

ESE 3030:
Stochastic Systems Analysis and Simulation
Lecture Notes

Notes prepared by Enzo Bergamo
based on Dr. Shirin Bidokhti's course at
the University of Pennsylvania

Spring 2023

Contents

1	Introduction	4
2	Probability Theory	5
2.1	Definitions & Axioms of Probability	5
2.2	Bayes's Theorem	6
2.3	Independence	6
2.4	Law of Total Probability	7
2.5	Random Variables	8
2.6	Probability mass function, probability density function, and cumulative distribution function	8
2.7	Distributions for derived random variables	9
2.8	Joint Distributions	10
2.9	Relevant Probability Distributions	10
2.9.1	Bernoulli distribution	11
2.9.2	Geometric distribution	11
2.9.3	Binomial distribution	12
2.9.4	Poisson distribution	13
2.9.5	Gaussian distribution	15
2.9.6	Exponential distribution	16
2.10	Expected Value & Variance	17
2.11	Notions of Convergence	20
2.11.1	Sure Convergence	20
2.11.2	Almost Sure Convergence	21
2.11.3	Convergence in Probability	22
2.11.4	Convergence in Distribution	23
3	Discrete-Time Markov Chains	25
3.1	Introduction	25
3.2	Graph Representation of a Markov Process	26

3.3	Matrix Representation of a Markov Process and the Transition Matrix . . .	27
3.4	The Chapman-Kolmogorov Equation	29
3.5	Definitions & Classifications for States and Classes	33
3.6	Limit Distributions	35
3.7	Limit Distributions in Markov Chains with More than One Class	38
3.8	Application: The Gambler’s Ruin Problem	42
3.9	Ergodicity	46
3.9.1	Definitions	46
3.9.2	Ergodic averages	48
3.10	Application: Discrete-Time Queueing Theory	48
3.11	Markov Chains with Rewards	51
3.11.1	Expected Reward until Specific Event	51
3.11.2	Expected Aggregate Reward over Multiple Transitions	57
3.12	Markov Decision Processes	58
4	Arrival Processes	64
4.1	Poisson Processes	64
4.2	Exponential Inter-Arrival Times	65
5	Continuous-Time Markov Chains	68
5.1	Chapman-Kolmogorov Equations in Continuous-Time	70
5.1.1	The Matrix Exponential	74
5.2	Application: Continuous-Time Queueing Theory	76
6	Gaussian Processes	82
6.1	Properties of the Gaussian Distribution	82
6.2	Gaussian Processes: General Aspects	83
6.3	Application: Brownian Motion	83

1. Introduction

These lecture notes are based on Dr. Shirin Bidokhti's ESE 3030: Stochastic Systems Analysis and Simulation during the Fall 2021 semester at the University of Pennsylvania. My personal notes from Dr. Bidokhti's lecture, slides from Dr. Alejandro Ribeiro's previous version of the course, Sheldon Ross's *Introduction to Probability Models*, and lecture notes from MIT's *Introduction to Probability* course by John Tsitsiklis and Dimitri P. Bertsekas were used in order to prepare these notes.

These notes may contain errors and do not substitute lecture attendance or consulting the references. This is especially true since these notes were prepared with the goal of helping students taking the course and thus may lack rigor at times. Note also that not all example problems have solutions as they will be often solved during recitation. Corrections are greatly appreciated — if a mistake/typo is found, please get in touch via e-mail.

2. Probability Theory

This course aims "to learn how to model, analyze and simulate stochastic systems." This inherently requires a solid understanding of probability theory — as such, this first module of the course focuses on reviewing the core components of this subject. Note that we assume familiarity with the subject and this serves as a review and opportunity to introduce some additional topics which are relevant to this course. For a more thorough discussion, see Chapters 1 through Chapter 3 of *Introduction to Probability Models* by Sheldon Ross or MIT's *Introduction to Probability* lecture notes by John Tsitsiklis and Dimitri P. Bertsekas, available here at no cost.

2.1. Definitions & Axioms of Probability

In order to illustrate these definitions, we make use of the archetypal probabilistic event: the roll of a single six-face die. We start by defining **sample space** Ω , the set that describes the possible results of a probabilistic event. In the context of our example, this corresponds to each possible outcome of the die roll (i.e. the die facing each of the six sides after being rolled). An **event** E is a subset of the sample space ($E \subseteq \Omega$) and represents specific results of a probabilistic experiment (E may be equal, for example, to a die showing the number four after being rolled).

While the definition of probability is an intrinsically complex topic, we skip the philosophical discussion and instead adopt the following definition: a **probability law** $\mathbf{P}(\cdot)$ is mapping from the space of events to the real numbers such that it represents the likelihood (or belief) of the occurrence of each event (e.g. the likelihood of a fair six-face die toss resulting in 3 is $1/6$). In order for $\mathbf{P}(\cdot)$ to be a probability law, it must satisfy the following three axioms.

- I. **Non-Negativity:** $\mathbf{P}(E) \geq 0; \forall E$ (i.e the probability law must only map to non-negative values)
- II. **Normalization:** $\mathbf{P}(\Omega) = 1$ (i.e. the probability of all events sums up to 1)

III. **Additivity:** $P(A \cup B) = P(A) + P(B)$ if $A \cap B = \emptyset$ (for two disjoint events, the probability of their union is equal to the sum of their probabilities)

Immediate results from these axioms include that $P(\emptyset) = 0$ and $P(\cup_i E_i) = \sum_i P(E_i)$ if $\cap_i E_i = \emptyset$, the probability of the empty event is zero and additivity generalizes for n events, respectively.

2.2. Bayes's Theorem

Using our definition of probability, we can also consider **conditional probability**. It allows us to consider the probability of a given event A given the occurrence of another event B . Mathematically, we have the following (where the last line is obtained using Bayes's Theorem with $P(B|A)$).

$$\begin{aligned} P(A|B) &= \frac{P(A; B)}{P(B)} \\ &= \frac{P(B|A)P(A)}{P(B)} \end{aligned}$$

2.3. Independence

We define **independence** $A \perp B$ for two events A, B . Intuitively, it follows that two events are independent if the occurrence of B does not affect the likelihood of the outcome of A (and vice-versa). Mathematically, events A and B are independent if and only if $P(A; B) = P(A)P(B)$. When applying this to Bayes's Theorem, we have the following.

$$\begin{aligned} P(A|B) &= \frac{P(A; B)}{P(B)} \\ &= \frac{P(A)P(B)}{P(B)} && \text{(independence)} \\ &= P(A) \end{aligned}$$

Intuition is thus correct here: if two events are independent, it follows that the occurrence of B does not affect the probability of A .

Note that for $n > 2$ events, pair-wise independence is insufficient for independence. Consider the set of events $E = \{E_1, E_2, E_3\}$; even if $E_1 \perp E_2$, $E_1 \perp E_3$, $E_2 \perp E_3$, it might be the case

that these events are not independent. For such, we require independence of all subsets (i.e. independence for all elements of 2^E).

2.4. Law of Total Probability

The **law of total probability** allows us to deconstruct events (which are often complex) into smaller, tractable components. It is given as the following (where $E_1; E_2; \dots; E_n$ are non-overlapping events that form a partition of the sample space).

$$\mathbf{P}(E) = \sum_i \mathbf{P}(E|E_i)\mathbf{P}(E_i)$$

Intuitively, we divide the sample space into a number of sections each with a probability associated with it; the probability of a given event, then, is the probability of the said event when in the section weighted by the probability of each section.

Example 1 Consider an airplane flying that enters a zone with probability $\mathbf{P}(P) = 0.05$. This area is monitored by a radar with false positive and false negative probabilities $\mathbf{P}(FP) = \mathbf{P}(+j\bar{P}) = 0.1$ and $\mathbf{P}(FN) = \mathbf{P}(-jP) = 0.01$, respectively. Given that the radar has a positive alert, what is the probability that the airplane did enter the zone?

$$\begin{aligned} \mathbf{P}(Pj+) &= \frac{\mathbf{P}(+jP)\mathbf{P}(P)}{\mathbf{P}(+)} \\ &= \frac{\mathbf{P}(+jP)\mathbf{P}(P)}{\mathbf{P}(+jP)\mathbf{P}(P) + \mathbf{P}(+j\bar{P})\mathbf{P}(\bar{P})} \\ &= \frac{\mathbf{P}(+jP)\mathbf{P}(P)}{\mathbf{P}(+jP)\mathbf{P}(P) + \mathbf{P}(+j\bar{P})\mathbf{P}(\bar{P})} \\ &= \frac{0.99 \cdot 0.05}{0.99 \cdot 0.05 + 0.1 \cdot 0.95} \\ &= 0.342 \end{aligned}$$

Example 2 A bin contains 3 types of flashlights. The probability that a type 1 flashlight will give more than 100 hours of light is 0.7, with the corresponding probabilities for type 2 and type 3 flashlights being 0.4 and 0.3, respectively. Suppose that 20% of the flashlights in the bin are type 1, 30% are type 2 and 50% are type 3. What is the probability that a randomly chosen flashlight will give more than 100 hours of use, and given that a flashlight lasted more than 100 hours, what is the probability that it was a type j flashlight for $j = 1; 2; 3$?

2.5. Random Variables

A random variable X is a mapping from the events to the real numbers $X : \Omega \rightarrow \mathbb{R}$. More simply, we assign a real number to each event in the sample space such that $X(\omega) \in \mathbb{R}$ for each event $\omega \in \Omega$. Returning to our original example, we may define X as the number on the die after being rolled (thus $X \in \{1; 2; 3; 4; 5; 6\}$) or the square of that number (thus $X \in \{1; 4; 9; 16; 25; 36\}$). For two rolls, it may take the value of the larger of the two rolls. Observe that while the difference is often hazy, note random variables map from abstract events (such as the die result) to actual numbers.

A commonly used type of random variable is an indicator random variable $\mathbb{1}_E$. It takes the value 1 if a given event E occurs and 0 otherwise. Considering again the die experiment, we may be interested in whether the roll results in a 3 or not. We can thus define the event E as whether the roll results in the number 3; it follows then that $\mathbb{1}_E$ is a random variable equal to 1 if the roll results in 3 and 0 otherwise.

2.6. Probability mass function, probability density function, and cumulative distribution function

Since random variables are associated with probabilistic events, we can obtain a probability associated with them.

$$P_X(x) = P(X = x) = P(\{\omega \in \Omega : X(\omega) = x\}) \quad (1)$$

When considering discrete random variables (i.e. countable set of possible values), we may define a probability mass function (p.m.f., $p_X(x)$) as given by Equation 1. Observe that the p.m.f. respects $p_X(x) \geq 0$ and $\sum_x p_X(x) = 1$. Likewise, for continuous random variables (i.e. uncountable set of possible values such as the real set) we may define the probability density function (p.d.f., $f_Y(y)$). It allows for the calculation of the probability of a random variable taking values in a given interval $[a; b]$. Similar to the p.m.f, the p.d.f. respects $f_Y(y) \geq 0$ and $\int_{-\infty}^{\infty} f_Y(y) dy = 1$. Note however that the probability that a continuous random variable takes any specific value (i.e. $P_Y(X = a)$) is always zero | the p.d.f. itself

is not the probability at a given value. Mathematically, we have the following.

$$P_Y(y \in [a; b]) = \int_a^b f_Y(y) dy$$

Likewise, we define the cumulative distribution function (c.d.f.), the probability that the random variable takes a value less than x . This is valid for both discrete and continuous random variables with the difference being whether the computation uses an integral or summation.

$$F_X(x) = P(X \leq x) = \sum_{i \leq x} p_X(i) \quad (\text{discrete random variable})$$

$$F_Y(y) = P(Y \leq y) = \int_{-\infty}^y f_Y(z) dz \quad (\text{continuous random variable})$$

Note that the c.d.f. is monotonically non-decreasing and approaches 1. In Figure 1, we exhibit the p.m.f. and c.d.f. for an arbitrary distribution.

Figure 1: Plot of p.m.f. and c.d.f. for Arbitrary Discrete Distribution

If a given random variable X follows a specific distribution with parameters $a; b$, we use the notation $X \sim (a; b)$. Finally, from the Fundamental Theorem of Calculus, we have that $\frac{d}{dx} F_X(x) = f_X(x)$.

2.7. Distributions for derived random variables

It is often of interest to find the distribution of a random variable which is a transformation of another, known random variable, for example $Y = 2X$ for a known X . The standard procedure to do this is using the c.d.f. of X to find the c.d.f. associated with Y ; once the c.d.f. is obtained, the p.d.f. can usually be trivially found using the Fundamental Theorem of Calculus.

Example 3 Consider a random variable X distributed uniformly between 0 and 1; therefore, its p.d.f. is equal to $f_X(x) = 1$ and its c.d.f. is equal to $F_X(x) = x$. Find the p.d.f. and c.d.f. of the random variables $Y = X^2$ and $Z = 2X$.

2.8. Joint Distributions

While we have focused our attention on individual random variables, it is often of interest to consider the joint distribution of two or more random variables.

Consider a random variable X corresponding to the face shown after rolling a die; likewise, define Y as the face shown on the opposite side in the same roll. Note that in a regular 6-side die, opposite sides of a die always add up to seven (e.g. 6 and 1, 3 and 4).

We can then consider the joint probability of $X; Y$ defined as $p(x; y)$. It should be clear, for example, that $p(3; 2)$ should be equal to zero since the probability of observing an event where one face shows 3 and the other 2 is zero. Likewise, $p(3; 4) = 1/6$, since this event corresponds simply to the probability of observing $X = 3$.

For two independent random variables X and Y , we have that $f(x; y) = f(x)f(y)$. This equality can be read as the probability of observing two events $X = x; Y = y$ is simply the product of the probability of each of the two events individually. Consider for example two dice being tossed with $X; Y$ being the number that results from the first and second die, respectively. Here, one toss does not affect the other $X(? Y)$, and thus $p(3; 2) = p(3; 4) = 1/36$.

2.9. Relevant Probability Distributions

There are several distributions that will be used throughout the course. A few of the should be familiar from previous coursework and they are reviewed below. While expectation and variance are discussed in Section 2.10, the values for these moments are also given below.

2.9.1. Bernoulli distribution

A single-parameter distribution, the Bernoulli distribution models a single binary experiment (i.e. either "yes" or "no", 1 or 0). The one parameter in the distribution is p and it corresponds to the probability of success (i.e. $X = 1$). Its p.m.f. is given below (for $x \in \{0, 1\}$) while the c.d.f. can only be written as a piece-wise function.

$$p_X(x) = p^x(1-p)^{1-x}$$

Figure 2: Plot of p.m.f. and c.d.f. for Bernoulli Distribution with $p = 0.3$

Its expectation is $E[X] = 1p + 0(1-p) = p$ and its variance is $\text{Var}[X] = p(1-p)$.

2.9.2. Geometric distribution

A single-parameter distribution, it models a sequence of binary experiments until the first success. It is thus parametrized by one success probability parameter p like the Bernoulli distribution. For example, the random variable that represents the number of tosses required until the first head follows a geometric distribution. Note that in the literature this distribution is often described as the shifted geometric distribution since it includes the "success" experiment in its count and thus has support in the positive integers while the geometric distribution does not (and thus has support in the non-negative integers). Its p.m.f. and c.d.f. are given below (for $x \in \{1, 2, \dots\}$).

$$p_X(x) = p(1-p)^{x-1} \quad (\text{p.m.f.})$$

$$F_X(x) = 1 - (1-p)^x \quad (\text{c.d.f.})$$

Figure 3: Plot of p.m.f. and c.d.f. for Geometric Distribution with $p = 0.3$

In order to calculate the expected value of a random variable that follows a geometric distribution, we first note the following.

$$\sum_{i=1}^{\infty} (1-p)^{i-1} = \frac{1-p}{1-(1-p)} = 1 \tag{1}$$

$$\sum_{i=1}^{\infty} i(1-p)^{i-1} = \frac{1}{p^2} \tag{2}$$

Observe then that the equation for the expected value $E[X] = \sum_{x=1}^{\infty} x(1-p)^{x-1}p = p \sum_{x=1}^{\infty} x(1-p)^{x-1}$; the last term is equal to the left-hand side of Equation 2 and thus $E[X] = p(1/p^2) = 1/p$. The corresponding variance is $Var[X] = (1-p)/p^2$.

2.9.3. Binomial distribution

Similar to the geometric distribution, the binomial distribution also focuses on modeling a sequence of binary experiments. With the binomial distribution, however, we have a fixed number of experiments and the total number of successes is accumulated. It is therefore a two-parameter distribution with n as the number of experiments and p as the success probability of each individual experiment. As an example, one can consider the case of tossing a coin n times each with a probability p of being heads. If X is a random variable counting the total number of heads, we have that X follows a binomial distribution

with parameters $n = 5; p = 0.3$. Its p.m.f. is given below (for $x = 0; 1; 2; \dots; n$) and there is no simple closed-form solution for the c.d.f. of this distribution.

$$p_X(x) = \binom{n}{x} p^x (1-p)^{n-x}$$

Figure 4: Plot of p.m.f. and c.d.f. for Binomial Distribution with $n = 12; p = 0.3$

For a binomial random variable, its expected value can be obtained using linearity of expectation and variance properties; we obtain thus $E[X] = np$ and its variance is $\text{Var}[X] = np(1-p)$.

