BUSI 573: STOCHASTIC MODELS IN OPERATIONS MANAGEMENT

Lecture Notes

Süleyman Kerimov
Jones Graduate School of Business
Rice University
Spring 2026

Contents

Acknowledgments
These notes adapt ideas from several expositional texts, including [RP23, Ros11, VDH24, Med02], and from lecture notes by Siddhartha Banerjee, Michel Goemans, Anupam Gupta, Karl Sigman, and János Sztrik. The primary goal of these notes is to serve as a study guide for Ph.D. students in preparation for their qualifying examination. They are not meant to be complete or fully rigorous; some proofs are omitted, incomplete, or only outlined, as they are discussed in class. This course introduces fundamental tools in probability and stochastic processes that are used to model various applications in operations management.

1 Limit Theorems

1.1 Why do we need measure theory?

It is unfortunate that we are starting this course with an informal example. Consider a circle with a radius of 1 meter. We say that two points (a and b) on the edge of the circle belong to the same family if you can go from a to b, or, b to a, by traveling 1 meter around the edge of the circle. Alternatively, you can consider an equivalence relation on the interval $I = [0, 2\pi)$, where $a, b \in I$ belong to the same equivalence class if the distance from a to b is 1, given that you are allowed to loop the interval.

Now each family will pick one of its members as a representative. What is the probability that a point a selected uniformly at random on the edge of the circle is a representative? At first glance, you may suspect that the answer is probably not 1, maybe it is 0.

Note that each family has infinitely many members: once you start from a point a, you will never visit point a again. This is because the circumference of the circle is 2π, which is irrational. Consider the following events

$$A = \{a \text{ is a representative}\},$$

$$B_i = \{a \text{ is } i \text{ steps clockwise from the representative of its family}\},$$

$$C_i = \{a \text{ is } i \text{ steps counter-clockwise from the representative of its family}\}.$$

Since a is chosen uniformly at random, we must have $\mathbb{P}(A) = \mathbb{P}(B_i) = \mathbb{P}(C_i)$ by symmetry. Moreover, since every family has a representative, we must have

$$\mathbb{P}(A) + \sum_{i=1}^{\infty} (\mathbb{P}(B_i) + \mathbb{P}(C_i)) = 1. \tag{1}$$

Let $x = \mathbb{P}(A)$. Then per (1), we get $x + \sum_{i=1}^{\infty} 2x = 1$, which has no solution for $x \in [0, 1]$. The event A is an example of a non-measurable event, because we cannot measure its probability. The reason why the example is not completely formal is that choosing exactly one representative from each family requires the axiom of choice, which we will not discuss.

Discussion 1.1. When $X \sim U[0, 1]$, the following looks contradictory: $1 = \mathbb{P}(0 \le X \le 1) = \sum_{x \in [0,1]} \mathbb{P}(X = x) = 0$?

Discussion 1.2. Argue that the set of rational numbers $\mathbb{Q}$ is countable (Cantor snake), and the set of irrational numbers $\mathbb{R} \setminus \mathbb{Q}$ is uncountable (Cantor's diagonal argument).

1.2 Probability spaces

Let Ω be an arbitrary set of points ω. For our purposes, Ω consists of all the possible results or outcomes ω of an experiment or observation. Next we define a collection of subsets of Ω, where these subsets can be viewed as events for which we can calculate a probability.

Definition 1.3. The collection of sets $\mathcal{F}$ is a sigma field (we also say σ-field), if it has the following properties:

1. $\Omega \in \mathcal{F}$,

2. If $A \in \mathcal{F}$, then $A^c \in \mathcal{F}$,

3. If $A_1, A_2, ... \in \mathcal{F}$, then $\cup_{i=1}^{\infty} A_i \in \mathcal{F}$.

We note that by DeMorgan's law, which states that $(\cup_{i=1}^{\infty} A_i)^c = \cap_{i=1}^{\infty} A_i^c$, (3) in Definition 1.3 can be replaced with: if $A_1, A_2, ... \in \mathcal{F}$, then $\cap_{i=1}^{\infty} A_i \in \mathcal{F}$. Therefore, σ-algebra is simply a non-empty collection of subsets of Ω, which is closed under countable unions, countable intersections, and complement. $(\Omega, \mathcal{F})$ is also referred as a measurable space.

Definition 1.4. A probability space is a measure space with total measure one. It is denoted by $(\Omega, \mathcal{F}, \mathbb{P})$, where

- Ω is a set (also known as sample space)
- $\mathcal{F}$ is a σ-field of subsets of Ω (the sets in $\mathcal{F}$ are also known as events)
- $\mathbb{P}$ is a function from $\mathcal{F}$ to $[0, 1]$ that satisfies $\mathbb{P}(\Omega) = 1$ and if $A_1, A_2, ... \in \mathcal{F}$ are pairwise disjoint, then
$$\mathbb{P}(\cup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mathbb{P}(A_i).$$

We write $\sigma(\mathcal{A})$ to represent the smallest σ-field that contains the collection of events $\mathcal{A}$. We also say that $\sigma(\mathcal{A})$ is the σ-field generated by $\mathcal{A}$. Let's say we want to calculate probabilities on the sample space $\Omega = [0, 1]$ (for example, we want to sample a uniform random number from this interval). One natural candidate for a σ-field $\mathcal{F}$ would be the collection of all possible subsets of Ω. But if you remember our informal example in the introduction, we will not be able to equip this σ-field with a probability measure $\mathbb{P}$, since sets like the set of representatives will belong to $\mathcal{F}$. What is the next natural try? Consider the σ-field generated by the set of all singletons: $\mathcal{F} = \sigma(\{x\}_{x \in [0,1]})$. But now, how can we calculate the probability that if a uniformly sampled random number belongs to the interval $[0, 0.5]$? We cannot represent this interval (which is an uncountable set) with a countable union of singletons. It turns out that the correct σ-field (which is called the Borel σ-field) is the smallest σ-field generated by all intervals of the form $[x, y)$: $\mathcal{B} = \sigma([x, y)_{x<y, x,y \in [0,1]})$. Finally, once you consider the Lebesgue measure, defined by $\mathbb{P}([x, y)) = y - x$ for $0 \leq x \leq y \leq 1$, we are basically good to go.

Discussion 1.5. Argue that singletons, set of rational and irrational numbers are in the Borel σ-field on $[0, 1]$.

Next, we discuss the continuity property of the probability function $\mathbb{P}$. Let $(A_n)_{n \geq 1}$ be a sequence of events, and let

$$\liminf A_n := \cup_{n=1}^{\infty} \cap_{i=n}^{\infty} A_i,$$

$$\limsup A_n := \cap_{n=1}^{\infty} \cup_{i=n}^{\infty} A_i.$$

Note that by definition, we have $\liminf A_n \subset \limsup A_n$: $\liminf A_n$ consists of all outcomes that are contained in all but a finite number of events $(A_n)_{n \geq 1}$, and $\limsup A_n$ consists of all outcomes that are contained in an infinite number of events $(A_n)_{n \geq 1}$. We say that $\lim_n A_n$ exists if $\limsup A_n = \liminf A_n$.

We say that $(A_n)_{n\geq 1}$ is an increasing sequence of events if $A_n \subset A_{n+1}$ for all $n \geq 1$. Note that $\cap_{i=n}^{\infty} A_i = A_n$ in this case, thus, $\liminf A_n := \cup_{n=1}^{\infty} \cap_{i=n}^{\infty} A_i = \cup_{n=1}^{\infty} A_n$. Also note that $\cup_{i=n}^{\infty} A_i = \cup_{i=1}^{\infty} A_i$. Thus, $\limsup A_n := \cap_{n=1}^{\infty} \cup_{i=n}^{\infty} A_i = \cap_{n=1}^{\infty} \cup_{i=1}^{\infty} A_i = \cup_{n=1}^{\infty} A_n$. Therefore, $\lim_n A_n = \cup_{n=1}^{\infty} A_n$.

We say that $(A_n)_{n\geq 1}$ is a decreasing sequence of events if $A_n \supset A_{n+1}$ for all $n \geq 1$. Via similar arguments, it follows that $\lim_n A_n = \cap_{n=1}^{\infty} A_n$ in this case.

Proposition 1.6. *If* $\lim_n A_n = A$, *then* $\lim_n \mathbb{P}(A_n) = \mathbb{P}(A)$.

Proof of Proposition 1.6. First, assume that $(A_n)_{n\geq 1}$ is an increasing sequence. Consider the sequence of events

$$B_{n+1} = A_{n+1} \cap A_n^c, \quad \forall n \geq 0,$$

where we define $A_0 = \emptyset$. First, note that B_n's are disjoint and that

$$\bigcup_{i=1}^{n} B_i = A_n \quad \text{and} \quad \bigcup_{i=1}^{\infty} B_i = A.$$

Then we can conclude that

$$\mathbb{P}(A) = \mathbb{P}\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(B_i) = \lim_n \sum_{i=1}^{n} \mathbb{P}(B_i) = \lim_n \mathbb{P}\left(\bigcup_{i=1}^{n} B_i\right) = \lim_n \mathbb{P}(A_n).$$

The proof when $(A_n)_{n\geq 1}$ is a decreasing sequence of events, i.e., $A_n \supset A_{n+1}$ for all $n \geq 1$, is similar (via De Morgan's law). Now we consider the general case. Let

$$C_n = \bigcup_{i=n}^{\infty} A_i.$$

Note that the C_n's are decreasing. Therefore,

$$\lim_n \mathbb{P}(C_n) = \mathbb{P}\Big(\lim_n C_n\Big) = \mathbb{P}\left(\bigcap_{n=1}^{\infty} C_n\right).$$

Now let

$$D_n = \bigcap_{i=n}^{\infty} A_i.$$

Note that the D_n's are increasing. Therefore,

$$\lim_n \mathbb{P}(D_n) = \mathbb{P}\Big(\lim_n D_n\Big) = \mathbb{P}\left(\bigcup_{n=1}^{\infty} D_n\right).$$

Note that

$$D_n = \bigcap_{i=n}^{\infty} A_i \subset A_n \subset \bigcup_{i=n}^{\infty} A_i = C_n,$$

which implies that

$$\mathbb{P}(D_n) \leq \mathbb{P}(A_n) \leq \mathbb{P}(C_n).$$

Since $\lim_n A_n = A$ exists, we have

$$\liminf A_n = \limsup A_n = A,$$

where

$$\lim_n \mathbb{P}(D_n) = \mathbb{P}\Big(\liminf A_n\Big) \quad \text{and} \quad \lim_n \mathbb{P}(C_n) = \mathbb{P}\Big(\limsup A_n\Big),$$

which concludes the proof. □

1.3 Intermezzo

To refresh your memory of probability theory, please refer to the notes on Canvas. What follows are some definitions included for completeness of this lecture.

Definition 1.7. A random variable X is a function that assigns a real number to each outcome in a sample space Ω. Formally, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Then a function $X : \Omega \to \mathbb{R}$ is called a random variable if it satisfies $\{\omega \in \Omega \,:\, X(\omega) \le x\} \in \mathcal{F}$ for all $x \in \mathbb{R}$. We also say that the random variable X is $\mathcal{F}$-measurable.

Given a random variable X, we define the σ-algebra generated by X, denoted by $\sigma(X)$, as the smallest σ-algebra with respect to which X is measurable, that is

$$\sigma(X) = \sigma\big(X^{-1}(B),\, B \in \mathcal{B}(\mathbb{R})\big) = \{X^{-1}(B) : B \in \mathcal{B}(\mathbb{R})\}.$$

Example 1.8. Consider an experiment, where we flip two coins, and let X be the number of heads. Note that $X^{-1}(\{0\}) = \{TT\}$, $X^{-1}(\{1\}) = \{HT, TH\}, X^{-1}(\{2\}) = \{HH\}$. The σ-algebra must contain the complements too, so that the σ-algebra generated by X is

$$\sigma(X) = \{\emptyset, \{TT\}, \{HT, TH\}, \{HH\}, \{HT, TH, HH\}, \{TT, HT, TH\}, \{HH, HT, TH, TT\}, \{HH, TT\}\}.$$

Definition 1.9. Let X be a random variable and g be any function. If X is discrete, then the expectation of $g(X)$ is defined as

$$\mathbb{E}[g(X)] = \sum_{x \in \Omega} g(x) f(x),$$

where f is the probability mass function of X. If X is continuous then the expectation of $g(X)$ is defined as

$$\mathbb{E}[g(X)] = \int_{-\infty}^{\infty} g(x) f(x)\, dx,$$

where f is the probability density function of X.

1.4 Lebesgue's dominated convergence theorem

We are now concerned with fundamental results on interchanging limits and expectations of random variables. We start with some preliminaries.

Definition 1.10. The sequence of random variables $X_n, n \ge 1$, is said to converge *almost* *surely* to a random variable X, written $X_n \to_{\text{a.s.}} X$, if

$$\mathbb{P}(\lim_{n\to\infty} X_n = X) = 1.$$

An equivalent definition is the following. We say that $X_n \to_{\text{a.s.}} X$ if and only if for any $\epsilon > 0$,

$$\lim_{m\to\infty} \mathbb{P}(|X_n - X| < \varepsilon \text{ for all } n \geq m) = 1.$$

Now consider the following example. Let $U \sim U(0,1)$ and $X_n = n\mathbf{1}_{\{n<1/U\}}$. Note that $X_n \to 0$ a.s., and therefore, $\mathbb{E}[\lim_{n\to\infty} X_n] = 0$. On the other hand, $\mathbb{E}[X_n] = n\mathbb{P}(U < 1/n) = 1$ for all n, and therefore, $\lim_{n\to\infty} \mathbb{E}[X_n] = 1$. One should be very careful when interchanging limits and expectations! Lebesgue's Dominated Convergence Theorem is a beautiful theorem that allows us to interchange limits and expectations safely, i.e., it tells us under what condition we can write $X_n \to X$ a.s. $\Rightarrow \mathbb{E}[\lim_{n\to\infty} X_n] = \lim_{n\to\infty} \mathbb{E}[X_n]$.

Theorem 1.11 (Monotone Convergence Theorem)**.** *If a sequence of non-negative random variables increasingly converge to a random variable (written* $0 \leq X_n \uparrow X$*), then* $\mathbb{E}[X_n] \uparrow \mathbb{E}[X]$.

Theorem 1.11 can be used to prove the following result.

Proposition 1.12 (Fatou's Lemma)**.** *Let* Y *be a random variable with* $\mathbb{E}[|Y|] < \infty$*. Then we have*

- *If* $Y \leq X_n$*, then* $\mathbb{E}[\liminf X_n] \leq \liminf \mathbb{E}[X_n]$.
- *If* $Y \geq X_n$*, then* $\mathbb{E}[\limsup X_n] \geq \limsup \mathbb{E}[X_n]$.

And, Proposition 1.12 can be used to prove the following result.

Theorem 1.13 (Lebesgue's Dominated Convergence Theorem)**.** *Assume that* $X_n \to X$ *a.s., and there is a random variable* Y *with* $\mathbb{E}[Y] < \infty$ *such that* $|X_n| < Y$ *for all* n*. Then* $\mathbb{E}[\lim_{n\to\infty} X_n] = \lim_{n\to\infty} \mathbb{E}[X_n]$.

Proof of Theorem 1.13. Note that $|X_n| < Y$ gives $-Y \leq X_n \leq Y$ for all n. Per Proposition 1.12, we have

$$\mathbb{E}[X] = \mathbb{E}[\liminf X_n] \leq \liminf \mathbb{E}[X_n] \leq \limsup \mathbb{E}[X_n] \leq \mathbb{E}[\limsup X_n] = \mathbb{E}[X],$$

Since $\mathbb{E}[X] = \liminf \mathbb{E}[X_n] = \limsup \mathbb{E}[X_n]$, the limit exists and $\lim_{n\to\infty} \mathbb{E}[X_n] = \mathbb{E}[X]$. □

Example 1.14. Using Proposition 1.11, let's prove that if $X_i \geq 0$ for all $i \geq 1$, then $\mathbb{E}[\sum_{i=1}^{\infty} X_i] = \sum_{i=1}^{\infty} \mathbb{E}[X_i]$. We have

$$\sum_{n=1}^{\infty} \mathbb{E}[X_n] = \lim_{k\to\infty} \sum_{n=1}^{k} \mathbb{E}[X_n] = \lim_{k\to\infty} \mathbb{E}\left[\sum_{n=1}^{k} X_n\right] = \mathbb{E}\left[\sum_{n=1}^{\infty} X_n\right],$$

where the last equality follows from applying the monotone convergence theorem:

$$\sum_{n=1}^{k} X_n \uparrow \sum_{n=1}^{\infty} X_n.$$

The assumption that $X_i \geq 0$ for all $i \geq 1$ is crucial. If we drop this assumption, the statement in Example 1.14 does not hold even if $\sum_{i=1}^{\infty} X_i$ is convergent. Consider $(\alpha_n)_{n=1}^{\infty}$ to be independent and identically distributed (i.i.d.) random variables with $\mathbb{P}(\alpha_1 = \pm 1) = 1/2$, and define a stopping time $\tau = \inf\{n \geq 1 : \sum_{k=1}^{n} \alpha_k = 1\}$. We will cover stopping times later in this lecture, but convince yourself that $\mathbb{P}(\tau < \infty) = 1$. Let $X_n = \alpha_n \mathbf{1}_{\{\tau \geq n\}}$. Then, we have that

$$\sum_{n=1}^{\infty} X_n = \sum_{n=1}^{\infty} \alpha_n \mathbf{1}_{\{\tau \geq n\}} = \alpha_1 + \cdots + \alpha_\tau = 1,$$

so that $\mathbb{E}\left[\sum_{n=1}^{\infty} X_n\right] = 1$.

Since the event $\{\tau \geq n\}$ belongs to $\sigma\{\alpha_1, \ldots, \alpha_{n-1}\}$ (we will discuss more about this later, but this basically means the occurrence of the event $\{\tau \geq n\}$ can be determined on the information available by all realizations $\{\alpha_1, \ldots, \alpha_{n-1}\}$), α_n and $\mathbf{1}_{\{\tau \geq n\}}$ are independent. Thus, we get

$$\mathbb{E}[X_n] = \mathbb{E}[\alpha_n]\,\mathbb{E}[\mathbf{1}_{\{\tau \geq n\}}] = 0, \quad n \geq 1.$$

Thus $\sum_{n=1}^{\infty} \mathbb{E}[X_n] = 0 \neq \mathbb{E}[\sum_{n=1}^{\infty} X_n]$.

1.5 Convergence

Here we discuss two types of convergence: convergence in probability and convergence in distribution. Before that, we present a useful result.

Proposition 1.15 (Borel-Cantelli Lemma). *Let $(A_n)_{n=1}^{\infty}$ be a sequence of events.*

1. *If $\sum_{i=1}^{\infty} \mathbb{P}(A_i) < \infty$, then $\mathbb{P}(\limsup A_n) = 0$.*
2. *If $\sum_{i=1}^{\infty} \mathbb{P}(A_i) = \infty$ and all events are independent, then $\mathbb{P}(\limsup A_n) = 1$.*

Proof of Proposition 1.15. We prove (1) first.

$$\mathbb{P}\left(\limsup A_n\right) = \mathbb{P}\left(\bigcap_{k=1}^{\infty} \bigcup_{j=k}^{\infty} A_j\right) \leq \mathbb{P}\left(\bigcup_{j=k}^{\infty} A_j\right) \quad \text{for any } k \geq 1.$$

Then the result follows since

$$\lim_{k \to \infty} \mathbb{P}\left(\bigcup_{j=k}^{\infty} A_j\right) = 0.$$

Now we prove (2). Let $B = \limsup A_n$. We will show that $\mathbb{P}(B^c) = 0$. Let

$$C_i = \bigcap_{n \geq i} A_n^c.$$

Then we have

$$B^c = \bigcup_{i=1}^{\infty} C_i.$$

Thus, we are done if $\mathbb{P}(C_i) = 0$ for all $i \geq 1$. For each i and $k \geq i$, we have

$$\mathbb{P}(C_i) = \mathbb{P}\left(\bigcap_{n=i}^{\infty} A_n^c\right) \leq \mathbb{P}\left(\bigcap_{n=i}^{k} A_n^c\right) = \prod_{n=i}^{k} (1 - \mathbb{P}(A_n)).$$

Now we utilize the fact that $\log(1-x) \le -x$ for all $x \in [0,1]$. This implies, for all $k \ge i$,

$$\log\bigl(\mathbb{P}(C_i)\bigr) \le \sum_{n=i}^{k} \log(1-\mathbb{P}(A_n)) \le -\sum_{n=i}^{k} \mathbb{P}(A_n).$$

If this is true for all $k \ge i$, then

$$\log\bigl(\mathbb{P}(C_i)\bigr) \le \lim_{k\to\infty} -\sum_{n=i}^{k} \mathbb{P}(A_n) = -\infty.$$

Hence, $\mathbb{P}(C_i) = 0$ for all $i \ge 1$. Note that (1) implies almost surely, only finitely many A_n's will occur, and (2) implies almost surely, infinitely many A_n's will occur. □

Discussion 1.16. Consider the following experiment. We toss a coin every minute. The probability that we get H on minute n is $1/n$. Argue that almost surely, infinitely many heads will occur. If the probability is $1/n^2$, then only finitely many times heads will occur, almost surely.

