4. Discrete-Time Markov Chains (*DTMC*)

A discrete-time stochastic process is a sequence of random variables $\{X_n\}$ indexed on a countable set $\{n\}$. Usually, the random variables X_n are dependent and hence stochastic processes are hard to analyze. Fortunately, we start with Discrete-Time Markov Chains (*DTMC*), a special type of stochastic processes. The simple dependence among X_n leads to nice results under very mild assumptions.

DTMC can be used to model a lot of real-life stochastic phenomena. For example, X_n can be the inventory on-hand of a warehouse at the nth period (which can be in any real time units), the amount of money that a taxi driver gets for his nth trip, or the potential profit of the nth investment opportunity faced by a fund manager. Frequently, we would like to know terms such as $\lim_{m\to\infty} \frac{\sum_{n=1}^{m} X_n}{m}$, which may be the average inventory on hand, the average reward per trip, and the average potential profit from an arriving opportunity. Similarly, nice physical meanings can be assigned to a term such as $\lim_{n\to\infty} P(X_n \geq k)$. It turns out that there are simple ways to calculate such limiting terms for *DTMC*.

4.0. Motivation to Study *DTMC*

The *Strong Law of Large Numbers* (*SLLN*) is one of the most widely applicable results to analyze the convergence of *i.i.d.* random variable $\{X_n\}$. Let $E(X) < \infty$. The *SLLN* says that $\frac{\sum_{i=1}^{n} X_i}{n} \to E(X)$ (*w.p.* 1).

<u>Example 4.0.1</u>. *SLLN*

(a). Suppose that one throws a fair dice repeatedly for n times, where n is “large”. Then, by the *SLLN*, the average value of all throws is close to 3.5, the expected value of a throw.

(b). To study the probability of an unfair dice, one can throw the dice for n times, where n is large. Then the probability of getting the face i in a throw is roughly equal to the fraction of times that face i appears. ✧✧

While *SLLN* gives nice results of *i.i.d.* random variables (under some other conditions, e.g., finite first moment), many random variables in real-life problems are in fact *dependent*.

<u>Example 4.0.2</u>. Sequence of Dependent Random Variables

(a). John does a part-time job to support his study. He gets \$700 every Sunday morning, and he uses up \$100 every day. Let X_n = the amount of money (in \$100) that John has at the morning of the nth day; the first day is arbitrarily set at March 11, 2002, Sunday. Then,

$$X_1 = 7; \quad X_{n+1} = \begin{cases} 7, & \text{if } n \bmod 7 = 0, \\ X_n - 1, & o.w., \end{cases}$$

and $\{X_n\}$ is a sequence of dependent (random) variables.

(b). A shop keeps at most 4 pieces of the brand *Accurate* watch. On every Sunday evening, the shopkeeper counts the number of Accurate watches that the shop has, and orders up to an inventory position of 4 watches if there is one or no Accurate watch in the shop. Any quantity ordered will be sent to the shop on Monday right before the shop is opened for business. Any customers who fail to find an Accurate watch will go away without coming back. Let D_n be the demand of Accurate watches in the nth week, and Y_n be the inventory position of Accurate watches in shop on the nth Sunday evening before any new order is made. Then

$$Y_{n+1} = \begin{cases} \max\{Y_n - D_{n+1}, 0\}, & \text{if } Y_n \geq 2, \\ \max\{4 - D_{n+1}, 0\}, & \text{if } Y_n = 0, 1. \end{cases}$$

Suppose that D_n are *i.i.d.* random variables. Then the distribution of Y_{n+1} is determined (completely) by Y_n and $\{Y_n\}$ is a sequence of dependent variables.

(c). Initially there are 5 red balls in Box 1 and 5 white balls in Box 2. At each time period, a ball is randomly selected from each of the two boxes and put into the other box. Let Z_n be the number of red balls in Box 1 after the nth switch; $Z_0 = 5$.

$$Z_{n+1} = \begin{cases} Z_n - 1, & \text{if the ball drawn from Box 1 is red while from Box 2 is white,} \\ Z_n, & \text{if the balls drawn are of the same colour,} \\ Z_n + 1, & \text{if the ball drawn from Box 1 is white while from Box 2 is red.} \end{cases}$$

The distribution of Z_{n+1} is determined completely by Z_n and $\{ Z_n \}$ is a sequence of dependent variables.

Given $Z_n = k$, the probabilities of the events {the ball drawn from Box 1 is red while from Box 2 is white}, {the balls drawn are of the same color}, and {the ball drawn from Box 1 is white while from Box 2 is red} can be completely determined. ✧✧

$\{X_n\}$, $\{Y_n\}$, and $\{Z_n\}$ in Example 4.0.2 are dependent variables. Will the mean of their partial sum converges to something? Clearly, for the simple case (a), $\lim_{m\to\infty} \frac{\sum_{n=1}^{m} X_n}{m}$ does exist. However, the situation is not clear for $\lim_{m\to\infty} \frac{\sum_{n=1}^{m} Y_n}{m}$ and $\lim_{m\to\infty} \frac{\sum_{n=1}^{m} Z_n}{m}$.

The above problem leads to an interesting extension: Suppose that $\{X_n\}$ is a sequence of dependent random variables modeling some real-life problems. Is it possible to deduce any long-term average such as $\lim_{m \to \infty} \frac{\sum_{n=1}^{m} X_n}{m}$? Can we say anything about $\lim_{n \to \infty} P(X_n = k)$, the limiting distribution of X_n? For a general system such that one gets a reward c_j when $X_n = j$, can we say anything about $\lim_{m \to \infty} \frac{\sum_{n=1}^{m} c_{X_n}}{m}$? If expected values exist, are they easy to find? Our interest is more than the limiting distributions. Suppose that there are three watches in the shop on the third Sunday, what is the probability that the shop needs to re-order watches on the 8*th* week? If the weather can be modeled by a *DTMC* and the weather now is cloudy, what is the probability that it will rain 10 days later? All these are interesting, useful, and challenging questions. It turns out that many of these questions have easy answers if $\{X_n\}$ is a *DTMC* (with minor structural conditions).