2.9.4. Poisson distribution

The Poisson distribution is similar to the binomial distribution but models specifically rare events (i.e. $p \ll 1; n \gg 1$). Note that even though the individual events are rare, the fact that n is large leads to a reasonable number of occurrences. The only parameter of the Poisson distribution is its rate parameter λ ; we will later see that it is the mean event rate per time period. Its p.m.f. is given below while its c.d.f. depends on the Gamma function and is thus omitted.

$$p_X(x) = \frac{\lambda^x e^{-\lambda}}{x!}$$

We show that the Poisson distribution can be derived from the Binomial distribution by setting $p = \frac{x}{n}$ and letting n approach infinity.

$$\begin{aligned}
 p_X(x) &= \lim_{n \rightarrow \infty} \binom{n}{x} p^x (1-p)^{n-x} \\
 &= \lim_{n \rightarrow \infty} \frac{n!}{x! (n-x)!} \left(\frac{x}{n}\right)^x \left(1 - \frac{x}{n}\right)^{n-x} \\
 &= \lim_{n \rightarrow \infty} \frac{n!}{(n-x)! x!} \frac{x^x}{n^x} \left(1 - \frac{x}{n}\right)^{n-x} \\
 &= \frac{x^x}{x!} \lim_{n \rightarrow \infty} \frac{n!}{(n-x)!} \frac{1}{n^x} \left(1 - \frac{x}{n}\right)^{n-x} \\
 &= \frac{x^x}{x!} \lim_{n \rightarrow \infty} O(n^x) O(1/n^x) \lim_{n \rightarrow \infty} \left(1 - \frac{x}{n}\right)^{n-x} \\
 &= \frac{x^x}{x!} \lim_{n \rightarrow \infty} \left(1 - \frac{x}{n}\right)^n \lim_{n \rightarrow \infty} O(n^x) O(1/n^x) \lim_{n \rightarrow \infty} \left(1 - \frac{x}{n}\right)^{-x} \\
 &= \frac{x^x e^{-x}}{x!} (1)(1) \\
 &= \frac{x^x e^{-x}}{x!}
 \end{aligned}$$

Figure 5: Plot of p.m.f. and c.d.f. for Poisson Distribution with $\lambda = 1$

The expected value of the Poisson can be obtained algebraically as follows.

$$\begin{aligned}
 E[X] &= \sum_{x=0}^{\infty} x \frac{e^{-\lambda} \lambda^x}{x!} \\
 &= \sum_{x=1}^{\infty} x \frac{e^{-\lambda} \lambda^x}{x!} \quad (\text{rst term is zero}) \\
 &= e^{-\lambda} \sum_{x=1}^{\infty} \lambda^x \frac{x}{x!} \\
 &= e^{-\lambda} \sum_{x=1}^{\infty} \lambda^x \frac{1}{(x-1)!} \\
 &= e^{-\lambda} \sum_{z=0}^{\infty} \lambda^{z+1} \frac{1}{z!} \\
 &= e^{-\lambda} \lambda \sum_{z=0}^{\infty} \frac{\lambda^z}{z!} \\
 &= \lambda e^{-\lambda} e^{\lambda} \quad (\text{Taylor Series expansion}) \\
 &= \lambda
 \end{aligned}$$

We thus have $E[X] = \lambda$ and likewise $\text{Var}[X] = \lambda$.

Example 4 Let $X \sim \text{Poisson}(\lambda_1)$ and $Y \sim \text{Poisson}(\lambda_2)$. Find the distribution associated with the random variable $Z = X + Y$. Considering that Poisson random variables may represent event occurrences, what is the interpretation of the distribution of Z ?

2.9.5. Gaussian distribution

The Gaussian (or Normal) distribution is one of the most important distributions in probability theory and statistics. It is particularly relevant when considering the phenomena associated with the Central Limit Theorem (see Section 2.11). It is a continuous distribution parameterized by its mean μ and variance σ^2 . Its p.d.f. is given below (for all $x \in \mathbb{R}$).

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

While there is no closed-form solution for the c.d.f. of the Gaussian distribution, due to its frequent use it is often denominated by $\mathcal{N}(\mu, \sigma^2)$. The standard Gaussian is a specific case of the Gaussian where $\mu = 0$; $\sigma^2 = 1$.

Figure 6: Plot of p.d.f. and c.d.f. for Gaussian Distribution with $\mu = 0$; $\sigma^2 = 1$

As previously mentioned, $E[X] = \mu$ and likewise $\text{Var}[X] = \sigma^2$.

Example 5 The time to pass through a queue to begin self-service at a cafeteria is modeled by a random variable $X \sim N(15, 9)$. Determine the probability that an arriving customer waits between 14 and 17 minutes for service analytically and through simulations.

2.9.6. Exponential distribution

The exponential distribution is a continuous distribution parameterized by a rate parameter λ and with support in the real non-negative numbers. Its p.d.f. and c.d.f. are given below (for all $x \geq 0$).

$$f_X(x) = \lambda e^{-\lambda x} \quad (\text{p.d.f.})$$

$$F_X(x) = 1 - e^{-\lambda x} \quad (\text{c.d.f.})$$

An important property of the exponential distribution is that it is the only continuous and memoryless distribution. This means that the current state of the system is not affected by its past (i.e. Markov property, which will be further discussed later). For example, if the random variable $X \sim \text{exp}(\lambda)$ represents waiting time until an event, its value does not depend on how much time has elapsed so far; mathematically, we have $P(X > s + t | X > s) = P(X > t) = P(X > s)$.

Figure 7: Plot of p.d.f. and c.d.f. for Exponential Distribution with $\lambda = 1$

For a random variable following an exponential distribution, its expectation $E[X] = 1 = \frac{1}{\lambda}$ can be obtained by integration by parts and $\text{Var}[X] = 1 = \frac{1}{\lambda^2}$.

Example 6 Show that the Markov (memoryless) property holds for the exponential distribution. BONUS: Sketch proof that the exponential distribution is the only continuous distribution with this property.

This can be proved directly using Bayes's rule.

$$\begin{aligned}
 P(X > a + b | X > a) &= \frac{P(X > a + b | X > a) P(X > a)}{P(X > a)} \\
 &= \frac{1 - e^{-(a+b)}}{1 - e^{-a}} \\
 &= \frac{e^{-(a+b)}}{e^{-a}} \\
 &= e^{-b} \\
 &= P(X > b)
 \end{aligned}$$

2.10. Expected Value & Variance

For a random variable X , we can define its k^{th} moment around its mean by the following formula:

$$\begin{aligned}
 E[X^k] &= \int_{-\infty}^{\infty} (x - E[X])^k f_X(X = x) dx \quad (\text{continuous r.v.}) \\
 &= \sum_x (x - E[X])^k p_X(X = x) \quad (\text{discrete r.v.})
 \end{aligned}$$

In the equation above, the expected value of X , $E[X]$, is the 1st moment around zero (not around $E[X]$); that is:

$$E[X] = \begin{cases} \int_{-\infty}^{\infty} xf_X(X = x)dx & \text{(continuous r.v.)} \\ \sum_x xp_X(X = x) & \text{(discrete r.v.)} \end{cases}$$

The expected value of a random variable gives a notion of the mean (a weighted average) to the distribution associated with the random variable. Important properties include Linearity of Expectation (expectation is a linear operator, $E[X + aY + b] = E[X] + aE[Y] + b$ for any two random variables X, Y and constants a, b), expectation of functions of random variables ($E[g(Y)] = \int_{-\infty}^{\infty} g(y)f_Y(y)dy$, $E[g(X)] = \sum_x g(x)p_X(x)$ for continuous and discrete random variables, respectively), Conditional Expectation, and Law of Iterated Expectation, the last two discussed below.

Example 7 Prove that Linearity of Expectation holds.

We present the concept of conditional expectation through an example. Let X be a random variable that corresponds to the height of an individual and Y their age. It can be shown that the best estimate (i.e. what should you guess) for X given no other information is $E[X]$. This should make sense: if you have no information about an individual, your best guess of their height should be the average height in the population. If we know one's age, however, our estimate would likely change; for example, if $Y = 2$ (i.e. the age is two years old), we should guess a value for X lower than if $Y = 20$. Clearly, the value of the mean of a random variable may be affected by information from others. This is mathematically encoded through a conditional expectation $E[X|Y = y]$, the mean of X given that $Y = y$, and it can be calculated as follows. Note that it is inherently a function of y .

$$E[X|Y = y] = \begin{cases} \int_{-\infty}^{\infty} xf_{X|Y}(X = x|Y = y)dx & \text{(continuous r.v.)} \\ \sum_x xp_{X|Y}(X = x|Y = y) & \text{(discrete r.v.)} \end{cases}$$

The notation $E[X|Y]$, however, is less clear. While $E[X|Y = y]$ is a value, $E[X|Y]$ is another random variable. If it is a random variable, however, this implies we can calculate its expected

value. In fact, the Law of Iterated Expectation tells us that $E[E[X|Y]] = E[X]$; the proof for the continuous case is given below.

$$\begin{aligned}
 E[E[X|Y]] &= \int_{Z^y} E[X|Y=y] f_Y(Y=y) dy \\
 &= \int_{Z^y} \int_{Z^x} x f_{X|Y}(X=x|Y=y) dx f_Y(Y=y) dy \\
 &= \int_{Z^x} x \int_{Z^y} f_{X|Y}(X=x|Y=y) f_Y(Y=y) dy dx \\
 &= \int_{Z^x} x f_{X,Y}(X=x; Y=y) dy dx \quad (f_{X|Y}(x;y) f_Y(y) = f_{X,Y}(x;y)) \\
 &= \int_{Z^x} x f_X(X=x) dx \\
 &= E[X]
 \end{aligned}$$

Intuition should help here. When we know $Y = y$, we change our estimate of $E[X]$; if we do not, however, we are back at the original problem of estimating one's height with no information of their age other than it comes from the distribution of Y . It should be the case, therefore, that our estimate will not change and therefore $E[E[X|Y]] = E[X]$.

Example 8 The number of goals scored by a soccer player is modeled by the random variable X goals per game and she expects to play Y games per season. Note that both X and Y are random variables since the number of goals scored varies as well as the number of games she will play (injuries, substitutions, etc.) are both stochastic. Let Z be the total number of goals scored in the season | what is $E[Z]$?

$$\begin{aligned}
 E[Z] &= E[X \cdot Y] \\
 &= E[Y]E[X]
 \end{aligned}$$

Variance corresponds to the 2nd moment and provides a measure of the spread of the distribution around the mean. It is therefore given by the following.

$$\begin{aligned}
 \text{Var}[X] &= \int_{Z^x} (x - E[X])^2 f_X(X=x) dx \quad (\text{continuous r.v.}) \\
 &= \sum_x (x - E[X])^2 p_X(X=x) \quad (\text{discrete r.v.})
 \end{aligned}$$

A simpler formulation for variance can be found through algebraic manipulation and is valid for both the continuous and discrete case; it is given below.

$$\begin{aligned} \text{Var}(X) &= \int_1^Z (x - E[X])^2 f_X(X = x) dx \\ &= E[(X - E[X])^2] \\ &= E[X^2 - 2X E[X] + E[X]^2] \\ &= E[X^2] - 2E[X]E[X] + E[X]^2 \\ &= E[X^2] - E[X]^2 \end{aligned}$$

Note that variance is not a linear operator, but $\text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y]$ if the random variables X, Y are uncorrelated (otherwise, $\text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y] + 2\text{Cov}[X; Y]$). Additionally, $\text{Var}[aX] = a^2\text{Var}[X]$ for a random variable X and constant a .

The third and fourth moments are called skewness and kurtosis | while extremely relevant in some contexts, they will not be used in this course.

2.11. Notions of Convergence

Here, we focus on studying the behavior of sequences of random variables X_1, X_2, \dots, X_n , also known as stochastic processes. Often, we will consider how different realizations x_1, x_2, \dots, x_n behave and establish different convergence types accordingly. This is further laid out in the following sections. As a preliminary, we define convergence: a sequence x_1, x_2, \dots converges to x if for any $\epsilon > 0$ there exists an n_0 such that for all $n > n_0$ the inequality $|x_n - x| < \epsilon$ holds. That is, for an arbitrarily small ϵ , we must be able to find a point n_0 in the sequence such that all points in the sequence after it (i.e. $x_{n_0+1}, x_{n_0+2}, \dots$) is no more than ϵ away from x .

2.11.1. Sure Convergence

For a given sequence of random variables X_1, X_2, \dots, X_n , we say that exhibits Sure Convergence if for all of its realizations we have that $\lim_{n \rightarrow \infty} x_n = x$ where x is a realization of the random variable X . Note that this is a very rigid requirement since all of its realizations must respect its conditions.

2.11.2. Almost Sure Convergence

A relaxation of the Sure Convergence is Almost Sure Convergence (a.s.). We say a sequence of random variables X_1, X_2, \dots, X_n converges almost surely $\mathbb{P}(\lim_{n \rightarrow \infty} X_n = X) = 1$. Note that the limit here is a random variable. Bertsekas and Tsitsiklis write that the "right way of interpreting this type of convergence is in terms of a sample space consisting of infinite sequences: all of the probability is concentrated on those sequences that converge to c . This does not mean that other sequences are impossible, only that they are extremely unlikely, in the sense that their total probability is zero." This should not be completely unfamiliar: in a continuous distribution, the probability of any given value x is zero while the actual probability of events/intervals is non-zero.

The Strong Law of Large Numbers can be seen as a statement regarding almost sure convergence. It states that the sample average \bar{X}_k converges a.s. to the real mean $\mu = E[X]$.

```

values = []
averages = []
n = 3000
for i in range(n):
    x = np.random.normal(loc = 3, scale = 4)
    values.append(x)
    averages.append(np.mean(values))

plt.plot(averages, label=r '$\bar{X}_k$ )
plt.plot(3*np.ones(n), label= '$\mu$ )
plt.xlabel( k )
plt.legend()
plt.grid()
plt.show()

plt.plot(np.abs(averages - real_mean), label=r '$|\bar{X}_k - \mu|$ )
plt.xlabel( k )
plt.legend()
plt.grid()
plt.show()

```

Figure 8: Different sequences $X_{0;k}, X_{1;k}, X_{2;k}, X_{3;k}, X_{0;k}$ all approach the real mean $= 3$ as $k \rightarrow 1$

2.11.3. Convergence in Probability

Relaxation of Almost Sure Convergence leads to Convergence in Probability. Convergence in Probability states that $\lim_{n \rightarrow \infty} P(|X_n - X| > \epsilon) = 0$. This does not imply that the realizations of the sequence converge, but that the probability does. Consider $X_n = X_0 + Z_n$ where $X_0 \sim N(0, 1)$ and $Z_n \sim \text{Bernoulli}(1/n)$. It can be shown that as $n \rightarrow \infty$ we have $P(Z_n = 1) = 0$, thus giving $X_n = X_0$ and showing convergence. However, it may still occur that $Z_n = 1$ (after all, $1/n$ is not zero), and thus the sequence only converges in probability.

The Weak Law of Large Numbers can be seen as a statement regarding almost Convergence in Probability. It states that the sample average \bar{X} converges in probability to the real mean $\mu = E[X]$, that is, $\lim_{n \rightarrow \infty} P(|\bar{X} - \mu| > \epsilon) = 0$

```

values = []
X_0 = np.random.uniform()
X = []

for n in range(1, 10000):
    X.append(X_0 + np.random.binomial(1, 1/n))

plt.plot(X, label=r '$X_n$')
plt.xlabel('$n$')

```

```
plt.legend()
plt.grid()
plt.show()
```

Figure 9: Sequence $X_n = X_0 + Z_n$ with $X_0 = 0.918$

2.11.4. Convergence in Distribution

A relaxation of Convergence in Probability and the weakest of the convergence types, convergence in Distribution states that in a stochastic process $X_1; X_2; \dots; X_n$ the c.d.f. of X_n converges to that of another random variable X ; that is, $\lim_{n \rightarrow \infty} F_{X_n}(x) = F_X(x)$. Consider $X_n = Y_i + Z_i/n$ where $Y_i \sim N(0; 1)$ and $Z_i \sim \text{Bernoulli}(0.5)$, it should be clear that the c.d.f. of X_n approaches that of a standard Gaussian as $n \rightarrow \infty$ since the the Bernoulli random variable vanishes.

```
values = []
N = 10000
for n in range(1, N+1):
    Y_i = np.random.normal(loc = 0, scale = 1)
    Z_i = 1 if random.random() > 0.5 else 0
    values.append(Y_i + Z_i / n)

x = np.linspace(norm.ppf(0.001),
                norm.ppf(0.999), 500)
plt.hist(values, bins=50, density=True, label= 'Empirical p.d.f. ')
plt.plot(x, norm.pdf(x, scale=1), label= '$F_Y(y)$ ')
plt.legend()
plt.show()
```

Figure 10: The p.d.f. of the sequence approaches that χ^2

3. Discrete-Time Markov Chains

This module provides an introduction to Markov Chains and introduces much of the modeling which will be used throughout this course. As you will see, this content makes extensive use of the probability content in the previous module. The topics discussed here correspond primarily to Chapter 4 of Introduction to Probability Models by Sheldon Ross.

3.1. Introduction

A discrete stochastic process is formally defined by a sequence of random variables X_0, X_1, \dots, X_n which is characterized by a joint p.m.f. $p(x_1; \dots; x_n)$ for $n \geq 1$. An important class of stochastic processes that finds applications in various systems is the class of Markov chain processes which are those that | as implied by its name | respect the Markovian property. This property greatly simplifies these joint distributions and implies that for all states $X_1; \dots; X_n$ and outcomes $i; j$, the following statement is true.

$$P(X_{n+1} = j | X_n = i; X_{n-1} = x_{n-1}; \dots; X_1 = x_1; X_0) = P(X_{n+1} = j | X_n = i):$$

Simply, this property states that given the present state $X_n = i$, the prediction of the next state X_{n+1} is independent of the past/history of the process $X_0; X_1; \dots; X_{n-1}$. Once we assume the Markov property, the entire description of the process is possible by knowing its "state space" | the set of values that X_i may take | and the "transition probabilities" which define the probability of transitioning from one state i to another state j in one time step. Throughout this module, we assume that the state space is finite, and thus X_i takes on values from a finite set. Note also that in this course we will assume that the Markov chains are time-invariant and thus the transition probabilities do not change with time. In other words, the probability of transitioning from state $X_n = i$ to $X_{n+1} = j$ is constant and independent of n . We use the notation P_{ij} to refer to the transition probabilities between i and j of the Markov chain. Note that P_{ij} are probabilities, and thus satisfy the two properties discussed previously, $P_{ij} \geq 0$; $\sum_j P_{ij} = 1$, non-negativity and normalization, respectively.