Definition 1.17. $(X_n)_{n=1}^{\infty}$ converges in probability to a random variable X (written $X_n \to_{\mathrm{p}} X$), if for any $\epsilon > 0$, $\mathbb{P}(|X_n - X| > \epsilon) \to 0$ as $n \to \infty$.

From the statement, it is immediate that almost sure convergence implies convergence in probability. The opposite is not true as the following example shows.

Example 1.18. Let $(X_n)_{n\ge1}$ be a sequence of random variables with

$$\mathbb{P}(X_n = 1) = \frac{1}{n}, \quad \text{and} \quad \mathbb{P}(X_n = 0) = 1 - \frac{1}{n}.$$

Note that for any $\varepsilon > 0$,

$$\mathbb{P}\{|X_n - 0| > \varepsilon\} = \frac{1}{n} \to 0 \quad \text{as } n \to \infty,$$

so that $X_n \to_{\mathrm{p}} 0$. But since $\sum_{n=1}^{\infty} \mathbb{P}(X_n = 1) = \infty$, we have $X_n = 1$ for infinitely many values of n, so we do not have almost sure convergence.

Theorem 1.19. *If $X_n \to_{\mathrm{p}} X$, then there is a subsequence $(X_{n_k})_{k\ge1}$ which converges to X almost surely.*

Proof of Theorem 1.19. Since for every $\varepsilon > 0$, $\mathbb{P}(|X_n - X| > \varepsilon) \to 0$, we can find an index n_1 such that

$$\mathbb{P}(|X_{n_1} - X| > 1/2) < 1/2.$$

Similarly, we can find an index $n_2 > n_1$ such that

$$\mathbb{P}(|X_{n_2} - X| > 1/4) < 1/4.$$

Repeating the argument above, we get a subsequence $(X_{n_k})_{k\ge1}$ such that for all $k \ge 1$,

$$\mathbb{P}(|X_{n_k} - X| > 1/2^k) < 1/2^k.$$

Since the series $\sum_{k=1}^{\infty} \mathbb{P}(|X_{n_k} - X| > 2^{-k})$ converges, per Borel-Cantelli Lemma (Proposition 1.15), only finitely many events

$$A_k = \{|X_{n_k} - X| > 2^{-k}\}$$

occur almost surely. Therefore, $X_{n_k} \to X$ almost surely. $\square$

Definition 1.20. Let F_n be the distribution function of X_n, and let F be the distribution function of X. We say that X_n converges in distribution to X if $\lim_{n\to\infty} F_n(x) = F(x)$ for all x at which F is continuous.

Proposition 1.21. *If $X_n \to_{\mathrm{p}} X$, then $X_n \to_{\mathrm{d}} X$. The converse is not true.*

Example 1.22. Let $(X_n)_{n\geq 1}$ be a sequence of Bernoulli random variables with $p = 1/2$. Also, let $X \sim$ Bernoulli(1/2). Then clearly $X_n \to_{\mathrm{d}} X$. But we don't have convergence in probability, since $\mathbb{P}(|X_n - X| \geq \epsilon) = 1/2$ for $\epsilon \in (0,1)$ and for any $n \geq 1$.

Example 1.23. Let $(X_n)_{n\geq 1}$ be a sequence of random variables with

$$F_n(x) = \begin{cases} 1 - \left(1 - \dfrac{1}{n+1}\right)^{(n+1)x}, & x > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Then $X_n \to_{\mathrm{d}} X$, where $X \sim \exp(1)$. Clearly, for $x \leq 0$, $F_n(x) = F_X(x)$. For $x \geq 0$, we also have

$$\lim_{n\to\infty} F_n(x) = 1 - \lim_{n\to\infty} \left(1 - \frac{1}{n+1}\right)^{(n+1)x} = 1 - e^{-x} = F_X(x).$$

1.6 Law of large numbers

We will discuss more about probability inequalities later in the course, but we need some of them now.

Proposition 1.24 (Markov's inequality)**.** *If X is a nonnegative random variable, then for any $a > 0$ we have*

$$\mathbb{P}(X \geq a) \leq \frac{\mathbb{E}[X]}{a}.$$

Proof of Proposition 1.24. Let $\mathbb{1}_{\{X\geq a\}}$ be the indicator function, which is 1 if $X \geq a$, and it is 0 otherwise. Since $X \geq 0$, clearly we have $a\mathbb{1}_{\{X\geq a\}} \leq X$. Taking expectations proves the result. $\square$

Discussion 1.25. Here is a stronger version of Markov's inequality. If X is a nonnegative random variable, then for any $a > 0$ we have

$$\mathbb{P}(X \geq U \cdot a) \leq \frac{\mathbb{E}[X]}{a},$$

where $U \sim U(0,1)$.

Proposition 1.26 (Chebyshev's inequality)**.** *If X is a random variable with $Var[X] < \infty$, then for any $b > 0$ we have*

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq b) \leq \frac{Var(X)}{b^2}.$$

Proof of Proposition 1.26. Since $(X-\mathbb{E}[X])^2$ is a nonnegative random variable, per Markov's inequality with $a = b^2$, we get

$$\mathbb{P}((X - \mathbb{E}[X])^2 \geq b^2) \leq \frac{\mathbb{E}[(X - \mathbb{E}[X])^2]}{b^2},$$

$$\Rightarrow \mathbb{P}(|X - \mathbb{E}[X]| \geq b) \leq \frac{Var(X)}{b^2}.$$

□

Theorem 1.27 (The Weak Law of Large Numbers). *If $(X_i)_{i=1}^{\infty}$ are i.i.d. with $\mu := \mathbb{E}[X_1] < \infty$, then for any $\epsilon > 0$,*

$$\lim_{n\to\infty} \mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^{n} X_i - \mu\right| > \epsilon\right) = 0$$

Proof of Theorem 1.27. Note that the expectation of $\frac{1}{n}\sum_{i=1}^{n} X_i$ is μ, and its variance is $\frac{\sigma^2}{n}$. Then per Chebyshev's inequality, we have

$$\lim_{n\to\infty} \mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^{n} X_i - \mu\right| > \epsilon\right) \leq \lim_{n\to\infty} \frac{\sigma^2}{n\epsilon^2} = 0.$$

□

Theorem 1.28 (The Strong Law of Large Numbers). *If $(X_i)_{i=1}^{\infty}$ are i.i.d. with $\mu = \mathbb{E}[X_1] < \infty$, then*

$$\mathbb{P}\left(\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} X_i = \mu\right) = 1.$$

Proof of Theorem 1.28. We will prove a weaker version of the statement, where we assume that $K := \mathbb{E}[X_1^4] < \infty$. Further assume that $\mu = 0$, and we generalize the proof in the end. Let $S_n = \sum_{i=1}^{n} X_i$ and consider

$$\mathbb{E}[S_n^4] = \mathbb{E}\left[(X_1 + \cdots + X_n)(X_1 + \cdots + X_n)(X_1 + \cdots + X_n)(X_1 + \cdots + X_n)\right].$$

Expanding the right-hand side, we get terms of the form

$$X_i^4, \quad X_i^3 X_j, \quad X_i^2 X_j^2, \quad X_i^2 X_j X_k, \quad \text{and} \quad X_i X_j X_k X_l,$$

where $i \neq j \neq k \neq l$. Thanks to our independence assumption, we have

$$\mathbb{E}[X_i^3 X_j] = \mathbb{E}[X_i^3]\mathbb{E}[X_j] = 0,$$

$$\mathbb{E}[X_i^2 X_j X_k] = \mathbb{E}[X_i^2]\mathbb{E}[X_j]\mathbb{E}[X_k] = 0,$$

$$\mathbb{E}[X_i X_j X_k X_l] = 0.$$

For given pair i and j, there will be $\binom{4}{2} = 6$ terms in the expansion in the form of $X_i^2 X_j^2$. Thus, we get

$$\mathbb{E}[S_n^4] = n\mathbb{E}[X_1^4] + 6\binom{n}{2}\mathbb{E}[X_1^2 X_2^2].$$

Using independence again, $\mathbb{E}[S_n^4] = nK + 3n(n-1)\mathbb{E}[X_1^2]^2$. Now, since $0 \leq \text{Var}(X_1) = \mathbb{E}[X_1^2] - (\mathbb{E}[X_1])^2$, we have

$$(\mathbb{E}[X_i^2])^2 \leq \mathbb{E}[X_i^4] = K.$$

Therefore, we have that

$$\mathbb{E}[S_n^4] \leq nK + 3n(n-1)K,$$

which implies that

$$\mathbb{E}\left[\frac{S_n^4}{n^4}\right] \leq K/n^3 + 3K/n^2.$$

Therefore, it follows that

$$\mathbb{E}\left[\sum_{n=1}^{\infty} \frac{S_n^4}{n^4}\right] = \sum_{n=1}^{\infty} \mathbb{E}\left[\frac{S_n^4}{n^4}\right] < \infty.$$

Now, for any $\epsilon > 0$, it follows from Markov's inequality that

$$\mathbb{P}\left(\frac{S_n^4}{n^4} > \epsilon\right) \leq \mathbb{E}\left[\frac{S_n^4}{n^4}\right] / \epsilon,$$

and therefore,

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\frac{S_n^4}{n^4} > \epsilon\right) < \infty,$$

which implies by the Borel–Cantelli lemma that $S_n^4/n^4 > \epsilon$ for only finitely many n's, almost surely. Since this is true for all $\epsilon > 0$, we can thus conclude that almost surely, we have

$$\lim_{n\to\infty} \frac{S_n^4}{n^4} = 0.$$

If $S_n^4/n^4 \to 0$, then we must also have $S_n/n \to 0$. Hence, we have proven that, almost surely,

$$\frac{S_n}{n} \to 0 \quad \text{as } n \to \infty.$$

For the case when $\mu \neq 0$, we can apply the same arguments to the random variables $X_i - \mu$ to obtain that, almost surely, we have

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (X_i - \mu) = 0.$$

□

2 Stein's Method

Stein's method is a powerful tool that helps to prove various central limit theorems and quantify the distance between two probability distributions. Recall the vanilla version of the central limit theorem.

Theorem 2.1 (Central Limit Theorem). *Let $(X_n)_{n\geq 1}$ be a sequence of independent and identically distributed random variables. Let $\mu = \mathbb{E}[X_1]$ and $\sigma^2 = Var(X_1) < \infty$. Let $S_n = \sum_{i=1}^n X_i$. Then we have*

$$\frac{S_n - n\mu}{\sigma\sqrt{n}} \to_{\mathrm{d}} \mathcal{N}(0,1),$$

where recall that $\mathcal{N}(0,1)$ is the standard normal distribution with density $f(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}}$.

But, under stronger assumptions, one can be more explicit to quantify the error in the approximation, and be more precise regarding the rate of convergence.

Theorem 2.2 (Berry-Esseen Theorem). *Let $(X_n)_{n\geq 1}$ be a sequence of independent and identically distributed random variables with $\mathbb{E}[X_1] = 0$ and $Var(X_1) = 1$, and assume that $\mathbb{E}[|X_1|^3] < \infty$. Let ϕ be the cumulative distribution function of $\mathcal{N}(0,1)$. Then we have*

$$\left|\mathbb{P}\left(\frac{S_n}{\sqrt{n}} \leq x\right) - \phi(x)\right| \leq 7.59\frac{\mathbb{E}[|X_1|^3]}{\sqrt{n}}.$$

This constant 7.59 was later improved by various papers. What central limit theorems suggest is that common events can be approximated by the normal distribution. But, in the context of the rare events, Poisson distribution provides a good approximation as well.

Theorem 2.3 (Poisson's Law of Small Numbers). *Let $X \sim \mathrm{Bin}(n, \lambda/n)$, $\lambda > 0$. Then for any $k \in \mathbb{N}$, we have*

$$\mathbb{P}(X = k) \to e^{-\lambda}\frac{\lambda^k}{k!} = \mathbb{P}(\mathrm{Poi}(\lambda) = k),$$

as $n \to \infty$.

Proof of Theorem 2.3. We have

$$\begin{aligned}\mathbb{P}(X = k) &= \binom{n}{k}(\lambda/n)^k(1-\lambda/n)^{n-k}\\ &= \frac{n(n-1)\cdots(n-k+1)}{k!}(\lambda/n)^k(1-\lambda/n)^{n-k}\\ &= (1-\lambda/n)^n \cdot \frac{\lambda^k}{k!} \cdot \frac{n}{n}\frac{n-1}{n}\cdots\frac{n-k+1}{n} \cdot (1-\lambda/n)^{-k}.\end{aligned}$$

For fixed k, as $n \to \infty$,

$$\frac{n}{n}\frac{n-1}{n}\cdots\frac{n-k+1}{n} \to 1, \quad (1-\lambda/n)^k \to 1.$$

Now we use the fact that $\exp(-p/(1-p)) \leq 1 - p \leq \exp(-p)$ for all $p \in (0,1)$.

$$\exp\left(-\frac{\lambda/n}{1-\lambda/n}\right) \leq 1 - \frac{\lambda}{n} \leq \exp(-\lambda/n),$$

so that

$$\exp\left(-\frac{\lambda}{1-\lambda/n}\right) \leq \left(1-\frac{\lambda}{n}\right)^n \leq \exp(-\lambda).$$

Therefore, we can conclude

$$\left(1-\frac{\lambda}{n}\right)^n \to \exp(-\lambda) \quad \text{as } n \to \infty.$$

□

Theorem 2.3 implies that $\text{Bin}(n, \lambda/n) \to_{\text{d}} \text{Poi}(\lambda)$. What follows is a discussion around bounding the distance between two probability distributions (e.g., distance between $\text{Bin}(n, \lambda/n)$ and $\text{Poi}(\lambda)$). Before introducing a powerful tool called *coupling*, let us first formalize what we mean by distance.

Definition 2.4. For two probability measures μ and ν, we define a probability metric as

$$d_{\mathcal{H}}(\mu, \nu) := \sup_{h \in \mathcal{H}} \left| \int h(x) d\mu(x) - \int h(x) d\nu(x) \right|,$$

where $h(\cdot)$ is called a test function, and $\mathcal{H}$ is the family of test functions. Similarly, for two random variables W and Z, the probability metric has the form

$$d_{\mathcal{H}}(W, Z) := \sup_{h \in \mathcal{H}} \left| \mathbb{E}[h(W)] - \mathbb{E}[h(Z)] \right|$$

Here are some examples of probability metrics. Let $X \sim \mu$ and $Y \sim \nu$.

1. If $\mathcal{H} = \{\mathbb{1}_{\{\cdot \leq x\}} : x \in \mathbb{R}\}$, then we get the Kolmogorov-Smirnov metric, which is denoted by d_K. Thus, $d_K(\mu, \nu) = \sup_{x \in \mathbb{R}} |F_\mu(x) - F_\nu(x)| = \sup_{x \in \mathbb{R}} |\mathbb{P}(X \leq x) - \mathbb{P}(Y \leq x)|$, and it can be interpreted as the maximum distance between distribution functions.

2. If $\mathcal{H} = \{\mathbb{1}_{\{\cdot \in A\}} : A \in \mathcal{B}(\mathbb{R})\}$, then we get the total variation metric, which is denoted by d_{TV}. Thus, $d_{TV}(\mu, \nu) = \sup_{A \in \mathcal{B}(\mathbb{R})} |\mu(A) - \nu(A)| = \sup_{A \in \mathcal{B}(\mathbb{R})} |\mathbb{P}(X \in A) - \mathbb{P}(Y \in A)|$. The total variation metric us the main metric we use for approximation by discrete distributions.

It is immediate that for two random variables X and Y, $d_K(X, Y) \leq d_{TV}(X, Y)$. The following lemma gives a nice characterization of the total variance distance.

Discussion 2.5. Show that if X and Y are two discrete random variables on Ω, then

$$d_{TV}(X, Y) = \frac{1}{2} \sum_{w \in \Omega} |\mathbb{P}(X = \omega) - \mathbb{P}(Y = \omega)|.$$

Discussion 2.6. Let F and G be the distribution functions with continuous densities f and g, respectively, i.e.,

$$\mu(A) = \int_A f(x)\, dx, \qquad \nu(A) = \int_A g(x)\, dx, \tag{2}$$

for all measurable sets $A \subseteq \mathbb{R}$. Then we have

$$d_{\mathrm{TV}}(f, g) = \frac{1}{2} \int_{-\infty}^{\infty} |f(x) - g(x)| \, dx. \tag{3}$$

2.1 Coupling

Definition 2.7 (Coupling of random variables)**.** The random variables $(\hat{X}_1, \ldots, \hat{X}_n)$ are a coupling of the random variables $(X_1, \ldots, X_n)$ when $(\hat{X}_1, \ldots, \hat{X}_n)$ are defined on the same probability space, and are such that the marginal distribution of $\hat{X}_i$ is the same as that of X_i for all $i = 1, \ldots, n$, i.e., for all measurable subsets $\mathcal{E}$ of $\mathbb{R}$,

$$\mathbb{P}(\hat{X}_i \in \mathcal{E}) = \mathbb{P}(X_i \in \mathcal{E}). \tag{4}$$

Note that the following is a trivial coupling: take $(\hat{X}_1, \ldots, \hat{X}_n)$ to be independent, with $\hat{X}_i$ having the same distribution as X_i. The following is another coupling: if $X, Y, U \sim U(0, 1)$, then $(U, 1 - U)$ is a coupling of (X, Y).

Now let X and Y be two discrete random variables with

$$\mathbb{P}(X = x) = p_x, \qquad \mathbb{P}(Y = y) = q_y, \qquad x \in \mathcal{X},\ y \in \mathcal{Y}.$$

The following result links the total variation distance between two discrete random variables and a coupling of them.

Theorem 2.8 (Maximal coupling)**.** *For any two discrete random variables X and Y, there exists a coupling $(\hat{X}, \hat{Y})$ of X and Y such that*

$$\mathbb{P}(\hat{X} \neq \hat{Y}) = d_{\mathrm{TV}}(p, q), \tag{5}$$

while, for any coupling $(\hat{X}, \hat{Y})$ of X and Y,

$$\mathbb{P}(\hat{X} \neq \hat{Y}) \geq d_{\mathrm{TV}}(p, q). \tag{6}$$

Moreover, the maximal coupling $(\hat{X}, \hat{Y})$ satisfies the following:

$$\mathbb{P}(\hat{X} = \hat{Y} = x) = \min(p_x, q_x), \tag{7}$$

$$\mathbb{P}(\hat{X} = x,\ \hat{Y} = y) = \frac{\max(p_x - q_x, 0) \cdot \max(q_y - p_y, 0)}{d_{TV}(p, q)}, \qquad x \neq y. \tag{8}$$

Theorem 2.9 (Poisson limit for binomial random variables)**.** *Let $(I_i)_{i=1}^n$ be independent with $I_i \sim \mathrm{Bernoulli}(p_i)$, and let $\lambda = \sum_{i=1}^n p_i$. Let $X = \sum_{i=1}^n I_i$, and let Y be a Poisson random variable with parameter λ. Then, there exists a coupling $(\hat{X}, \hat{Y})$ of X and Y such that*

$$\mathbb{P}(\hat{X} \neq \hat{Y}) \leq \sum_{i=1}^n p_i^2. \tag{9}$$

Proof of Theorem 2.9. Let $J_i \sim \mathrm{Poi}(p_i)$ and assume that $(J_i)_{i=1}^n$ are independent. Note that the respective mass functions are

$$p_{i,x} = \mathbb{P}(I_i = x) = p_i^x (1 - p_i)^{1-x}, \qquad q_{i,x} = \mathbb{P}(J_i = x) = e^{-p_i} \frac{p_i^x}{x!} \tag{10}$$

Let $(\hat{I}_i, \hat{J}_i)$ be a coupling of I_i, J_i, where $(\hat{I}_i, \hat{J}_i)$ are independent for all i. Per Theorem 2.8, for each pair I_i, J_i, the maximal coupling $(\hat{I}_i, \hat{J}_i)$ satisfies

$$\mathbb{P}(\hat{I}_i = \hat{J}_i = x) = \min(p_{i,x}, q_{i,x}) = \begin{cases} 1 - p_i, & x = 0 \\ p_i e^{-p_i}, & x = 1 \\ 0, & x \geq 2 \end{cases} \tag{11}$$

since $1 - p_i \leq e^{-p_i}$ for $x = 0$. Since $1 - e^{-p_i} \leq p_i$, we have

$$\mathbb{P}(\hat{I}_i \neq \hat{J}_i) = 1 - \mathbb{P}(\hat{I}_i = \hat{J}_i) = 1 - (1 - p_i) - p_i e^{-p_i} = p_i(1 - e^{-p_i}) \leq p_i^2. \tag{12}$$

Next, let $\hat{X} = \sum_{i=1}^n \hat{I}_i$ and $\hat{Y} = \sum_{i=1}^n \hat{J}_i$. Then, $\hat{X}$ has the same distribution as $X = \sum_{i=1}^n I_i$, and $\hat{Y}$ has the same distribution as $Y = \sum_{i=1}^n J_i \sim \mathrm{Poi}(p_1 + \cdots + p_n)$. Per Boole's inequality[1] and (12), we have

$$\mathbb{P}(\hat{X} \neq \hat{Y}) \leq \mathbb{P}\left(\bigcup_{i=1}^n \{\hat{I}_i \neq \hat{J}_i\}\right) \leq \sum_{i=1}^n \mathbb{P}(\hat{I}_i \neq \hat{J}_i) \leq \sum_{i=1}^n p_i^2. \tag{13}$$

□

2.2 Stein-Chen Method

Now we discuss the Stein-Chen Method, which upper bounds the total variation metric between W and Z, where W is some random variable and Z is a Poisson random variable. That is, we want to show that

$$d_{\mathrm{TV}}(W, \mathrm{Poi}(\lambda)) := \sup_{A \subset \mathbb{Z}_{\geq 0}} |\mathbb{P}(W \in A) - \mathbb{P}(\mathrm{Poi}(\lambda) \in A)|$$

is small.