4.1. States and Time of Markov Chains

DTMC are sequences of *dependent* random variables $\{X_n\}$ with nice conditional probability structure. They have been used widely to model real-life problems, and they are building blocks for more complex, advanced applications.

Obviously the process $\{X_n\}$ is of discrete-time: the index n takes up discrete values. In this course, we take $n \in \{0, 1, 2, \dots\}$.

The *state* of $\{X_n\}$, S, is the set of all possible values (or status) to be taken up by the process $\{X_n\}$. Usually, we take the state to be the whole set or a subset of non-negative integers. In fact, for many cases, we can assign physical meanings to each state. For example, in modeling the condition of a machine that has four states, states 0, 1, 2, and 3 can mean, respectively, "break down", "idle", "under repair", and "working", and S can either be taken as {0, 1, 2, 3} or {"break down", "idle", "under repair", "working"}; the latter can also be shortened to $\{b, i, r, w\}$.

4.2. Markov Property

A *DTMC* $\{X_n\}$ possesses the *Markov property*, i.e.,

$$P(X_{n+1} = j \mid X_n, X_{n-1}, \dots, X_1) = P(X_{n+1} = j \mid X_n), \text{ for any } n \geq 0 \text{ and for any } j \in S.$$

This random-variable version is equivalent to

$$P(X_{n+1} = j \mid X_n = i, X_{n-1}, \dots, X_1) = P(X_{n+1} = j \mid X_n = i), \text{ for any } n \geq 0 \text{ and for any } i, j \in S.$$

This property shows that $\{X_n\}$ has very peculiar dependence: given the current state X_n, the distribution of the future states of the *DTMC* is completely determined and is independent of the past $(X_{n-1}, \dots, X_1)$. [More technically, given the current state (X_n), the future $(\{X_{n+1}, \dots\})$ is conditionally independent of the past $(\{X_1, \dots, X_{n-1}\})$.]

4.3. Stationary Transition

We only consider $\{X_n\}$ with *stationary* transition, i.e., $P(X_{n+1} = j \mid X_n) = P(X_1 = j \mid X_0)$ for all j and n. With this assumption, the analysis of *DTMC* becomes very simple.

4.4. Defining a *DTMC*

The procedure to define a *DTMC* is to:

(a) specify the states,
(b) demonstrate the Markov property, and
(c) find the (stationary) probability transition matrix, which is in the form

$$\mathbf{P} = \begin{pmatrix} P(X_1 = 0 \mid X_0 = 0) & P(X_1 = 1 \mid X_0 = 0) & \cdots & \cdots \\ P(X_1 = 0 \mid X_0 = 1) & P(X_1 = 1 \mid X_0 = 1) & \cdots & \cdots \\ \vdots & \vdots & \ddots & \vdots \\ \vdots & \vdots & \cdots & \ddots \end{pmatrix}.$$

Remark 4.4.1. Each row of $\mathbf{P}$ specifies a conditional distribution, e.g, the first row is the conditional distribution of $(X_1 \mid X_0 = 0)$, and the second row is the conditional distribution of $(X_1 \mid X_0 = 1)$. Clearly, each row of $\mathbf{P}$ sums up to 1.

Example 4.4.2.

(a). *i.i.d.* random variables

Let $X_n = 1$ if the nth flip is a head, and $X_n = 0$ otherwise; $n = 0, 1, 2, \ldots$. Suppose that the coin gives a head with probability p, $0 < p < 1$. Is $\{X_n\}$ a *DTMC*?

Sol. Let the state space be $\{0, 1\}$. We need to
(i) show that the Markov property holds;
(ii) show that the transitions are stationary;
(iii) define the (stationary) transition probability matrix.

By the independence of the flips,

$$P(X_{n+1} = 1 \mid X_n, X_{n-1}, \ldots, X_1) = P(X_{n+1} = 1) = P(X_{n+1} = 1 \mid X_n) = p,$$
$$P(X_{n+1} = 0 \mid X_n, X_{n-1}, \ldots, X_1) = P(X_{n+1} = 0) = P(X_{n+1} = 0 \mid X_n) = 1-p,$$

which complete (i)–(iii) simultaneously. The transition probability matrix is

$$\mathbf{P} = \begin{pmatrix} 1-p & p \\ 1-p & p \end{pmatrix}.$$

The transition matrix can easily be represented by a transition diagram:

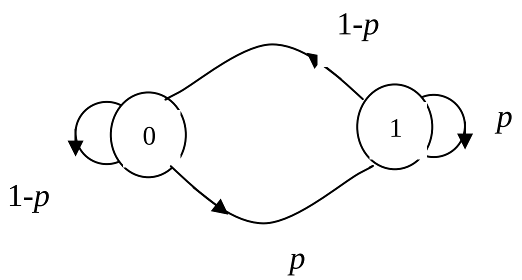

(b). Two Gamblers

Peter and Sam bet against each other. Each time one of them flips a coin, which gives a head with probability p, $0 < p < 1$. If the coin lands head, Peter wins one dollar from Sam; otherwise Sam wins a dollar from Peter. Initially Peter has $\$c_1$ (> 0) and Sam $\$c_2$ (> 0); the game ends when one of them is out of money. Given that all flips are independent from each other, can this be modeled as a *DTMC*?