Leveraging the fact that the state space and the transition probabilities make up a complete description of a Markov chain, we develop two critical ways to represent Markov chains:

the graph representation and the matrix representation, both described in the following sections.

3.2. Graph Representation of a Markov Process


Formally, we define a directed weighted graph G as the tuple $G = (V; E; w)$ where $V = \{v_1; v_2; \dots; v_n\}$ is a finite set of n nodes (or vertices), E as the set of ordered tuples $(v_i; v_j)$ which indicate a connection from node v_i to v_j , and w as a function $E \rightarrow \mathbb{R}$ which assign a value (or weight) to each edge.

As mentioned before, a time-invariant Markov chain may be defined fully by its state space and transition probabilities. Using a graph representation, each state from the state space corresponds to a vertex $v_i \in V$ while the transition probabilities P_{ij} between states $i; j$ is represented by an edge $e = (v_i; v_j)$ with corresponding weight $w(e) = P_{ij}$.

Consider a simple Markov chain with state space with cardinality 2. The two states; 1 are such that with probability 0:6 we stay at the same state in the next time step and with probability 0:4 transition to the other state. In graph representation, we have the following.

$$\begin{aligned}
 V &= \{1; 2\} \\
 E &= \{(1; 1); (1; 2); (2; 1); (2; 2)\} \\
 w(i; j) &= \begin{cases} 0:6 & \text{if } i = j \\ 0:4 & \text{otherwise} \end{cases}
 \end{aligned}$$

We can represent this through set notation (as given above) or visually as follows.



3.3. Matrix Representation of a Markov Process and the Transition Matrix

Under the assumption that the system is time-invariant, a Markov chain may also be fully represented by one transition matrix P . Here, each element P_{ij} (where i, j are the row and column index, respectively) corresponds to the probability of transitioning from state i to state j . Considering the example given in the previous section, we have the following.

$$P = \begin{matrix} & \begin{matrix} 2 & 3 \end{matrix} \\ \begin{matrix} 4 & 5 \end{matrix} & \begin{pmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{pmatrix} \end{matrix}$$

Here, each row i corresponds to the probability of transitioning from state i to each other state j . As such, the sum of the elements of each row should add up to 1, a fact that is readily confirmed in the example above.

Example 9 Consider a computer server that processes jobs one at a time. Each job takes some random amount of time to be completed, following a geometric distribution with probability q . This implies that a job's processing is completed at each time step with probability q . Likewise, at each time step, a new job arrives at each time step with probability p . While the server is busy completing a job, other jobs that have arrived wait in a queue with maximum capacity of 3 packets. If more packets arrive, they get dropped and do not join the queue.

In this example, the random process $(X_0; X_1; X_2; \dots)$ tracks the length of the queue which is stochastic in nature. For one realization of the process, it might take the form of the sequence $(0; 1; 1; 2; 3; 3; 2; 1; 2; 2; \dots)$ while for another it might be $(0; 0; 0; 1; 0; 1; 2; 1; 0; 1; \dots)$. Note that here each X_i can be in one of the states $\{0; 1; 2; 3\}$. Markov chains provide a tractable framework to model this process as the Markovian property is a natural assumption: the length of a queue on a given time step (i.e., X_n) is a function of the length at the previous time step $n - 1$ (i.e., X_{n-1}) but it should not depend on the state at the time step before that, X_{n-2} .

As previously mentioned, we can represent a Markov chain through either a matrix or graph representation. For the process described in Example 9, the matrix representation is as

3 Discrete-Time Markov Chains

follows.

$$P = \begin{matrix} & \begin{matrix} 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1-p & p & 0 \\ q(1-p) & (1-q)(1-p)+qp & (1-q)p \\ 0 & q(1-p) & (1-q)(1-p)+qp \\ 0 & 0 & q \end{bmatrix} \end{matrix} \quad (3)$$

Likewise, the graph representation is given below.




Figure 11: Markov chain for Example 9

This system can be simulated using the graph or the matrix form. Using arbitrary parameters $p = 0:3$ and $q = 0:7$, the code below shows one way in which we can obtain multiple realizations of this stochastic system.

```
n_sims = 3
for sims in range(n_sims):
    N = 10
    curr = 0
    n_states = 4

    p, q = 0.3, 0.7

    X = [curr]

    for n in range(N):
        if curr == 0:
            probs = [1 - p, p, 0, 0]
            curr = np.random.choice(n_states, p=probs)
        elif curr == 1:
            probs = [q*(1-p), (1-q)*(1-p)+q*p, (1-q)*p, 0]
            curr = np.random.choice(n_states, p=probs)
```

```

elif curr == 2:
    probs = [0, q*(1-p), (1-q)*(1-p)+q*p, (1-q)*p]
    curr = np.random.choice(n_states, p=probs)
elif curr == 3:
    probs = [0, 0, q, 1-q]
    curr = np.random.choice(n_states, p=probs)

X.append(curr)

print(X)

```

The code yields n_{sims} realizations of the system; for example:

```

[0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0]
[0, 1, 0, 0, 0, 1, 1, 0, 0, 1, 0]
[0, 0, 1, 1, 1, 2, 3, 3, 2, 2, 1]

```

For completeness, when using the parameters $p = 0.3; q = 0.7$, the transition matrix from Equation 3 takes the following form.

$$P = \begin{matrix} & \begin{matrix} 2 & 3 \end{matrix} \\ \begin{matrix} 2 \\ 3 \\ 4 \end{matrix} & \begin{bmatrix} 0.7 & 0.3 & 0 & 0 \\ 0.49 & 0.42 & 0.09 & 0 \\ 0 & 0.49 & 0.42 & 0.09 \\ 0 & 0 & 0.7 & 0.3 \end{bmatrix} \end{matrix}$$

3.4. The Chapman-Kolmogorov Equation

The transition probabilities allow us to immediately quantify the probability of transitioning from state i to a state j in one time step. Nonetheless, in order to understand the long-term behavior of a Markov chain process, it is of interest to obtain the probabilities of going from a given state i to another state j in a number of steps greater than one.

Formally, we define n step transition probability $P_{ij}^{(n)}$ as the probability of going from state i to state j in n time steps for $n \geq 0$:

$$P_{ij}^{(n)} = \Pr(X_{m+n} = j | X_m = i) \stackrel{(a)}{=} \Pr(X_n = j | X_0 = i) \tag{4}$$

3 Discrete-Time Markov Chains

where m is arbitrary and (a) holds because we assumed that the Markov chain is time-invariant.

Likewise, we define the multi-step transition matrix $P^{(n)}$ where for each index $i(j)$ we have $P_{ij}^{(n)}$, the probability of transitioning from state i to state j in n time steps. Consider again Example 9, whose graph representation is reproduced below.




Figure 12: Markov chain for Example 9

It might be of interest, for example, to calculate the probability of having 1 person in the queue at time $t+2$ given that we have 0 people at time t ; that is, $P(X_{t+2} = 1 | X_t = 0)$.

In order to calculate the probability of reaching state 1 from state 0 in 2 time steps, we must consider all the different ways this may happen. In this example, this may occur in two ways: (I.) we stay in state 0 for one time step and then transition to state 1 $(0; 1)$, or (II.) we transition to state 1 in one time step and then stay in state 1 for another time step $(0; 1; 1)$. Mathematically, we can calculate this through the Law of Total Probability.

$$\begin{aligned}
 P(X_{t+2} = 1 | X_t = 0) &= P(X_{t+2} = 1 | X_{t+1} = 0; X_t = 0) P(X_{t+1} = 0 | X_t = 0) \\
 &\quad + P(X_{t+2} = 1 | X_{t+1} = 1; X_t = 0) P(X_{t+1} = 1 | X_t = 0) \\
 &\hspace{15em} \text{(Total Probability)} \\
 &= P(X_{t+2} = 1 | X_{t+1} = 0) P(X_{t+1} = 0 | X_t = 0) \\
 &\quad + P(X_{t+2} = 1 | X_{t+1} = 1) P(X_{t+1} = 1 | X_t = 0) \quad \text{(Markov Property)} \\
 &= p(1-p) + ((1-q)(1-p) + qp)p
 \end{aligned}$$

For more general problems (i.e. to obtain the probability of reaching a state at time $m+n$ conditioned on being at state i at time 0), we can apply the Chapman-Kolmogorov Equa-

3 Discrete-Time Markov Chains

tion . This equation is obtained directly by applying the Theorem of Total Probability.

$$\begin{aligned}
 P(X_{m+n} = j | X_0 = i) &= \sum_{k \in S} P(X_{m+n} = j | X_m = k; X_0 = i) P(X_m = k | X_0 = i) && \text{(Total Probability)} \\
 &= \sum_{k \in S} P(X_{m+n} = j | X_m = k) P(X_m = k | X_0 = i) && \text{(Markov property)} \\
 &= \sum_{k \in S} P(X_n = kj | X_0 = i) P(X_m = j | X_0 = k) && \text{(Time-invariance)}
 \end{aligned}$$

This equation represents a two-step process, the first with length m and the second n . In order to get to state j in $m + n$ steps, we go through each possible intermediary state k and consider the probability of reaching such state in m time steps. For each of these intermediary states k , we compute the probability of then reaching the final state j in the remaining n steps. Consider now the process of calculating the probability of reaching state $j = 2$ from state $i = 0$ in three time steps, for example. Letting $m = 1$ and $n = 2$:

$$\begin{aligned}
 P(X_{1+2} = 2 | X_0 = 0) &= \sum_{k \in S} P(X_2 = kj | X_0 = 0) P(X_1 = 2j | X_0 = k) \\
 &= P(X_2 = 1j | X_0 = 0) P(X_1 = 2j | X_0 = 1) && \text{(Impossible)} \\
 &\quad + P(X_2 = 1j | X_0 = 0) P(X_1 = 2j | X_0 = 1) && (0! \ 0! \ 1! \ 2 \text{ or } 0! \ 1! \ 1! \ 2) \\
 &\quad + P(X_2 = 2j | X_0 = 0) P(X_1 = 2j | X_0 = 2) && (0! \ 1! \ 2! \ 2) \\
 &\quad + P(X_2 = 3j | X_0 = 0) P(X_1 = 2j | X_0 = 3) && \text{(Impossible)} \\
 &= 0 \\
 &\quad + (p(1-p) + ((1-q)(1-p) + qp)p)(1-q)p \\
 &\quad + p(1-q)p((1-q)(1-p) + qp) \\
 &\quad + 0 \\
 &= (p(1-p) + ((1-q)(1-p) + qp)p)(1-q)p + p(1-q)p((1-q)(1-p) + qp)
 \end{aligned}$$

Using the same arbitrary values $p = 0.3$; $q = 0.7$, for example, we obtain that the probability is equal to values 0.4158. It is worth going through the work to convince yourself that the Chapman-Kolmogorov partitions the sample space completely; you may also verify that the equation remains valid for different choices of m ; n .

3 Discrete-Time Markov Chains

A critical property of the transition matrix follows from the Chapman-Kolmogorov equation. Define as $P^{(m)}$ and $P^{(n)}$ the matrices containing the respective m - and n -step transition probabilities from each state i to j , $P_{ij}^{(m)}$ and $P_{ij}^{(n)}$, respectively. Note that $P(X_{m+n} = j | X_0 = i) = \sum_k P(X_n = k | X_0 = i) P(X_m = j | X_0 = k)$, is by definition the $(i; j)$ -index of the product of the matrices $P^{(m)}$ and $P^{(n)}$. From this, we obtain that $P^{(m+n)} = P^{(m)}P^{(n)}$.

For the the case where $m = n = 1$, we have that $P^{(m+n)} = P^{(2)} = P^{(m)}P^{(n)} = P^{(1)}P^{(1)} = PP = P^2$. More generally, we have that $P^{(k)} = P^k$. We prove this relationship via induction.

Base Case: for $n = 2$, we have that $P^{(2)} = P^{(1)}P^{(1)} = P^2$.

Induction Hypothesis: we assume that for $m = k$ the relationship $P^{(k)} = P^k$ holds.

Induction Step: We prove that this relationship holds for $m = k + 1$.

$$\begin{aligned} P^{(k+1)} &= P^{(k)}P^{(1)} && \text{(since } P^{(m+n)} = P^{(m)}P^{(n)} \text{)} \\ &= P^k P && \text{(Induction Hypothesis: } P^{(k)} = P^k \text{)} \\ &= P^{k+1} \end{aligned}$$

Therefore, in order to calculate the k -step transition probabilities from state i to j , one must simply compute the $(i; j)$ -entry of the matrix P^k . We can then obtain, for example $P^3 = P^2P^1$.

$$\begin{aligned} P^3 &= P^2P^1 \\ &= \begin{pmatrix} 0.637 & 0.336 & 0.027 & 0 \\ 0.5488 & 0.3675 & 0.0756 & 0.0081 \\ 0.2401 & 0.4116 & 0.2835 & 0.0648 \\ 0 & 0.343 & 0.504 & 0.153 \end{pmatrix} \begin{pmatrix} 0.7 & 0.3 & 0 & 0 \\ 0.49 & 0.42 & 0.09 & 0 \\ 0 & 0.49 & 0.42 & 0.09 \\ 0 & 0 & 0.7 & 0.3 \end{pmatrix} \\ &= \begin{pmatrix} 0.61054 & 0.34545 & 0.04158 & 0.00244 \\ 0.5642 & 0.3560 & 0.070497 & 0.009234 \\ 0.369754 & 0.383817 & 0.201474 & 0.081367 \\ 0.16807 & 0.39102 & 0.34965 & 0.09126 \end{pmatrix} \end{aligned}$$

As a sanity-check, we can observe that the entry $(2, 2)$ of P^3 is equal to 0.04158, the same value we had obtained previously using Chapman-Kolmogorov's equation directly. Moreover,

assuming we have some estimates of the probability of being at each state at time t (call it $x^t \in \mathbb{R}^n$), we can easily obtain the probabilities at the next time step through the recursive equation $x^{t+1} = P^T x^t$.

3.5. Definitions & Classifications for States and Classes

In the context of Markov chains, we have a series of definitions and classifications which are central to the study of this topic.

We first consider the classification of the relationship between two arbitrary states i, j . We say a state j is accessible ($i \rightarrow j$) by i if there exists n such that $P_{ij}^n > 0$; that is, the probability of transitioning from i to j for some number of steps n is non-zero (and thus it is possible to reach j from i in some number of steps). Moreover, we say two states i, j communicate (and we write $i \leftrightarrow j$) if $i \rightarrow j$ and $j \rightarrow i$. For example, in the Markov chain represented in Figure 13, state 5 is accessible from state 2 but they do not communicate; states 4, 5, however, do communicate.




Figure 13: A Markov chain with multiple classes

A state i may be transient or recurrent. Intuitively, a state is classified as transient if visits to this state are temporary (i.e. we may visit it a few times then possibly never again) and recurrent otherwise. Formally, define the probability f_i of visiting a state i after starting

at state i as follows.

$$f_i = P \left[\bigcap_{n=1}^{\infty} X_n = i \mid X_0 = i \right]$$

Then, a state is classified as transient if $f_i < 1$ and recurrent iff $f_i = 1$. In other words, if the probability at some point of returning to a given state is equal to 1, we have a recurrent state and a transient one otherwise. Consider again the Markov Chain in Figure 13. If we start at state 2, for example, we may transition to state 1 at which point it will be impossible to ever return to state 2; as such $f_2 < 1$ and state 2 is a transient state. Compare that with state 4: while we may transition to state 5, it is the case that $f_4 = 1$, we will revisit state 4 since we are limited in transitioning to the states in red. Therefore $f_4 = 1$ and state 4 is a recurrent state.

We also define the period for a state. Formally, a state i has period d if $d = \max\{d : P_{ii}^n = 0 \text{ for all } n \notin D\}$ where $D = \{kd \mid k \in \mathbb{Z}^+\}$. While the mathematical definition is a bit nebulous, it effectively means that state i has period d if the probability of returning to state i is non-zero (i.e. $P_{ii}^n > 0$) only when n is multiple of d and d is the largest number with this property. For example, if state i is such that $d = 3$, it must be the case that the probability of returning to state i after starting from state i in 4 or 5 time-steps is zero. Since 6 is a multiple of 3, however, it must be that the probability of returning in 6 time-steps is non-zero (and the same is true for 9 time-steps, for example). If $d = 1$, we say that the state is aperiodic. Consider the Markov Chain in Figure 14. If we start the process at state 2 (i.e. $X_1 = 2$), we will be at either state 1 or state 3 on the next time step; in the third time step, however, we are guaranteed to be at state 2 again. More generally, we have that on even time steps $(X_2; X_4; X_6; \dots; X_{2k})$ we will be either on state 1 or state 3. On odd time steps $(X_1; X_3; X_5; \dots; X_{2k+1})$, we are guaranteed to be on state 2. In other words, we have that $P_{22}^2 = 0; P_{22}^4 = 0; P_{22}^{2k} = 0$ and that

$$P_{22} = 0; P_{22}^3 = 0; P_{22}^{2k+1} = 0$$

. By the definition of periodicity, state 2 has period 2.




Figure 14: A Periodic Markov Chain

Finally, we define classes. A class of states T is defined such that for all states $i, j \in T$ we have that $i \leftrightarrow j$ and there is no pair of states $i \in T, k \notin T$ such that $i \leftrightarrow k$. More directly, it's the largest possible subset of states where all states communicate with each other. Interestingly, if a member from a class is transient or recurrent, so is the class as a whole. Therefore, the transient and recurrent properties can be used to classify classes as a whole. In the example in Figure 13, it follows that the classes highlighted in blue and red are recurrent while the one in green is transient. Similarly, it can be shown that periodicity is also a class-wide property.