Proposition 2.10 (Characterizing operator of Poisson)**.** *For $\lambda > 0$, define the functional operator $\mathcal{A}$ by*

$$\mathcal{A}f(k) = \lambda f(k+1) - kf(k).$$

1. *If the random variable Z has the Poisson distribution with mean λ, then $\mathbb{E}\mathcal{A}f(Z) = 0$ for all bounded f.*
2. *If for some non-negative integer-valued random variable W, $\mathbb{E}\mathcal{A}f(W) = 0$ for all bounded functions f, then W has the Poisson distribution with mean λ.*

Proof of Proposition 2.10. We only prove the first part. Note that

$$\lambda \mathbb{E}[f(Z+1)] = e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^{k+1}}{k!} f(k+1) = e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^{k+1}}{(k+1)!}(k+1)f(k+1) = \mathbb{E}[Zf(Z)].$$

□

Having Proposition 2.10, the following two results are very intuitive.

[1]Let $(A_i)_{i=1}^{\infty}$ be a sequence of events. Then, we have $\mathbb{P}(\cup_{i=1}^{\infty} A_i \leq \sum_{i=1}^{\infty} \mathbb{P}(A_i)$.

Proposition 2.11. *Let $\mathcal{P}_\lambda(A) := \mathbb{P}(\mathrm{Poi}(\lambda) \in A)$, $A \subseteq \mathbb{Z}_{\geq 0}$. The unique solution f_A of*

$$\lambda f_A(k+1) - k f_A(k) = \mathbf{1}[k \in A] - \mathcal{P}_\lambda(A) \tag{14}$$

with $f_A(0) = 0$ is given by

$$f_A(k) = \lambda^{-k} e^{\lambda} (k-1)! \big[\mathcal{P}_\lambda(A \cap U_k) - \mathcal{P}_\lambda(A)\mathcal{P}_\lambda(U_k)\big],$$

where $U_k = \{0, 1, \ldots, k-1\}$.

Exercise 2.12. Prove Proposition 2.11.

Thanks to Proposition 2.11, we have the following immediately.

Proposition 2.13. *If $W \geq 0$ is an integer-valued random variable with mean λ, then*

$$\big|\mathbb{P}(W \in A) - \mathcal{P}_\lambda(A)\big| = \big|\mathbb{E}[\lambda f_A(W+1) - W f_A(W)]\big|.$$

One last result is needed to present the main theorem.

Proposition 2.14. *If f_A solves* (14)*, then*

$$\|f_A\| \leq \min\{1, \lambda^{-1/2}\} \qquad \text{and} \qquad \Delta f(k) := \|f(k+1) - f(k)\| \leq \frac{1 - e^{-\lambda}}{\lambda} \leq \min\{1, \lambda^{-1}\},$$

where $||f|| := \sup_{x \in D} |f(x)|$ and D is the domain of f.

Theorem 2.15 (Poisson Approximation Theorem)**.** *Let $\mathcal{F}$ be the set of functions satisfying the conditions in Proposition 2.14. If $W \geq 0$ is an integer-valued random variable with mean λ and $Z \sim \mathrm{Po}(\lambda)$, then*

$$d_{\mathrm{TV}}(W, Z) \leq \sup_{f \in \mathcal{F}} \big|\mathbb{E}[\lambda f(W+1) - W f(W)]\big|.$$

Let's now apply Theorem 2.15 to generalize Theorem 2.3: recall we have already shown that $X_n \sim \mathrm{Bin}(n, \lambda/n)$ and $Z \sim \mathrm{Poi}(\lambda)$ then $d_{\mathrm{TV}}(W_n, Z) \to 0$ as $n \to \infty$.

2.3 Law of small numbers

It is well known that if $X_n \sim \mathrm{Bin}(n, \lambda/n)$ and $Z \sim \mathrm{Poi}(\lambda)$ then $d_{\mathrm{TV}}(W_n, Z) \to 0$ as $n \to \infty$,

Theorem 2.16. *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables with $\mathbb{P}(X_i = 1) = p_i$, $W = \sum_{i=1}^n X_i$, and $\lambda = \mathbb{E}[W] = \sum_{i=1}^n p_i$. If $Z \sim \mathrm{Poi}(\lambda)$, then*

$$d_{\mathrm{TV}}(W, Z) \leq \min\{1, \lambda^{-1}\} \sum_{i=1}^n p_i^2.$$

Proof of Theorem 2.16. We apply Theorem 2.15. Let f satisfy the conditions of Proposition

2.14 and note that

$$\begin{aligned}
\mathbb{E}[Wf(W)] &= \sum_{i=1}^{n} \mathbb{E}[X_i f(W)] \\
&= \sum_{i=1}^{n} \mathbb{E}[f(W) \mid X_i = 1]\mathbb{P}(X_i = 1) \\
&= \sum_{i=1}^{n} p_i \mathbb{E}[f(W_i + 1)], \qquad (15)
\end{aligned}$$

where $W_i = W - X_i$ and (15) follows since X_i is independent of W_i. Since $\lambda f(W+1) = \sum_i p_i f(W+1)$, we obtain

$$\big|\mathbb{E}[\lambda f(W+1) - Wf(W)]\big| = \left|\sum_{i=1}^{n} p_i \, \mathbb{E}[f(W+1) - f(W_i+1)]\right| \leq \sum_{i=1}^{n} p_i \|\Delta f\| \, \mathbb{E}[|W - W_i|].$$

To see why the inequality holds, note that $f(W+1) - f(W_i+1) = \sum_{k=W_i+1}^{W} f(k+1) - f(k)$ so that by the triangle inequality $|f(W+1) - f(W_i+1)| \leq \sum_{k=W_i+1}^{W} ||\Delta f|| = ||\Delta f|| |W - W_i|$, and we just take the expectation.

Since $|W - W_i| = X_i$, we get

$$\big|\mathbb{E}[\lambda f(W+1) - Wf(W)]\big| \leq \|\Delta f\| \sum_{i=1}^{n} p_i \mathbb{E}[X_i] = \|\Delta f\| \sum_{i=1}^{n} p_i^2.$$

Using $\|\Delta f\| \leq \min\{1, \lambda^{-1}\}$ from Proposition 2.14, we conclude that

$$d_{\mathrm{TV}}(W, Z) \leq \min\{1, \lambda^{-1}\} \sum_{i=1}^{n} p_i^2.$$

By Theorem 2.15, we are done. □

3 Concentration Inequalities

We are interested in concentration inequalities, which help us understand how close random variables are to their expected values (or to other values). So far, we have studied Markov's and Chebyshev's inequalities, which apply to a single random variable, and the law of large numbers, which characterizes the behavior of sums of many random variables. We start with classical bounds.

3.1 Chernoff Bound

The generic Chernoff bound is inspired by applying Markov's inequality to the exponential of a random variable. Note that for any random variable $X \geq 0$ and $a, t > 0$, we have $\mathbb{P}(X \geq a) = \mathbb{P}(e^{tX} \geq e^{ta}) \leq \frac{\mathbb{E}[e^{tX}]}{e^{ta}}$, since e^{tX} is monotonically increasing. If $S_n = \sum_{k=1}^n X_k$, then we also have $\mathbb{P}(S_n \geq a) \leq e^{-ta}\mathbb{E}[\prod_{i=1}^n e^{tX_i}]$, which looks useful when the random variables are independent. One gets useful (sometimes tight) bounds when optimizing the right-hand side over t.

Example 3.1. Let $(X_i)_{i=1}^n$ be a sequence of i.i.d. Bernoulli random variables with $\mathbb{P}(X_k = 1) = p$, and let $X = \sum_{i=1}^n X_i$. First, note that

$$\begin{aligned} \mathbb{E}\left[e^{tX_i}\right] &= pe^t + (1-p)e^0 \\ &= 1 + p(e^t - 1) \\ &\leq e^{p(e^t-1)}, \end{aligned} \tag{16}$$

where we used $1 + x \leq e^x$ with $x = p(e^t - 1)$.

$$\begin{aligned} \mathbb{P}(X \geq a) &\leq \frac{\mathbb{E}\left[e^{tX}\right]}{e^{at}} \\ &\leq e^{-at}\,\mathbb{E}\left[e^{t\sum_i X_i}\right] \\ &\leq e^{-at}\,\mathbb{E}\left[e^{tX_1}\right]\mathbb{E}\left[e^{tX_2}\right]\cdots\mathbb{E}\left[e^{tX_n}\right] \\ &\leq e^{-at}\,e^{\sum_{i=1}^n p(e^t-1)}, \end{aligned} \tag{17}$$

where (17) follows from (16). Now taking $a = (1+\delta)\mathbb{E}[X] = (1+\delta)np$ and $t = \log(1+\delta)$, we get

$$\begin{aligned} \mathbb{P}(X \geq (1+\delta)np) &\leq \frac{e^{np(1+\delta-1)}}{(1+\delta)^{(1+\delta)np}} \\ &\leq \left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{np}. \end{aligned} \tag{18}$$

Using arguments analogous to the ones in Example 3.1, we can get the following two general results.

Theorem 3.2 (Chernoff bound for Bernoulli trials)**.** *Let $X = \sum_{i=1}^n X_i$, where $X_i = 1$ with probability p_i and $X_i = 0$ with probability $1 - p_i$, and suppose that $X_1, \ldots, X_n$ are* *independent. Let* $\mu = \mathbb{E}[X] = \sum_{i=1}^n p_i$. *Then we have*

(i) ***Upper Tail:***

$$\mathbb{P}(X \geq (1+\delta)\mu) \leq \exp\left(-\frac{\delta^2}{2+\delta}\,\mu\right), \qquad \textit{for all } \delta > 0.$$

(ii) ***Lower Tail:***

$$\mathbb{P}(X \leq (1-\delta)\mu) \leq \exp\left(-\frac{\mu\delta^2}{2}\right), \qquad \textit{for all } 0 < \delta < 1.$$

Combining both bounds for $\delta \in (0,1)$ *yields*

$$\mathbb{P}(|X-\mu| \geq \delta\mu) \leq 2\exp\left(-\frac{\mu\delta^2}{3}\right).$$

Theorem 3.3 (Chernoff bound for bounded random variables)**.** *Let* $X_1, X_2, \ldots, X_n$ *be random variables such that*

$$a \leq X_i \leq b \quad \textit{for all } i.$$

Let

$$X = \sum_{i=1}^n X_i \quad \textit{and} \quad \mu = \mathbb{E}[X].$$

Then, for all $\delta > 0$*:*

(i) ***Upper Tail:***

$$\mathbb{P}(X \geq (1+\delta)\mu) \leq \exp\left(-\frac{2\delta^2\mu^2}{n(b-a)^2}\right).$$

(ii) ***Lower Tail:***

$$\mathbb{P}(X \leq (1-\delta)\mu) \leq \exp\left(-\frac{2\delta^2\mu^2}{n(b-a)^2}\right).$$

Example 3.4. Suppose that we are repeatedly tossing a fair coin. Let S_n be the number of heads we observe from the first n tosses. Per Chebyshev's inequality, we get $\mathbb{P}(|S_n/n - 1/2| \geq \epsilon) \leq 1/4n\epsilon^2$. For example, when $\epsilon = 1/2$, we get $\mathbb{P}(|S_n/n - 1/2| \geq 1/2) \leq 4/n$. If we use Chernoff bound instead, per Theorem 3.2, we have

$$\mathbb{P}\left(\left|S_n - \frac{n}{2}\right| \geq \delta\frac{n}{2}\right) \leq 2\exp\left(-\frac{n\delta^2}{6}\right).$$

When $\delta = 1/2$, we obtain $\mathbb{P}(|S_n - 1/2| \geq 1/4) \leq 2\exp(-n/24)$, which is a much better bound.

3.2 Hoeffding's Inequality

Theorem 3.5 (Hoeffding's Inequality)**.** *Let* $(X_i)_{i=1}^n$ *be a sequence of independent bounded random variables with* $a_i \leq X_i \leq b_i$ *for all* $i = 1, ..., n$ *with probability* 1. *Let* $S_n = \sum_{i=1}^n X_i$.

Then for any $t > 0$, we have

$$\mathbb{P}(|S_n - \mathbb{E}[S_n]| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

Proof of Theorem 3.5. We start with a useful result (Hoeffding's Lemma), which you will prove in your homework: if Y is a random variable with $\mathbb{E}[Y] = 0$ and $a \leq Y \leq b$, then we have $\mathbb{E}[\exp(tY)] \leq \exp(\frac{t^2(b-a)^2}{8})$. The proof simply uses the convexity of the exponential function, i.e., $e^{ty} \leq \frac{y-a}{b-a}e^{tb} + \frac{b-y}{b-a}e^{ta}$ for all $a \leq y \leq b$.

Now, note that by Markov's inequality, for all $s > 0$ we have

$$\Pr[S_n - \mathbb{E}[S_n] \geq t] = \Pr\left[e^{s(S_n - \mathbb{E}[S_n])} \geq e^{st}\right] \leq \frac{\mathbb{E}\left[e^{s(S_n - \mathbb{E}[S_n])}\right]}{e^{st}} \tag{19}$$

Letting $Y_i = X_i - \mathbb{E}[X_i]$, we get

$$\mathbb{E}\left[e^{s(S_n - \mathbb{E}[S_n])}\right] = \mathbb{E}\left[e^{s\sum_{i=1}^n (X_i - \mathbb{E}[X_i])}\right] = \mathbb{E}\left[\prod_{i=1}^n e^{s(X_i - \mathbb{E}[X_i])}\right] = \mathbb{E}\left[\prod_{i=1}^n e^{sY_i}\right]$$

Per Hoeffding's Lemma, we get

$$\mathbb{E}\left[\prod_{i=1}^n e^{sY_i}\right] \leq \prod_{i=1}^n e^{\frac{s^2(b_i - a_i)^2}{8}} \tag{20}$$

Using the bound (41) in (40), we get

$$\Pr[S_n - \mathbb{E}[S_n] \geq t] \leq \frac{\exp(\frac{s^2}{8}\sum_{i=1}^n (b_i - a_i)^2)}{\exp(st)}. \tag{21}$$

Now it is time to optimize the right-hand side. It is easy to see that $s = \frac{4t}{\sum_{i=1}^n (b_i - a_i)^2}$ is the minimizer of the right-hand side. Hence, we have

$$\Pr[S_n - \mathbb{E}[S_n] \geq t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right). \tag{22}$$

In a similar fashion, one can prove that for all $t > 0$

$$\Pr[S_n - \mathbb{E}[S_n] \leq -t] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right). \tag{17}$$

Combining both inequalities, we get

$$\mathbb{P}(|S_n - \mathbb{E}[S_n]| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

□

3.3 Martingale Inequalities

The family of random variables $\{X(t) : t \in T\}$ is called a stochastic process, where the parameter t is interpreted as time, and $X(t)$ is interpreted as the state of the process at time t. Throughout the semester, we will study numerous stochastic processes to model queueing systems, dynamic matching markets, online resource allocation settings, etc. We finish this lecture with an important family of stochastic processes: martingales.

Definition 3.6. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Then a filtration on the probability space is an increasing family of sub-σ-fields $(\mathcal{F}_n)_{n\geq 0}$ of $\mathcal{F}$ such that $\mathcal{F}_n \subset \mathcal{F}_{n+1} \subset \mathcal{F}$ for all $n \geq 1$. We further say that a stochastic process $X = (X_n)_{n\geq 1}$ is adapted to the filtration $(\mathcal{F}_n)_{n\geq 1}$ if X_n is $\mathcal{F}_n-$measurable.

What are we trying to achieve via filtration here? If you consider the parameter n as time, then intuitively you can interpret $\mathcal{F}_n$ as all historical information that is available to us up to (and including) time n. In other words, the value of X_n only depends on what has already happened up to time n. The sigma fields are increasing over time, because we do not forget the history.

Definition 3.7. Let $(\mathcal{F}_n)_{n\geq 1}$ be a filtration, and X is adapted to the filtration. Assume that $\mathbb{E}[|X_n|] < \infty$ for all $n \geq 1$. Then,

1. X is called a martingale if $\mathbb{E}[X_n|\mathcal{F}_{n-1}] = X_{n-1}$ a.s. for all $n \geq 2$.
2. X is called a supermartingale if $\mathbb{E}[X_n|\mathcal{F}_{n-1}] \leq X_{n-1}$ a.s. for all $n \geq 2$.
3. X is called a submartingale if $\mathbb{E}[X_n|\mathcal{F}_{n-1}] \geq X_{n-1}$ a.s. for all $n \geq 2$.

If you are gambling in Las Vegas, then you are playing supermartingale games, while the casino is playing submartingale games. This course is a fair game so we have a martingale. Note that in Definition 3.7[1] by the tower property, $\mathbb{E}[\mathbb{E}[X_n|\mathcal{F}_{n-1}]] = \mathbb{E}[X_n] = \mathbb{E}[X_{n-1}]$, and by recursively we get $\mathbb{E}[X_n] = \mathbb{E}[X_1]$. Similarly, one can also show that $\mathbb{E}[X_{n+m}|\mathcal{F}_n] = X_n$ for any $m \geq 1$.

Example 3.8. Let $X_1, X_2, ...$ be a sequence of independent random variables with $\mathbb{E}[|X_n|] < \infty$ for all $n \geq 1$, and $\mathbb{E}[X_n] = 0$ for all $n \geq 1$. Let $S_n = \sum_{i=1}^n X_i$ and $\mathcal{F}_n = \sigma(X_1, X_2, ..., X_n)$. Then almost surely we have $\mathbb{E}[S_n|\mathcal{F}_{n-1}] = \mathbb{E}[X_n|\mathcal{F}_{n-1}] + \mathbb{E}[S_{n-1}|\mathcal{F}_{n-1}] = \mathbb{E}[X_n] + S_{n-1} = S_{n-1}$. Thus, S_n is a martingale.

Example 3.9. Let $X_1, X_2, ...$ be a sequence of independent random variables with $\mathbb{E}[X_n] = 1$ for all $n \geq 1$. Let $M_0 = 1$, $\mathcal{F}_0 = \{\emptyset, \Omega\}$, $M_n = \prod_{k=1}^n X_k$, and $\mathcal{F}_n = \sigma(X_1, X_2, ..., X_n)$. Then $M = (M_n)_{n\geq 0}$ is a martingale, since $\mathbb{E}[M_n|\mathcal{F}_{n-1}] = \mathbb{E}[M_{n-1}X_n|\mathcal{F}_{n-1}] = M_{n-1}\mathbb{E}[X_n|\mathcal{F}_{n-1}] = M_{n-1}\mathbb{E}[X_n] = M_{n-1}$.

Example 3.10. Let $X_1, X_2, ...$ be a sequence of independent and identically distributed random variables with $\mathbb{E}[X_1] = 0$ and $Var(X_1) = \sigma^2$. Let $S_n = \sum_{i=1}^n X_i$ and define

$Z_n := S_n^2 - n\sigma^2$ for $n \geq 1$. Let $\mathcal{F}_n = \sigma(X_1, X_2, ..., X_n)$. Then Z_n is a martingale, since

$$\begin{aligned}
\mathbb{E}\big[S_{n+1}^2 \mid \mathcal{F}_n\big] &= \mathbb{E}\big[(S_n + X_{n+1})^2 \mid \mathcal{F}_n\big] \\
&= \mathbb{E}\big[S_n^2 \mid \mathcal{F}_n\big] + \mathbb{E}[2S_n X_{n+1} \mid \mathcal{F}_n] + \mathbb{E}\big[X_{n+1}^2 \mid \mathcal{F}_n\big] \\
&= S_n^2 + 2S_n\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] + \mathbb{E}\big[X_{n+1}^2\big] \\
&= S_n^2 + 2S_n\mathbb{E}[X_{n+1}] + \sigma^2 \\
&= S_n^2 + \sigma^2.
\end{aligned}$$

Definition 3.11. A map $T : \Omega \to \mathbb{Z}_{\geq 0} \cup \{\infty\}$ is called a stopping time if $\{T \leq n\} = \{\omega : T(w) \leq n\} \in \mathcal{F}_n$ for all $n \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$.

It is a simple exercise to show that in Definition 3.11, $\{T \leq n\} \in \mathcal{F}_n$ is equivalent to $\{T = n\} \in \mathcal{F}_n$ or $\{T \geq n\} \in \mathcal{F}_{n-1}$. Intuitively, whether the process will stop at time n according to your stopping time T depends only on the history up to (and including) time n.

Example 3.12. Given a sequence of independent random variables $(X_n)_{n\geq 1}$, the first time X_n hits an arbitrary set A, i.e., $T_A = \min\{n : X_n \in A\}$, is clearly a stopping time. But the last time that X_n visits an arbitrary set A, i.e., $T_A = \max\{n : X_n \in A\}$, or, the time when X_n reaches its maximum, i.e., $T = \min\{n : X_n = \max_{k\geq 1} X_k\}$, are not stopping times.