Sol. Define $c = c_1+c_2$. Let X_n be the amount of money that Peter has after the nth flip; $X_n \in \{0, 1, \ldots, c\}$; $X_0 = c_1$.

$$X_{n+1} = \begin{cases} c, & \text{if } X_n = c, \\ X_n + 1, & \text{if } 0 < X_n < c \text{ and Peter wins in the } n\text{th game}, \\ X_n - 1, & \text{if } 0 < X_n < c \text{ and Peter losses in the } n\text{th game}, \\ 0, & \text{if } X_n = 0. \end{cases}$$

Clearly, given the amount of money that Peter has after the nth flip, the $(n+1)$st flip completely determines the amount of money that Peter owns after the $(n+1)$st flip; the history of prior flips and of prior amount of money owned by Peter play no part. Thus, the process has Markov property. The independence of flips shows that the process has stationary transition. The question is to find the transition probability matrix.

Given that $X_n = c$, the game ends, and $X_{n+1} = X_{n+2} = \ldots = c$. Similarly, given $X_n = 0$, the game ends, and $X_{n+1} = X_{n+2} = \ldots = 0$. For $X_n = i$, $1 \le i \le c-1$, $P(X_{n+1} = i+1 \mid X_n = i) = p$, $P(X_{n+1} = i-1 \mid X_n = i) = 1-p$. The probability transition matrix is:

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & \cdots & \cdots & 0 \\ 1-p & 0 & p & \ddots & \ddots & \vdots \\ 0 & 1-p & 0 & p & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \vdots & 1-p & 0 & p \\ 0 & \cdots & \cdots & \cdots & 0 & 1 \end{pmatrix}.$$

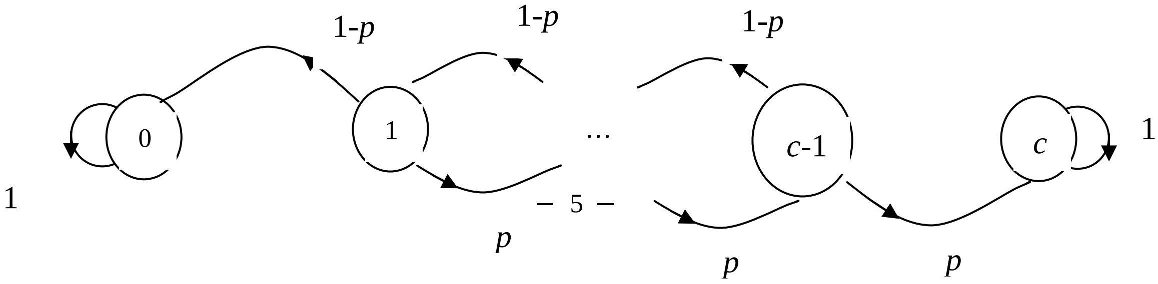

Can we deduce the probability that Peter eventually wins or deduce the expected length of the game before it stops?

(c). Inventory [Cont'n of Example 4.0.2(b)]

Is $\{Y_n\}$ in Example 4.0.2(b) a *DTMC*?

Sol. The sample space of $\{Y_n\}$ is dependent on that of D_n. For example, if D_n can only be 1 or 0, then the state space of $\{Y_n\}$ is {1, 2, 3, 4}; if D_n can only be zero or two, the state space is {0, 2, 4} if $Y_0 \in \{0, 2, 4\}$, and is {0, 1, 2, 4} or {0, 1, 2, 3, 4} dependent on whether Y_0 = 1 or Y_0 = 3.

Here, we consider the general case in which D_n can take up any non-negative integer values. The state space of $\{Y_n\}$ will then be {0, 1, 2, 3, 4}. From the problem statement, we know that

$$Y_{n+1} = \begin{cases} \max\{Y_n - D_{n+1}, 0\}, & \text{if } Y_n \geq 2, \\ \max\{4 - D_{n+1}, 0\}, & \text{if } Y_n = 0 \text{ or } 1. \end{cases}$$

Such an expression shows that given Y_n the distribution of Y_{n+1} does not depend on $\{Y_0, \ldots, Y_n\}$ and hence $\{Y_n\}$ is a Markov process. It has stationary distribution since D_n's are *i.i.d.* To find probability transition function, let $P(D_n = i) = \alpha_i$, $i = 0, 1, 2, \ldots$. Then

$$\mathbf{P} = \begin{pmatrix} \sum_{i=4}^{\infty}\alpha_i & \alpha_3 & \alpha_2 & \alpha_1 & \alpha_0 \\ \sum_{i=4}^{\infty}\alpha_i & \alpha_3 & \alpha_2 & \alpha_1 & \alpha_0 \\ \sum_{i=2}^{\infty}\alpha_i & \alpha_1 & \alpha_0 & 0 & 0 \\ \sum_{i=3}^{\infty}\alpha_i & \alpha_2 & \alpha_1 & \alpha_0 & 0 \\ \sum_{i=4}^{\infty}\alpha_i & \alpha_3 & \alpha_2 & \alpha_1 & \alpha_0 \end{pmatrix}.$$

You may draw the transition diagram by yourselves.

(d). Weather

The weather tomorrow is dependent on today's weather. If today is rainy, half of the times that tomorrow will also be rainy; 40% will be cloudy; 10% will be sunny. If today is cloudy, 20% of the times that tomorrow will be rainy; 40% will be cloudy; 40% will be sunny. If today is sunny, 10% of the times that tomorrow will be rainy; 20% will be cloudy; 70% will be sunny. Can this be modeled as a *DTMC*?

Sol. The state space of a *DTMC* needs not be integer. It can be any discrete quantity so long as physically such a construction makes sense. In this case, let X_n be the random quantity denoting the weather of the *n*th day; $X_n \in$ {rainy, cloudy, sunny}. The problem statement says that the weather tomorrow just depends on today's weather, which is Markov. The statement also shows that the transition probability is stationary.

$$\mathbf{P} = \begin{pmatrix} 0.5 & 0.4 & 0.1 \\ 0.2 & 0.4 & 0.4 \\ 0.1 & 0.2 & 0.7 \end{pmatrix}.$$ ✧✧

<u>Exercise 4.4.3</u>. Show that the processes in (a) and (c) of Example 4.0.2 are indeed *DTMC*. Find their transition probability matrices.