3.6. Limit Distributions

Much of the discussion involving Markov chains so far made use of conditional expectations; that is, we have been focusing on questions such as the probability of reaching state j given that we are in state i . After all, from the Law of Total Probability, $P(X_n = j) = \sum_i P(X_n = j | X_0 = i)P(X_0 = i)$. For illustration purposes, consider again the Markov chain defined in Figure 13: the probability of reaching state 1 from state 4 is zero for any n while from states 2,3 it is non-zero. As we will see, however, under specific circumstances, the initial condition does not affect the probability of reaching a state in the long-run.

Define an irreducible Markov chain as a Markov chain with one (and only one) communication class; in graph theory, this corresponds to a strongly connected component. For a finite, aperiodic irreducible Markov chain, the limiting probability $\lim_{n \rightarrow \infty} P(X_n = j | X_0 = i)$ exists and it is independent of the starting state i (that is, $\lim_{n \rightarrow \infty} P(X_n = j | X_0 = i) = \lim_{n \rightarrow \infty} P(X_n = j)$). Moreover, the limiting probability $\lim_{n \rightarrow \infty} P(X_n = j)$ is such that $\lim_{n \rightarrow \infty} P(X_n = j) = \pi_j$ where $\pi = [\pi_1; \pi_2; \dots; \pi_n]^T$ is the unique solution to the system of

equations below.

$$\sum_i P_{ij} = 1 \quad \text{Balance Equations}$$

In matrix representation, the balance equations can be written as follows:

$$\pi = \pi P^T \tag{5}$$

The proof of the existence of the above limit and its independence from the initial state rely on the chain being aperiodic and irreducible and it is beyond the scope of this course. But the balance equations can be seen (without formal rigor) through simple algebraic manipulation.

$$\begin{aligned} P(X_{n+1} = j) &= \sum_i P(X_{n+1} = j | X_n = i) P(X_n = i) && \text{(Total Probability)} \\ \lim_{n \rightarrow \infty} P(X_{n+1} = j) &= \lim_{n \rightarrow \infty} \sum_i P(X_{n+1} = j | X_n = i) P(X_n = i) \\ &= \sum_i P(X_{n+1} = j | X_n = i) \lim_{n \rightarrow \infty} P(X_n = i) \\ &= \sum_i P_{ij} \pi_i \end{aligned}$$

These results can be seen numerically as well. Note first that $\pi_j^{(n)} = \lim_{n \rightarrow \infty} P(X_n = j) = \sum_i \pi_i P_{ij}$. This implies the following.

$$\lim_{n \rightarrow \infty} P^n = \begin{matrix} & \begin{matrix} 1 & 2 & \dots & j & \dots & N \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ \vdots \\ j \\ \vdots \\ N \end{matrix} & \begin{matrix} \pi_1 & \pi_2 & \dots & \pi_j & \dots & \pi_N \\ \pi_1 & \pi_2 & \dots & \pi_j & \dots & \pi_N \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \pi_1 & \pi_2 & \dots & \pi_j & \dots & \pi_N \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \pi_1 & \pi_2 & \dots & \pi_j & \dots & \pi_N \end{matrix} \end{matrix}$$

Consider our motivating example in Figure 12. We can calculate π^{100} as an approximation of the limiting distribution. It can be seen that the resulting matrix has the same structure

3 Discrete-Time Markov Chains

as the one above.

$$P = \begin{pmatrix} 0.7 & 0.3 & 0 & 0 \\ 0.49 & 0.42 & 0.09 & 0 \\ 0 & 0.49 & 0.42 & 0.09 \\ 0 & 0 & 0.7 & 0.3 \end{pmatrix} \quad \lim_{n \rightarrow \infty} P^n = P^{100} = \begin{pmatrix} 0.575 & 0.352 & 0.065 & 0.008 \\ 0.575 & 0.352 & 0.065 & 0.008 \\ 0.575 & 0.352 & 0.065 & 0.008 \\ 0.575 & 0.352 & 0.065 & 0.008 \end{pmatrix}$$

Estimating limiting probabilities by matrix multiplication, however, can be very expensive (especially for large Markov chains) and potentially inaccurate. The theoretical method leverages the balance equations described in (5). The balance equations imply the fact that the limit distribution is a stationary distribution, which means that the probabilities of being at each state do not change with time. To see this, let π^t be the (vector consisting of the) probabilities of being in different states at time t and recall that the one step evolution of the probabilities is captured by $\pi^{t+1} = P \pi^t$ (see Section 3.4). Now the balance equations imply that $\pi^t = \pi^{t+1} = \pi$ for large t , meaning that after running the Markov chain for a very long time, the probabilities of being at each state converge and no longer change with time.

Another view on (5) reveals from Linear Algebra that finding π is related to an Eigenvalue Problem: The limiting distribution can be obtained by the (properly normalized, $\sum_i \pi_i = 1$) eigenvector of P^T associated with the eigenvalue $\lambda = 1$. Importantly, $\lambda = 1$ is the eigenvalue largest in absolute value (otherwise P^n would not converge).

```

n = 10
dist = np.zeros((n, 15))

for iter in range(15):
    x = np.random.random(size=4)
    x = x / np.sum(x)
    print(x[0])
    dist[0, iter] = x[1]
    for i in range(1, n):
        x = P.T @ x
        dist[i, iter] = x[1] / np.sum(x)

plt.plot(dist)
plt.xlabel('$n$')


```

```
plt.ylabel( '$x_1$' )
plt.grid()
plt.show()
```

Figure 15: Convergence of several randomly initialized x_1 to $x_1 = 0.352$ after 10 iterations of $x_1 = P^T x_1$

3.7. Limit Distributions in Markov Chains with More than One Class

The discussion in Section 3.6 focuses on limit distributions with a single class. Here, we consider cases involving more than one class. This is the case, for example, of the Markov chain from Figure 13 which is reproduced below.




The main difference we observe when we have more than one class is that the solution to the system of two equations $\pi_j = \sum_i \pi_i P_{ij}$ and $\sum_i \pi_i = 1$ (or alternatively the number of

3 Discrete-Time Markov Chains

eigenvectors with associated eigenvalue = 1) will not be unique. In fact, we will obtain one solution for each recurrent class. This should make sense: if we start in state i belonging to a recurrent class C_1 the long-term probability of being at a state $j \notin C_1$ will be zero (and vice-versa). Consider the following simple Markov chain.



The transition matrix for the Markov chain above is simply the 2 x 2 identity matrix I with eigenvectors $v_1 = [1; 0]^T$; $v_2 = [0; 1]^T$ and eigenvalues $\lambda_1 = \lambda_2 = 1$. This should make sense: if we start at state 0, the limit probability is simply being at state 0 all the time and analogously for state 1. Consider now the following Markov chain and corresponding transition matrix.



Here, the eigenvectors with $\lambda = 1$ are $v_1 = [1; 0; 0]^T$ and $v_2 = [0; 0; 1]^T$. For this Markov chain, we have two possible cases. In Case (I.), the system is initialized in state 0 or 2 and it stays in the corresponding state; in Case (II.), it is initialized in State 1 (a transient state) and eventually transitions to either state 0 and 1 and we are back to Case (I.). This explains why in both cases the limit probability is such that $p_{11} = 0$. Moreover, consider the value of $\lim_{n \rightarrow \infty} P^n$.

$$\lim_{n \rightarrow \infty} P^n = \begin{bmatrix} 1 & 0 & 0 \\ 0.5 & 0 & 0.5 \\ 0 & 0 & 1 \end{bmatrix}$$

3 Discrete-Time Markov Chains

That matrix tells us that in the long-run the probability of transitioning to state 1 is zero and with probability 1/2 each we transition from state 1 to states 2. Consider the following starting probability $x^0 = [0; 1; 0]$. The result $\lim_{n \rightarrow \infty} P^n x^0 = [0; 0.5; 0; 0.5]^T$ implies that after $n \gg 1$ steps, there is zero probability of being at state 1 and probability 1/2 of being at each state 2, 3. While it may continue at state 1 for a while, the probability of staying there forever is zero ($\lim_{n \rightarrow \infty} (0.2)^n = 0$). Likewise, with $x^0 = [1; 0; 0]$, we have $\lim_{n \rightarrow \infty} P^n x^0 = [1; 0; 0]^T$, a result that should match our intuition (starting at state 0, we remain there with probability 1).

With the results so far, we can develop a more systematic approach to analyze multi-class Markov chains. Consider once again the three-class Markov chain from Figure 13. The corresponding transition matrix is given below.

$$P = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.2 & 0 & 0.6 & 0.2 & 0 \\ 0.2 & 0.6 & 0 & 0 & 0.2 \\ 0 & 0 & 0 & 0.3 & 0.7 \\ 0 & 0 & 0 & 0.6 & 0.4 \end{bmatrix} \end{matrix}$$

Let's try to find $\lim_{n \rightarrow \infty} P^n$ without directly computing it | a process that should give further insights into the workings of Markov chains. We know states 2, 3 consist of a transient class, call it C_2 . After an infinite number of steps, we should expect that the probability of transitioning to those states should be zero. Therefore:

$$\lim_{n \rightarrow \infty} P^n = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{bmatrix} ? & 0 & 0 & ? & ? \\ ? & 0 & 0 & ? & ? \\ ? & 0 & 0 & ? & ? \\ ? & 0 & 0 & ? & ? \\ ? & 0 & 0 & ? & ? \end{bmatrix} \end{matrix}$$

Additionally, observe that the probabilities of transitioning out of C_2 to C_1 (consisting of state 1) and to C_3 (consisting of states 4, 5) are the same. This is the case since we can transition to C_1 via state 2 ($P_{21} = 0.2$) or via state 3 ($P_{31} = 0.2$) and we can transition to C_3 via state 2 ($P_{24} = 0.2$) or via state 3 ($P_{35} = 0.2$); therefore, $P_{C_2 C_1} = P_{C_2 C_3} = 0.5$.

3 Discrete-Time Markov Chains

We also know that class C_3 (composed by states 4,5) is an aperiodic, recurrent class. We can therefore treat this class independently from the rest of the chain to calculate the long-term behavior when the initial state is there. From the bottom-right corner of the original matrix, we have that the behavior of that class is ruled by the following transition matrix.

$$\lim_{n \rightarrow \infty} P_{C_3}^n = \begin{matrix} & \begin{matrix} 2 & 3 \end{matrix} \\ \begin{matrix} 4 \\ 5 \end{matrix} & \begin{bmatrix} 0.3 & 0.7 \\ 0.6 & 0.4 \end{bmatrix} \end{matrix}$$

The corresponding normalized eigenvector for the transpose of this matrix is $\mathbf{v}_3 = [0.4615 \ 0.5385]^T$. This gives us the long-term probabilities of being at states 4,5 assuming that we start in one of those two states. If we start in state 1, the probability is 0; if we start in state 2,3, however, we know that the probability of transitioning to class C_3 is 0.5. We weigh the values obtained in $\mathbf{v}_3 = [0.4615 \ 0.5385]^T$ by 0.5 in order to obtain the long-term probability of ending up at states 4,5 assuming we start at C_2 . As a result, we obtain $[0.5 \ 0.4615 \ 0.5 \ 0.5385] = [0.231 \ 0.269]$. Therefore:

$$\lim_{n \rightarrow \infty} P^n = \begin{matrix} & \begin{matrix} 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{bmatrix} ? & 0 & 0 & 0 & 0 \\ ? & 0 & 0 & 0.2308 & 0.2693 \\ ? & 0 & 0 & 0.2308 & 0.2693 \\ ? & 0 & 0 & 0.4615 & 0.5385 \\ ? & 0 & 0 & 0.4615 & 0.5385 \end{bmatrix} \end{matrix}$$

Finally, we repeat the same procedure for class C_1 , likewise weighting the probabilities as we did before. Note that since C_1 is composed of only one state, it follows that $\mathbf{c}_1 = [1]^T$ and $\lim_{n \rightarrow \infty} P_{C_1}^n = [1]$. This yields the following matrix.

$$\lim_{n \rightarrow \infty} P^n = \begin{matrix} & \begin{matrix} 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0.2308 & 0.2693 \\ 0.5 & 0 & 0 & 0.2308 & 0.2693 \\ 0 & 0 & 0 & 0.4615 & 0.5385 \\ 0 & 0 & 0 & 0.4615 & 0.5385 \end{bmatrix} \end{matrix}$$

Indeed, we obtain computationally the following:


$$\lim_{n \rightarrow \infty} P^n = P^{100} = \begin{matrix} & \begin{matrix} 0 & 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0.2308 & 0.2693 \\ 0.5 & 0 & 0 & 0.2308 & 0.2693 \\ 0 & 0 & 0 & 0.4615 & 0.5385 \\ 0 & 0 & 0 & 0.4615 & 0.5385 \end{bmatrix} \end{matrix}$$

3.8. Application: The Gambler's Ruin Problem

In this section, we introduce Discrete-Time Markov chains as a framework to analyze quantitatively and qualitatively a stochastic process. The description of the problem is given below.

Example 10 A gambler begins with an initial wealth w_0 . She places bets on a simple game where with probability p she gets an extra dollar ($w_t = w_{t-1} + 1$) and with probability $q = (1 - p)$ she loses a dollar ($w_t = w_{t-1} - 1$). She continues to play until she is broke ($w_t = 0$) or she has a total of N dollars ($w_t = N$).

For this problem, we can construct a Markov chain where each state corresponds to an amount of wealth (i.e. state k corresponds to w_k). Observe that we have three classes here: two recurrent ($\{0, N\}$) and one transient ($\{i; 0 < i < N\}$). From before, we know that the probability of being at any transient state as $n \rightarrow \infty$ is equal to zero, and thus we will either reach states 0 or N .



3 Discrete-Time Markov Chains

Furthermore, define S_i as the probability of success (i.e. reaching N) conditioned on being at state i . Therefore, for all $i \in \{0; N\}$, we have the following (note that $S_0 = 0; S_N = 1$):

$$S_i = \sum_k P(X_{t+1} = k | X_t = i) S_k \quad (\text{Total Probability})$$

$$S_i = S_{i+1}p + S_{i-1}q$$

$$(p+q)S_i = S_{i+1}p + S_{i-1}q \quad (p+q=1)$$

$$p(S_{i+1} - S_i) = q(S_i - S_{i-1})$$

$$S_{i+1} - S_i = \frac{q}{p}(S_i - S_{i-1}) \quad (\frac{q}{p} = r)$$

At this point, we observe that $S_0 = 0$ (probability of success when broke is zero) and therefore we have that $S_2 - S_1 = r(S_1 - S_0)$. By substitution $S_3 - S_2 = r(S_2 - S_1) = r^2(S_1 - S_0)$ and more generally, by induction, it follows that $S_i - S_{i-1} = r^{i-1}(S_1 - S_0)$.

Next, we can sum up the probabilities up to i as follows to obtain a closed-form equation for S_i as a function of S_1 . Here, we assume $r < 1$.

$$(S_i - S_{i-1}) + (S_{i-1} - S_{i-2}) + \dots + (S_2 - S_1) = r^{i-1}S_1 + r^{i-2}S_1 + \dots + S_1$$

$$S_i - S_1 = S_1(r^{i-1} + r^{i-2} + \dots + 1)$$

$$S_i = S_1(1 + r + r^2 + \dots + r^{i-1})$$

$$S_i = S_1 \sum_{k=0}^{i-1} r^k$$

$$S_i = S_1 \frac{1 - r^i}{1 - r}$$

Here, we use the fact that $S_N = 1$ (the probability of success with wealth N is 1); then.

$$1 = S_N$$

$$1 = S_1 \frac{1 - r^N}{1 - r}$$

$$S_1 = \frac{1 - r}{1 - r^N}$$

Finally, we obtain the following.

$$S_i = S_1 \frac{1 - p^i}{1 - p}$$

$$S_i = \frac{1 - p^i}{1 - p} \frac{1 - p^N}{1 - p^N}$$

$$S_i = \frac{1 - p^i}{1 - p} \frac{1 - p^N}{1 - p^N}$$

With a closed-form equation for S_i , we can see several insights into this process.

- I. If $p = 1$ (i.e. the probabilities of winning or losing are equal), it can be shown that $S_i = i/N$. That is, the probability of success is simply a function of how close you are to the target wealth N (i.e. $\lim_{i \rightarrow N} i/N = 1$).
- II. For a large Markov chain (large N) and $p > 1$ (i.e. the probability of losing in each round is greater than that of winning), $\lim_{N \rightarrow \infty} (1 - p^i)/(1 - p^N) = 0$. In the long-run, the probability of winning approaches zero.
- III. For a large Markov chain (large N) and $p < 1$ (i.e. the probability of winning in each round is less than that of losing), $\lim_{N \rightarrow \infty} (1 - p^i)/(1 - p^N) > 0$. In the long-run, the probability of winning is non-zero.

Figure 16: Probabilities S_i for $p = 1/4$; $p = 3/4$; $p = 2$; $p = 1/2$ (left to right)

3 Discrete-Time Markov Chains

While the values in Figure 16 were obtained analytically, it is possible to also obtain information via simulations. Consider a similar situation where we do not have a wealth limit and instead the gambler plays until she is broke M rounds have occurred.