Theorem 3.13. *Let X be a martingale, and T be a stopping time. Then $X_{T\wedge n}$, where $T \wedge n = \min\{T, n\}$, is a martingale.*

Proof of Theorem 3.13. It is easy to verify that $X_{T\wedge n} = X_{T\wedge(n-1)} + \mathbf{1}_{\{T\geq n\}}(X_n - X_{n-1})$. Thus,

$$\begin{aligned}
\mathbb{E}[X_{T\wedge n} \mid \mathcal{F}_{n-1}] &= \mathbb{E}\big[X_{T\wedge(n-1)} \mid \mathcal{F}_{n-1}\big] + \mathbb{E}\big[\mathbf{1}_{\{T\geq n\}}(X_n - X_{n-1}) \mid \mathcal{F}_{n-1}\big] \\
&= X_{T\wedge(n-1)} + \mathbf{1}_{\{T\geq n\}}\,\mathbb{E}[X_n - X_{n-1} \mid \mathcal{F}_{n-1}] \\
&= X_{T\wedge(n-1)}.
\end{aligned}$$

□

Theorem 3.14 (Martingale Stopping Theorem)**.** *Let $(X_n)_{n\geq 1}$ be a martingale adapted to the filtration $(\mathcal{F}_n)_{n\geq 1}$, and suppose that T is a stopping time for this filtration. Then*

$$\mathbb{E}[X_T] = \mathbb{E}[X_1],$$

if any of the following three sufficient conditions hold:

1. *T is bounded, i.e., there exists a constant C such that $T(w) \leq C$ for all $w \in \Omega$;*
2. *X_n is bounded for all n and $\mathbb{P}(T < \infty) = 1$;*
3. *$\mathbb{E}[T] < \infty$ and X has bounded increments, i.e., there exists $M < \infty$ such that for all $n \geq 1$,*
$$\mathbb{E}\big[\,|X_{n+1} - X_n| \mid \mathcal{F}_n\big] < M.$$

In the homework, you will prove the following result, which is a corollary of Theorem 3.14.

Theorem 3.15 (Wald's Identity)**.** *Let $X_1, X_2, ...$ be a sequence of independent and identically distributed random variables, and let T be a stopping time for the filtration $\mathcal{F}_n = \sigma(X_1, X_2, ..., X_n)$ for all $n \geq 1$. If $\mathbb{E}[T] < \infty$, then*

$$\mathbb{E}[\sum_{i=1}^{T} X_i] = \mathbb{E}[X_1] \cdot \mathbb{E}[T].$$

Example 3.16 (Ballot counting problem)**.** Assume that we have two candidates A and B, and let N_A and N_B be the number of votes for each of the candidates A and B, respectively, with $N_A + N_B = N$. Assume that $N_A > N_B$. We start counting votes one by one in a uniformly random ordering. We want to find the probability that candidate A is always ahead when counting votes. Let Y_k be the difference between the number of votes for candidates A and B after counting k votes. Let $X_k = \frac{Y_{N-k}}{N-k}$. First show that $X = (X_k)_{k\geq 0}$ is a martingale. Then define a stopping time $T = \min\{k \in [0, N] : X_k = 0\}$, or $T = N - 1$ if there is no such k and apply Theorem 3.14.

The following result provides a concentration bound for martingales, where note that we are not imposing independence (martingales can have dependent increments!).

Theorem 3.17 (Azuma–Hoeffding Inequality)**.** *Let $(X_n)_{n\geq 1}$ be a martingale adapted to the filtration $(\mathcal{F}_n)_{n\geq 1}$ and assume that almost surely $|X_n - X_{n-1}| \leq c_i$ for all n, where define $X_0 = \mathbb{E}[X_1]$. Then, for all $t > 0$,*

$$\Pr(X_n - X_0 \geq t) \leq \exp\left(-\frac{t^2}{2\sum_{i=1}^{n} c_i^2}\right),$$

and

$$\Pr(X_n - X_0 \leq -t) \leq \exp\left(-\frac{t^2}{2\sum_{i=1}^{n} c_i^2}\right).$$

Combining both bounds yields

$$\Pr(|X_n - X_0| \leq t) \leq 2\exp\left(-\frac{t^2}{2\sum_{i=1}^{n} c_i^2}\right).$$

Note that one can view Theorem 3.17 as providing a concentration bound on the sum function of random variables: $X_n - X_0 = \sum_{i=1}^{n} Y_i$, where $Y_i = X_1 - X_0$. The following result lets us getting other bounds when the functions are relatively more general.

Theorem 3.18 (McDiarmid's Inequality)**.** *Consider n independent random variables $X_1, \ldots, X_n$ taking values in A_i for each i, and a function $f : \prod_{i=1}^{n} A_i \to \mathbb{R}$ satisfying*

$$|f(x) - f(x')| \leq c_i \quad \textit{whenever } x \textit{ and } x' \textit{ differ only in the } i\textit{th coordinate.}$$

Let

$$\mu = \mathbb{E}[f(X_1, \ldots, X_n)]$$

be the expected value of the random variable $f(X)$. Then for any $\beta > 0$,

$$\Pr(f(X) \geq \mu + \beta) \leq \exp\left(-\frac{2\beta^2}{\sum_{i=1}^{n} c_i^2}\right),$$

and

$$\Pr(f(X) \le \mu - \beta) \le \exp\left(-\frac{2\beta^2}{\sum_{i=1}^n c_i^2}\right).$$

Example 3.19. Consider the following balls and bins problem, where we throw n balls uniformly at random and independently into n bins. Let B_i denote the number of balls in bin i. Note that B_i is a $\text{Bin}(n, 1/n)$ random variable. Per Markov's inequality, we have

$$\Pr[B_i \ge 1 + \lambda] \le \frac{n \cdot 1/n}{1+\lambda} = \frac{1}{1+\lambda} \approx \frac{1}{\lambda}.$$

However, Chebyshev's inequality gives a much better bound, since

$$\Pr[|B_i - 1| \ge \lambda] \le \frac{n \cdot 1/n \cdot (1 - 1/n)}{\lambda^2} = \frac{(1-1/n)}{\lambda^2} \approx \frac{1}{\lambda^2}.$$

If we let $\lambda = 2\sqrt{n}$, the probability of any fixed bin having more than $2\sqrt{n}+1$ balls is at most $\frac{1}{4n}$. Taking a union bound over all bins, we have that with probability at least $1 - n \cdot \frac{1}{4n} \le \frac{1}{4}$, the load on every bin is at most $1 + 2\sqrt{n}$.

Now, if we apply Chernoff bound for Bernoulli trials, we get

$$\Pr[B_i \ge 1 + \lambda] \le \exp\left(-\frac{\lambda^2}{2+\lambda}\right).$$

If we set $\lambda = \Theta(\log n)$, the probability that bin i has more than $1 + \lambda$ balls is at most $1/n^2$. Taking a union bound over all bins, the probability that any bin has at least $1 + \lambda$ balls is at most $1/n$, i.e., the maximum load is $O(\log n)$ balls with high probability.

Now define $f(X) = f(X_1, X_2, ..., X_n)$ to be the number of empty bins, where X_i denotes the location of ball i, i.e., ball i is placed in bin number X_i. Note that $\mathbb{E}[f(X)] = n(1-1/n)^n$ by linearity of expectation. Noting that changing the assignment of one ball can only change the number of empty bins by at most 1, we get the following bound by McDiarmid's inequality:

$$\mathbb{P}(|f(X) - n(1-1/n)^n| \ge \beta) \le 2\exp\left(\frac{-2\beta^2}{n}\right).$$

Discussion 3.20. Suppose that you are drawing samples from a distribution X with $\mathbb{E}[X] = \mu$ and $Var(X) = \sigma^2$. Let's say you have a bound on the variance, i.e., $\sigma \le C$. How large should the sample size be to ensure that with probability $1-p$, the sample average is 2 away from the mean μ? What if you have a bound on the distribution instead, i.e., $|X| \le C$?

Discussion 3.21. Consider the symmetric random walk. Denote the position of the random walk after n steps by $S_n = \sum_{i=1}^n X_i$, where $\mathbb{P}(X_i = 1) = \mathbb{P}(X_i = -1) = 1/2$. What are the tail bounds under Chebyshev's and Hoeffding's inequalities?

Discussion 3.22. Recall the coupon collector problem: given n different types of coupons, how many coupons in expectation do we need to draw with replacement before having drawn each coupon at least once? Using linearity of expectation, one can easily show that $\mathbb{E}[X] = \sum_{i=1}^n \frac{n}{n-i+1} = nH(n) \approx n\log(n)$, where $H(n)$ is the harmonic number. What does Markov's inequality, Chebyshev's inequality imply on tail bounds?

4 Queueing Theory

We start with some preliminaries.

Definition 4.1 (Markov chain)**.** A stochastic process $\{X_n, n \geq 0\}$ is called a *Markov chain* if, for every $x_i \in \mathcal{S}$,

$$\Pr\{X_n = x_n \mid X_{n-1} = x_{n-1}, \ldots, X_0 = x_0\} = \Pr\{X_n = x_n \mid X_{n-1} = x_{n-1}\}, \tag{23}$$

The definition implies that given the present state of the system, the future is independent of the past. The conditional probability

$$p_{jk}(n) := \Pr\{X_n = k \mid X_{n-1} = j\}, \qquad j, k \in S,$$

is called the *transition probability* from state j to state k. We say that the chain is homogeneous if $p_{jk}(n)$ does not depend on n, i.e.,

$$p_{jk} := \Pr\{X_n = k \mid X_{n-1} = j\} = \Pr\{X_{n+m} = k \mid X_{n+m-1} = j\}$$

for all $m \in \mathbb{Z}$. Let $P = (p_{ij})_{i,j \in \mathcal{S}}$ be the transition matrix.

Given an irreducible Markov chain (there is a single communicating class), then we have seen (sometime in the past :)) that there is a unique probability distribution π on $\mathcal{S}$ such that $\pi P = \pi$.

Theorem 4.2 (Ergodic theorem for Markov chains)**.** *If $\{X_t, t \geq 0\}$ is a Markov chain on the state space S with unique invariant distribution π, then for any initial condition, we have*

$$\lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} \mathbf{1}\{X_t = x\} = \pi(x) \qquad \forall\, x \in S, \text{a.s.}$$

In order to calculate π, we use the global balance equations for the Markov chain, which states that $\pi_j = \sum_{i \in \mathcal{S}} \pi_i p_{ij}$, or equivalently $\pi_i \sum_{j \in \mathcal{S} \setminus \{i\}} p_{ij} = \sum_{j \in \mathcal{S} \setminus \{i\}} \pi_j p_{ji}$.

Definition 4.3 (Reversibility)**.** We say that a Markov chain is reversible if

$$P\big(X_{t_1} = x_1, X_{t_2} = x_2, \ldots, X_{t_k} = x_k\big) = P\big(X_{s-t_1} = x_1, X_{s-t_2} = x_2, \ldots, X_{s-t_k} = x_k\big),$$

for all $k \in \mathbb{N}$, $s, t_1, \ldots, t_k \in \mathbb{Z}$, and $x_1, \ldots, x_k \in S$.

Discussion 4.4. If we have reversibility, then we can calculate π via detailed balance equations, which states that $\pi_j p_{ji} = \pi_i p_{ij}$ for all $i, j \in \mathcal{S}$. We can check whether a Markov chain is reversible via Kolmogorov's closed loop criterion: an ergodic Markov chain is reversible if and only if

$$p_{j_0 j_1}\, p_{j_1 j_2} \cdots p_{j_{k-1} j_k} = p_{j_0 j_k}\, p_{j_k j_{k-1}} \cdots p_{j_2 j_1}\, p_{j_1 j_0}, \tag{24}$$

for every finite sequence of distinct states $j_0, j_1, j_2, \ldots, j_k$.

Now we turn our focus on queueing theory. A queueing system is characterized by

1. Arrival pattern of customers: whether arrivals occur singly or in batches, what distribution governs the interrarrival times?

2. Service pattern of customers: what is the average time required to serve a customer?
3. The number of servers.
4. The capacity of system: infinite or finite capacity?
5. The queue discipline: first-in-first-out (FIFO), last-in-first-out (LIFO), priority queues, ...

Theorem 4.5 (Little's Law)**.** *Given a queueing system, in the steady-state, let L be the average number of customers in the system, let λ be the average arrival rate, and let W be the average waiting time (waiting in the queue plus waiting while getting service). Then $L = \lambda W$.*

Discussion 4.6. PASTA property.

4.1 Elementary queueing systems: exponential models

4.1.1 The $M/M/1$ model

We start with the simplest queueing system, the M/M/1 queue. Here, arrivals follow a Poisson process with parameter λ, i.e., the inter-arrival times are independent and exponential with mean $\frac{1}{\lambda}$, and the service times are independent and exponential with mean $\frac{1}{\mu}$. The utilization is defined as $\rho = \frac{\lambda}{N\mu}$, where $N = 1$.

Let $L(t)$ be the number of customers (both waiting in the queue and receiving service) at time t and let $p_n = \lim_{t\to\infty} \mathbb{P}(L(t) = n)$ for all $n \geq 0$. Then, we have

$$\begin{aligned} \lambda p_n &= \mu p_{n+1}, \qquad (n \geq 0) \\ \text{or} \quad p_{n+1} &= \frac{\lambda}{\mu} p_n = a p_n = a^2 p_{n-1} \\ &\vdots \\ &= a^{n+1} p_0 \end{aligned}$$

or

$$p_n = a^n p_0, \qquad n \geq 0.$$

Using the fact that $\sum_{n=0}^{\infty} p_n = 1$, for $a < 1$, we have

$$p_n = (1-a)a^n, \qquad n = 0, 1, 2, \ldots$$

Since $a = \rho$, we get

$$p_0 = (1-a) = 1-\rho$$

and

$$p_n = (1-\rho)\rho^n, \qquad n = 1, 2, \ldots$$

Note that the distribution is geometric and is memoryless. Let N be the number of customers in the system and W be the waiting time in the system in steady-state. Thus, we have

$$\begin{aligned}\mathbb{E}[N] &= \sum_{n=0}^{\infty} np_n = \sum_{n=1}^{\infty} n(1-\rho)\rho^n \\ &= \rho(1-\rho)\sum_{n=1}^{\infty} n\rho^{n-1} = \frac{\rho(1-\rho)}{(1-\rho)^2} = \frac{\rho}{1-\rho},\end{aligned} \tag{25}$$

and

$$\begin{aligned}\mathbb{E}[N^2] &= \sum_{n=0}^{\infty} n^2 p_n = \sum_{n=1}^{\infty} n^2(1-\rho)\rho^n \\ &= (1-\rho)\sum_{n=1}^{\infty}\big[(n^2-n)+n\big]\rho^n \\ &= (1-\rho)\left(\frac{2\rho^2}{(1-\rho)^3} + \frac{(1-\rho)\rho}{(1-\rho)^2}\right) = \frac{2\rho^2}{(1-\rho)^2} + \frac{\rho}{1-\rho} \\ &= \frac{\rho+\rho^2}{(1-\rho)^2}.\end{aligned} \tag{26}$$

Therefore, we have

$$Var(N) = \mathbb{E}[N^2] - \big(\mathbb{E}[N]\big)^2 = \frac{\rho}{(1-\rho)^2}. \tag{27}$$

Using Little's formula $L = \lambda W$, we get

$$\mathbb{E}[W] = \frac{\mathbb{E}[N]}{\lambda} = \frac{1}{\lambda}\frac{\rho}{1-\rho} = \frac{1}{\mu(1-\rho)}. \tag{28}$$

4.1.2 $M/M/1/K$ model

Now we assume that there is a bound on the maximum queue-length, i.e., when there are K customers waiting in the queue, any arrival leaves the system without getting a service. Analagous calculations yield

$$\lambda p_n = \mu p_{n+1}, \qquad n = 0, 1, 2, \ldots, K-1. \tag{29}$$

$$p_n = p_0 a^n, \qquad a = \frac{\lambda}{\mu}, \quad n = 0, 1, 2, \ldots, K. \tag{30}$$

Using the fact that

$$\sum_{n=0}^{K} p_n = 1,$$

we have

$$p_0 \sum_{n=0}^{K} a^n = 1.$$

Therefore,

$$p_0 = \begin{cases} \left[\sum_{n=0}^{K} a^n\right]^{-1} = \dfrac{1-a}{1-a^{K+1}}, & \lambda \neq \mu, \\ \dfrac{1}{K+1}, & \lambda = \mu. \end{cases}$$

We get for any $n = 0, 1, \ldots, K$, that

$$p_n = p_0 a^n = \begin{cases} \dfrac{(1-a)a^n}{1-a^{K+1}}, & \lambda \neq \mu, \\ \dfrac{1}{K+1}, & \lambda = \mu. \end{cases} \tag{31}$$

We can find the expected number of customers in the system as follows. If $\lambda = \mu$, then

$$L_K = \sum_{n=0}^{K} n p_n = \sum_{n=0}^{K} \frac{n}{K+1} = \frac{K}{2},$$

and if $\lambda \neq \mu$,

$$\begin{aligned} L_K &= \frac{(1-a)a}{1-a^{K+1}} \sum_{n=0}^{K} n a^{n-1} \\ &= \frac{(1-a)a}{1-a^{K+1}} \frac{1-(K+1)a^K + K a^{K+1}}{(1-a)^2} \\ &= \frac{a}{1-a} - \frac{(K+1)a^{K+1}}{1-a^{K+1}}. \end{aligned}$$

where we used the geometric stair sum formula.

4.2 Birth and death process

Now consider the following generalization, where arrival and service rates are state-dependent. That is, when there are n customers in the system, the arrival rate is λ_n and the service rate is μ_n. Not much will change as now we have

$$\lambda_n p_n = \mu_{n+1} p_{n+1}, \qquad n = 0, 1, 2, \ldots \tag{32}$$

Thus,

$$p_{n+1} = \frac{\lambda_n}{\mu_{n+1}} p_n = \frac{\lambda_n}{\mu_{n+1}} \frac{\lambda_{n-1}}{\mu_n} p_{n-1} = \cdots = \prod_{k=0}^{n} \frac{\lambda_k}{\mu_{k+1}} p_0, \qquad n = 0, 1, 2, \ldots,$$

or

$$p_n = \prod_{k=0}^{n-1} \frac{\lambda_k}{\mu_{k+1}} p_0, \qquad n = 1, 2, \ldots \tag{33}$$

Using $\sum_{n=0}^{\infty} p_n = 1$, we get

$$p_0 = \frac{1}{1 + \sum_{n=1}^{\infty} \prod_{k=0}^{n-1} \frac{\lambda_k}{\mu_{k+1}}}. \tag{34}$$

The necessary and sufficient condition for the existence of a steady state is the convergence of

$$\sum_{n=1}^{\infty} \prod_{k=0}^{n-1} \frac{\lambda_k}{\mu_{k+1}},$$

Note that when $\lambda_n = \lambda$ and $\mu_n = \mu$ for all $n = 0, 1, 2, \ldots$, we recover the M/M/1 system.

4.3 Continuous-Time Markov Chains

Following our discussions so far, we start with another set of preliminaries. Let $\{X(t), 0 \leq t < \infty\}$ be a Markov process with countable state space $S = \{0, 1, 2, \ldots\}$. Assume that the process is time homogeneous. Then the transition probability function given by

$$p_{ij}(t) = \Pr\{X(t+u) = j \mid X(u) = i\}, \qquad t > 0, \; i, j \in S, \tag{35}$$

is then independent of $u \geq 0$. Then for all $t > 0$, we have

$$0 \leq p_{ij}(t) \leq 1, \qquad \sum_j p_{ij}(t) = 1, \quad \text{for all } i \in S.$$

Denote the matrix of transition probabilities by

$$P(t) = \big(p_{ij}(t)\big), \qquad i, j \in S.$$

Set $p_{ij}(0) = \delta_{ij}$ (Kronecker delta function). Then the initial condition can be written as

$$P(0) = I.$$

Denote the probability that the system is at state j at time t by

$$\pi_j(t) = \Pr\{X(t) = j\};$$

the vector $\pi(t) = \{\pi_1(t), \pi_2(t), \ldots\}$ is the probability vector of the state of the system at time t, and $\pi(0)$ is the initial probability vector. We get

$$\begin{aligned} \pi_j(t) &= \sum_i \Pr\{X(t+u) = j \mid X(u) = i\} \Pr\{X(u) = i\} \\ &= \sum_i p_{ij}(t) \Pr\{X(0) = i\} \\ &= \sum_i p_{ij}(t) \pi_i(0). \end{aligned}$$

Thus, once we are given an initial probability vector $\pi(0)$ and the transition functions $p_{ij}(t)$, the state probabilities can be calculated as follows:

$$\pi(t) = \pi(0) P(t).$$

Definition 4.7 (Sojourn time)**.** The waiting time for change of state from state i is a random variable denoted by τ_i, and it is called the sojourn time at state i.