4.5. Chapman-Kolmogorov equations

Let $p_{ij}^{(n)} = P(X_n = j \mid X_0 = i)$ be probability that $\{X_n\}$ moves from state i to state j in n steps. The Chapman-Kolmogorov equations are recursive relation relating $\left\{p_{ij}^{(n)}\right\}$ to $\left\{p_{ij}^{(n-1)}\right\}$, and hence one can find $p_{ij}^{(n)}$ from the one-step transition probabilities $p_{ij}^{(1)} = p_{ij} = P(X_1 = j|X_0 = i)$.

$$\begin{aligned}
p_{ij}^{(n)} &= P(X_n = j|X_0 = i) = \sum_{k=0}^{\infty} P(X_n = j, X_1 = k \mid X_0 = i) \\
&= \sum_{k=0}^{\infty} P(X_n = j, X_1 = k \mid X_0 = i) = \sum_{k=0}^{\infty} \frac{P(X_n = j, X_1 = k, X_0 = i)}{P(X_0 = i)} \\
&= \sum_{k=0}^{\infty} \frac{P(X_n = j|X_1 = k, X_0 = i)P(X_1 = k, X_0 = i)}{P(X_0 = i)} \\
&= \sum_{k=0}^{\infty} P(X_n = j \mid X_1 = k, X_0 = i)P(X_1 = k \mid X_0 = i) \\
&= \sum_{k=0}^{\infty} P(X_n = j \mid X_1 = k)P(X_1 = k \mid X_0 = i) \\
&= \sum_{k=0}^{\infty} P(X_{n-1} = j \mid X_0 = k)P(X_1 = k \mid X_0 = i) = \sum_{k=0}^{\infty} p_{ik}\, p_{kj}^{(n-1)}.
\end{aligned}$$

Similarly, one can show that $p_{ij}^{(n)} = \sum_{k=0}^{\infty} p_{ik}^{(n-1)} p_{kj}$.

From the expression of $p_{ij}^{(2)}$, one can see that $\mathbf{P}^{(2)} = \mathbf{P}^2$; in general, we have $\mathbf{P}^{(n)} = \mathbf{P}^n$.

Example 4.5.1. [Cont'n of Example 4.4.2 (d)] *n*-step transition probability

Suppose that the weather across days is a *DTMC* following the transition probability in Example 4.4.2 (d). Given that today is cloudy, find

(a) the probability that it is sunny two days later;
(b) the probability that it is sunny ten days later.

Sol. To simplify notation, we use *r*, *c*, and *s* to represent the states, where *r* stands for rainy; *c* for cloudy; *s* for sunny. We are asked to find $p_{cs}^{(2)} = P(X_2 = s|X_0 = c)$. We may apply Chapman-Kolmogorov equation to get the answer. Since the problem is simple, we solve it from the first principle.

$$
\begin{aligned}
& P(X_2 = s|X_0 = c) = P(X_2 = s \text{ and } \{X_1 = r \text{ or } X_1 = c \text{ or } X_1 = s\}|X_0 = c) \\
= \quad & P(X_2 = s, X_1 = r|X_0 = c) + P(X_2 = s, X_1 = c|X_0 = c) + P(X_2 = s, X_1 = s|X_0 = c)
\end{aligned}
$$

Let us digress a bit to see a generic way to simplify $P(A \cap B|C)$ for events A, B, and C. $P(A\cap B|C) = \frac{P(A\cap B\cap C)}{P(C)} = \frac{P(A\cap B\cap C)}{P(B\cap C)}\frac{P(B\cap C)}{P(C)} = P(A|B\cap C)P(B|C)$. Applying this result to

$$
\begin{aligned}
& P(X_2 = s, X_1 = r|X_0 = c) \\
= \quad & P(X_2 = s|X_1 = r, X_0 = c)\, P(X_1 = r|X_0 = c) \\
= \quad & P(X_2 = s|X_1 = r)\, P(X_1 = r|X_0 = c) && \text{(Markov Property)} \\
= \quad & P(X_1 = s|X_0 = r)\, P(X_1 = r|X_0 = c) && \text{(stationary transition)} \\
= \quad & p_{cr}p_{rs}.
\end{aligned}
$$

$$
\begin{aligned}
& P(X_2 = s|X_0 = c) \\
= \quad & p_{cr}p_{rs} + p_{cc}p_{cs} + p_{cs}p_{ss} = (0.2)(0.1) + (0.4)(0.4) + (0.4)(0.7) = 0.46.
\end{aligned}
$$

(b). We are asked to find $p_{cs}^{(10)}$. One may apply Chapman-Kolmogorov equation recursively to get $p_{\cdot,\cdot}^{(2)}$, $p_{\cdot,\cdot}^{(3)}$, .., till $p_{cs}^{(10)}$. However, a neater way is to find

$$
\mathbf{P}^{(10)} = \mathbf{P}^{10} = \mathbf{P}^2\mathbf{P}^8 = (\mathbf{P}\cdot\mathbf{P})(\mathbf{P}\cdot\mathbf{P})^4 = \begin{pmatrix} 0.34 & 0.38 & 0.28 \\ 0.22 & 0.32 & 0.46 \\ 0.16 & 0.26 & 0.58 \end{pmatrix} \begin{pmatrix} 0.34 & 0.38 & 0.28 \\ 0.22 & 0.32 & 0.46 \\ 0.16 & 0.26 & 0.58 \end{pmatrix}^4
$$

$$
= \begin{pmatrix} 0.34 & 0.38 & 0.28 \\ 0.22 & 0.32 & 0.46 \\ 0.16 & 0.26 & 0.58 \end{pmatrix} \begin{pmatrix} 0.2187 & 0.3053 & 0.4760 \\ 0.2175 & 0.3044 & 0.4781 \\ 0.2167 & 0.3039 & 0.4794 \end{pmatrix} = \begin{pmatrix} 0.2177 & 0.3046 & 0.4777 \\ 0.2174 & 0.3044 & 0.4782 \\ 0.2172 & 0.3042 & 0.4785 \end{pmatrix}.
$$

