ from each transient start. (c) Verify that, for each transient start $i \in \{1, 2, 3, 4\}$, the two absorption probabilities $\phi_{i,5}$ and $\phi_{i,6}$ sum to 1.

3. (**Stationarity**) Compute the stationary distribution for the 8-state "circle" chain where each step moves clockwise by the sum of two independent $\{1, 2, 3, 4\}$ dice.

4. (**Detailed balance**) Show that the random walk on an undirected graph (Example 24) is reversible and find π.

5. (**Hitting probability**) Solve the gambler's ruin hitting probability $\phi_{i,n} = \Pr(H_n < H_0 \mid X_0 = i)$ and expected hitting time $\eta_{i,A}$ to $A = \{1, n\}$ via linear equations; verify the closed forms.

6. (**Fundamental matrix**) For an absorbing chain with transient block Q, prove $N = (I - Q)^{-1} = \sum_{k \geq 0} Q^k$ and interpret N_{ij} probabilistically.

7. (**Absorption probabilities**) Build an absorbing chain for a simple board game with ladders and chutes. Compute $B = NR$ and $\underline{t} = N\mathbb{1}$.

8. (**Time reversal**) For a birth–death chain, compute the time-reversed transition probabilities and verify reversibility.

9. (**Ergodic averages**) Simulate a small irreducible aperiodic chain and empirically verify the ergodic theorem for several test functions f.