Note that

$$\Pr\{\tau_i > s{+}t \mid X(0) = i\} = \Pr\{\tau_i > s{+}t \mid X(0) = i,\, \tau_i > s\} \Pr\{\tau_i > s \mid X(0) = i\}, \qquad t \geq 0. \tag{36}$$

Denote

$$\bar{F}_i(u) := \Pr\{\tau_i > u \mid X(0) = i\}, \qquad u \geq 0.$$

Then (70) can be written as follows:

$$\bar{F}_i(t+s) = \bar{F}_i(t)\, \bar{F}_i(s), \qquad s, t \geq 0.$$

The only right continuous solution of this functional equation is (why?)

$$\bar{F}_i(u) = e^{-a_i u}, \qquad u \geq 0, \quad a_i > 0 \text{ is a constant.} \tag{37}$$

This implies that the sojourn time τ_i at state i is distributed exponentially with parameter a_i. Moreover, the sojourn times τ_i and τ_j are independent. Finally, we have $t \geq 0$, $T \geq 0$,

$$p_{ij}(T+t) = \sum_k p_{ik}(T)\, p_{kj}(t), \qquad i, j, k \in S. \tag{38}$$

or, in matrix form,

$$P(T+t) = P(T)P(t). \tag{39}$$

(39) is called the Chapman-Kolmogorov equation.

4.4 Transition density matrix

Now we discuss the transition density matrix, which is also known as infinitesimal generator or rate matrix. Consider

$$q_{ij} = \lim_{h \to 0} \frac{p_{ij}(h) - p_{ij}(0)}{h} = \lim_{h \to 0} \frac{p_{ij}(h)}{h}, \qquad i \neq j, \tag{40}$$

and

$$q_{ii} = \lim_{h \to 0} \frac{p_{ii}(h) - p_{ii}(0)}{h} = \lim_{h \to 0} \frac{p_{ii}(h) - 1}{h}. \tag{41}$$

Let $-q_i := q_{ii}$. We only bother with the cases when q_i and $q_i j$ are finite. Letting $Q = (q_i j)_{i,j \in \mathcal{S}}$, we have the following matrix notation

$$Q = \lim_{h \to 0} \frac{P(h) - I}{h}.$$

From (40) and (41), it follows that, when h is small,

$$p_{ij}(h) = h q_{ij} + o(h), \qquad i \neq j, \tag{42}$$

and

$$p_{ii}(h) = 1 - h q_i + o(h), \tag{43}$$

where $o(h)$ is a function of h that tends to zero more rapidly than h, i.e., $\frac{o(h)}{h} \to 0$ as $h \to 0$.

Now note that $\sum_j p_{ij}(h) = 1$, which implies $\sum_{j\neq i} p_{ij}(h) + p_{ii}(h) - 1 = 0$. Thus, we get $\sum_{j\neq i} q_{ij} + q_{ii} = 0$, or, $\sum_{j\neq i} q_{ij} = q_i$.

The Q-matrix $Q = (q_{ij})$ satisfies: (i) its diagonal elements are negative and off-diagonal elements are positive, and (ii) the sum of each row is 0. If we have a finite set $S = \{0, 1, 2, \ldots, m\}$ then the matrix looks like

$$Q = \begin{pmatrix} -q_0 & q_{01} & \cdots & q_{0m} \\ q_{10} & -q_1 & \cdots & q_{1m} \\ \vdots & \vdots & \ddots & \vdots \\ q_{m0} & q_{m1} & \cdots & -q_m \end{pmatrix}.$$

4.5 Chapman-Kolmogorov Backward and Forward Equations

From (39), we have

$$p_{ij}(h+t) = \sum_k p_{ik}(h)p_{kj}(t) = \sum_{k\neq i} p_{ik}(h)p_{kj}(t) + p_{ii}(h)p_{ij}(t),$$

so that

$$\frac{p_{ij}(h+t) - p_{ij}(t)}{h} = \sum_{k\neq i} \frac{p_{ik}(h)}{h} p_{kj}(t) + \left(\frac{p_{ii}(h) - 1}{h}\right) p_{ij}(t).$$

Now if we take the limit $h \to 0$ and interchange the limit and summation operations (if the state space is finite, this interchange is justified clearly, if the state space is countable, this interchange is again justified if we assume that $\sup_i q_i < \infty$, that is if we have uniformly bounded jump rates), we have

$$\lim_{h\to 0} \frac{p_{ij}(h+t) - p_{ij}(t)}{h} = \sum_{k\neq i} \left[\lim_{h\to 0} \frac{p_{ik}(h)}{h}\right] p_{kj}(t) + \left[\lim_{h\to 0} \frac{p_{ii}(h) - 1}{h}\right] p_{ij}(t),$$

or

$$p'_{ij}(t) = \sum_{k\neq i} q_{ik}p_{kj}(t) + q_i p_{ij}(t), \tag{44}$$

which is another form of the Chapman–Kolmogorov (backward) equation: in matrix notation, we have $P'(t) = QP(t)$.

We can also use (39) as

$$p_{ij}(t+h) = \sum_k p_{ik}(t)p_{kj}(h) = \sum_{k\neq i} p_{ik}(t)p_{kj}(h) + p_{ij}(t)p_{jj}(h).$$

Again, if we take the limit and interchange it with the summation operation, we get

$$p'_{ij}(t) = \sum_{k\neq j} p_{ik}(t)q_{kj} + q_j p_{ij}(t), \tag{45}$$

which is the Chapman–Kolmogorov *forward* equation: in matrix notation, we have $P'(t) = P(t)Q$.

Recall that $\pi(t) = \pi(0)P(t)$ so that both equations yield

$$\frac{d}{dt}\pi(t) = Q\pi(t) = \pi(t)Q. \tag{46}$$

Discussion 4.8. Here is an alternative definition of continuous-time Markov chains. Consider a stochastic process such that when you enter to state i, the time you spend at state i before you transition to another state $j \neq i$ is an exponential random variable with parameter a_i (with mean $\frac{1}{a_i}$). The parameter a_i only depends on state i and it is independent of other states j's. When you leave state i, you immediately transition to another state (to state j with probability p_{ij}). Thus, we have

$$p_{ii} = 0, \qquad 0 \leq p_{ij} \leq 1,$$

$$\sum_j p_{ij} = 1, \qquad j \in S.$$

Therefore, a continuous-time Markov chain is a stochastic process such that (i) its transition from one state to another state of the state space S is as in a discrete-time Markov chain and (ii) the sojourn time τ_i is an exponential random variable with some parameter a_i. The sojourn times in different states must be independent exponential random variables.

To see the relationship between p_{ij} and $p_{ij}(t)$, note that

$$p_{ij}(h) = h a_i p_{ij} + o(h),$$

since $p_{ij}(h)$ is the probability that the state of the process changes from i to j in an infinitesimal interval h. Thus,

$$q_{ij} = \lim_{h \to 0} \frac{p_{ij}(h)}{h} = a_i p_{ij},$$

Similarly, $1 - p_{ii}(h)$ is the probability that the state of the system changes from state i to some other state in the interval h, so that

$$1 - p_{ii}(h) = a_i h \sum_j p_{ij} + o(h) = a_i h + o(h).$$

Thus,

$$q_i = \lim_{h \to 0} \frac{1 - p_{ii}(h)}{h} = a_i;$$

Thus, the Q-matrix can also be written as

$$Q = \begin{pmatrix} -a_0 & a_0 p_{01} & \cdots & a_0 p_{0m} \\ a_1 p_{10} & -a_1 & \cdots & a_1 p_{1m} \\ \vdots & \vdots & \ddots & \vdots \\ a_m p_{m0} & a_m p_{m1} & \cdots & -a_m \end{pmatrix}. \qquad (47)$$

4.6 Birth and death processes revisited

A birth-and-death process (which we already discussed) is a continuous-time Markov chain $\{X(t), t \in T\}$ with state space $S = \{0, 1, 2, \ldots\}$ and with rates

$$\begin{aligned} q_{i,i+1} &= \lambda_i \quad \text{(say)}, \qquad i = 0, 1, \ldots, \\ q_{i,i-1} &= \mu_i \quad \text{(say)}, \qquad i = 1, 2, \ldots, \\ q_{i,j} &= 0, \quad j \neq i \pm 1,\ j \neq i, \quad i = 0, 1, \ldots, \\ q_i &= (\lambda_i + \mu_i), \qquad i = 0, 1, \ldots, \quad \mu_0 = 0, \end{aligned}$$

From the Chapman-Kolmogorov forward equations, we get
For $i, j = 1, 2, \ldots,$

$$p'_{i,j}(t) = -(\lambda_j + \mu_j)p_{i,j}(t) + \lambda_{j-1}p_{i,j-1}(t) + \mu_{j+1}p_{i,j+1}(t). \tag{48}$$

and

$$p'_{i,0}(t) = -\lambda_0 p_{i,0}(t) + \mu_1 p_{i,1}(t). \tag{49}$$

Set the boundary conditions as

$$p_{i,j}(0+) = \delta_{ij}, \qquad i, j = 0, 1, \ldots \tag{50}$$

Let

$$p_j(t) = \Pr\{X(t) = j\}, \qquad j = 0, 1, \ldots,\ t > 0.$$

and assume that at time $t = 0$, our initial condition starts at state i. Therefore,

$$P_j(0) = \Pr\{X(0) = j\} = \delta_{ij}, \tag{51}$$

and

$$P_j(t) = p_{ij}(t),$$

The forward equations become

$$P'_j(t) = -(\lambda_j + \mu_j)P_j(t) + \lambda_{j-1}P_{j-1}(t) + \mu_{j+1}P_{j+1}(t), \qquad j = 1, 2, \ldots, \tag{52}$$

$$P'_0(t) = -\lambda_0 P_0(t) + \mu_1 P_1(t). \tag{53}$$

Suppose that all the λ_i's and μ_i's are nonzero to ensure that the Markov chain is irreducible (single communicating class). Since we have ergodicity, the following limit

$$\lim_{t \to \infty} p_{ij}(t) = p_j$$

exist, and they are independent of the initial state i. Then per (52) and (53), we obtain

$$0 = -(\lambda_j + \mu_j)p_j + \lambda_{j-1}p_{j-1} + \mu_{j+1}p_{j+1}, \qquad j \geq 1, \tag{54}$$

$$0 = -\lambda_0 p_0 + \mu_1 p_1. \tag{55}$$

Finally, one can show by solving the above equations inductively that if $\sum_{k=0}^{\infty} \pi_k < \infty$, where

$$\pi_j = \frac{\lambda_0 \lambda_1 \cdots \lambda_{j-1}}{\mu_1 \mu_2 \cdots \mu_j}, \qquad j \geq 1, \qquad \pi_0 = 1, \tag{56}$$

then we have

$$p_j = \frac{\pi_j}{\sum_k \pi_k}, \qquad j \geq 0. \tag{57}$$

4.7 The $M/M/c$ model

Now let's revisit our discussion on the $M/M/c$ model. Assume that we have a single queue having Poisson arrivals with rate λ, and there are $1 < c < \infty$ parallel servers, each having an i.i.d. exponential service time with mean $\frac{1}{\mu}$. We can capture this model with a suitable birth-and-death process.

Note that if there are n customers in the system, where $n < c$, then first n servers are busy and the time between two consecutive service completions is the minimum of n i.i.d. exponential random variables with each parameter being μ, where the minimum is exponential with rate $n\mu$. If there are at least c many customers in the system (that is, $n \geq c$) then all c servers are busy and the time between two consecutive service completions is exponential with rate $c\mu$. Thus, we have a birth-and-death process with birth rate λ and death rates

$$\mu_n = n\mu, \qquad n = 0, 1, 2, \ldots, c,$$

$$\mu_n = c\mu, \qquad n = c+1, c+2, \ldots.$$

Denote the utilization $\rho = \lambda/(c\mu)$. Assume that steady state exists and that the system is in steady state. From the previous section, we get for $1 \leq n \leq c$,

$$p_n = \frac{\lambda\lambda \cdots \lambda}{(\mu)(2\mu)\cdots(n\mu)}\, p_0 = \frac{(\lambda/\mu)^n}{n!}\, p_0 = \frac{\lambda}{n\mu}\, p_{n-1}, \tag{58}$$

and for $n \geq c$, we get

$$\begin{aligned} p_n &= \frac{(\lambda)(\lambda)\cdots(\lambda)}{[(\mu)(2\mu)\cdots(c\mu)][(c\mu)(c\mu)\cdots(c\mu)]}\, p_0 \\ &= \frac{\lambda^n}{c!\, \mu^c\, c^{\,n-c} \mu^{\,n-c}}\, p_0 = \frac{(\lambda/\mu)^n}{c!\, c^{\,n-c}}\, p_0 \\ &= \frac{\lambda}{c\mu}\, p_{n-1} = \rho^{\,n-c} p_c. \end{aligned} \tag{59}$$

In a more compact form, for all $n \geq 1$, we can write

$$\big[\min(n,c)\big]\mu\, p_n = \lambda p_{n-1}.$$

Using the fact that $\sum_{n=0}^{\infty} p_n = 1$, we have

$$\begin{aligned} p_0^{-1} &= 1 + \sum_{n=1}^{c-1} \frac{(\lambda/\mu)^n}{n!} + \sum_{n=c}^{\infty} \frac{(\lambda/\mu)^n}{c!\, c^{\,n-c}} \\ &= \sum_{n=0}^{c-1} \frac{(\lambda/\mu)^n}{n!} + \frac{1}{c!\, c^{-c}} \sum_{n=c}^{\infty} \left(\frac{\lambda}{c\mu}\right)^n. \end{aligned} \tag{60}$$

To guarantee the existence of a steady-state, the series $\sum_{n=c}^{\infty}(\lambda/(c\mu))^n$ must be convergent, which implies that $\rho < 1$. Finally we get,

$$p_0 = \left[\sum_{n=0}^{c-1} \frac{(\lambda/\mu)^n}{n!} + \frac{(\lambda/\mu)^c}{c!\,(1-\lambda/(c\mu))}\right]^{-1}. \tag{61}$$

Note that we can analyze the $M/M/\infty$ model in a similar fashion.

4.8 The $M/m/c/c$ system: Erlang loss model

Now consider the $M/m/c$ model with an additional twist: if all the c servers are busy, any arrival leaves the system without getting a service. Such systems where arrivals are rejected are called a loss system. This is a birth-and-death process with

$$\lambda_n = \lambda, \quad \mu_n = n\mu, \quad n = 0, 1, 2, \ldots, c-1, \text{ and}$$

$$\lambda_n = 0, \quad \mu_n = c\mu, \quad n \geq c.$$

We just follow the same analysis as before:

$$p_n = \frac{\left(\frac{\lambda}{\mu}\right)^n}{n!} p_0, \quad n = 1, \ldots, c, \tag{62}$$

$$p_n = 0, \quad n > c, \tag{63}$$

and

$$p_0 = \left[\sum_{k=0}^{c} \frac{\left(\frac{\lambda}{\mu}\right)^k}{k!}\right]^{-1}. \tag{64}$$

Thus,

$$p_n = \frac{\left(\frac{\lambda}{\mu}\right)^n / n!}{\sum_{k=0}^{c}\left(\frac{\lambda}{\mu}\right)^k / k!}, \quad n = 0, 1, 2, \ldots, c. \tag{65}$$

The distribution of $\{p_n\}$ is also called truncated Poisson.

4.9 Priorities

Discussion 4.9. Which system is more efficient?: (i) n $M/M/1$ queues with arrival rates λ and service rates μ, or, (ii) a single $M/M/1$ queue with arrival rate λn and service rate μn. Pooling reduces congestion in general, but not always?

Discussion 4.10. Now let's consider an $M/M/1$ queue with two types of customers: type 1 and type 2. Type i customers arrive independently according to a Poisson process with rate λ_i, $i = 1, 2$. For simplicity, let's assume that the service times of all customers are exponentially distributed with parameter μ. For stability, we must assume $\rho_1 + \rho_2 < 1$, where $\rho_i = \frac{\lambda_i}{\mu}$. Assume that type 1 customers have a strict priority over type 2 customers. That is, if there is an arrival of type 1 customer to the system, and type 2 customer is receiving a service, then we interrupt this service and we start serving the arriving type

1 customer, where type 2 customer rejoins the queue (it actually does not matter where exactly this customers rejoins: top of the line, end of the line, etc.)

We first note that type 1 customers can completely neglect type 2 customers. Therefore, the average number of type 1 customers in the system $L_1 = \frac{\rho_1}{1-\rho_1}$, and the average waiting time of type 1 customers is $W_1 = \frac{L_1}{\lambda_1}$ by Little's Law.

Because of the memoryless property, and the fact that the service times are distributed with the same mean, the (average) total number of customers in the system $L = L_1 + L_2$ does not depend on the priority rule we impose. The crucial assumption we are making here is that the servers are not idling intentionally, i.e., the server only idles when the system is completely empty. Therefore, $L = \frac{\rho_1+\rho_2}{1-(\rho_1+\rho_2)}$. Therefore, we have $L_2 = L - L_1 = \frac{\rho_2}{(1-\rho_1)(1-\rho_1-\rho_2)}$. By Little's Law again, we can also find the average waiting time of type 2 customers: $W_2 = \frac{L_2}{\lambda_2}$.

Discussion 4.11. Continuing our discussion from Discussion 4.10, now let's consider a non-preemptive case, where we have an M/M/1 queue with two types of customers, but now the arrival of a type 1 customer does not disrupt the service of a type 2 customer. After the server is done serving a type 2 customer, we start serving type 1 customers if there are any. Find the average number of customers and average waiting times of both types.

4.10 Jackson Networks

Jackson's network model is defined as follows. Assume that customers from one node (queueing system) i proceed to an arbitrary node, and new customers may arrive to a node from outside (say customers arrive to node i according to a Poisson process with rate λ_i). Suppose that there are k nodes, where the ith node ($i = 1, \ldots, k$) consists of c_i exponential servers with parameter μ_i (that is, each node contains a $M/M/c$ queueing system). Customers after receiving service at the ith node proceed to the jth node with probability p_{ij}.

Customers at node i depart from the system with probability

$$q_i = 1 - \sum_{j=1}^{k} p_{ij}.$$

Consider Jackson's general network model with k nodes. The arrivals can be categorized into two groups: the external arrivals (with rate λ_i) and internal arrivals (with rate $\sum_{j=1}^{k} p_{ji}\lambda_j$). Therefore, the effective arrival rate to node i (or the effective rate of flow through node i) is

$$\alpha_i = \lambda_i + \sum_{j=1}^{k} p_{ji}\alpha_j, \qquad i = 1, 2, \ldots, k; \tag{66}$$

where these equations are also referred as traffic (flow balance, conservation, etc.) equations.

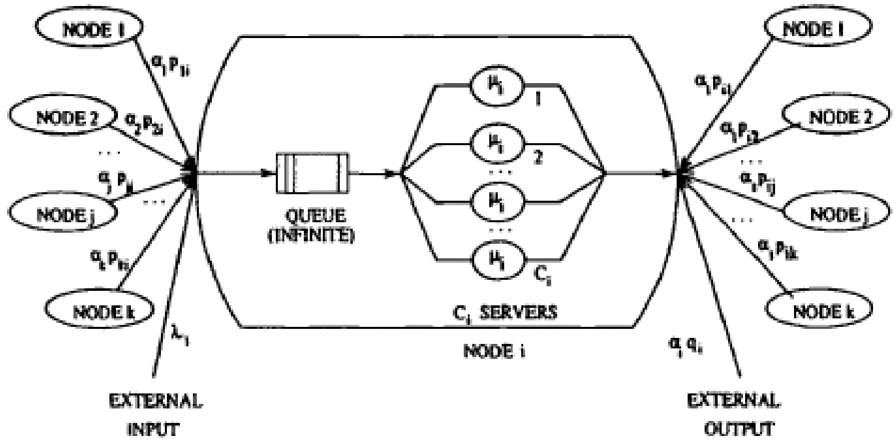

Figure 5.1 Node i in a Jackson Network.

Theorem 4.12 (Jackson's Theorem). *Let $(n_1, n_2, \ldots, n_k)$ denote the state of the complete system in which there are n_i (in the queue and in service) at node i in a Jackson network of Markovian queues in equilibrium, and let $p(n_1, \ldots, n_k)$ be the probability that the system is in the state $(n_1, \ldots, n_k)$.*

Assume that

$$\rho_i = \frac{\alpha_i}{\mu_i} < 1, \qquad i = 1, 2, \ldots, k,$$

where $\{\alpha_i\}$ are given by the balance equations

$$\alpha_i = \lambda_i + \sum_{j=1}^{k} \alpha_j p_{ji}, \qquad i = 1, 2, \ldots, k. \tag{67}$$

If $p_i(n)$ denotes the probability that there are n customers in the system (in queue plus service) for the $M/M/c_i$ queue with input rate α_i and service rate μ_i for each of the c_i servers, i.e.,

$$p_i(n) = p_i(0)\, \frac{\left(\frac{\alpha_i}{\mu_i}\right)^n}{n!}, \qquad n = 0, 1, 2, \ldots, c_i, \tag{68}$$

$$= p_i(0)\, \frac{\left(\frac{\alpha_i}{\mu_i}\right)^n}{c_i!\, c_i^{n-c_i}}, \qquad n = c_i + 1, \ldots, \tag{69}$$

then we have

$$p(n_1, n_2, \ldots, n_k) = p_1(n_1)\, p_2(n_2) \cdots p_k(n_k). \tag{70}$$

Proof Sketch. Let $p_t(n_1, \ldots, n_k)$ be the probability that the complete system is in state $(n_1, \ldots, n_k)$ at time t.