Hence, $p_{cs}^{(10)} = 0.4782$. ✧✧

4.6. Long-Term Averages and Limiting Distribution

For *i.i.d.* random variables X_k, under very mild conditions (e.g., $E(X) < \infty$), $\lim_{n\to\infty} \frac{\sum_{k=1}^{n} X_k}{n} = E(X)$. We would like to know whether a similar result hold for *DTMC*, and if it indeed exists, we would like to find a quick way to calculate it. Such results are useful in daily operations. For example, if X_n is the on-hand inventory of the n time unit, $\lim_{n\to\infty} \frac{\sum_{k=1}^{n} X_k}{n}$ will be the long-run average inventory (of the chosen inventory policy), which can be used to determine the long-run average inventory cost. In the following, we first argue heuristically the value of such a limit, if it exists. Then we stay the theorem formally to justify the results.

Suppose that for any j, $\pi_j = \lim_{n\to\infty} P(X_n = j \mid X_0)$, i.e., the limiting probability of state j exists, and *is independent of the initial state*. Naturally, $\Sigma_j \pi_j = 1$. The assumption *suggests* the following results.

(a) Intuitively, this result says that for "large" n, X_n's behave as if they are *i.i.d*. So π_j should also be the fraction of time the chain in state j.

(b) When n is large, the proportion of time in (periods in, visits to) state j (roughly) π_j. The number of times (periods, visits) that the chain is in state $j \approx \pi_j n$, each time contributing a value of j to $\Sigma_j X_j$. Then $\frac{\sum_{k=1}^{n} X_k}{n} \approx \frac{\sum_j j\pi_j n}{n} = \sum_j j\pi_j$ for large n.

(c) Suppose that a cost of c_j is incurred for each time visit of state i. By the same argument as in (b), the long-term average reward $\approx \sum_j c_j \pi_j$.

(d) (*An intuitive argument to get* π_j) This is established by a conservation law: the rate (number of occurrences per unit time) into any state must be equal to the rate out of the state. Consider again "large" n. The *DTMC* (roughly) visits state j for $\pi_j n$ times (periods), which is also the number of departures from state j. Thus, the rate of departure from state $j \approx \pi_j$ (which, later, is shown to be exact). Now consider the rate into state j. The *DTMC* visits state i for $\pi_i n$ times. For each visit of state i, it will next visit state j with probability p_{ij}. Hence for large n, the number of i to j transitions $\approx \pi_i p_{ij} n$. The total number of visits to state j $\approx \sum_i \pi_i p_{ij} n$, and the rate of into state $j \approx \sum_i \pi_i p_{ij}$ (again, this rate is actually exact). Equating the into and out of rates,

$$\pi_j = \sum_i \pi_i p_{ij}, \qquad \text{for any } j. \qquad \text{(balance equation)}$$

For an M-state *DTMC*, there are M equations, one for each state, and M unknowns. However, the equations are linearly dependent and there are only M-1 degree of freedom. Fortunately, there is one more equation

$$\sum_j \pi_j = 1. \qquad \text{(normalization equation)}$$

The above results are nice. However, it is under the assumption that for any j, $\pi_j = \lim_{n\to\infty} P(X_n = j \mid X_0)$. Moreover, even if the limiting probabilities actually exist, we still need to justify our heuristic argument.

Does $\pi_j = \lim_{n\to\infty} P(X_n = j \mid X_0)$ for any j? Arbitrarily consider $\mathbf{P} = \begin{pmatrix} 0.4 & 0.5 & 0.1 \\ 0.1 & 0.2 & 0.7 \\ 0.01 & 0 & 0.99 \end{pmatrix}$. We find that $\mathbf{P}^{(5)} = \begin{pmatrix} 0.0462 & 0.0543 & 0.8995 \\ 0.0254 & 0.0224 & 0.9522 \\ 0.0175 & 0.0104 & 0.9722 \end{pmatrix}$, $\mathbf{P}^{(10)} = \begin{pmatrix} 0.0192 & 0.013 & 0.9678 \\ 0.0184 & 0.0117 & 0.9699 \\ 0.0180 & 0.0112 & 0.9707 \end{pmatrix}$, $\mathbf{P}^{(20)} = \begin{pmatrix} 0.0181 & 0.0113 & 0.9706 \\ 0.0181 & 0.0113 & 0.9707 \\ 0.0181 & 0.0113 & 0.9707 \end{pmatrix}$, and $\mathbf{P}^{(25)} = \begin{pmatrix} 0.0181 & 0.0113 & 0.9707 \\ 0.0181 & 0.0113 & 0.9707 \\ 0.0181 & 0.0113 & 0.9707 \end{pmatrix}$; for $n \geq 25$, $\mathbf{P}^{(n)}$ becomes a matrix such that all entities within a column have the same value. This shows that indeed $\lim_{n\to\infty} P(X_n = i|X_0)$ exists for this $\mathbf{P}$ (i.e., $\mathbf{P}^{(\infty)}$ exists). Based on our heuristic argument, $\pi_0 = 0.0181$, $\pi_1 = 0.0113$, and $\pi_2 = 0.9707$; $\lim_{n\to\infty} \frac{\sum_{k=1}^{n} X_k}{n} = \pi_1 + 2\pi_2 = 1.9527$. If a cost of c_i, $i = 0, 1, 2$, is incurred for each time that state i is visited, then the average reward will be $\pi_0 c_0 + \pi_1 c_1 + \pi_2 c_2 = 0.0181c_0 + 0.0113c_1 + 0.9707c_2$.