Define the random variables W_t which corresponds to the wealth at time step t ; define also the indicator random variable X_t which indicates whether the gamble was won at time t . It follows that $W_{t+1} = W_t + (2X_t - 1)$ and recursively we can obtain $E[W_{t+1}] = w_0 + (2p - 1)t$. Consider now 50 simulations with $p = 0.55$, $w_0 = 20$ and $M = 1000$.

```
N, t, w, T_max, p, n_sims = 50, 0, 20, 1000, 0.55, 50
results = np.zeros((n_sims, T_max))

for sim in range(n_sims):
    t = 0
    w = 20
    while t < T_max and w != 0:
        x = 1 if random.random() < p else 0
        if x == 1:
            w += 1
            results[sim, t] = w
        else:
            w -= 1
            results[sim, t] = w
        t += 1

X = np.zeros((n_sims, T_max))
for sim in range(n_sims):
    X[sim] = np.arange(0, 1000)
lr = LinearRegression().fit(X.reshape(-1, 1), results.reshape(-1, 1))

exp = lr.coef_ * np.arange(0, 1000).T + lr.intercept_
mean = np.mean(results, axis=0)

plt.plot(results.T, color= 'blue', linewidth=0.5, alpha=0.5)
plt.plot(mean, color= 'red')
plt.plot(exp, color= 'green')
plt.xlabel( '$t$')
plt.ylabel( '$i$')
plt.show()
```

Figure 17: Simulation results for 50 realizations of the stochastic process. Green is linear fitted model, red is average, blue are individual realizations

Indeed, fitting a linear model $\hat{y}_i = \hat{\alpha}_0 + \hat{\alpha}_1 x_i$, we obtain $\hat{\alpha}_0 = 22.138$ and $\hat{\alpha}_1 = 0.0985$, both close to the actual values $\alpha_0 = 22$ and $(2p - 1) = (2 \cdot 0.55 - 1) = 0.1$. This shows the process of estimating parameters from a stochastic process through a combination of analytical steps and a number of realizations of said process.

3.9. Ergodicity

3.9.1. Definitions


We first define the random variable N_i which represents the number of times that we visit state i such that $N_i = \sum_{j=1}^{\infty} 1(X_j = i)$. Recall we defined f_i as the probability of revisiting state i after starting at state i in any number of steps. Therefore, the probability of the total number of visits to a state i being equal to n should be equal to the probability of visiting it n times and never visiting it again. It therefore follows a geometric distribution with parameter f_i (i.e. $P(N_i = n) = f_i^n(1 - f_i)$). With this formulation, one sees that when $f_i = 1$, i.e., when the state is recurrent, we have $E[N_i] = \infty$ and when state i is transient, $E[N_i] < \infty$.

Now define the minimum number of steps to revisit a state, $\tau_i = \min_n(X_n = i | X_0 = i)$. State i is positive recurrent if $E[\tau_i] < \infty$ and is null recurrent if $E[\tau_i] = \infty$; $E[N_i] = \infty$. Similar to our previous state classifications, positive recurrence and null recurrence are also class properties. If a class is positive recurrent, and a-periodic, it is called ergodic.

3 Discrete-Time Markov Chains

A finite state Markov chain that is irreducible is necessarily positive recurrent because it has one and only one class and all the finite states cannot be transient or null-recurrent. Irreducible Markov chains of infinite (countable) states can, however, be null recurrent.

Example 11 The following infinite Markov chain exemplifies a case of null recurrence.



The probability of returning in 2 time steps is $P(X_0 = 2) = 1 \cdot \frac{1}{2} = \frac{1}{2}$ and for 3 time steps is $P(X_0 = 3) = 1 \cdot \frac{1}{2} \cdot \frac{1}{3} = \frac{1}{6}$. More generally, $P(X_0 = t) = \frac{1}{t!}$. We can then calculate the expected value of the random variable

$$E[N_0] = \sum_{t=2}^{\infty} t \cdot \frac{1}{t!} = \sum_{t=2}^{\infty} \frac{1}{(t-1)!} = 1 + \sum_{t=1}^{\infty} \frac{1}{t!} = 1 + e - 1 = e \quad (\text{diverges})$$

We can also calculate the probability of returning to state 0.

$$f_0 = \sum_{t=2}^{\infty} \frac{1}{t!} = e - 1 = 1$$

Since $f_0 = 1$, we have that $E[N_0] = \infty$. We therefore have that for state 0, $E[N_0] = \infty$; $E[N_0] = 1$ and it is classified as null recurrent state.

3.9.2. Ergodic averages

With ergodic classes, $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} P(X_j = i | X_0 = i) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} P(X_j = i)$ applies. Define $T_i^{(n)}$ as the amount of time spent at state i up to time instant n . Then:

$$T_i^{(n)} = \sum_{j=0}^{n-1} 1_{[X_j = i]}$$

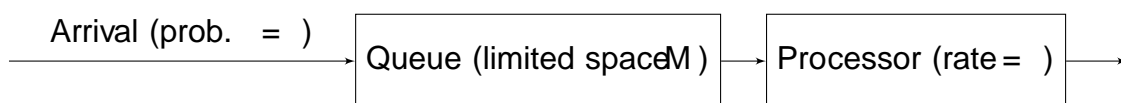
$$E[T_i^{(n)}] = \sum_{j=0}^{n-1} E[1_{[X_j = i]}] \quad \text{(Linearity of Expectation)}$$

$$= \sum_{j=0}^{n-1} P(X_j = i) \quad \text{(Expectation of Indicator r.v.)}$$

In the limit, we have that $\lim_{n \rightarrow \infty} \frac{1}{n} E[T_i^{(n)}] = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} P(X_j = i) = \pi_i$. In the case of ergodic Markov chains, a single run of the Markov chain is representative of the entire stochastic process and thus the expected value is not required (i.e. $\lim_{n \rightarrow \infty} \frac{1}{n} T_i^{(n)} = \pi_i$). This means that in order to find the limit probabilities in ergodic Markov chain, we just need to calculate the average fraction of time spent at each state. Moreover, let $f: S \rightarrow R$; for example, it might be a function associated with a reward for each state $s \in S$. Then, for an ergodic Markov chain, $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} f(X_j) = \sum_{i \in S} \pi_i f(i)$.

3.10. Application: Discrete-Time Queueing Theory

Example 12 Consider the discrete-time communication queue system represented in the block diagram below. In this system, we have a queue with limited length M . Packets enter the queue at each time with probability λ and they leave the queue at each time step with probability μ . Since the events of packet arrival and departure are stochastic, it may happen that a packet that arrives has to wait in the queue until it is its turn to be processed/transmitted. Any packet received when the queue is full is immediately dropped. Determine the event where a packet is dropped. What should be the value of μ such that $P(D) < 10^{-6}$? We assume that arrival and departure from the queue do not take place at the same time step.



In order to be able to analyze the packet drop probability, we need to know when the queue




Figure 18: MC representing the evolution of the number of people in the queue.

is full and for this purpose we have to keep track of the length of the queue. For this purpose, we use Markov chains. The number of states M , each state i , $i = 1; 2; \dots; M$ representing the event that there are i packets in the queue. Transitions happen between consequent states $i - 1, i, i + 1$ because we move from state i to $i + 1$ whenever a packet arrives, and we move from state i to $i - 1$ when a packet leaves. Otherwise, we stay in state

Now we have to quantify the transition probabilities. The probability of a new packet arriving is λ and therefore the probability of transitioning from state i to $i + 1$ is λ for every state $i = 1; \dots; M - 1$. For state $i = M$, we would stay in state i whether a packet arrives or not. Similarly, the probability of transitioning from state i to $i - 1$ is μ for all states $i = 1; 2; \dots; M$. For state 0 , we would stay in state 0 whether a packet leaves or not. The probability of no packet arriving or leaving is $1 - \lambda - \mu$, and for $i = 1; 2; \dots; M - 1$, this determines the probability of staying in state i , for state 0 the probability of staying in state 0 is $1 - \mu$ and for state M it is $1 - \lambda$. The Markov chain is then given as follows. We now consider this Markov process in its steady state where the probability of being in state i is π_i .

Observe that a packet drop occurs with probability $P(D) = \pi_M \lambda$, the probability of being at state M and receiving a new packet.

We thus need to find the steady-state probabilities $\pi_1; \pi_2; \dots; \pi_M$. We do so using the balance

3 Discrete-Time Markov Chains

equations (and the fact $\sum_i P_{ij} = 1$). The balance equations are given below

$$\begin{aligned} p_0 &= p_0(1 - p) + p_1 q \\ p_i &= p_i(1 - p) + p_{i-1} p + p_{i+1} q \\ p_M &= p_M(1 - p) + p_{M-1} p \end{aligned}$$

In order to find values for p_i , we first manipulate the above equations to get the following.

$$\begin{aligned} p_1 &= p_0 p \\ p_{i+1} - p_i &= -p(p_i - p_{i-1}) \\ p_M &= p_{M-1} p \end{aligned}$$

Therefore,

$$\begin{aligned} p_{i+1} - p_i &= -p(p_i - p_{i-1}) \\ &= -p^i (-p_0) \end{aligned}$$

We then obtain the following.

$$p_i = p^i p_0$$

It is only left to find p_0 . We do this by using the fact that the sum of p_i should be equal to 1; that is,

$$\sum_{i=0}^M p^i p_0 = 1$$

Using the fact that

$$\sum_{i=0}^M p^i = \frac{1 - p^{M+1}}{1 - p}$$

Let $x = p$. We then find the following.

$$\begin{aligned} \frac{1 - p^{M+1}}{1 - p} p_0 &= 1 \\ p_0 &= \frac{1 - p}{1 - p^{M+1}} \end{aligned}$$

Plugging the equation for ρ_0 above back into the general equation for ρ_i , we find the following.

$$\rho_i = \frac{1 - \rho}{1 - \rho^M} \rho^i \quad \text{for } i = 0; 1; \dots; M-1$$

If $\lambda = 0.3$; $\mu = 0.5$, for example, we obtain $M = 25$. Note that as λ and μ get closer, the value of M should increase. This makes sense: if the arrival rate is a lot smaller than the processing rate, we should not worry about overflow. With closer values $\lambda = 0.75$; $\mu = 0.8$, we obtain a larger value for M than before with $M = 171$.


3.11. Markov Chains with Rewards

In applications such as statistical decision-making, we model a Markov chain where each state i has a reward r_i associated with it. This means that as the process $X_1; X_2; \dots$ evolves, depending on the states that are visited, a sequence of rewards $\{R(X_i)\}$ is collected. Looking at $\{R(X_1); R(X_2); \dots\}$ as another random process, it should be clear that (i) rewards are functions of the states, and (ii) the reward sequence is not i.i.d.. While we use the term reward, this may capture other values associated with states such as cost or time.


It is not difficult to see that the time-average reward until time n would be $\frac{1}{n} \sum_{i=1}^n R(X_i)$ which would converge, for an ergodic MC, to $\sum_i p_i r_i$ when n is very large. This is called the long-term stationary reward. In many problems, however, we are interested in the expected aggregate reward up to a fixed time n , or until a predetermined stopping time; Examples 13, 14 deal with this type of situation.

3.11.1. Expected Reward until Specific Event

Example 13 We consider a problem where we want to find the expected amount of time until we visit state 1 in the Markov chain given below.



For this problem, we first observe that analyzing the Markov chain below is equivalent. The only change between the two is that we remove the transitions out of state 1; since we are only interested in what happens before we reach state 1, it should be clear that they will yield the same results. Since we cannot transition out of state 1, we call it a trapping state.



Let v_i be the expected reward (or more specifically, in this case, wait time) until we reach state 1 starting from state i . Consider if we had an additional state 5 that with probability 1 transitions to state 1 and that no state transitions into it; in that case, it is straightforward that $v_5 = 1$. For cases such as in this matrix where we do not transition with probability 1 to state 1, the value of v_i is equal to 1 (the reward/cost associated with the state) plus a weighted average of each v_j (where j are the states that i can transition to and the weights are the corresponding probabilities). For states 2, 3, 4, we then have the following.

$$v_1 = 0$$

$$v_2 = 1 + P_{23}v_3$$

$$v_3 = 1 + P_{32}v_2 + P_{34}v_4$$

$$v_4 = 1 + P_{43}v_3$$

3 Discrete-Time Markov Chains

We can write this in matrix form as follows (with the condition that $v_1 = 0$).

$$\begin{aligned}
 \begin{matrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{matrix} &= \begin{matrix} 0 \\ 1 \\ 1 \\ 0 \end{matrix} + \begin{matrix} 0 & 1 & 0 & 0 \\ 0.1 & 0 & 0.9 & 0 \\ 0 & 0.4 & 0 & 0.6 \\ 0.3 & 0 & 0.7 & 0 \end{matrix} \begin{matrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{matrix} \\
 &= \begin{matrix} 1 + 0.1v_1 + 0.9v_3 \\ 1 + 0.4v_2 + 0.6v_4 \\ 1 + 0.3v_1 + 0.7v_4 \\ 0 \end{matrix} \\
 &= \begin{matrix} 0 \\ 1 + 0.9v_3 \\ 1 + 0.4v_2 + 0.6v_4 \\ 1 + 0.7v_3 \end{matrix} \\
 &= \begin{matrix} 0 \\ 9:18 \\ 9:09 \\ 7:36 \end{matrix}
 \end{aligned}$$

Note that since the corresponding Markov chain is ergodic, the solution is unique. This would not be the case in case this property did not hold. The code below simulates this Markov chain to find these statistics.

```

P = np.array([[1, 0, 0, 0],
              [0.1, 0, 0.9, 0],
              [0, 0.4, 0, 0.6],
              [0.3, 0, 0.7, 0]])

n_sims = 5000
n_steps = 80
count = np.zeros((4, n_sims))

```

```
for state in [1, 2, 3]:
    for sim in range(n_sims):
        steps = 0
        x = np.zeros(4)
        x[state] = 1

        for step in range(n_steps):
            probs = P.T @ x
            # Obtain the next state and update state vector
            next_state = np.random.choice(4, p=probs)

            x = np.zeros(4)
            x[next_state] = 1

            # Remember Python is zero-indexed
            if next_state == 0:
                count[state, sim] = steps + 1
                break
            else:
                steps += 1

print(Average time to reach state {}: .format(state),
      np.mean(count[state]))
```

The output is 9:12, 9:17, and 7:28 for states 2, 3, and 4, respectively. This matches our theoretical result nicely.

Example 14 We next consider an application in queueing theory that can best be understood using Markov Chains with rewards. Suppose that we have a queueing system with arrival rate λ , server processing rate μ , and buffer capacity M . As discussed before, the number of customers in the queue can be modeled by the Markov chain shown in Fig. 18. We are now interested in the expected sum of the customer waiting times, assuming that we have i customers waiting, until the system becomes idle again (i.e. reaches state 0 for the first time).

In order to do this, we must (i.) count the amount of time it takes for each existing customer to leave the system and (ii.) count the amount of time customers that arrive in the system to leave. The first step to model this problem is to modify the Markov of Fig. 18 by making

state 0 (the idle state) recurrent so there is no transition out of state 0 when it is visited. This captures the fact that we stop as soon as we arrive to state 0. We refer to this new Markov chain as the modified Markov chain.

Next, to understand the expected sum waiting times, look at Figure 19 (left). The x -axis shows the time of arrival and departure. The y -axis looks at different customers. For example at time t , customer 1 is in the system and so is customer 2. There are then 2 customers in total in the system (represented by the y value).

We are interested in the sum of the wait times from time t where there are $i = 2$ customers in the queue until time t^0 where the system is idle for the first time. This is equal to the total sum of the red intervals in Figure 19 (center). We approach this problem using a double counting trick and translate it into a Markov chain problem with rewards. More precisely, one can see that the total sum of the red x -intervals is the same as the total sum of the green y -intervals in Figure 19 (right). Each green interval at time t has a value exactly equal to the number of customers in the queue at that time step. Let reward $r_i = i$ for each state i . Then, the expected total sum of the green y -intervals is equal to the expected aggregate reward of the MC until being trapped in state 0.

Figure 19: Sum of customer wait times

As before, let v_i be the expected aggregate reward until being trapped in state 0, assuming that we start from state i . Then, can find v_i as the unique solution of the matrix equation below (with the condition $v_0 = 0$).

$$v = r + Pv$$


Here, v is the vector stacking all v_i 's, r is the reward vector $[0; 1; 2; \dots; M]^T$, and P is the

3 Discrete-Time Markov Chains

transition probability matrix of the modified MC (in which state 0 has no outgoing edges and is a trapping state).

Example 15 A drunken man has arrived at a small airport with a main area that leads to three gates. Going from the main area of the airport to gate 1 takes 5 minutes (one way), going to gate 2 takes 2 minutes (one way), and going to gate 3 takes 2 minutes (one way). His flight leaves from gate 3. If he walks to any of the other two gates, he will need to walk back to the main area of the airport. He is so drunk that he chooses each gate uniformly at random with probability $\frac{1}{3}$.

We construct the following Markov chain where state 0 corresponds to the main area and states 1, 2, 3 corresponds to the respective gates.



Then, we follow the same framework from Example 13. Here, however, we use the time to walk to and back each gate. We also know that $t_3 = 0$, since the time to get to state 3 while at state 3 is zero. A simulation similar to that in Example 13 provides empirical evidence of the formulation given above.

3 Discrete-Time Markov Chains

up in state i after n transitions. It could also capture the "value" of state i . Including it in our formulation leads to the following.

$$v_i(n) = r_i + \sum_{j} [P]_{ij} r_j + \sum_{j} P^2_{ij} r_j + \dots + \sum_{j} P^{n-1}_{ij} r_j + \sum_{j} P^n_{ij} u_j$$

We illustrate this with an example.

Example 16 Consider the following two-state Markov chain with rewards $r_1 = 0$ for state 1 and $r_2 = 1$ for state 2. Calculate $v_1(n); v_2(n)$ for $n = 1; 2$.

Here, we see that for $n = 1$ we have

$$v_1(1) = r_1 = 0$$

$$v_2(1) = r_2 = 1$$

and for $n = 2$, we have

$$v_1(2) = r_1 + P_{11}r_1 + P_{12}r_2 = 0$$

$$v_2(2) = r_2 + P_{21}r_1 + P_{22}r_2 = 1$$

Moreover, the long-term (stationary) reward for this chain is 0.5 because for this completely symmetric MC, the stationary distribution is given by $\pi_1 = \pi_2 = 0.5$.

So far, we have seen how to compute expected aggregate reward over a Markov chain with rewards. Markov chains with rewards are key elements in Markov decision processes and here we briefly discuss Markov decision processes and dynamic programming.

3.12. Markov Decision Processes

In many applications, an agent faces sequential decision-making in a stochastic environment that can be modeled by a Markov chain. The agent's actions lead to some rewards but also interact with the environment.