Let

$$q_i = 1 - \sum_j p_{ij}, \qquad a_i(n) = \min\{n_i, c_i\} = \begin{cases} n_i, & \text{if } n < c_i, \\ c_i, & \text{if } n \geq c_i, \end{cases}$$

$$\delta_i = \min\{n_i, 1\} = \begin{cases} 1, & n_i \geq 1, \\ 0, & n_i = 0. \end{cases}$$

Our goal is to write the differential equations satisfied by p_t. Therefore, we will consider the state changes in an infinitesimal interval $(t, t+h)$ following the interval $(0, t)$. p_t. Consider the following four mutually exclusive ways to move from t to $t+h$:

(A) State at t is $(n_1, \ldots, n_k)$ and there are no arrivals or departures occur to or from any node externally. We get

$$\Pr(A) = p_t(n_1, \ldots, n_k)\left[1 - \left(\sum_i \lambda_i\right) h - \sum_i a_i(n_i)\mu_i h\right] + o(h). \tag{71}$$

(B) State at t is $(n_1, \ldots, n_i + 1, \ldots, n_k)$ and there is one service completion at i in $(t, t+h)$, and this completion departs from the system (with probability q_i). We get

$$\Pr(B) = \sum_{i=1}^{k} p_t(n_1, \ldots, n_i + 1, \ldots, n_k)\left[a_i(n_i + 1)\mu_i q_i\right] h + o(h). \tag{72}$$

(C) State at t is $(n_1, \ldots, n_i - 1, \ldots, n_k)$ and there is one arrival from the external source to node i in the interval $(t, t+h)$. We get

$$\Pr(C) = \sum_{i=1}^{k} p_t(n_1, \ldots, n_i - 1, \ldots, n_k)\left[\lambda_i h \delta_i\right] + o(h). \tag{73}$$

(D) State at t is $(n_1, \ldots, n_i + 1, \ldots, n_j - 1, \ldots, n_k)$: there is one service completion at node i in $(t, t+h)$, and the one whose service is completed moves to node j with probability p_{ij}. Thus,

$$\Pr(D) = \sum_i \sum_j p_t(n_1, \ldots, n_i + 1, \ldots, n_j - 1, \ldots, n_k)\left[a_i(n_i + 1)\mu_i h p_{ij}\right] + o(h). \tag{74}$$

Merging all cases, we have

$$p_{t+h}(n_1, \ldots, n_k) = \Pr(A) + \Pr(B) + \Pr(C) + \Pr(D). \tag{75}$$

(75) can be written as

$$p_{t+h}(n_1, \ldots, n_k) - \Pr(A) = \Pr(B) + \Pr(C) + \Pr(D). \tag{76}$$

$$p_{t+h}(n_1, \ldots, n_k) - p_t(n_1, \ldots, n_k)\left[1 - \left(\sum_i \lambda_i\right) h - \sum_i a_i(n_i)\mu_i h\right] - o(h) = \Pr(B) + \Pr(C) + \Pr(D). \tag{77}$$

Taking the limit as $h \to 0$ and solving for $p'(t) = 0$ gives the equations satisfied by steady-state probabilities:

$$\left[\sum_i \lambda_i + \sum_i a_i(n_i)\mu_i\right] p(n_1,\ldots,n_k) = \sum_i a_i(n_i+1)\mu_i q_i\, p(n_1,\ldots,n_i+1,\ldots,n_k)$$
$$+ \sum_i \lambda_i \delta_i\, p(n_1,\ldots,n_i-1,\ldots,n_k)$$
$$+ \sum_i \sum_j a_i(n_i+1)\mu_i p_{ij}\, p(n_1,\ldots,n_i+1,\ldots,n_j-1,\ldots,n_k).$$

Finally, one can show that (70) uniquely satisfies the equations above.

□

4.11 The M/G/1 model

Now assume that we have a Poisson arrival process with rate λ, the service times are i.i.d. and follow a general distribution with mean $\mathbb{E}[S] = \frac{1}{\mu}$, and there is a single server. Again, for stability, we assume that $\rho = \frac{\lambda}{\mu} < 1$.

Let R be the residual service time and let P_k denote the probability that there are k customers in the system in the steady-state. By PASTA property, we have

$$\begin{aligned} W_q &= \sum_{k=1}^{\infty} \left(E(R) + (k-1)E(S)\right) P_k \\ &= \sum_{k=1}^{\infty} E(R)P_k + \left(\sum_{k=1}^{\infty} (k-1)P_k\right) E(S) \\ &= E(R)\rho + L_q E(S). \end{aligned}$$

where the equation follows since $1-\rho = P_0$. By Little's Law, we get

$$W_q = \frac{\rho\mathbb{E}[R]}{1-\rho} \tag{78}$$

which is known as the Pollaczek-Khintchine mean value formula. So now we have to characterize $\mathbb{E}[R]$.

Proposition 4.13. *We have* $\mathbb{E}[R] = \frac{\mathbb{E}[S^2]}{2\mathbb{E}[S]}$.

There are various ways to prove this, but the most pedagogical one is via renewal reward theorem (see the next section). From Proposition 4.13, we have

$$E(R) = \frac{E(S^2)}{2E(S)} = \frac{\mathrm{Var}(S) + E[S]^2}{2E(S)} = \frac{1}{2}\left(C_S^2 + 1\right) E(S), \tag{79}$$

C_S^2 is the squared coefficient of the service time S. Thus for the mean waiting time we have

$$W_q = \frac{\rho E(R)}{1-\rho} = \frac{\rho}{2(1-\rho)}\left(C_S^2 + 1\right) E(S).$$

And by Little's law, we get

$$L_q = \frac{\rho^2}{1-\rho} \frac{C_S^2 + 1}{2}.$$

Discussion 4.14. Kingman's G/G/N formula:

$$W \approx \frac{1}{\mu N} \cdot \frac{\rho^{\sqrt{2(N+1)}-1}}{1-\rho} \cdot \frac{C_A^2 + C_S^2}{2} + \frac{1}{\mu}.$$

4.12 The renewal reward theorem

Definition 4.15. A random point process $\psi = \{t_n\}$ for which the (non-negative) interarrival times $X_n = t_n - t_{n-1}$, $n \geq 1$, form an i.i.d. sequence is called a *renewal process*.

Following the definition, t_n is called the nth *renewal epoch* and $F(x) := P(X \leq x)$, $x \geq 0$, denotes the common inter-arrival time distribution. $t_n = X_1 + \cdots + X_n$, and $N(t) = \max\{n : t_n \leq t\}$ is the counting process. The rate of the renewal process is denoted by $\lambda \triangleq 1/\mathbb{E}[X]$.

Theorem 4.16 (Elementary renewal theorem)**.** *For a renewal process,*

$$\lim_{t\to\infty} \frac{N(t)}{t} = \lambda \quad a.s.$$

and

$$\lim_{t\to\infty} \frac{\mathbb{E}[N(t)]}{t} = \lambda.$$

Proof of Theorem 4.16. We start with the first statement. Note that $t_n = X_1 + \cdots + X_n$, $n \geq 1$. Consider the time

$$t_{N(t)} \leq t < t_{N(t)+1}, \tag{80}$$

Then we have

$$t_{N(t)} = \sum_{j=1}^{N(t)} X_j, \qquad t_{N(t)+1} = \sum_{j=1}^{N(t)+1} X_j,$$

which can be written as

$$\frac{1}{N(t)} \sum_{j=1}^{N(t)} X_j \;\leq\; \frac{t}{N(t)} \;<\; \frac{1}{N(t)} \sum_{j=1}^{N(t)+1} X_j.$$

By the strong law of large numbers, the left and right hand sides converge to $E(X)$ as $t \to \infty$ almost surely. One can prove the second part via Wald's identity (which you will prove in your homework with some provided hints). □

Now let $R(t) = \sum_{j=1}^{N(t)} R_j$ be the total amount of reward collected by time t, where $N(t)$ is the counting process for the renewal process. We want to calculate our long-run reward rate

$$\lim_{t\to\infty} \frac{R(t)}{t}.$$

Theorem 4.17 (Renewal reward theorem)**.** *For a positive recurrent renewal process in which a reward R_j is earned during cycle length X_j and such that $\{(X_j, R_j) : j \geq 1\}$ is i.i.d. with $\mathbb{E}[|R_j|] < \infty$, the long run rate at which rewards are earned is given by*

$$\lim_{t\to\infty} \frac{R(t)}{t} = \frac{\mathbb{E}[R]}{\mathbb{E}[X]} = \lambda \mathbb{E}[R] \quad a.s., \tag{81}$$

where (X, R) *denotes a typical "cycle"* (X_j, R_j)*;* $\lambda = \{\mathbb{E}[X]\}^{-1}$ *is the arrival rate for the renewal process.*

Moreover,

$$\lim_{t\to\infty} \frac{\mathbb{E}[R(t)]}{t} = \frac{\mathbb{E}[R]}{\mathbb{E}[X]}. \tag{82}$$

Discussion 4.18. Now let's prove Proposition 4.13 via the renewal reward theorem. Consider a renewal point process $\{t_n : n \geq 1\}$ with i.i.d. interarrival times $X_n = t_n - t_{n-1}$, $n \geq 1$. Define

$$A(t) = t_{N(t)+1} - t, \qquad t \geq 0. \tag{83}$$

$A(t)$ is called the *excess at time* t, or *remaining lifetime*. If $t_{n-1} \leq t < t_n$, then

$$A(t) = t_n - t \leq X_n.$$

Note that if $\{t_n\}$ is a Poisson process at rate λ, then by the memoryless property we have $A(t) \sim \exp(\lambda)$, $t \geq 0$. But for a general renewal process (as in residual service time in $M/G/1$ queue), we need to be smarter. We want to show that

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t A(s)\,ds = \frac{\mathbb{E}[X^2]}{2\mathbb{E}[X]} \quad \text{a.s.}.$$

Note that we can view the i.i.d. X_j as cycle lengths (service times), and $r(t) = A(t)$ as the generated reward rate at time t. Let R_1 be the generated reward in the first cycle. Then we have

$$R_1 = \int_0^{X_1} A(s)\,ds = \int_0^{X_1} (X_1 - s)\,ds = \frac{X_1^2}{2}.$$

Since $\{(X_j, R_j)\}$'s are i.i.d., by the renewal reward theorem, we almost surely have

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t A(s)\,ds = \frac{\mathbb{E}[R]}{\mathbb{E}[X]} = \frac{\mathbb{E}[X^2]}{2\mathbb{E}[X]}$$

Discussion 4.19 (Inspection paradox)**.** Let $S(t) = t_{N(t)+1} - t_{N(t)}$ be the length of the interarrival time covering time t. If $t_{j-1} \leq t < t_j$, then we have $S(t) = X_j$. Define the reward rate as $r(t) = S(t)$. Then we get

$$R_j = \int_{t_{j-1}}^{t_j} S(s)\,ds = \int_{t_{j-1}}^{t_j} X_j\,ds = X_j \int_{t_{j-1}}^{t_j} ds = X_j^2.$$

By the renewal reward theorem, we have almost surely that

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t S(s)\,ds = \frac{\mathbb{E}[R]}{\mathbb{E}[X]} = \frac{\mathbb{E}[X^2]}{\mathbb{E}[X]},$$

where the fact that $\frac{E[X^2]}{E[X]} \geq \mathbb{E}[X]$ yields the inspection paradox.

5 Foster-Lyapunov Techniques & Dynamic Matching Models

We will attempt to study Foster-Lyapunov techniques via dynamic matching models (also known as matching queues). Before we proceed with the dynamic matching model, here is one informal version of the Foster-Lyapunov criteria. Let X be an irreducible discrete-time Markov chain with a countable state space $\mathcal{S}$. Let $\mathbb{P}(x, A) = \mathbb{P}(X(t) \in A | X(t-1) = x)$ be a transition operator. Let $h : S \to \mathbb{R}_{\geq 0}$ be some function. Then we denote the drift of $h(\cdot)$ at $x \in \mathcal{S}$ by

$$\Delta h(x) = \int P(x, dy) h(y) - h(x).$$

It follows that X is positive recurrent if for all $x \in S$, there exists some non-negative function f, a finite set C, and constants $\delta, b > 0$ such that

$$\Delta h(x) \leq -\epsilon f(x) + b \mathbb{1}_{\{x \in C\}}.$$

This is a powerful result that we can use to establish ergodicity for various stochastic processes. One can further show that $\mathbb{E}_\pi[f(X)] \leq \frac{b}{\epsilon}$. We will be more formal soon.

5.1 Two-way matching model

Consider the following two-way matching model. There is a finite set of *agent types* $\mathcal{A} = \{1, 2, \ldots, n\}$, a finite set of *matches* $\mathcal{M} = \{1, \ldots, d\}$, and a *match value* $r_m > 0$ for each match $m \in \mathcal{M}$. Each match $m \in \mathcal{M}$ is characterized by *two* participating agent types, denoted by the set $\mathcal{A}(m)$. The *network topology* is specified by a *matching matrix* $M \in \{0, 1\}^{n \times d}$, where $M_{im} = 1$ if and only if $i \in \mathcal{A}(m)$. There is no harm in assuming that each agent type participates in at least one match. Each agent type $i \in \mathcal{A}$ is associated with an *arrival probability* $\lambda_i > 0$; $\sum_{i \in \mathcal{A}} \lambda_i = 1$. We refer to the tuple $\mathcal{G} = (M, \lambda, r)$ as the *matching network*.

The matching network induces a weighted undirected simple graph, where the set of vertices is $\mathcal{A}$ and the set of edges is $\mathcal{M}$: there is an edge between $i, j \in \mathcal{A}$ with weight r_m if and only if there exists $m \in \mathcal{M}$ such that $\mathcal{A}(m) = \{i, j\}$. We can assume without of loss generality that $\mathcal{G}$ is connected.

Dynamics. Time is discrete, and there is a single agent arrival every period. The arriving agent is of type $i \in \mathcal{A}$ with probability λ_i. We maintain a separate queue for each agent type, and agents join their type-dedicated queues upon arrival. All queues are empty at time $t = 0$.

Match $m \in \mathcal{M}$ is *available* at time t if and only if the queues of both agent types in $\mathcal{A}(m)$ are non-empty at that time. Performing $m \in \mathcal{M}$ once requires one agent from each type in $\mathcal{A}(m)$ and generates a value of r_m. Matched agents leave the market immediately.

The process A_i^t counts the number of arrivals to queue $i \in \mathcal{A}$ until (and including) time t. The sequence of events in a time period is: an agent arrival is realized, then matches are performed, and queue-lengths are updated. The process Q_i^t tracks the number of agents waiting in queue $i \in \mathcal{A}$ at time t, *after* all matches for this period have been performed.

Matching policy. A matching *policy* is a mapping from histories of arrivals and performed matches to a (possibly empty) set of matches. Given the history, the matching policy determines how many times each match is performed at each time period. An *admissible* matching policy is an increasing non-anticipative process $D^t := (D^t_m : m \in \mathcal{M}, t \geq 0)$, where D^t_m is the number of times match $m \in \mathcal{M}$ is performed by time t; D^t must satisfy

$$Q^t = A^t - MD^t \text{ for all } t \geq 0. \tag{84}$$

We assume that D^t is right-continuous with left limits (RCLL). $\Delta D^t_m := D^t_m - D^{t-1}_m$ is then the number of times match $m \in \mathcal{M}$ is performed at time $t > 0$. We add the superscript D on expectations to make explicit the dependence on the policy, where the superscript is omitted when the context is clear. The family of all admissible matching policies is denoted by Π.

Greedy policies are a large family of admissible policies. These policies perform, whenever possible, a match among those available within a prespecified set. The reason of defining a prespecified set will be clear later.

Definition 5.1 (greedy policy)**.** Given a matching network $\mathcal{G}$ and a subset $\mathcal{S} \subseteq \mathcal{M}$ (not necessarily strict), we say that a policy D is a *greedy policy* with respect to $\mathcal{S}$, if

(i) a match is performed whenever at least one match becomes available to perform in $\mathcal{S}$, and

(ii) matches in $\mathcal{M} \backslash \mathcal{S}$ are never performed, i.e., $D^t_m = 0$ for all $m \in \mathcal{M} \backslash \mathcal{S}$ and for all $t \geq 0$.

Optimality criterion. The expected *total value* generated by time t under a policy D is given by

$$\mathcal{R}^{D,t} := \mathbb{E}^D[r \cdot D^t].$$

For any *fixed* t, the *optimal value* $\mathcal{R}^{*,t} := \max_{D \in \Pi} \mathcal{R}^{D,t}$ is trivially attained by the policy, which takes no action until time t and follows an optimal (static) weighted matching at time t. That is,

$$\mathcal{R}^{*,t} := \mathbb{E} \begin{bmatrix} \max & r \cdot y \\ \text{s.t.} & My \leq A^t \\ & y \in \mathbb{Z}^d_{\geq 0} \end{bmatrix},$$

where the expectation is taken over all realizations of A^t.

The function $\mathcal{R}^{*,t}$ can be interpreted as the *hindsight upper bound* at time t, i.e., the decision maker is allowed to correct past decisions so that previously performed matches may be revoked to perform new ones at all times. A matching policy is *hindsight optimal* if it is, *at all times*, *almost* as good as the optimal value.

Definition 5.2 (hindsight optimality)**.** A matching policy D is *hindsight optimal* if

$$\mathcal{R}^{*,t} - \mathcal{R}^{D,t} = \mathcal{O}(1) \text{ for all } t > 0,$$

which implies, in particular, $\mathcal{R}^{D,t}/\mathcal{R}^{*,t} = 1 - \mathcal{O}(1/t)$ for all $t > 0$.

The existence of a hindsight optimal matching policy means that the tension between short- and long-term objectives is essentially moot; a good performance at time t_0 does not necessitate a significant compromise at time $t_1 > t_0$. Observe that a hindsight optimal matching policy is also optimal in the long-run average sense:

$$\frac{\mathcal{R}^{*,T} - \mathcal{R}^{D,T}}{\mathcal{R}^{*,T}} = \mathcal{O}(1/T) \to 0 \text{ as } T \to \infty. \tag{85}$$

5.2 Static planning problem and general position condition

Relaxing the integrality constraints and applying Jensen's inequality gives the following upper bound on $\mathcal{R}^{*,t}$:

$$\mathcal{R}^{*,t} = \mathbb{E}\begin{bmatrix} \max & r \cdot y \\ \text{s.t.} & My \le A^t \\ & y \in \mathbb{Z}^d_{\ge 0} \end{bmatrix} \le \begin{array}{ll} \max & r \cdot x \\ \text{s.t.} & Mx \le \lambda t \\ & x \in \mathbb{R}^d_{\ge 0}. \end{array}$$

With the change of variables $z = x/t$, we can write the upper bound in standard form as follows:

$$\begin{array}{ll} \max & r \cdot z \\ \text{s.t.} & Mz + s = \lambda \\ & z \in \mathbb{R}^d_{\ge 0}, s \in \mathbb{R}^n_{\ge 0}. \end{array} \tag{SPP}$$

We refer to this formulation as the *static-planning problem* (SPP). The following definition introduces the notion of *general position* that captures the level of stability in a matching network and plays a crucial role in our main results. In fact, general position is a necessary condition to achieve hindsight optimality (next lecture).

Definition 5.3 (general position)**.** A matching network $\mathcal{G}$ satisfies the *general position* condition (**GP**) if (SPP) has a unique non-degenerate optimal solution (z^*, s^*), i.e., all n basic variables in this solution are strictly positive. Define the sets

$$\mathcal{M}_+ := \{m \in \mathcal{M} : z^*_m > 0\},\ \mathcal{M}_0 := \mathcal{M} \backslash \mathcal{M}_+,\ \mathcal{Q}_+ := \{j \in \mathcal{A} : s^*_j > 0\} \text{ and } \mathcal{Q}_0 := \mathcal{A} \backslash \mathcal{Q}_+,$$

where $\mathcal{M}_+$ is the set of *active* matches, $\mathcal{M}_0$ is the set of *redundant* matches, $\mathcal{Q}_+$ is the set of *under-demanded (non-empty)* queues, and $\mathcal{Q}_0$ is the set of *over-demanded (empty)* queues. The *general position gap* is defined as

$$\epsilon := \min_{m \in \mathcal{M}_+} z^*_m \wedge \min_{j \in \mathcal{Q}_+} s^*_j.$$

Residual graph. To achieve hindsight optimality, any matching policy must mostly avoid performing redundant matches. Accordingly, the policies that we will propose are greedy with respect to the set $\mathcal{S} = \mathcal{M}_+ \subsetneq \mathcal{M}$. Let $\mathcal{G}' := \mathcal{G} - \mathcal{M}_0$ be the (SPP)-*residual graph*, which is obtained from $\mathcal{G}$ by removing all redundant matches (every $m \in \mathcal{M}$ with $z^*_m = 0$). The (SPP)-residual graph $\mathcal{G}'$ is then a union of (possibly) multiple components, and we write $\mathcal{G}' = \cup_{k \in [K]} \mathcal{C}_k$, where $\mathcal{C}_k$ is the k^{th} component of $\mathcal{G}'$. Since $\mathcal{G}$ is a simple graph, any edge (match) removal can increase the number of components at most by 1; $K \le |\mathcal{M}_0| + 1$. Let $\mathcal{A}(\mathcal{C}_k)$ be the set of all vertices (queues) in $\mathcal{C}_k$, and let $\mathcal{M}(\mathcal{C}_k)$ be the set of all edges (matches) in $\mathcal{C}_k$ for all $k \in [K]$.
The (SPP)-residual graph $\mathcal{G}'$ has some useful properties, which will be crucial in the design and analysis of our policies.