Everything looks nice, doesn't it?

Example 4.6.1.

(a). ($\mathbf{P}^{(\infty)}$ does not exist.) A guard takes alternate weekly day and night shifts. The pattern goes on forever. Let $X_n = 1$ if the guard has day shift in the nth week; $X_n = 0$, otherwise. It readily seen that $\{X_n\}$ is a *DTMC* with the transition probability matrix $\mathbf{P} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. For this $\mathbf{P}$, $\mathbf{P}^n = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ when n is even, and $\mathbf{P}^n = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ when n is odd. So $\lim_{n\to\infty} P(X_n = j \mid X_0)$ does not exist for all j.

In retrospect, the above result is not surprising at all. After all, the change of states follows a fixed pattern, 1 (day shift), 0 (night shift), 1 (day shift), 0 (night shift), … forever and hence $\lim_{n\to\infty} P(X_n = i|X_0)$ cannot exist.

(b). [Cont'n of Example 4.2.2 (b); $\mathbf{P}^{(\infty)}$ exists but the rows are different.] Suppose that Peter and Sam totally has 3 dollars and they bet with a fair coin. Then the transition probability matrix $\mathbf{P}$ is $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 1 \end{pmatrix}$. Then $\mathbf{P}^{(\infty)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 2/3 & 0 & 0 & 1/3 \\ 1/3 & 0 & 0 & 2/3 \\ 0 & 0 & 0 & 1 \end{pmatrix}$. Note that unlike the previous case, while the limit exists, now the rows of $\mathbf{P}^{\infty}$ are *different*. ✧✧

The example shows that there are loopholes in our heuristics argument. Here we will give simple conditions such that the nice results (a)-(d) and the existence of limiting probability indeed exist. It turns out that it is related to the *connectivity* and the *periodicity* of the states (chains).

State j is *accessible* from state i (state i can reach state j), $i \to j$, if $p_{ij}^{(n)} > 0$ for some n. Graphically, it just means that in the transition diagram, there is a directed path (sequence of transitions with non-zero probabilities) starting from state i and ending at state j with n line segments (n transitions). States i and j communicate with each other, denoted by $i \leftrightarrow j$, if $i \to j$ and $j \to i$. If all states of a *DTMC* communicate with each other, the chain is said to be *irreducible*. Mathematically, it means that there is an n (> 0) such that $p_{ij}^{(n)} > 0$ for all i, j. The simplest way to check whether a small *DTMC* is irreducible or not is by drawing out its transition diagram.

Exercise 4.6.2. Connectivity of *DTMC*

Which of the following *DTMC*'s are irreducible? The missing numbers of the transition probability matrices are all zero.

(a). $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ (b). $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ (c). $\begin{bmatrix} 0.2 & 0.8 \\ 1 & 0 \end{bmatrix}$

(d). $\begin{bmatrix} 0.3 & 0.7 & & & \\ & 0.1 & & & 0.9 \\ & & & 1 & \\ & & 0.6 & 0.4 & \\ 0.2 & 0.8 & & & \end{bmatrix}$ (e). $\begin{bmatrix} 0.3 & 0.7 & & & \\ & 0.1 & & & 0.9 \\ & & & 1 & \\ 0.5 & & 0.1 & 0.4 & \\ 0.2 & 0.8 & & & \end{bmatrix}$ (f). $\begin{bmatrix} 0.3 & 0.7 & & & \\ & 0.1 & & & 0.9 \\ & & & 1 & \\ 0.5 & & 0.1 & 0.4 & \\ 0.2 & 0.7 & & 0.1 & \end{bmatrix}$

✧✧

State i is of *period* d if d = HCF$\{n > 0 \mid p_{ii}^{(n)} > 0\}$. A state is *aperiodic* if $d = 1$.

Fact 4.6.4. States in an irreducible chain have the same period. An irreducible chain is aperiodic if d (of any state) = 1. ✧✧

Example 4.6.5. Periodicity of Chains

Find the periodicity of the following chains.

(a).

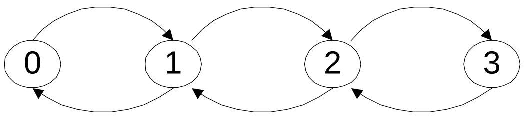

(b).

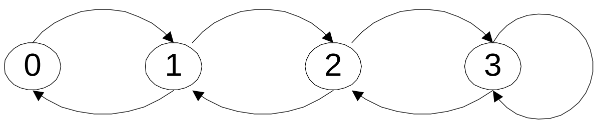

(c).

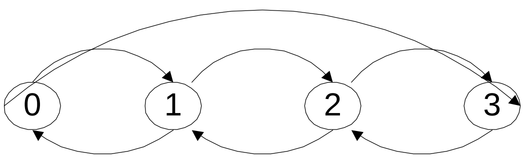

(d).

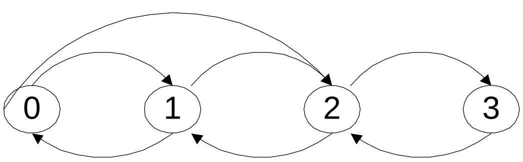

(e).

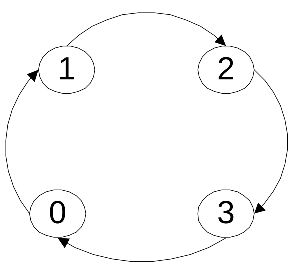

(f).