For example, the environment may be an island being explored by a robot. The states may represent the robot's location (e.g. in a grid manner) and the robot's actions may be to go north, east, west, or south. Once actions are chosen, however, the robot may not end up in the expected state (because of the non-smooth paths, or stochasticity); for example, if the robot chooses to go north, he may go north 80% of the time, go right 20% of the time, or may remain in its location if there is an obstacle that the robot is not aware of which may happen with some probability in our model. These probabilities would depend on the current location of the robot (the state) as well as its action.

Also, depending on the action and the state, the agent collects rewards (or costs). These may take the form of a reward for not being damaged by bumping into obstacles or falling into holes. For costs, we may consider the time or distance the robot had to move, for example. Finally, as previously discussed, depending on the action, there may also be a final-state reward, for example, a large reward if the location of a treasure is found.

In general, we have an environment that has state X_t at time t , an agent that takes action A_t at time t and in return receives a reward that is dependent on X_t and the action A_t . Depending on the action and the state, the environment's state consequently changes to X_{t+1} in the next time step. The goal is for the agent to maximize his reward. This reward may be over a finite time horizon, looking at n steps ahead.

In previous sections, without the need to do decision-making, we modeled our problems by a Markov chain (with rewards) that was composed of M states, a transition probability matrix P , and a set of rewards $\{r_i\}$. When modeling decision-making, we consider the following setting: in each state i , the agent has a set of K_i possible actions. Each action k , $k = 1, 2, \dots, K_i$, has a corresponding transition probability matrix $P^{(1)}; P^{(2)}; \dots; P^{(K_i)}$ and reward $r_i^{(k)}$. If the agent chooses action k , reward $r_i^{(k)}$ is collected from state i and the system moves from state i to another state j with probability $P_{ij}^{(k)}$. For multiple time-steps, we consider this process repeatedly and find the maximum aggregated reward. This framework is exemplified below.

Example 17 Consider a Markov chain with states 1 and 2 where we can make decisions (Decision 1 or Decision 2) when in state 2. In state 1, the corresponding reward is $r_1 = 0$

and we would remain in state 1 with probability $P_{11} = 0.99$ or go to state 2 with probability $P_{12} = 0.01$. In state 2, we have two possible actions. Under Decision 1, we get reward $r_2^{(1)} = 1$ and transition out of state 2 to state 1 with probability $P_{21}^{(1)} = 0.01$ and to state 2 with probability $P_{22}^{(1)} = 0.99$. Under Decision 2, the reward is $r_2^{(2)} = 50$ and we move to state 2 with probability $P_{21}^{(2)} = 1$. Figure 20 illustrates this.

Figure 20: Two-state MC with decisions in state 2.

We are now interested in finding optimal policies that would maximize the expected aggregate reward. If we just look at the two Markov chains separately and look at the overall long-term stationary rewards, we see that the Markov chain corresponding to Decision 1 has steady-state probabilities (where the superscript indicates the decision associated with it) equal to $\pi_1^{(1)} = \pi_2^{(1)} = 0.5$ and therefore the long-term (stationary) reward would be $0.5 \cdot 0 + 0.5 \cdot 1 = 0.5$. For the Markov chain of Decision 2, however, we have $\pi_1^{(2)} = 0$ and $\pi_2^{(2)} = 1$ and the long-term (stationary) reward would be $0 \cdot 0 + 1 \cdot 50 = 50$. In this sense, Decision 2 offers higher rewards in the long run. Nonetheless, if we look at the problem over a finite time horizon, this is not the case. For example, if we are deciding between Decisions 1 and Decisions 2 to maximize the expected reward in a single time step while at state 2, then choosing Decision 2 is optimal since the corresponding reward is 50 compared to 1 under the Markov chain of Decision 1. More generally, we are interested in optimal dynamic policies in order to maximize the expected aggregate reward for time steps.

The optimal policy may depend on the state, the duration of time over which rewards are collected, and time. More precisely, consider decision-making at time t , for maximizing aggregate reward over $m+1, \dots, m+n-1$. The agent's decision may depend on m, n ,

3 Discrete-Time Markov Chains

and X_{m+h} , but not on the state or action of the past (i.e. $t = m; m + 1; \dots; m + h - 1$).

This justifies the name Markov Decision Processes, as once again the history of the process given the present does not affect the decision-making | another appearance of the Markov property.

Let's start by considering $n = 1$, i.e., the optimal policy for one time-step. Let's further allow a final state reward u_j for each state j , the value of ending up in that state. In this case, the expected aggregate reward in state i under action k is $r_i^{(k)} + \sum_j P_{ij}^{(k)} u_j$. Let u be the vector of final state rewards; since we are conducting decision-making in one time step only, the optimal action would be the action k such that $v_i^*(1; u) = r_i^{(k)} + \sum_j P_{ij}^{(k)} u_j$ is maximized.

We now move on to $n = 2$. Suppose that at time m , we aim to look for the optimal policy considering time steps m and $m + 1$ (i.e. two time steps). The key is that decision-making at time $m + 1$ is independent of time m (given the state X_{m+1}). So at time $m + 1$, the last step of decision-making for $n = 2$, the optimal decision and the maximum reward can be found following the strategy we discussed for $n = 1$. For time m , however, it is more interesting because we have to consider the immediate reward following the action at time m as well as the next state and its expected aggregate reward moving forward. In particular, the expected maximum aggregate reward in two steps, starting in state i is given by $v_i^*(2; u) = \max_k [r_i^{(k)} + \sum_j P_{ij}^{(k)} v_j^*(1; u)]$. More generally, we have the following for n steps (where $v_i^*(0; u) = u_i$) and a set K_i actions:

$$v_i^*(n; u) = \max_{k \in K_i} [r_i^{(k)} + \sum_j P_{ij}^{(k)} v_j^*(n-1; u)]$$

This recursive equation is known as Bellman's equation, and it was developed by Richard Bellman in the 1960s. This equation plays a central role in a machine learning framework called reinforcement learning. Using it, we may use dynamic programming techniques to break down the decision-making process into smaller problems.

We finish this section by going back to Example 17. For $n = 2$ and $u = 0$, we obtain the

3 Discrete-Time Markov Chains

following. We first compute the values for $n = 1$.

$$\begin{aligned}v_1^1(1; u) &= \max_{k \in K_1} \left(r_1^{(k)} + \sum_j P_{1j}^{(k)} u_j \right) g \\ &= r_1 + \sum_j P_{1j} u_j \quad \text{(no decision to make)}\end{aligned}$$

$$= r_1$$

$$= 0$$

$$v_2^1(1; u) = \max_k \left(r_2^{(1)} + \sum_j P_{2j}^{(1)} u_j ; r_2^{(2)} + \sum_j P_{2j}^{(2)} u_j \right) g$$

$$= \max \{ r_2^{(1)} ; r_2^{(2)} \} g$$

$$= \max \{ 1 ; 50 \} g$$

$$= 50$$

3 Discrete-Time Markov Chains

Then looking at $n = 2$.

$$\begin{aligned}
 v_1^?(2; u) &= \max_{k \in K_1} f r_1^{(k)} + \sum_j P_{1j}^{(k)} u_j g \\
 &= r_1 + \sum_j P_{1j} v_j^?(1; u) && \text{(no decision to make)} \\
 &= 0 + 0.99 v_1^?(0; u) + 0.01 v_2^?(0; u) \\
 &= 0 + 0.99 + 0.01 \cdot 50 = 0.5 \\
 &= 0.5
 \end{aligned}$$

$$\begin{aligned}
 v_2^?(2; u) &= \max_{k \in K_1} f r_1^{(k)} + \sum_j P_{1j}^{(k)} u_j g \\
 &= \max_k f r_2^{(1)} + \sum_j P_{2j}^{(1)} v_j^?(1; u); r_2^{(2)} + \sum_j P_{2j}^{(2)} v_j^?(1; u) g \\
 &= \max_k f r_2^{(1)} + \sum_j P_{2j}^{(1)} v_j^?(1; u); r_2^{(2)} + \sum_j P_{2j}^{(2)} v_j^?(1; u) g \\
 &= \max f 1 + 0.01 v_1^?(1; u) + 0.99 v_2^?(1; u); 50 + 1 v_1^?(1; u) g \\
 &= \max f 1 + 0.01 \cdot 0 + 0.99 \cdot 50; 50 + 1 \cdot 0 g \\
 &= \max f 1 + 0.99 \cdot 50; 50 g \\
 &= 50.5
 \end{aligned}$$

This indicates that in order to maximize the expected aggregate reward in two time steps, the policy is as follows: in state 2, the first time, we choose Decision 1 and the second time, we choose Decision 2. One can show that more generally, the optimal policy is to choose Decision 1 except in the last step. In the last step, the optimal decision would be Decision 2.

4. Arrival Processes

So far, we have modeled time in discrete steps. We now discuss the continuous nature of time and occurrence of events through arrival processes. This topic corresponds to Chapter 5 of Introduction to Probability Models by Sheldon Ross.

4.1. Poisson Processes

Arrival processes model occurrence of events: arrival of a communication packet, arrival of a customer, occurrence of a crisis, arrival of visitors on a webpage, etc. An arrival process may be described by the arrival times (the point in time when each arrival occurred, $S_1; S_2; S_3; \dots$), the interarrival times (the time between events, $T_1; T_2; T_3; \dots$), or the number of arrivals up to a point in time ($N(t); t \geq 0$). In this section, we focus on Poisson arrival processes.

In order to introduce this topic, we first illustrate these quantities through the formulation of this problem in discrete time. At each discrete time index, an arrival occurs with probability p . Therefore, the arrival process as a whole can be modeled as a sequence of binary experiments and thus follows a binomial distribution with probability p . In other words, $N(t) \sim \text{Binom}(t; p)$. Additionally, note that the interarrival time consists of the number of time steps where we do not have an arrival; thus, it follows a geometric distribution and $T_i \sim \text{Geom}(p)$. Since these interarrivals follow a geometric distribution, the memoryless property, $P(T \leq s + t | T \leq t) = P(T \leq s)$ applies: in coin tosses, the number of sequential "heads" does not affect the probability of seeing "tails".




Figure 21: A discrete-time arrival process; here, for example,

$$N(t_1) = 1; N((t_1 + t_2) = 2) = 1; N(t_3) = 3$$

In continuous time, we do not have the notion of a time index. Nonetheless, if we consider the limit of the above Bernoulli process, we will have a Poisson process which is defined on

continuous time domain. Formally, a Poisson process is an arrival process with the following three properties. This forms our first definition of Poisson processes.

- I. Stationary increments: the number of arrivals in an interval depends only on the length of the interval { not when it takes place; that is:

$$P(N(t_2) - N(t_1) = k) = P(N(t_2 - t_1) = k)$$

$$P(N(s+t) - N(t) = k) = P(N(s) = k)$$

- II. Independent increments: the number of arrivals in disjoint time intervals is independent; simply, the number of arrivals in one interval does not affect the number of arrivals in another disjoint interval. Thus, for times $t_1 < t_2 < t_1 < t_2$ we have the following.

$$P(N(t_2) - N(t_1) = j; N(t_1) = k) = P(N(t_2) - N(t_1) = j)P(N(t_1) = k)$$

- III. Poisson pmf of the number of arrivals: the number of arrivals in an interval of length t is Poisson with rate t . I.e., $N(t) \sim \text{Poisson}(t)$.

As an immediate result, for two times s, t with $s < t$, we have the following.

$$P(N(t) = j | N(s) = k) = P(N(t) - N(s) = j - k | N(s) = k) \quad \text{(Property I.)}$$

$$= P(N(t) - N(s) = j - k) \quad \text{(Property II.)}$$

$$= P(N(t - s) = j - k) \quad \text{(Property I.)}$$

$$= \frac{(t - s)^{j - k} e^{-(t - s)}}{(j - k)!} \quad \text{(Property III.)}$$

4.2. Exponential Inter-Arrival Times

One of the properties of Poisson processes is that the inter-arrival times T_1, T_2, \dots are i.i.d. following an Exponential distribution with mean $1/\lambda$. This can be seen as follows. Time until the next arrival is at least t (i.e., $T > t$), if there is no arrival in the interval of length t . So

the corresponding probability is

$$\Pr(T > t) = \Pr(N(t) = 0) = e^{-\lambda t}$$

and consequently

$$\Pr(T \leq t) = 1 - e^{-\lambda t}$$

which points to the cdf of an exponential random variable with mean $\frac{1}{\lambda}$.

Interestingly, an equivalent alternative definition for Poisson processes is based on this property: An arrival process is Poisson with rate λ if its inter-arrival times are i.i.d. distributed according to Exponential with mean $\frac{1}{\lambda}$.

Recall that exponential distributions have the memoryless property, i.e.,

$$\Pr(T > s + t | T > s) = \Pr(T > t):$$

So no matter how long we have waiting for an event, the remaining time until the occurrence of the event is still Exponential with the mean $\frac{1}{\lambda}$. Exponential RVs are the only continuous RVs with the memoryless property. In discrete time RVs, Geometric RVs are the only RVs satisfying the memoryless property.

Consider now the probability of an arrival in an infinitesimally small time h . This means that time from $t = 0$ to $t = S_1$ (i.e., T_1) is less than h .

$$\begin{aligned} P(T < h) &= \int_0^h f_T(t) dt \\ &= \int_0^h \lambda e^{-\lambda t} dt \\ &= \int_0^h \lambda e^{-\lambda t} dt \\ &= 1 - e^{-\lambda h} \\ &= 1 - \left(1 - \frac{\lambda h}{1!} + \frac{\lambda^2 h^2}{2!} - \dots \right) \quad (\text{Taylor Series}) \\ &= \lambda h + o(h) \end{aligned}$$

The notation $o(h)$ comes from asymptotic analysis notation and indicates a class of function whose behavior as $h \rightarrow 0$ approaches zero quicker than a linear function (and thus is negligible for small h).

The above results suggests that at each time, the probability of having an event within the next in nitesimally small amount of time h is λh . For this reason, λ is referred to as the rate of the arrival process. We next argue that the probability of having more than one arrival in an in nitesimal time-interval is negligible. So, the probability of having no events is $1 - \lambda h$

Consider the event of having two arrivals within the small time interval of length h . This implies that $T_1 < h; T_2 < h$. The probability of such events can be computed as follows because $T_1; T_2$ are iid in a Poisson process:

$$\Pr(T_1 < h \& T_2 < h) = \Pr(T_1 < h) \Pr(T_2 < h) = \lambda^2 h^2 = o(h):$$

A third de nition of Poisson processes is based on this property. An arrival process is Poisson with rate λ if it satisfies Properties I and II (stationary and independent increments) and further $\Pr(N(t) = 1) = \lambda t + o(t)$ and $\Pr(N(t) > 1) = o(t)$.

Suppose that we have two independent Poisson processes with rates λ_1 and λ_2 , respectively (e.g. arrival of passengers and arrival of buses). We may be interested in the event that one occurs before the other. Let the time to the arrival of a passenger be T_1 and the time to the arrival of a bus be T_2 . The passenger arriving before the bus corresponds to the event that $T_1 < T_2$ and can be computed as follows:

$$\begin{aligned} P(T_1 < T_2) &= \int_0^{\infty} P(T_1 < t | T_2 = t) f_{T_2}(t) dt \\ &= \int_0^{\infty} F_{T_1}(t) f_{T_2}(t) dt \\ &= \int_0^{\infty} (1 - e^{-\lambda_1 t}) \lambda_2 e^{-\lambda_2 t} dt \\ &= \lambda_2 \int_0^{\infty} e^{-\lambda_2 t} - e^{-(\lambda_1 + \lambda_2)t} dt \\ &= \lambda_2 \left[\frac{1}{\lambda_2} - \frac{1}{\lambda_1 + \lambda_2} \right] \\ &= \frac{\lambda_1}{\lambda_1 + \lambda_2} \end{aligned}$$

The result above should be compatible with your intuition. The probability of seeing one event rather than the other has a chance proportional to the rate of the arrival of the events.

In the homework, you will establish that the minimum of two Exponential RVs (of rates λ and μ) is an Exponential with rate $(\lambda + \mu)$. This comes in handy when we have multiple Poisson arrival processes and we are interested in the combined Poisson process. The time until the first event (of any kind) is then modelled by an Exponential RV and its rate is the sum of the rate of the involved Poisson processes.

5. Continuous-Time Markov Chains

With discrete Markov chains, transitions (to a different or same state) took place at each successive time index n . With continuous-time Markov chains, we use the notation $X(t) = i; X[t : s] = i$ to indicate that we are at state i at time instant $t \in \mathbb{R}^+$ and that we were in state i in the interval from t to s , with $t, s \in \mathbb{R}^+$, respectively. Note that in this context we do not have a notion of a time index/step for transition to occur as it is the case with discrete-time Markov chains. Instead, we remain at state i for some (random) time $T_i \in \mathbb{R}^+$ at which point a transition to a different state occurs. As the name implies, with continuous-time Markov chains we require that the distribution of transition times is memoryless (e.g. follows the Markov property). Since it cannot depend on the past, it cannot be affected by the holding time (the amount of time spent at a given state) and thus T_i must be a random variable such that $P(T_i > t) = P(X[0; t] = i | X(0) = i)$.

Observe that the condition $P(T_i > t) = P(X[0; t] = i | X(0) = i)$ corresponds exactly to the Markov property as shown below.

$$\begin{aligned}
 P(T_i > s + t | T_i > s) &= P(X[0 : s + t] = i | X[0 : s] = i) \\
 &= P(X[s : s + t] = i | X[0 : s] = i) \\
 &= P(X[s : s + t] = i | X(s) = i) \\
 &= P(X[0 : s + t] = i | X(0) = i)
 \end{aligned}$$

From previous sections, we know that the only memoryless continuous distribution (i.e. follows the Markov property) is the exponential distribution; as such, it must be the case that with continuous-time Markov chains we have that the holding time T_i will always follow an exponential distribution. Given this fact, we can define a continuous-time Markov chain as a set of states $i \in S$ with holding times $T_i \sim \exp(-\lambda_i)$ and once transitions occur, they take place according to the transition probabilities of embedded discrete-time Markov chain for each $P_{ij}; i \neq j$.