Lemma 5.4. Assume that $\mathcal{G}$ satisfies **GP**. Then each component $\mathcal{C}_k$, $k \in [K]$, of the (SPP)–residual graph $\mathcal{G}'$ satisfies the following properties: (i) $\mathcal{C}_k$ contains at most one cycle, (ii) if $\mathcal{C}_k$ does not contain a cycle, then $\mathcal{C}_k$ is a tree and $|\mathcal{A}(\mathcal{C}_k) \cap \mathcal{Q}_+| = 1$, and (iii) if $\mathcal{C}_k$ contains a cycle, then the cycle is of odd length and $|\mathcal{A}(\mathcal{C}_k) \cap \mathcal{Q}_+| = 0$.

As an important consequence of the general position condition, bounding the all-time regret of a policy can be boiled down to analyzing the total length of the over-demanded queues, provided that the policy is restricted to active matches.

Lemma 5.5. Suppose that $\mathcal{G}$ satisfies the general position condition, and let (z^*, s^*) be a non-degenerate optimal solution of (λ). Suppose that the following conditions hold under a policy D:

1. Only matches in $\mathcal{M}_+$ are performed, and
2. $\sum_{i\in\mathcal{Q}_0} \mathbb{E}[Q_i(t)] \leq B$ for every $t > 0$, where $B > 0$ does not depend on t.

Then, $\mathcal{R}^{*,t} - \mathcal{R}^{D,t} \leq r_{\max} nB$, where $r_{\max} \triangleq \max_{m\in\mathcal{M}_+} r_m$.

The optimality test lemma should already hint you that Foster-Lyapunov techniques will be very useful to bound stationary expectations of queue-lengths so that we can establish constant regret bounds. The following is one way to bound stationary expectations.

Lemma 5.6. [GZ08][Corollary 4] Let $X = (X_t : t \geq 0)$ be a discrete-time S-valued Markov chain with transition kernel P, and suppose $f : S \to \mathbb{R}$ is nonnegative. If there exists a nonnegative function $g : S \to \mathbb{R}$ and a constant c for which

$$\int_S P(x, dy)g(y) - g(x) \leq -f(x) + c \quad \text{for all } x \in S, \tag{86}$$

then

$$\int_S \pi(dx)f(x) \leq c, \tag{87}$$

for any stationary distribution π of X.

5.3 Candidate matching policies

Definition 5.7 (longest-queue policy)**.** Given a matching network $\mathcal{G}$, the longest-queue policy, denoted by $LQ(\mathcal{M}_+)$, is a greedy policy with respect to $\mathcal{M}_+$ such that

(i) At any time $t > 0$, upon arrival of an agent (say type$-i$), perform the available match $m \in \mathcal{M}_+$ such that $A(m) = \{i, j\}$ and $j \in \{Q_k^t : A(m') = \{i, k\} \textit{ for some } m' \in \mathcal{M}_+\}$, where ties are broken arbitrarily, and

(ii) at the end of each time period (after a match is performed), all agents of types $i \in \mathcal{Q}_+$ leave the market unmatched.

Definition 5.8 (static priority policy)**.** Given a matching network $\mathcal{G}$, the static priority policy, denoted by $SP(\mathcal{M}_+, p)$, is a greedy policy with respect to $\mathcal{M}_+$ such that

(i) $p : \mathcal{M}_+ \to \{1, \ldots, |\mathcal{M}_+|\}$ is a bijective static priority order. We say that $m \in \mathcal{M}_+$ has a higher priority than $m' \in \mathcal{M}_+$ if and only if $p(m) < p(m')$,

(ii) at any time $t > 0$, upon arrival of an agent (say type$-i$), perform the highest priority match $m \in \mathcal{M}_+$ among those available, where $m \in \{p(m') : i \in A(m')\}$, and

(iii) at the end of each time period (after a match is performed), all agents of type$-i$, $i \in \mathcal{Q}_+$, leave the market unmatched.

Discussion 5.9. What are other candidate matching policies? How do they differ in terms of operational costs?

Discussion 5.10. One appropriate Lyapunov function to show that longest-queue policy is hindisight optimal is $g(Q^t) = \sum_{i \in \mathcal{Q}_0} (Q_i^t)^2$. However, it is not appropriate for the static priority policy.

6 Mean field theory

6.1 The power of two in token systems

Consider the following trading favors model. There is a finite set of agents $\mathcal{A} = \{1, 2, \ldots, n\}$, $n \geq 2$. The time $t \in \mathbb{Z}_{\geq 0}$ is discrete. The number of tokens agent $i \in \mathcal{A}$ has at time t is denoted by $s_i^t \in \mathbb{Z}$. We assume that $s_i^0 = 0$ for all $i \in \mathcal{A}$. Let $s^t \in \mathbb{Z}^n$ track the number of tokens agents have at time t.

Let $P = (p_i)_{i \in \mathcal{A}}, Q = (q_i)_{i \in \mathcal{A}}$ be full-support probability measures over the set of agents. At each time period, nature picks one agent to become a *service requester* according to P. Let d be a positive integer, which we call *service availability density.* At each time period, nature picks *available service providers* by selecting d agents according to Q independently and with replacement. Thus, at most d agents are available to provide service at each time period. We say that agents are *symmetric* if $p_i = q_i = \frac{1}{n}$ for all $i \in \mathcal{A}$.

We refer the tuple (n, P, Q, d) as the token system. We will analyze the behavior of the token system under a natural matching policy called the *minimum token selection rule.* This policy, at each time period, selects the available provider with the lowest number of tokens as the *service provider* (ties are broken by choosing uniformly at random). At each time t, if agent i is the service requester and agent j is the service provider, i pays one token to j. Note that an agent can provide service to herself.[2] In this case, $s^{t+1} = s^t$. Otherwise, $s_i^{t+1} = s_i^t - 1$, $s_j^{t+1} = s_j^t + 1$, and $s_k^{t+1} = s_k^t$ for all $k \in \mathcal{A} \backslash \{i, j\}$. In either case, we have $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all $t \geq 0$. The case $d = 1$ is the degenerate case, where the system simply selects one service requester and one service provider independently at random.

Stability. Under a token system (n, P, Q, d), the state of amount of tokens s^t evolves according to a Markov chain defined on the state space $\left\{s \in \mathbb{Z}^n \, : \, \sum_{i \in \mathcal{A}} s_i = 0\right\}$. Our assumptions that P and Q are full-support probability measures over $\mathcal{A}$ ensure that this Markov chain is irreducible. Furthermore, since there is a positive probability that the service requester is the service provider herself at each time period, the Markov chain is aperiodic.

We say that a token system (n, P, Q, d) is *stable* if this Markov chain has a stationary probability distribution. The reason we associate the existence of a stationary distribution with stability is the fact that it is equivalent to each of the following conditions:

- **(C1)** There is a uniformly small probability that the number of tokens owned or owed by any agent is large. Formally, there is a function $f \colon \mathbb{Z}_{\geq 0} \to [0, 1]$ such that $\lim_{M \to \infty} f(M) = 0$, and for all times t large enough and all agents $i \in \mathcal{A}$ it holds that $\mathbb{P}(|s_i^t| > M) < f(M)$.
- **(C2)** The expected time for the token system to clear is finite. Formally, let T_0 be the first time the system returns to 0, i.e., $T_0 = \min\{t > 0 \, : \, s^t = 0\}$. Then $\mathbb{E}(T_0)$ is finite. Note that by the Markov property, this is also the expected time to return to 0 after a later visit to 0.

[2]This technical assumption is useful for the proofs. Moreover, it is motivated by the kidney exchange setting, where patient-donor pairs from the same hospital can be matched internally in a centralized setting when hospitals merge their patient-donor pools. Our results hold with minor modifications if an agent cannot serve herself.

6.1.1 The case $d = 1$

Note that when $d = 1$, s_i^t is a lazy random walk in one dimension for all $i \in \mathcal{A}$, and so by standard arguments the token system is not stable. Moreover, s^t is a random walk on $\mathbb{Z}^{n-1}$ ($n-1$ dimensions), and hence, the Markov chain will not be recurrent by Pólya's Recurrence Theorem for all $n \geq 4$, so that the market will eventually stop clearing at all.

An immediate corollary is that the token system is also not stable under the *random tie-breaking selection rule*, where the rule selects the service provider, even if all providers are available ($d = n$), uniformly at random.

6.1.2 Two agents

Here we analyze the token distribution when there are only 2 agents ($n = 2$). Although understanding the token distribution for this case is simple, the analysis will provide insights regarding the token distribution for the general case. Informally, the best concentration around the initial point is achieved when $n = 2$; as the number of agents increases, the "distance" of the token distribution from the initial point increases as well.

Proposition 6.1. *For $d \geq 2$, the token system with 2 agents is stable if and only if $q_1^d < p_1$, $q_2^d < p_2$. In this case, let π be the the steady-state distribution of the Markov chain $(s^t : t \geq 0)$. Then for all $M \in \mathbb{Z}_+$, we have*

$$\mathbb{P}_\pi(|s_i^t| > M) = \frac{\left(\frac{p_2 q_1}{p_1 - q_1^d}\left(\frac{p_2 q_1^d}{p_1(1-q_1^d)}\right)^M + \frac{p_1 q_2}{p_2 - q_2^d}\left(\frac{p_1 q_2^d}{p_2(1-q_2^d)}\right)^M\right)}{\left(1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}\right)},$$

for $i = 1, 2$. Moreover, the expected time between two successive occurrences of the initial state $(0, 0)$ is given by $1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}$.

The proof is straightforward via studying an appropriate birth-and-death process. One implication of Proposition 6.1 is that the probability of owning or owing a large number of tokens decays exponentially, i.e., $f(M) = O(a^M)$, where the constant $a \in (0, 1)$ can be found via simple algebra.

Proposition 6.1 identifies the level of asymmetry that can be tolerated between service request and service provision (within and across agents) rates. Moving forward, we focus on the symmetric case. Note that for symmetric agents and $d \geq 2$, Proposition 6.1 implies that the system is stable. When furthermore $d = 2$, we get that

$$\mathbb{P}_\pi(|s_i^t| > M) = \frac{2}{3}\left(\frac{1}{3}\right)^M,$$

and that $\mathbb{E}(T_0) = 3$. In the general case, we will argue that $a = 1/3$ is the best rate one can hope for.

6.1.3 Stability

In general, it seems difficult to determine whether a given system (n, P, Q, d) is stable. Indeed, the results of the previous section show that already for $n = 2$, this can be highly sensitive to the precise values of P and Q. The next result shows that in the symmetric case, stability is achieved for any $n \geq 2$, assuming $d \geq 2$.

Theorem 6.2. *The token system is stable for any $d \geq 2$ when the agents are symmetric. Furthermore, (C1) holds with $f(M) = 5/M$.*

Proof. Denote by k^t the agent chosen to request service at time t, and denote by I^t the set of agents chosen to be available to provide service at time t. Let $|I^t|$ denote the size of I^t, and note that $|I^t|$ takes values in $\{1, 2, \ldots, d\}$. Let $j^t \in I^t$ be the agent chosen to provide service, i.e., an agent chosen uniformly from the agents in $i \in I^t$ that minimizes s_i^t. Hence $s_{j^t}^t = \min_{i \in I^t} s_i^t$, so that in particular, the service provider j^t is an agent with the minimum number of tokens among the available agents in I^t.

Let $V^t = \sum_{i=1}^n (s_i^t)^2$, and let

$$v^t := \mathbb{E}[V^t] = \sum_{i=1}^n \mathbb{E}[(s_i^t)^2].$$

Note that by symmetry, we have $v^t = n\mathbb{E}[(s_1^t)^2]$. Let E_t be the event $\{k^t = j^t\}$, i.e., the event that the service provider and requester are the same agent. Let E_t^c be the event $\{k^t \neq j^t\}$. Since k^t is uniformly distributed and k^t and j^t are independent, the probability of E_t is $1/n$. Note that conditioned on E_t, it holds that $s^{t+1} = s^t$. Then we have

$$\begin{aligned} v^{t+1} &= \frac{1}{n}\mathbb{E}[V^{t+1}|E_t] + \frac{n-1}{n}\mathbb{E}[V^{t+1}|E_t^c] \\ &= \frac{1}{n}\mathbb{E}[V^t] + \frac{n-1}{n}\mathbb{E}[V^{t+1}|E_t^c] \\ &= \frac{1}{n}v^t + \frac{n-1}{n}\mathbb{E}[V^{t+1}|E_t^c]. \end{aligned}$$

Now,

$$\begin{aligned} \mathbb{E}[V^{t+1}|E_t^c] &= \sum_{i \in \mathcal{A}} \mathbb{E}[(s_i^{t+1})^2|E_t^c] \\ &= \mathbb{E}[(s_{k^t}^{t+1})^2|E_t^c] + \mathbb{E}[(s_{j^t}^{t+1})^2|E_t^c] + \sum_{i \in I^t \setminus \{j^t\}} \mathbb{E}[(s_i^{t+1})^2|E_t^c] + \sum_{i \in \mathcal{A} \setminus I^t \setminus \{k^t\}} \mathbb{E}[(s_i^{t+1})^2|E_t^c] \\ &= \mathbb{E}[(s_{k^t}^t - 1)^2 + (s_{j^t}^t + 1)^2|E_t^c] + \sum_{i \in I^t \setminus \{j^t\}} \mathbb{E}[(s_i^t)^2] + \sum_{i \in \mathcal{A} \setminus I^t \setminus \{k^t\}} \mathbb{E}[(s_i^t)^2] \\ &= -2\mathbb{E}[s_{k^t}^t|E_t^c] + 1 + 2\mathbb{E}[s_{j^t}^t|E_t^c] + 1 + \sum_{i \in \mathcal{A}} \mathbb{E}[(s_i^t)^2]. \end{aligned}$$

Since $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all $t \geq 0$, we must have that $\mathbb{E}[s_{k^t}^t|E_t^c] \geq 0$. Since k^t and j^t are independent, we have $\mathbb{E}[s_{j^t}^t|E_t^c] = \mathbb{E}[s_{j^t}^t]$. Hence,

$$\mathbb{E}[V^{t+1}|E_t^c] \leq v^t + 2\mathbb{E}[s_{j^t}^t] + 2,$$

and so

$$v^{t+1} \leq v^t + \frac{2(n-1)}{n}\mathbb{E}[s_{j^t}^t] + \frac{2(n-1)}{n}.$$

Without loss of generality, assume that $I^t = \{1, 2, \ldots, |I^t|\}$. Now with probability

$n^{-(d-1)}$, we have that $|I^t| = 1$, in which case $I^t = \{j^t\} = \{1\}$. Hence $\mathbb{E}[s^t_{j^t} \mid |I^t| = 1] = 0$. With probability $1-n^{-(d-1)}$, we have that $|I_t| \geq 2$, in which case $s^t_{j^t} = \min\{s^t_1, s^t_2, \dots, s^t_{|I^t|}\} \leq \min\{s^t_1, s^t_2\}$. Using the fact that $\min\{a, b\} = \frac{1}{2}(a + b - |a - b|)$ for all $a, b \in \mathbb{R}$, we get $2\mathbb{E}[\min\{s^t_1, s^t_2\}] = -\mathbb{E}[|s^t_1 - s^t_2|]$, and we get

$$\mathbb{E}[s^t_{j^t} \mid |I^t| \geq 2] \leq -\frac{1}{2}\mathbb{E}[|s^t_1 - s^t_2|].$$

Thus,

$$\mathbb{E}[s^t_{j^t} \mid |I^t| \geq 2] \leq -\frac{1}{2}(1 - n^{-(d-1)})\mathbb{E}[|s^t_1 - s^t_2|],$$

and we get

$$v^{t+1} \leq v^t - \frac{n-1}{n}(1 - n^{-(d-1)})\mathbb{E}[|s^t_1 - s^t_2|] + 2.$$

Now we make use of the following claim. For real random variables $Y, Z_1, \dots, Z_n$, we have

$$\mathbb{E}\left[\left|Y - \frac{1}{n}\sum_{i=1}^{n} Z_i\right|\right] \leq \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[|Y - Z_i|], \tag{88}$$

where the claim follows immediately from convexity and Jensen's inequality. Again by the symmetry of the problem, we have

$$\mathbb{E}[|s^t_1 - s^t_2|] = \frac{1}{n-1}\sum_{i=2}^{n}\mathbb{E}[|s^t_1 - s^t_i|].$$

Therefore by (88) and using the fact that $\sum_{i=2}^{n} s^t_i = -s^t_1$ for all $t \geq 0$, we have

$$\mathbb{E}[|s^t_1 - s^t_2|] \geq \mathbb{E}\left[\left|s^t_1 - \frac{1}{n-1}\sum_{i=2}^{n} s^t_i\right|\right] = \mathbb{E}\left[\left|s^t_1 + \frac{s^t_1}{n-1}\right|\right] = \frac{n}{n-1}\mathbb{E}[|s^t_1|],$$

Therefore, we have

$$v^{t+1} \leq v_t - (1 - n^{-(d-1)})\mathbb{E}[|s^t_1|] + 2,$$

and via recursion, we get

$$v^t \leq 2t - (1 - n^{-(d-1)})\sum_{u=0}^{t-1}\mathbb{E}[|s^u_1|].$$

Since $v^t \geq 0$, we have

$$\frac{1}{t}\sum_{u=0}^{t-1}\mathbb{E}[|s^u_1|] \leq \frac{2}{1 - n^{-(d-1)}} \leq 4, \tag{89}$$

where the second inequality follows from the fact that for $d, n \geq 2$, the denominator $1 - n^{-(d-1)}$ is at least $1/2$.

Denote the distribution of s^t by $\mu^t \in \Delta(\mathbb{Z}^n)$, and let $\nu^t = \frac{1}{t}\sum_{u=0}^{t-1}\mu_t$. Let Y^t be a random variable with distribution ν^t for all $t \geq 1$, and denote by Y^t_i the i'th index of Y^t for all $i \in \mathcal{A}$. Then per (89), we have

$$\mathbb{E}[|Y_i^t|] = \frac{1}{t}\sum_{u=0}^{t-1}\mathbb{E}[|s_i^u|] \le 4, \tag{90}$$

for all $i \in \mathcal{A}$.

Suppose towards a contradiction that the Markov chain has no stationary probability measure. Then for any finite subset $E \subset \mathbb{Z}^n$, it holds that $\lim_{t\to\infty}\mathbb{P}(s_i^t \in E) = 0$ by standard arguments about Markov chains. In particular, $\lim_{u\to\infty}\mathbb{E}[|s_i^u|] = \infty$, which contradicts (90). Thus, we have stability.

Since the Markov chain has a stationary distribution, and since it is irreducible and aperiodic, the distribution of s_i^t will converge to the stationary distribution. It follows from (90) that $\lim_{t\to\infty}\mathbb{E}[|s_i^t|] \le 4$, and so $\mathbb{E}[|s_i^t|] \le 5$ for all t large enough. Hence, it follows from Markov's inequality that for all t large enough and for all $i \in \mathcal{A}$, we have

$$\mathbb{P}(|s_i^t| \ge M) \le \frac{5}{M}.$$

□

6.2 Density dependent Markov chains

We begin with the definition of density dependent Markov chains. Let $\mathbb{Z}^*$ be either $\mathbb{Z}^m$ for some finite dimension m, or $\mathbb{Z}^{\mathbb{N}}$, and similarly define $\mathbb{R}^*$. Let $L \subseteq \mathbb{Z}^*$ be the set of possible non-zero transitions of the system. For each $\vec{l} \in L$, define a nonnegative function $\beta_{\vec{l}} : \mathbb{R}^* \to [0,1]$.

Definition 6.3. A sequence (indexed by n) of continuous time Markov chains $(X_n(t) : t \ge 0)$ on the state spaces $S_n = \left\{\vec{k}/n : \vec{k} \in \mathbb{Z}^*\right\}$ is a density dependent Markov chain if there exists a $\beta_{\vec{l}}$: $\mathbb{R}^* \to [0,1]$ such that for all n the transition rate of X_n is given by $q_{x,y}^{(n)} = n\beta_{n(y-x)}(x)$, $x, y \in S_n$.