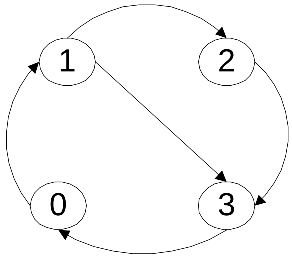

✧✧

A state can be classified according to how the chain returns to it. State i is *recurrent* if P(return to state $i|X_0 = i) = 1$; else it is transient. A recurrent state is *positive* (recurrent) if E(return to state $i\ |X_0 = i) < \infty$, and is *null* recurrent if E(return to state $i\ |X_0 = i) = \infty$. States that communicate with each other are of the same type.

One of the most important results in the theory of *DTMC* is that π, in terms of fraction of time, indeed exists for positive, irreducible chains.

Theorem 4.6.6. An irreducible *DTMC* $\{X_n\}$ is positive *iff* there exists unique nonnegative solution to

$$\sum_{j=0}^{\infty} \pi_j = 1, \qquad \text{(normalization equation)}$$

and $$\pi_j = \sum_{i=0}^{\infty} \pi_i\, p_{ij}, \text{ for all } j. \qquad \text{(balance equations)}$$

$\{\pi_j\}$ are called stationary (steady-state) distribution of $\{X_n\}$. ✧✧

Exercise 4.6.7. Let π be a nonnegative probability distribution that satisfies balance equations of a positive irreducible chain. Show that π is the stationary distribution, i.e., if $X_0 \sim \pi$, then $X_n \sim \pi$. ✧✧

Remarks 4.6.8. The stationary distribution π of an irreducible positive *DTMC* has some nice properties:

(a). π_j is also the long-run proportion of time that state j is visited, i.e.,

$$\pi_j = \lim_{n\to\infty} \frac{\sum_{k=1}^{n} 1_{\{X_k = j | X_0 = i\}}}{n}.$$

(b). A third interpretation of π_j is that it is the long-run proportion of *expected* time that state j is visited, i.e.,

$$\pi_j = \lim_{n\to\infty} \frac{\sum_{k=1}^{n} E[1_{\{X_k=j\}} \mid X_0 = i]}{n} = \lim_{n\to\infty} \frac{\sum_{k=1}^{n} p_{ij}^{(k)}}{n}.$$

(c). If a cost c_j is incurred whenever state j is visited, then the long-run average cost is given by $\lim_{n\to\infty} \frac{\sum_{k=1}^{n} c_{X_k}}{n} = \sum_{i=0}^{\infty} \pi_j c_j$. The long-run average expected cost $\lim_{n\to\infty} \frac{E\left[\sum_{k=1}^{n} c_{X_k}\right]}{n}$ is also given by $\sum_{j=0}^{\infty} \pi_j c_j$.

(d). If a random cost C_j is incurred whenever state j is visited, then the long-run average cost $\lim_{n\to\infty} \frac{\sum_{k=1}^{n} C_{X_k}}{n} = \sum_{j=0}^{\infty} \pi_j E(C_j)$. Similarly, the long-run average expected cost $\lim_{n\to\infty} \frac{E\left[\sum_{k=1}^{n} C_{X_k}\right]}{n}$ is also given by $\sum_{j=0}^{\infty} \pi_j E(C_j)$.

(e). If the chain is aperiodic, $\lim_{n\to\infty} P(X_n = i \mid X_0) = \pi_i$. ✧✧

<u>Example 4.6.9</u>. [Cont'n of Example 4.4.2(d)]

Denote the rainy, cloudy, and sunny state of the chain in Example 4.4.2(d) by r, c, and s. Let π_r, π_c, and π_s be the fraction of time that the chain is in state r, c, and s, respectively. Then the balance equations give

$\pi_r + \pi_c + \pi_s = 1$,
$0.5\pi_r + 0.2\pi_c + 0.1\pi_s = \pi_r$,
and $0.4\pi_r + 0.4\pi_c + 0.2\pi_s = \pi_s$.

Solving, $\pi_r = 5/23$, $\pi_c = 7/23$, and $\pi_s = 11/23$.

(a) Fraction of rainy days = 5/23; fraction of cloudy days = 7/23; fraction of sunny days = 11/23.

(b) A subcontractor can earn \$100 in a sunny day, \$30 in a cloudy day, and -\$10 in a rainy day. The long-term average money earned per day = $100\pi_r + 30\pi_c - 10\pi_s = \frac{600}{23}$.

(c) Instead of (b), suppose that the amount earned by the subcontractor in a sunny, cloudy, and rainy day distributes, respectively, as unif[50, 150], exp(1/30), and –10 dollars. Then the long-term average money earned per day is still $100\pi_r + 30\pi_c - 10\pi_s = \frac{600}{23}$.

(d) Suppose that a penalty cost \$10 is induced every time the weather changes from sunny to rainy. To find the long-run average penalty cost due to such a weather change, re-define

$$C_s = \begin{cases} 10, & \text{if } X_{n+1} = r \text{ given that } X_n = s, \\ 0 & o.w., \end{cases} \quad \text{and } C_r = C_c = 0.$$

From (d) of Remark 4.6.8, the long-run average penalty cost = $10\pi_s p_{sr} = \pi_s$. ✧✧

Exercise 4.6.10. (Exercise 14.6-8 of Hillier and Lieberman)

A production process contains a machine that deteriorates rapidly in both quality and output under heavy use, so that it is inspected at the end of each day. Immediately after inspection, the condition of the machine is noted and classified into one of the four possible states:

State	Condition
0	Good as new
1	Operable – minimum deterioration
2	Operable – major deterioration
3	Inoperable and replaced by a good-as-new machine

The process can be modeled as a Markov chain with its (one-step) transition matrix **P** given by

$$\begin{pmatrix} 0 & \frac{7}{8} & \frac{1}{16} & \frac{1}{16} \\ 0 & \frac{3}{4} & \frac{1}{8} & \frac{1}{8} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