Given the above, we define some quantities. Let λ_i be the transition rate out of state i (i.e. the holding time for state i is $T_i \sim \exp(-\lambda_i)$) and $q_{ij} = \lambda_i P_{ij}$ be the transition rate from state i to state j . Note that the latter is simply the transition rate of state i weighted by the probability of going from state i to state j and thus $\sum_j q_{ij} = \sum_j \lambda_i P_{ij} = \lambda_i \sum_j P_{ij} = \lambda_i$.

The following results further clarify the underlying meaning of $\lambda_i; q_{ij}$. Note that $P_{ii}(0) = 1$ and $P_{ij}(0) = 0; j \neq i$, that is, the probability of being at state i assuming we are at state i after a time interval of length 0 is equal to 1 and the probability of being at state j assuming we are at state i in a time interval of length 0 is equal to 0, respectively. Assuming we are at state i , the probabilities of being at state i after an infinitesimally small time interval h (i.e. $P_{ii}(h)$) and the probabilities of being at state $j; j \neq i$ after an infinitesimally small time interval h (i.e. $P_{ij}(h)$). We also remind you that as discussed in Section 4.1, for small interval $h \neq 0$, the probability of more than one event taking place is negligible ($o(h)$). Therefore:

$$\begin{aligned} P_{ii}(h) &= \mathbf{P}(T_i > h) \\ &= 1 - \mathbf{P}(T_i \leq h) \\ &= 1 - \lambda_i h + o(h) \end{aligned}$$

$$\begin{aligned} P_{ij}(h) &= \mathbf{P}(X(h) = j | X(0) = i) \\ &= \mathbf{P}(T_i \leq h) P_{ij} + o(h) \\ &= \lambda_i P_{ij} + o(h) \\ &= h q_{ij} + o(h) \end{aligned}$$

Then:

$$\begin{aligned} \lim_{h \downarrow 0} P_{ii}(h) &= \lim_{h \downarrow 0} 1 - \lambda_i h + o(h) \\ \lambda_i &= \lim_{h \downarrow 0} \frac{P_{ii}(h) - 1}{h} \\ \lambda_i &= \lim_{h \downarrow 0} \frac{P_{ii}(h) - P_{ii}(0)}{h - 0} && (P_{ii}(0) = 1) \\ \lambda_i &= \frac{\partial P_{ii}(t)}{\partial t} \Big|_{t=0} && (\text{limit definition of the derivative}) \end{aligned}$$

Likewise:

$$\begin{aligned} \lim_{h \downarrow 0} P_{ij}(h) &= \lim_{h \downarrow 0} \lambda_{ij} h + o(h) \\ q_{ij} &= \lim_{h \downarrow 0} \frac{P_{ij}(h)}{h} \\ q_{ij} &= \lim_{h \downarrow 0} \frac{P_{ij}(h) - 0}{h - 0} \\ q_{ij} &= \lim_{h \downarrow 0} \frac{P_{ij}(h) - P_{ij}(0)}{h - 0} && (P_{ij}(0) = 0) \\ q_{ij} &= \frac{\partial P_{ij}(t)}{\partial t} \Big|_{t=0} && (\text{limit definition of the derivative}) \end{aligned}$$

This further re-enforces the notion of rates for λ_i and q_{ij} . Indeed, λ_i is the instantaneous rate at which the probability of being at state i is decreasing at $t = 0$; similarly, q_{ij} is the rate at which the probability of being at state j (starting at state i) is increasing at $t = 0$. They are, therefore, the rates of transition.

5.1. Chapman-Kolmogorov Equations in Continuous-Time


In Section 3.4, the Chapman-Kolmogorov equation was introduced to calculate multiple-step transition probabilities. When considering continuous-time Markov Chains, those equations transfer almost directly; instead of considering time steps, however, we look at two time

intervals, one with length s and another with length t .

$$\begin{aligned}
 P_{ij}(s+t) &= \mathbf{P}(X(t+s) = j | X(0) = i) \\
 &= \sum_k \mathbf{P}(X(t+s) = j | X(0) = i; X(t) = k) \mathbf{P}(X(t) = k | X(0) = i) \\
 &= \sum_k \mathbf{P}(X(t+s) = j | X(t) = k) \mathbf{P}(X(t) = k | X(0) = i) \quad (\text{Markov property}) \\
 &= \sum_k \mathbf{P}(X(s) = j | X(0) = k) \mathbf{P}(X(t) = k | X(0) = i) \quad (\text{Time invariance}) \\
 &= \sum_k P_{kj}(s) P_{ik}(t)
 \end{aligned}$$

Consider now the following simple Markov chain.

Example 18 A Markov chain with states $0; 1$ is such that the transition from state 0 to 1 occurs with rate q_{01} and the transition from state 1 to 0 occurs with rate q_{10} . What are the limit probabilities for this Markov chain?



Observe first that $\sum_j q_{0j} = q_{01}$; $\sum_j q_{1j} = q_{10}$. We apply the Chapman-Kolmogorov equations with $s = h; h \neq 0$.

$$\begin{aligned}
 P_{00}(t+h) &= \sum_k P_{k0}(h) P_{0k}(t) \\
 &= P_{00}(h) P_{00}(t) + P_{01}(h) P_{10}(t)
 \end{aligned}$$

Note we have all two possibilities here: we either (1) stay in state 0 for a time h and then continue to stay there for another time t or (2) transition from 0 to 1 in time h and the

transition back in time h . Then:

$$\begin{aligned}
 P_{00}(t+h) &= P_{00}(h)P_{00}(t) + P_{01}(t)P_{10}(h) \\
 P_{00}(t+h) &= (1 - q_{01}h)P_{00}(t) + (1 - P_{00}(t))q_{10}h \\
 P_{00}(t+h) &= P_{00}(t) - q_{01}hP_{00}(t) + q_{10}h - P_{00}(t)q_{10}h \\
 P_{00}(t+h) &= P_{00}(t) + q_{10}h - P_{00}(t)(q_{01} + q_{10})h \\
 \frac{P_{00}(t+h) - P_{00}(t)}{h} &= \frac{P_{00}(t)}{h} + q_{10} - P_{00}(t)(q_{01} + q_{10}) \\
 \frac{P_{00}(t+h) - P_{00}(t)}{h} &= q_{10} - P_{00}(t)(q_{01} + q_{10})
 \end{aligned}$$

Applying the limit as $h \rightarrow 0$.

$$\begin{aligned}
 \lim_{h \rightarrow 0} \frac{P_{00}(t+h) - P_{00}(t)}{h} &= \lim_{h \rightarrow 0} q_{10} - P_{00}(t)(q_{01} + q_{10}) \\
 \frac{d}{dt}P_{00}(t) &= q_{10} - P_{00}(t)(q_{01} + q_{10}) \\
 P_{00}'(t) &= q_{10} - P_{00}(t)(q_{01} + q_{10})
 \end{aligned}$$

Note that this is a first-order ordinary differential equation. It can be shown that a solution to this problem is of the form $P_{00}(t) = \frac{q_{10}}{q_{10} + q_{01}} + \frac{q_{01}}{q_{10} + q_{01}} e^{-(q_{10} + q_{01})t}$. Using the initial condition that $P_{00}(0) = 1$, the probability of being at state 0 starting at state 0 after an interval of length 0 is equal to 1, we obtain that $c = \frac{q_{01}}{q_{10} + q_{01}}$; therefore.

$$P_{00}(t) = \frac{q_{10}}{q_{10} + q_{01}} + \frac{q_{01}}{q_{10} + q_{01}} e^{-(q_{10} + q_{01})t}$$

Applying a similar procedure to $P_{11}(t)$, we obtain the following equation after solving the ordinary differential equation.

$$P_{11}(t) = \frac{q_{01}}{q_{10} + q_{01}} + \frac{q_{10}}{q_{10} + q_{01}} e^{-(q_{10} + q_{01})t}$$

Using the fact that $P_{01}(t) = 1 - P_{00}(t)$ and $P_{10}(t) = 1 - P_{11}(t)$:

$$\begin{aligned}
 P_{01}(t) &= \frac{q_{01}}{q_{10} + q_{01}} - \frac{q_{10}}{q_{10} + q_{01}} e^{-(q_{10} + q_{01})t} \\
 P_{10}(t) &= \frac{q_{10}}{q_{10} + q_{01}} - \frac{q_{01}}{q_{10} + q_{01}} e^{-(q_{10} + q_{01})t}
 \end{aligned}$$

With these closed-form equations, we can find the long-term probabilities simply by taking the limit as $t \rightarrow \infty$; then.

$$\begin{aligned}\lim_{t \rightarrow \infty} P_{00}(t) &= \frac{q_{10}}{q_{10} + q_{01}} \\ \lim_{t \rightarrow \infty} P_{01}(t) &= \frac{q_{01}}{q_{10} + q_{01}} \\ \lim_{t \rightarrow \infty} P_{11}(t) &= \frac{q_{01}}{q_{10} + q_{01}} \\ \lim_{t \rightarrow \infty} P_{10}(t) &= \frac{q_{10}}{q_{10} + q_{01}}\end{aligned}$$

This example offers insights into the framework for more complex continuous-time Markov chains. More generally, using our previous results, we have that the following holds.

$$\begin{aligned}\lim_{h \rightarrow 0} P_{ij}(t+h) &= \lim_{h \rightarrow 0} \sum_{k \in S} P_{kj}(h) P_{ik}(t) \\ \lim_{h \rightarrow 0} P_{ij}(t+h) &= \lim_{h \rightarrow 0} P_{jj}(h) P_{ij}(t) + \sum_{k \in S; k \neq j} P_{kj}(h) P_{ik}(t) \\ \lim_{h \rightarrow 0} P_{ij}(t+h) &= \lim_{h \rightarrow 0} P_{jj}(h) P_{ij}(t) + \sum_{k \in S; k \neq j} P_{kj}(h) P_{ik}(t) \\ \lim_{h \rightarrow 0} P_{ij}(t+h) &= \lim_{h \rightarrow 0} (1 - q_j h) P_{ij}(t) + \sum_{k \in S; k \neq j} q_{kj} h P_{ik}(t) \\ \lim_{h \rightarrow 0} P_{ij}(t+h) &= \lim_{h \rightarrow 0} P_{ij}(t) - q_j h P_{ij}(t) + \sum_{k \in S; k \neq j} q_{kj} h P_{ik}(t) \\ \lim_{h \rightarrow 0} \frac{P_{ij}(t+h) - P_{ij}(t)}{h} &= \lim_{h \rightarrow 0} -q_j P_{ij}(t) + \sum_{k \in S; k \neq j} q_{kj} P_{ik}(t) \\ \frac{\partial}{\partial t} P_{ij}(t) &= -q_j P_{ij}(t) + \sum_{k \in S; k \neq j} q_{kj} P_{ik}(t) \\ P_{ij}^{\prime}(t) &= \sum_{k \in S; k \neq j} q_{kj} P_{ik}(t) - q_j P_{ij}(t)\end{aligned}$$

The differential equation above is denominated in the **Kolmogorov forward equation**.

Note that, unlike the previous simple case, there is no simple solution to this differential equation. Observe it corresponds to the sum of the probabilities of going into state j (in time t go from i to k , i.e. $P_{ik}(t)$, and instantly go from k to j , i.e. q_{kj}) minus the probability of out of it (in time t going from i to j , i.e. $P_{ij}(t)$, and instantly out of j , i.e. q_j).

Note that in that equation we set it up such that we go from i to k in time t and instantly to state j . Alternatively, we can set the equations up such that we consider the case where

we transition immediately from state i to k (i.e. q_{ik}) and in t time go from k to j (i.e. $P_{kj}(t)$).

$$\begin{aligned} \lim_{h \downarrow 0} P_{ij}(t+h) &= \lim_{h \downarrow 0} \sum_{k \in S} P_{kj}(t) P_{ik}(h) \\ \lim_{h \downarrow 0} P_{ij}(t+h) &= \lim_{h \downarrow 0} P_{ij}(h) P_{ij}(t) + \sum_{k \in S; k \neq i} P_{kj}(t) P_{ik}(h) \\ \lim_{h \downarrow 0} P_{ij}(t+h) &= \lim_{h \downarrow 0} (1 - q_i h) P_{ij}(t) + \sum_{k \in S; k \neq i} P_{kj}(t) q_{ik} h \\ \lim_{h \downarrow 0} \frac{P_{ij}(t+h) - P_{ij}(t)}{h} &= \lim_{h \downarrow 0} (-q_i P_{ij}(t) + \sum_{k \in S; k \neq i} q_{ik} P_{kj}(t)) \\ \frac{\partial}{\partial t} P_{ij}(t) &= -q_i P_{ij}(t) + \sum_{k \in S; k \neq i} q_{ik} P_{kj}(t) \\ P_{ij}^0(t) &= \sum_{k \in S; k \neq i} q_{ik} P_{kj}(t) - q_i P_{ij}(t) \end{aligned}$$

The result we obtain is what is known as the **Kolmogorov backward equation**. While interpretation is slightly harder than the forward equations, we likewise have the sum of the probabilities of transitioning immediately from i to k multiplied by the probability of reaching k from j in the remaining time t minus the probability of out of i then transitioning in the remaining time t into j .

5.1.1. The Matrix Exponential

As previously mentioned, while we can solve these equations for simple cases such as the two-state Markov chain in Example 18, this is generally not the case. This is primarily due to the fact that we have coupled differential equations (i.e. for each $P_{ij}(t)$ depends on k functions $P_{kj}(t)$). An efficient way of solving coupled differential equations is using the matrix exponential by making use of the fact that $\frac{\partial}{\partial t} e^{\mathbf{A}t} = \mathbf{A}e^{\mathbf{A}t}$.

We first define the meaning of the matrix exponential $e^{\mathbf{A}t}$. While taking the matrix power of a constant may not immediately make sense, we can make use of the Taylor expansion of the exponential to justify the notation. Namely, we have the following.

$$e^{\mathbf{A}t} = \sum_{k=0}^{\infty} \frac{(\mathbf{A}t)^k}{k!} = \mathbf{I} + \mathbf{A}t + \frac{\mathbf{A}^2 t^2}{2!} + \frac{\mathbf{A}^3 t^3}{3!} + \dots$$

Indeed, this definition provides a clear explanation of the fact that $\frac{\partial}{\partial t} e^{\mathbf{A}t} = \mathbf{A}e^{\mathbf{A}t}$.

$$\begin{aligned} \frac{\partial}{\partial t} e^{\mathbf{A}t} &= \frac{\partial}{\partial t} \mathbf{I} + \mathbf{A}t + \frac{\mathbf{A}^2 t^2}{2!} + \frac{\mathbf{A}^3 t^3}{3!} + \dots \\ &= \mathbf{A} + \mathbf{A}^2 t + \frac{\mathbf{A}^3 t^2}{2!} + \dots \\ &= \mathbf{A} \left(\mathbf{I} + \mathbf{A}t + \frac{\mathbf{A}^2 t^2}{2!} + \dots \right) \\ &= \mathbf{A} \sum_{k=0}^{\infty} \frac{(\mathbf{A}t)^k}{k!} \\ &= \mathbf{A}e^{\mathbf{A}t} \end{aligned}$$

Example 19 Find the general solution to the system of equations below.

$$\begin{aligned} \dot{f}(x) &= f(x) + g(x) \\ \dot{g}(x) &= 4f(x) + g(x) \end{aligned}$$

We can write this in matrix form as follows.

$$\begin{aligned} \frac{\partial}{\partial x} \begin{pmatrix} f(x) \\ g(x) \end{pmatrix} &= \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} f(x) \\ g(x) \end{pmatrix} \\ \frac{\partial}{\partial x} \begin{pmatrix} f(x) \\ g(x) \end{pmatrix} &= \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} f(x) \\ g(x) \end{pmatrix} \end{aligned}$$

Such system admits a solution of the form $e^{\mathbf{A}x}$. We can write it as follows by leveraging the spectral decomposition of \mathbf{A} .

$$\begin{aligned} \exp \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} x &= \exp \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} x \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \\ &= \begin{pmatrix} e^{2x} & 0 \\ 0 & e^{2x} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} e^{x} \end{aligned}$$

Then, $e^{A^t}\mathbf{x}$ with \mathbf{x} being a vector of constants such that initial conditions are satisfied is a solution to this system. Consider for example $\mathbf{x} = [1; 0]$, then.

$$\begin{aligned}
 f(x) &= e^{-x/2} + e^{3x/2} \\
 g(x) &= e^{3x} - e^{-x} \\
 &= e^{-x/2} + e^{3x/2} \\
 &\quad e^{3x} - e^{-x}
 \end{aligned}$$

Indeed.

$$\begin{aligned}
 & \frac{d}{dx} \begin{pmatrix} e^{-x/2} + e^{3x/2} \\ e^{3x} - e^{-x} \end{pmatrix} = \begin{pmatrix} -\frac{1}{2}e^{-x/2} + \frac{3}{2}e^{3x/2} \\ 3e^{3x} + e^{-x} \end{pmatrix} \\
 &= \begin{pmatrix} -\frac{1}{2} & \frac{3}{2} \\ 3 & -1 \end{pmatrix} \begin{pmatrix} e^{-x/2} + e^{3x/2} \\ e^{3x} - e^{-x} \end{pmatrix}
 \end{aligned}$$

Returning to the problem at hand, define first the transition rate matrix \mathbf{R} and the continuous-time transition matrix as follows.

$$\mathbf{R} = \begin{pmatrix} q_{11} & q_{12} & \dots & q_{1n} \\ q_{21} & q_{22} & \dots & q_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ q_{n1} & q_{n2} & \dots & q_{nn} \end{pmatrix} \quad \mathbf{P}(t) = \begin{pmatrix} P_{11}(t) & P_{12}(t) & \dots & P_{1n}(t) \\ P_{21}(t) & P_{22}(t) & \dots & P_{2n}(t) \\ \vdots & \vdots & \ddots & \vdots \\ P_{n1}(t) & P_{n2}(t) & \dots & P_{nn}(t) \end{pmatrix}$$

Note then that the matrix products $\mathbf{R}\mathbf{P}(t); \mathbf{P}(t)\mathbf{R}$ correspond, respectively, to the system of differential equations that follows from the Kolmogorov backward and forward equations. In other words, we have $\mathbf{P}'(t) = \mathbf{R}\mathbf{P}(t)$ and $\mathbf{P}'(t) = \mathbf{P}(t)\mathbf{R}$. These are linear systems of differential equations, and therefore we have that their solutions are given by $\mathbf{P}(t) = e^{\mathbf{R}t}$. This means that in order to find the probability of transitioning from state i to state j in an interval of length t , we can find the $(i; j)$ entry of the matrix $e^{\mathbf{R}t}$.