In Definition 6.3, the index n can be interpreted as the total population or volume of the system, and the components of $\vec{k}/n$ can be interpreted as the densities of different types present in the system. The $\beta_{\vec{l}}\,(x)$ can be interpreted as the probability of transition $\vec{l}$ from $x \in S_n$ to $y \in S_n$, where $nx + \vec{l} = ny$. Given a density dependent Markov chain X_n with transition rates $q_{x,y}^{(n)} = q_{\vec{k},\vec{k}+\vec{l}}^{(n)} = n\beta_{\vec{l}}\,(\vec{k}/n)$, define $F(x) = \sum_{\vec{l}\in L}\vec{l}\beta_{\vec{l}}\,(x)$. The following theorem is key in our analysis:

Theorem 6.4 (Kurtz's theorem). *Suppose we have a density dependent Markov chain X_n (of possibly countably infinite dimension) satisfying the Lipschitz condition $|F(x) - F(y)| \le M|x-y|$ for some constant M. Further suppose $\lim_{n\to\infty} X_n(0) = x_0$, and let X be the deterministic process:*

$$X(t) = x_0 + \int_0^t F(X(u))du,\ t \ge 0. \tag{91}$$

Consider the path $\{X(u) : u \le T\}$ for some fixed $T \ge 0$, and assume that there exists a neighborhood K around this path satisfying

$$\sum_{\vec{l}\in L}|\vec{l}|\sup_{x\in K}\beta_{\vec{l}}\,(x) < \infty. \tag{92}$$

Then $\lim_{n\to\infty} \sup_{u\le T} |X_n(u) - X(u)| = 0$ *almost surely.*

The Lipschitz condition ensures the uniqueness of the solution for the differential equation $\dot{X} = F(X)$, which follows by taking the derivative of (91) with respect to t. Condition (92) ensures that the jump rate is bounded in the process. Kurtz's theorem implies that as $n \to \infty$, the behavior of a density dependent Markov chain can be characterized by the deterministic process given in (91), where the convergence holds on a finite time interval $[0, T]$ for an arbitrary T. We next model and study our system as a density dependent Markov chain, which we refer to as the *finite model*.

The finite and infinite models. Let us model the system with n symmetric agents as a density dependent Markov chain and denote it by $(X_n(t), t \ge 0)$. Note that in Definition 6.3, the $\beta_{\vec{l}}$'s are independent of n, and the transition rates are linear in n. In order to fit our system to this definition, we assume that each agent has an exponential clock with rate 1. The ticking of agent i's clock corresponds to a service request by i, and the service provider is selected immediately using the minimum token selection rule. Note that because of the memoryless property and the continuity of the distribution that governs the clocks, agents request service uniformly, and exactly one agent requests service at a time. As a slight abuse of notation, let $n_i(t)$ be the number of agents with i tokens at time t, $m_i(t)$ be the number of agents with at least i tokens at time t, and $z_i(t) := m_i(t)/n$ be the fraction of agents with at least i tokens at time t. Let $\vec{z}(t) = (..., z_{-2}(t), z_{-1}(t), z_0(t), z_1(t), z_2(t), ...)$, and we drop the time index t when the meaning is clear. We represent the state of X_n by $\vec{z} = \vec{k}/n \in \mathbb{Z}^{\mathbb{N}}/n$. Let us call this process the *finite model*. Note that the initial state of X_n is $\vec{z}(0) = (..., 1, 1, 1, 0, 0, ...)$, where $z_i(0) = 1$ for all $i \le 0$, and $z_i(0) = 0$ for all $i \ge 1$.

Next, we describe the transition probabilities $\beta_{\vec{l}}$'s. The set of possible non-zero transitions from $\vec{k} = n\vec{z}$ is $L = \{e_{ij} : i, j \in \mathbb{Z}, i \ne j\}$, where e_{ij} is an infinite dimensional vector of all zeros except the i'th index (which corresponds to the index of z_i) is -1 and the j'th index (which corresponds to the index of z_j) is 1. Note that after transition e_{ij} occurs, nz_i decreases by 1 and nz_j increases by 1 simultaneously. Hence, the transition e_{ij} corresponds to the event when an agent with i many tokens requests service and an agent with $j-1$ many tokens provides service. Since the probability that the service requester has i many tokens is $z_i - z_{i+1}$, and the probability that the service provider has j many tokens is $z_j^d - z_{j+1}^d$, we have $\beta_{e_{ij}}(\vec{z}) = (z_i - z_{i+1})(z_{j-1}^d - z_j^d)$.[3] Denote the *infinite model* by X, which is the limit of the finite model X_n, i.e., $X = \lim_{n\to\infty} X_n$. Since X is characterized by the deterministic process (91), we need to analyze the components of $F(x)$. Note that the i'th component of $F(x) = \sum_{\vec{l}\in L} \vec{l}\beta_{\vec{l}}(x)$ (which corresponds to z_i) is $\sum_{j\in\mathcal{A}\setminus\{i\}}(z_j - z_{j+1})(z_{i-1}^d - z_i^d) - \sum_{j\in\mathcal{A}\setminus\{i\}}(z_i - z_{i+1})(z_{j-1}^d - z_j^d)$, which simplifies to

$$(1 - z_i + z_{i+1})(z_{i-1}^d - z_i^d) - (1 - z_{i-1}^d + z_i^d)(z_i - z_{i+1}) = (z_{i-1}^d - z_i^d) - (z_i - z_{i+1}). \quad (93)$$

Now let's recall the main result we want to prove. Let π be the steady-state distribution of the Markov chain $(s^t : t \ge 0)$. For any $M \in \mathbb{Z}_+$, define

$$p_{n,M} := \mathbb{P}_\pi(|s_1| \le M)$$

[3]The probability that all available agents have at least j many tokens is z_j^d, and we subtract the probability that all available agents have at least $j+1$ many tokens, which is z_{j+1}^d.

when there are n symmetric agents.

Theorem 6.5.

$$\lim_{n\to\infty} p_{n,M} \geq 1 - \left(\frac{1}{2}\right)^M \qquad \textit{for all } M \in \mathbb{Z}_+.$$

Now that we have represented our system using the finite model, we are ready to prove Theorem 6.5. The proof is organized as follows. We first show that the conditions of Theorem 6.4 hold. Then using Theorem 6.4, we obtain the system of ordinary differential equations that characterize the infinite model. This characterization lets us represent the probability of interest in (C1) as n grows large using π_0 (the fraction of agents that have at least 0 tokens in the long-run). Finally, we find lower and upper bounds for π_0 to conclude.

Condition (92) is clearly satisfied since the magnitude of any jump is bounded, and the jump rate is bounded above by 1 for any state.

Now let's show that the Lipschitz condition of Theorem 6.4 holds with $M = 2 + 2d$.

Lemma 6.6 (Lipschitz condition)**.** The finite model satisfies the Lipschitz condition in L_1-distance.

Proof. Let $x = (x_i)_{i\in\mathbb{Z}}$ and $y = (y_i)_{i\in\mathbb{Z}}$ be two states of the finite model. By (93), we have

$$\begin{aligned}
|F(x) - F(y)| &= \sum_{i=-\infty}^{\infty} |(x_{i-1}^d - x_i^d) - (x_i - x_{i+1}) - (y_{i-1}^d - y_i^d) + (y_i - y_{i+1})| \\
&\leq 2 \sum_{i=-\infty}^{\infty} |x_i^d - y_i^d| + 2 \sum_{i=-\infty}^{\infty} |x_i - y_i| \\
&\leq (2 + 2d) \sum_{i=-\infty}^{\infty} |x_i - y_i| \\
&= (2 + 2d)|x - y|,
\end{aligned}$$

where in the first inequality we used the triangle inequality, and in the second inequality we used the expansion $(a - b)^n = (a - b)(a^{n-1} + a^{n-2}b + ... + ab^{n-2} + b^{n-1})$ and the fact that $0 \leq x_i, y_i \leq 1$ for all $i \in \mathbb{Z}$. ☐

By differentiating (91) with respect to t and using (93), we get the following system of ordinary differential equations that characterizes the infinite model:

$$\frac{dz_i}{dt} = (z_{i-1}^d - z_i^d) - (z_i - z_{i+1}) \text{ for all } i \in \mathbb{Z}. \tag{94}$$

Intuitively, (94) can be interpreted as follows. Let us consider the expected change in m_i (the number of agents with at least i many tokens) over a small time interval dt. First note that a transition occurs whenever one of the agent's exponential clock ticks, which happens with rate ndt. Under such transition, m_i increases by 1 if an agent with $i - 1$ many tokens is selected as the service provider, which happens with probability $z_{i-1}^d - z_i^d$. m_i decreases by 1 if an agent with i many tokens is selected as the service requester, which happens with probability $z_i - z_{i+1}$. Hence, the expected increase in m_i is $(z_{i-1}^d - z_i^d)ndt$, and the expected decrease in m_i is $(z_i - z_{i+1})ndt$, which gives $dm_i = (z_{i-1}^d - z_i^d)ndt - (z_i - z_{i+1})ndt$, and since $m_i/n = z_i$, dividing both sides by ndt gives (94).

Define an *equilibrium point*, which is a point $\vec{a}$ such that if $\vec{z}(t') = \vec{a}$, then $\vec{z}(t) = \vec{a}$ for all $t \geq t'$. Denote the equilibrium point of the infinite model by $\vec{\pi}$, and assume $d = 2$ for simplicity from now on (the following arguments can be easily generalized for $d > 2$). Clearly $\vec{\pi}$ is an equilibrium point of the infinite model if and only if $\frac{d\pi_i}{dt} = 0$ for all $i \in \mathbb{Z}$. Moreover, since agents start with 0 tokens and exchange one token at each transition, the expected number of tokens agents have is 0, and it can be written as follows:

$$\sum_{i \in Z} i \cdot \frac{n_i}{n} = \sum_{i \geq 1} i \cdot \frac{n_i}{n} + \sum_{i \leq 0} i \cdot \frac{n_i}{n} = \sum_{i \geq 1} \frac{m_i}{n} - \sum_{i \leq 0} \frac{n - m_i}{n} = \sum_{i \geq 1} z_i - \sum_{i \leq 0} (1 - z_i) = 0. \quad (95)$$

Using (94) and (95), $\vec{\pi}$ can be found by solving the following system of equations:

$$(\pi_{i-1}^2 - \pi_i^2) - (\pi_i - \pi_{i+1}) = 0 \text{ for all } i \in \mathbb{Z}, \quad (96)$$

$$\sum_{i \geq 1} \pi_i - \sum_{i \leq 0} (1 - \pi_i) = 0. \quad (97)$$

Note that (96) implies $\pi_{i+1} - \pi_i^2 = \pi_0 - \pi_{-1}^2$ for all $i \in \mathbb{Z}$. Since $\lim_{i \to \infty} \pi_i = 0$, we have $\pi_0 = \pi_{-1}^2$, and inductively we have the following relation:

$$\pi_{i+1} = \pi_i^2 \quad \text{for all } i \in \mathbb{Z}. \quad (98)$$

Using (98), (97) becomes

$$\sum_{i \geq 1} \pi_0^{2^i} - \sum_{i \geq 0} \left(1 - \pi_0^{2^{-i}}\right) = 0.$$

Such series are known as *lacunary series*, where the function has no analytic continuation across its disc of convergence (see Hadamard's Gap Theorem). There is no closed form expression for such series to the best of our knowledge and thus, we are unable to find the equilibrium point explicitly. Note that in the long-run, the probability that $|s_i^t| \leq M$ for any $i \in \mathcal{A}$ is equal to $\pi_{-M} - \pi_{M+1}$. Using (98), proving that for all $M \in \mathbb{Z}_+$,

$$\lim_{n \to \infty} p_{n,M} \geq 1 - a^M \quad \text{for some } a \in (0, 1),$$

is equivalent to proving that the following inequalities hold:

$$\pi_0^{2^{-M}} - \pi_0^{2^{M+1}} \geq 1 - a^M \quad \text{for all } M \in \mathbb{Z}_+. \quad (99)$$

Lemma 6.7. We have $\frac{1}{2} < \pi_0 < \frac{3}{4}$.

Proof. We first prove that $\pi_0 > \frac{1}{2}$. Note that (97) can be written as

$$\sum_{M \geq 1} (\pi_M + \pi_{-M+1} - 1) = 0. \quad (100)$$

We claim that if $\pi_1 + \pi_0 = \pi_0^2 + \pi_0 - 1 < 0$, which implies $\pi_0 \leq \frac{1}{2}$, then all the terms in the summation (100) is negative, which is a contradiction. We have $\pi_M + \pi_{-M+1} - 1 = \pi_0^{2^M} + \pi_0^{2^{-M+1}} - 1$ by (98). Let $f(M) = 1 - \pi_0^{2^M} - \pi_0^{2^{-M+1}}$. Assume to the contrary that $\pi_0 \leq \frac{1}{2}$. Then $f(1) > 0$, and clearly we have $\lim_{M \to \infty} f(M) = 0$. We will show that

$f(M) \geq f(M+1)$ for all $M \in \mathbb{Z}_+$. The derivative of $f(M)$ with respect to M is

$$\frac{df(M)}{dM} = -2^M \cdot \log(2) \cdot \pi_0^{2^M} \cdot \log(\pi_0) + 2^{-M+1} \cdot \log(2) \cdot \pi_0^{2^{-M+1}} \cdot \log(\pi_0), \tag{101}$$

where log is the natural logarithm. Since $\pi_0 \leq \frac{1}{2}$ by assumption, we have $\log(2) \cdot \log(\pi_0) < 0$. Thus, (101) has the same sign with

$$2^M \cdot \pi_0^{2^M} - 2^{-M+1} \cdot \pi_0^{2^{-M+1}} = \frac{2^{2M-1} \cdot \pi_0^{2^M} - \pi_0^{2^{-M+1}}}{2^{M-1}}. \tag{102}$$

Since $2^M > 2M-1$ for all $M \geq 1$ and $2\pi_0 \leq 1$, we have

$$2^{2M-1} \cdot \pi_0^{2^M} - \pi_0^{2^{-M+1}} = (2\pi_0)^{2M-1} \cdot \pi_0^{2^M - 2M+1} - \pi_0^{2^{-M+1}} \leq \pi_0^{2^M-2M+1} - \pi_0^{2^{-M+1}} < 0 \tag{103}$$

for all $M > 1$, where the last inequality follows from the fact that $2^M - 2M + 1 > 0$, $-M+1 < 0$, and $0 < \pi_0 < 1$.

Now we prove that $\pi_0 < \frac{3}{4}$. Assume to the contrary that $\pi_0 \geq \frac{3}{4}$. Then clearly we have

$$\sum_{i \geq 1}^{\infty} \pi_i = \sum_{i \geq 1}^{\infty} \pi_0^{2^i} \geq 0.8. \tag{104}$$

Now we want to find an upper bound for

$$\sum_{i \leq 0}^{\infty} (1 - \pi_i) = \sum_{i \geq 0}^{\infty} (1 - \pi_0^{2^{-i}}). \tag{105}$$

Since $\frac{1-\pi_0^{2^{-i}}}{1-\pi_0^{2^{-i-1}}} = 1 + \pi_0^{2^{-i-1}}$ is an increasing function for $i \geq 0$, we can upper bound (105) by the following geometric series

$$(1-\pi_0) + (1 - \pi_0^{2^{-1}}) + (1 - \pi_0^{2^{-2}})(1 + r + r^2 + \ldots) \leq 0.6, \tag{106}$$

where $r = \frac{1-\pi_0^{2^{-2}}}{1-\pi_0^{2^{-3}}}$. But, (104) and (106) contradict to (97). ☐

Now we are ready to conclude the proof of Theorem 6.5. We will show that $g(M) := \pi_0^{2^{-M}} - \pi_0^{2^{M+1}} - 1 + 2^{-M} \geq 0$ for all positive integers M. Note that $\lim_{M \to \infty} g(M) = 0$. Hence, we will show that $g(M) \geq g(M+1)$ for all $M \in \mathbb{Z}_+$. It is easy to check that $g(M) \geq 0$ for all $M \leq 6$ using Lemma 6.7. The derivative of g with respect to M is $\frac{dg(M)}{dM} = \log(\pi_0) \cdot \log(0.5) \cdot \pi_0^{2^{-M}} \cdot 2^{-M} + \log(\pi_0) \cdot \log(0.5) \cdot \pi_0^{2^{M+1}} \cdot 2^{2M+1} \cdot 2^{-M} - \log(2) \cdot 2^{-M}$, where log is the natural logarithm. By Lemma 6.7, we have $\frac{1}{2} < \pi_0 < \frac{3}{4}$, and thus $0.19 < \log \pi_0 \cdot \log(0.5) < 0.5$. Since $\pi_0^{2^{-M}} \leq 1$, the first term in $\frac{dg(M)}{dM}$ is upper bounded by $\frac{1}{2} \cdot 2^{-M}$. For the second term, note that $\pi_0^8 \cdot 2 < \frac{1}{3}$. Since $2^{M+1} > 8(2M+1)$ for all $M \geq 6$, the second term is upper bounded by $\frac{1}{2} \cdot \frac{1}{3} \cdot 2^{-M}$, and $\frac{dg(M)}{dM}$ is upper bounded by $\frac{1}{2} \cdot 2^{-M} + \frac{1}{6} \cdot 2^{-M} - \log(2) \cdot 2^{-M}$, which is negative since $\log(2) > 2/3$. Hence, we have shown that $g(M)$ is a decreasing function on $[6, \infty]$, which concludes the proof.

Example 6.8 (The supermarket model)**.** Recall the definition of the supermarket model:

customers arrive as a Poisson stream of rate λn, where $\lambda < 1$, at a collection of n FIFO servers. Each customer chooses some constant $d \geq 2$ servers independently and uniformly at random with replacement and queues at the server currently containing the fewest customers. The service time for a customer is exponentially distributed with mean 1.

Define $n_i(t)$ to be the number of queues with i customers at time t; $m_i(t)$ to be the number of queues with at least i customers at time t; $p_i(t) = n_i(t)/n$ to be the fraction of queues of size i; and $s_i(t) = \sum_{k=i}^{\infty} p_k(t) = m_i(t)/n$ to be the tails of the $p_i(t)$. Note that $s_0 = 1$ always, and that the s_i are non-increasing. In an empty system, which corresponds to one with no customers, $s_0 = 1$ and $s_i = 0$ for $i \geq 1$.

We can represent the state of the system at any given time by an infinite dimensional vector $\mathbf{s} = (s_0, s_1, s_2, \ldots)$. The time evolution of the infinite system is specified by

$$\begin{cases} \dfrac{ds_i}{dt} = \lambda(s_{i-1}^d - s_i^d) - (s_i - s_{i+1}) \quad \text{for } i \geq 1, \\ s_0 = 1. \end{cases} \tag{107}$$

The intuition goes as follows. Consider a supermarket system with n queues, and determine the expected change in the number of servers with at least i customers over a small period of time of length dt. The probability a customer arrives during this period is $\lambda n\, dt$, and the probability an arriving customer joins a queue of size $i-1$ is $s_{i-1}^d - s_i^d$. (This is the probability that all d servers chosen by the new customer are of size at least $i-1$ but not all are of size at least i.) Thus the expected change in m_i due to arrivals is exactly $\lambda n(s_{i-1}^d - s_i^d)\, dt$. Similarly, the probability a customer leaves a server of size i in this period is $n_i\, dt = n(s_i - s_{i+1})\, dt$. Hence, if the system behaved according to these expectations, we would have

$$\frac{dm_i}{dt} = \lambda n(s_{i-1}^d - s_i^d) - n(s_i - s_{i+1}).$$

Removing a factor of n from the equations yields the system. One can show that when $d \geq 2$ and given that $\sum_{i=1}^{\infty} s_i < \infty$, $s_i = \lambda^{\frac{d^i - 1}{d-1}}$ yields the unique solution.

Example 6.9 (The closed model)**.** Assume that there are n servers serve according to the FIFO rule and the total number of customers in the system is Mn which is fixed throughout the process for some constant M. Start with a system where all the queues have size M. At each time period, select one non-empty queue uniformly at random and complete the service of the customer in that queue. Before leaving the queue, the customer selects some constant number d of servers independently and uniformly at random from the n servers with replacement and joins to the queue that contains the fewest customers where the ties are broken arbitrarily.[4] Using the same definitions in Example 5.2, let $\vec{s}(t) = (s_0(t), s_1(t), s_2(t), \ldots)$ and we drop the time index t when it is clear from the context. Again, we can represent the state of the system by $\vec{s} = k/n$. Now the set of possible transitions from k is $L = \{e_{ij} : i, j \geq 1, i \neq j\}$ where e_{ij} is a vector of all zeros except the i'th index is -1 and the j'th index is 1. For simplicity, assume that $d = 2$. For $|i-j| > 1$ and $i - j = 1$, we have $\beta_{e_{ij}}(\vec{s}) = \dfrac{(s_i - s_{i+1})}{s_1}(s_{j-1}^2 - s_j^2)$. For $i - j = -1$, since the probability that the customer will rejoin to the queue which she currently belongs to is $1/n^2 + 2 \cdot (1/n) \cdot s_{i+1} + 2 \cdot (1/n) \cdot ((n_i - 1)/n) \cdot 1/2 = (s_i + s_{i+1})/n$, we have $\beta_{e_{ij}}(\vec{s}) = \dfrac{(s_i - s_{i+1})}{s_1}(s_{j-1}^2 - s_j^2 - (s_i + s_{i+1})/n)$.

[4]Note that the customer can rejoin to the queue which she currently belongs to.

References

[GZ08] Peter W Glynn and Assaf Zeevi. Bounding stationary expectations of markov processes. In *Markov processes and related topics: a Festschrift for Thomas G. Kurtz*, volume 4, pages 195–215. Institute of Mathematical Statistics, 2008.

[Med02] Jyotiprasad Medhi. *Stochastic models in queueing theory*. Elsevier, 2002.

[Ros11] Nathan Ross. Fundamentals of stein's method. 2011.

[RP23] Sheldon M Ross and Erol A Peköz. *A second course in probability*. Cambridge University Press, 2023.

[VDH24] Remco Van Der Hofstad. *Random graphs and complex networks*, volume 2. Cambridge university press, 2024.