(a). Find the steady-state probabilities.
(b). If the costs of being in states 0, 1, 2, 3 are 0, \$1,000, \$3,000, and \$6,000, respectively, what is the long-run expected average cost per day? ✧✧

4.7. First Passage Times

The *first passage time* in going from state i to state j, T_{ij}, is the number of transitions taken to visit state j for the first time given that $X_0 = i$, i.e., T_{ij} is the n such that $\{X_n = j, X_{n-1} \neq j, \ldots, X_1 \neq j | X_0 = i\}$. T_{ii} is called the *recurrence time* for state i. Let $\mu_{ij} = E[T_{ij}]$. There is a simple way to calculate μ_{ij} for a *fixed* j if $P(\text{visit } j \mid X_0 = i) = 1$ for all i. In that case, μ_{ij}'s satisfy the set of equations

$$\mu_{ij} = 1 + \sum_{k \neq j} p_{ik}\mu_{kj}; \text{ and for } i \neq j,$$

and $\mu_{jj} = 1 + \sum_{k \neq j} p_{jk} \mu_{kj}$.

For a positive recurrent chain, μ_{jj} can be found easily from $\pi_j = \frac{1}{\mu_{jj}}$ (why?).

Example 4.7.1. Consider the transition probability matrix

$$\mathbf{P} = \begin{pmatrix} 0.08 & 0.184 & 0.368 & 0.368 \\ 0.632 & 0.368 & 0 & 0 \\ 0.264 & 0.368 & 0.368 & 0 \\ 0.08 & 0.184 & 0.368 & 0.368 \end{pmatrix}.$$

Suppose that there are $X_0 = 3$. Find the expected time until state 0 is visited for the first time.

Sol. The chain is finite-state irreducible and hence we can apply the equations:

$$\mu_{30} = 1 + p_{31}\mu_{10} + p_{32}\mu_{20} + p_{33}\mu_{30} \quad \Leftrightarrow \quad \mu_{30} = 1 + 0.184\mu_{10} + 0.368\mu_{20} + 0.368\mu_{30}$$
$$\mu_{20} = 1 + p_{21}\mu_{10} + p_{22}\mu_{20} + p_{23}\mu_{30} \quad \Leftrightarrow \quad \mu_{20} = 1 + 0.368\mu_{10} + 0.368\mu_{20}$$
$$\mu_{10} = 1 + p_{11}\mu_{10} + p_{12}\mu_{20} + p_{13}\mu_{30} \quad \Leftrightarrow \quad \mu_{10} = 1 + 0.368\mu_{10}$$

Solving, $\mu_{10} = 1.58$, $\mu_{20} = 2.51$, and $\mu_{30} = 3.50$. μ_{00} can be found from μ_{10}, μ_{20}, and μ_{30}. A quicker way is that

π_0 = proportion of time that state 0 is visited = $\frac{1}{\mu_{00}}$, since on average the chain visits state 0 once for a duration of μ_{00} and the chain stays there for one period of time. ✧✧

4.8. Absorption States

In the gambler's ruin problem, there are two absorption states, one denoting that Peter wins and the one denoting that Sam wins. To find the probability that Peter wins is equivalent to determine the probability of absorption by the absorbing state 0.

Let f_i = the probability of absorption by the absorbing state 0 given that $X_0 = i$. If state 0 is the only absorbing state, then $f_i = 1$ for all i. Otherwise, f_i can be found from the following set of equations:

$f_i = \sum_{j=0}^{\infty} p_{ij} f_j$, for $i = 0, 1, \ldots$; where $f_k = 1$, and $f_i = 0$ if state i is another absorbing state.

Example 4.8.1. Gambler's ruin problem.

Suppose that the probability transition matrix of a gamblers' ruin problem is given by

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 1-p & 0 & p & 0 & 0 \\ 0 & 1-p & 0 & p & 0 \\ 0 & 0 & 1-p & 0 & p \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$ Find f_i for i = 1, 2, and 3.

Sol. The equations are

$f_1 = (1-p) + pf_2$,
$f_2 = (1-p)\, f_1 + pf_3$,
and $f_3 = (1-p)\, f_2$.

We get f_i by solving these equations. In general, for a Gambler's ruin problem with state space {0, 1, …, M}, $1-f_i = \dfrac{\sum_{m=0}^{i-1} \rho^m}{\sum_{m=0}^{M-1} \rho^m} = \begin{cases} \dfrac{1-\rho^i}{1-\rho^M}, & \text{for } p \neq 0.5, \\ \dfrac{i}{M}, & p = 0.5, \end{cases}$ for i = 1, 2, ..., M, where $\rho = \frac{1-p}{p}$.

✧✧

Exercise 4.8.2. Suppose that the weather can be modeled as a four-state *DTMC* with state space {rainy, cloudy, windy, sunny} and the transition probability matrix $\begin{pmatrix} 0.5 & 0.2 & 0.2 & 0.1 \\ 0.2 & 0.3 & 0.3 & 0.2 \\ 0.4 & 0.1 & 0.3 & 0.2 \\ 0.1 & 0.2 & 0.1 & 0.6 \end{pmatrix}$. Given that today is cloudy, find the probability that the weather changes into a sunny day at an earlier time than a rainy day. ✧✧

Remarks on Chapter 4 *of Ross*. We will cover Sections 4.1, 4.2, 4.4, 4.5.1, and 4.6 of the chapter. While we don't go deep in Section 4.3, we will state results deduced from the material in the section. Examples 4.19 and 4.20 look hard, though they are very educational. Roughly speaking, problems #3.1 to #3.56 are useful for understanding *DTMC*. Students may like to try those with probabilistic rather than mathematical flavor.