5.2. Application: Continuous-Time Queueing Theory

When discussing discrete-time Markov chains, we introduced queueing systems under the assumption that at each time step the objects under study (packets, customers, etc.) are

processed with some fixed probability. With continuous-time Markov chains, we can further our model by introducing specific (exponential) processing times or equivalently arrivals following a Poisson distribution. We develop these ideas through an example.

Example 20 *A single-server queue is such that arrival and service times are distributed following an exponential distribution with rates λ and μ , respectively. The quantities of interest are the limit probabilities, the expected length of the queue, and the average wait time.*

We model this using a Markov chain where each state i , $i \geq 0$, corresponds to the length of the queue. First note that $q_{i,i-1} = \lambda P_{i,i-1}$ for all $i \geq 1$. It should be clear that $P_{i,i-1} = \frac{\lambda}{\lambda + \mu} P_{i+1}$ and thus the following is true for all $i \geq 1$.

$$q_{i,i-1} = \lambda P_{i,i-1} = \frac{\lambda}{\lambda + \mu} P_{i+1} = \lambda P_{i+1} \frac{1}{\lambda + \mu}$$

$$q_{i,i+1} = \mu P_{i,i+1} = \frac{\mu}{\lambda + \mu} P_i = \mu P_i \frac{1}{\lambda + \mu}$$

Note that throughout this analysis we assume $\lambda > \mu$; otherwise, the queue will grow indefinitely. We then have the following balance equation.

$$P_i = P_{i+1} \frac{\lambda}{\lambda + \mu} + P_{i-1} \frac{\mu}{\lambda + \mu}$$

$$P_i = \frac{\lambda}{\lambda + \mu} P_{i+1} + \frac{\mu}{\lambda + \mu} P_{i-1}$$

$$(\lambda + \mu) P_i = \lambda P_{i+1} + \mu P_{i-1}$$

Note also that $P_0 = P_1$. Following the usual substitutions and the fact that $\sum_{i=0}^{\infty} P_i = 1$, we obtain that $P_i = \left(\frac{\mu}{\lambda}\right)^i P_0$ and $P_0 = 1 - \frac{\mu}{\lambda} = \frac{\lambda - \mu}{\lambda}$.

At this point, we may calculate the other quantities of interest. We first calculate the

expected length of the queue; this corresponds to the expected value of $X(t)$, call it L .

$$\begin{aligned}
 E[X(t)] &= E[L] \\
 &= \sum_i i P_i \\
 &= \sum_i i \left(\frac{\lambda}{\lambda + \mu} \right)^i P_0 \\
 &= \sum_i i \left(\frac{\lambda}{\lambda + \mu} \right)^i (1 - \frac{\lambda}{\lambda + \mu}) \\
 &= (1 - \frac{\lambda}{\lambda + \mu}) \sum_i i \left(\frac{\lambda}{\lambda + \mu} \right)^i \\
 &= (1 - \frac{\lambda}{\lambda + \mu}) \frac{\lambda}{(\lambda + \mu)^2} \\
 &= \frac{\lambda \mu}{\lambda + \mu}
 \end{aligned}$$

This is a sensible result: it's the ratio between the service and service and arrival rates. Note if $\lambda < \mu$ (arrivals are slower than processing), the expected length is shorter; in the same vein, if $\lambda > \mu$ (remember that $\lambda > \mu$), we have that the expected value approaches L .

We may also calculate the wait time W . Let L be the queue length; for a new customer joining this queue, the total time in the system is $T_1 + T_2 + \dots + T_{L+1}$ where T_i is the service time of person i . Note that we have $L + 1$ terms since we need to include that time for serving the new customer. Then:

$$\begin{aligned}
 E[W] &= E \left[\sum_{i=1}^{L+1} T_i \right] \\
 &= E[(L+1)E[T_i]] \\
 &= E[L+1]E[T_i] \\
 &= (E[L] + E[1])E[T_i] \\
 &= \frac{\lambda \mu}{\lambda + \mu} + 1 \cdot \frac{1}{\mu} \\
 &= \frac{1}{\mu} \\
 &= \frac{1}{\mu}
 \end{aligned}$$

As a side-note, one can write the equation above as $E[W] = 1/\mu - E[L]$. Changing notation, we have that $L = W$. This relationship is known as Little's Law and interestingly it holds for any queueing system independent of its structure or distribution.

A simulation of the system with $\lambda = 5$; $\mu = 2$ is provided below. Note that this system is somewhat more complicated to simulate and analyze than previous examples.

```
#####
# Simulation #
#####

Lambda, Mu = 2, 5
N = 5000

interarrivals = np.random.exponential(1/Lambda, size = N)
service_time = np.random.exponential(1/Mu, size = N)

arrivals = np.cumsum(interarrivals)

enter_service_time = np.zeros(N)
leave_service_time = np.zeros(N)

enter_service_time[0] = arrivals[0]
leave_service_time[0] = arrivals[0] + service_time[0]

for customer in range(1, N):
    # Case 1: Previous customer finishes service before next arrival
    if leave_service_time[customer - 1] < arrivals[customer]:
        enter_service_time[customer] = arrivals[customer]

    # Case 2: Previous customer finishes service after next arrival;
    # in this case, the next only starts once the previous is done
    else:
        enter_service_time[customer] = leave_service_time[customer - 1]

    # Each customer always leaves after entering service + service
    leave_service_time[customer] = enter_service_time[customer] + \
        service_time[customer]

#####
# Analysis #
#####
```

```
L = np.zeros(N)

# Assumes maximum length is 10...
# Might have to change depending on Mu, Lambda
count = np.zeros(11)
probs = np.zeros((N, 11))

for customer in range(N):
    # Define t as the arrival time of each person
    t = arrivals[customer]

    # If no wait, queue length is zero
    if arrivals[customer] == enter_service_time[customer]:
        L[customer] = 0

    # If wait, find who was being serviced at the time and
    # calculate difference
    else:
        for other_customer in range(customer):
            if enter_service_time[other_customer] < t < \
                leave_service_time[other_customer]:
                L[customer] = customer - other_customer
                break

    count[int(L[customer])] = count[int(L[customer])] + 1
    curr_prob = count / np.sum(count)
    probs[customer] = curr_prob

print( Average wait time W: , str(np.mean(leave_service_time - arrivals)))
print( Average queue length L: , str(np.mean(L)))
```

The output of this script provides an average wait time of 0.343 time units and an average queue length of 0.672. This matches nicely with the theoretical values of $E[W] = 1/3 = 0.333$ and $E[L] = 2/3 = 0.666$. Plots for the simulation are provided below.

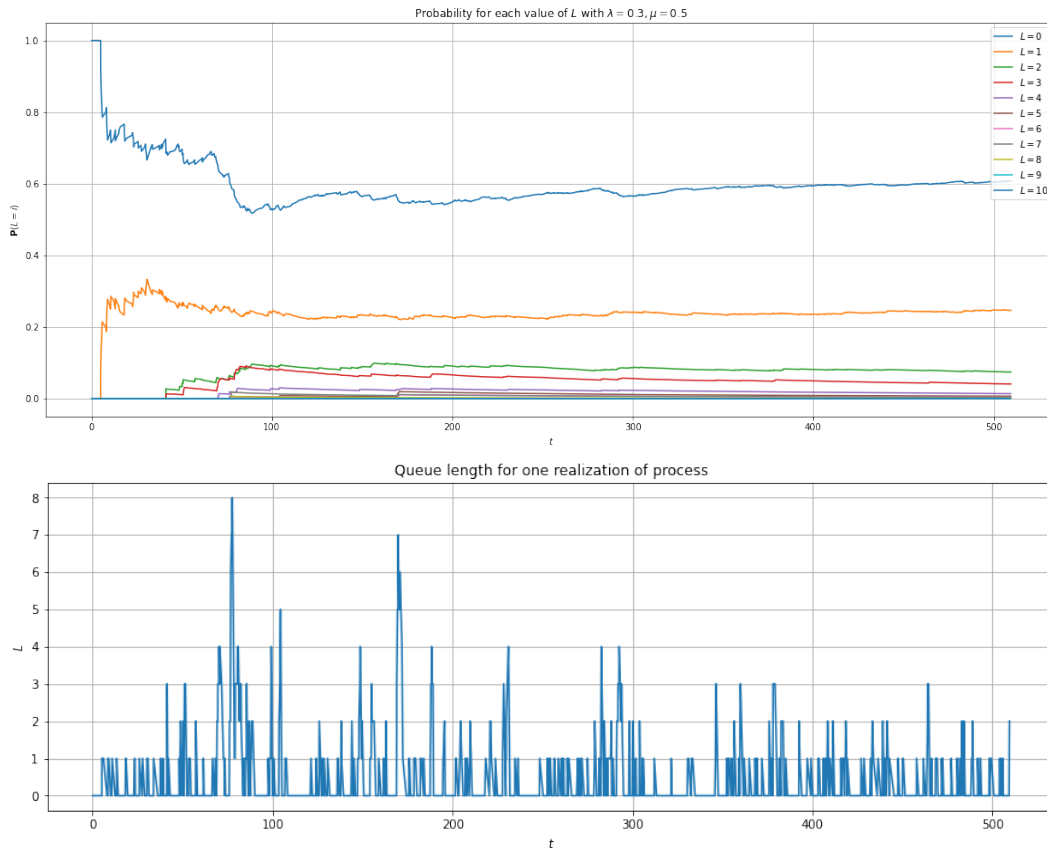


Figure 22: (1) Probabilities for each value of L , (2) one realization of the process

Observe here that $P_0; P_1; P_2$ stabilize around 0.60; 0.25; 0.10, respectively. This also matches theory, with the limit probabilities $P_0 = 1 - \rho = 0.6; P_1 = (\rho)P_0 = 0.24$ and $P_2 = (\rho)^2 P_0 = 0.096$.

Additionally, note that this analysis generalizes (albeit with more complex algebra) to more complex systems (e.g. N servers).

6. Gaussian Processes

The Gaussian (or Normal) distribution is ubiquitous in different fields, from engineering to economics to the natural sciences. This relates directly to the many analytical properties of the Gaussian and its close connection with the Central Limit Theorem. In this section, we introduce a class of stochastic processes governed by the Gaussian distribution and covers the content from Chapter 10 of Sheldon Ross's *Introduction to Probability Models*.

6.1. Properties of the Gaussian Distribution

The univariate Gaussian is a distribution fully described by its mean μ and variance σ^2 . Let X be a random variable such that $X \sim N(\mu; \sigma^2)$; its p.d.f. is given as follows.

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right)$$

The plot of the p.d.f. gives rise to the well-known bell curve. Consider then k normal random variables $X_1; X_2; \dots; X_k$ with corresponding means $\mu_1; \mu_2; \dots; \mu_k$ and covariances $\Sigma_{ij} = E[(X_i - \mu_i)(X_j - \mu_j)]$ for each $i; j$ (note that Σ_{ii} is simply the variance of X_i). The joint distribution of these random variables corresponds to the multivariate formulation of the Gaussian distribution with p.d.f. given below.

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\det(\Sigma)^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\right)$$

Here, \mathbf{X} is the random vector with entries $[X_1; X_2; \dots; X_k]^T; \mathbf{x} \in \mathbb{R}^k$, and covariance matrix Σ such that $[\Sigma]_{ij} = \Sigma_{ij}$. Note that if the random variables are independent, Σ is a diagonal matrix and its determinant is equal to the product of the individual variances. It is worth emphasizing again the fact that the probability distribution of this sequence of random variables is fully described by the mean vector and the covariance matrix.

Important properties of Gaussian random variables (proof omitted) include the following.

- I. Let $X_1 \sim N(\mu_1; \sigma_1^2); X_2 \sim N(\mu_2; \sigma_2^2)$. Then $aX_1 + bX_2 \sim N(a\mu_1 + b\mu_2; a^2\sigma_1^2 + b^2\sigma_2^2)$ for $a; b \in \mathbb{R}$.

- II. If $X \sim N(\mu; \sigma^2)$, the random variable $Y = (X - \mu)/\sigma$ follows the a standard Gaussian distribution, i.e. $Y \sim N(0;1)$.
- III. Under somewhat mild conditions, given an i.i.d. set of observations $\{Z_i\}_{i=1}^n$, the sample mean $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n Z_i$ approaches a Gaussian distribution with mean $E[Z_i]$ and variance $\text{Var}[Z_i]/n$.

6.2. Gaussian Processes: General Aspects

While we have considered so far discrete sequences of Gaussian random variables in the form of vectors \mathbf{X} , we can generalize these to infinite dimensions. This gives rise to **Gaussian processes**; here, we use the notation $X(t)$ to identify the Gaussian variable at time t .

As discussed previously, the distribution of a finite sequence of Gaussian random variables is described by their mean and covariance matrix. This also holds for the case of Gaussian processes, albeit with a slight change of notation to account for its continuous nature. We define the mean value function $\mu(t) = E[X(t)]$ as the expectation of the random process at time t and the autocorrelation function $R(t_i; t_j) = E[X(t_i)X(t_j)]$; note that from the former we can obtain the autocovariance through the formula $C(t_i; t_j) = R(t_i; t_j) - \mu(t_i)\mu(t_j)$ and thus the whole process is explained by the mean and autocorrelation functions.

Note that given the distribution of $X(t)$, we can make use of the tools from the probability toolkit. If, for example, we know the value of $X(t_1)$, this will affect our estimate of $X(t_2)$ for $t_2 > t_1$. Analytically, this can be computed using Bayes's Theorem.

$$f_{X(t_2)|X(t_1)}(t_2|t_1) = \frac{f_{X(t_1)X(t_2)}(t_1,t_2)f_{X(t_1)}(t_1)}{f_{X(t_2)}(t_2)}$$

6.3. Application: Brownian Motion

A direct application of Gaussian processes is in the context of Brownian motion. The name comes from Robert Brown, a botanist who observed that small particles in fluids behaved according to the system described here. Due to the major contributions to this topic by Norbert Wiener, this type of process is also known as Wiener process. A stochastic process $X(t); t \geq 0$ is said to be a Brownian process if the following conditions are respected.

- I. $X(0) = 0$
- II. $fX(t) : t \geq 0$ has stationary and independent increments
- III. $X(t) \sim N(0; \sigma^2 t)$

A simple example of Brownian motion is seen in the continuous-time version of the Gambler's Ruin problem (i.e. time is not indexed by integers and the wealth is a real number). Specifically, we have the following framework: a particle's position at time t is described by a random variable $X(t)$, and at each instant $h > 0$ with equal probabilities the particle moves to the right or to the left by an amount $\pm \sqrt{h}$. Define an index $i = 0; 1; \dots; N$ such that the process ends at time $T = Nh$ and any time step has a corresponding index $t = ih$. Let $Y_i \sim \text{Bernoulli}(1/2)$, then for all time indices $i > 0$, we have the following:

$$\begin{aligned}
 X(t) &= X(ih) \\
 &= X((i-1)h) + (\pm \sqrt{h}) Y_{i-1} \\
 &= X(0) + \sum_{j=1}^{i-1} (\pm \sqrt{h}) Y_j \quad (\text{by induction})
 \end{aligned}$$

We have that $\sum_{j=1}^{i-1} (\pm \sqrt{h}) Y_j \sim N(0; \sigma^2(ih))$ using the Central Limit Theorem. It follows then that $X(t) \sim N(0; \sigma^2(ih))$. Note that this is the case since we have that $X(0) = 0$. Simply, this implies that as time progresses, the variance (i.e. uncertainty) of the particle's position increases with time. The simulation below provides empirical evidence of these results.

```

sigma = 1
h = 0.001
t = np.arange(0, 5, h)
N = len(t)
T = N * h

n_sims = 1000
X = np.zeros((n_sims, N))

upper_interval = np.array([1.95*(np.sqrt(sigma**2 * i * h)) \
    for i in range(N)])
lower_interval = np.array([-1.95*(np.sqrt(sigma**2 * i * h)) \

```

```

for i in range(N)]

for sim in range(n_sims):
    for i in range(1, N):
        Y_i = 1 if random.random() < 0.5 else -1
        X[sim, i] = X[sim, i-1] + sigma * np.sqrt(h) * Y_i

plt.figure(figsize=(10, 5))
plt.plot(X[:25, :].T, linewidth=0.1)
plt.fill_between(np.arange(N), upper_interval, -lower_interval, \
                alpha=0.4, label='95% CI')
plt.legend()
plt.xlabel('$t$')
plt.ylabel('$X(t)$')
plt.show()

```




Figure 23: $N = 25$ realization of Brownian process and 95% confidence interval with
 $\mu = 0$; $\sigma^2 = 1$

Note that there is no requirement for equiprobabilities in the problem formulation. If we adopt a probability $p = \frac{1}{2}(1 + \frac{\rho}{h})$ for the object to move to the right and $1 - p$ to the left, we obtain that $X(t) \sim N(0, t; \frac{\rho}{2}(h))$. That is, the object has a tendency to drift away as t increases; note however that the variance remains the same as before. Figure 24 provides the simulation results of this type of process.

