# Modern Discrete Probability An Essential Toolkit 

Sébastien Roch



December 20, 2023

To Betsy

## Contents

Preface ..... v
Notation ..... ix
1 Introduction ..... 1
1.1 Background ..... 1
1.1.1 Review of graph theory ..... 2
1.1.2 Review of Markov chain theory ..... 9
1.2 Some discrete probability models ..... 18
2 Moments and tails ..... 27
2.1 Background ..... 27
2.1.1 Definitions ..... 28
2.1.2 Basic inequalities ..... 29
2.2 First moment method ..... 33
2.2.1 The probabilistic method ..... 33
2.2.2 Boole's inequality ..... 36
2.2.3 $\triangleright$ Random permutations: longest increasing subsequence . ..... 38
2.2.4 $\triangleright$ Percolation: existence of a non-trivial critical value on $\mathbb{Z}^{2}$ ..... 40
2.3 Second moment method ..... 44
2.3.1 Paley-Zygmund inequality ..... 45
2.3.2 $\triangleright$ Random graphs: subgraph containment and connectivity in the Erdö́s-Rényi model ..... 47
2.3.3 $\triangleright$ Percolation: critical value on trees and branching number ..... 55
2.4 Chernoff-Cramér method ..... 63
2.4.1 Tail bounds via the moment-generating function ..... 64
2.4.2 Sub-Gaussian and sub-exponential random variables ..... 69
2.4.3 $\triangleright$ Probabilistic analysis of algorithms: knapsack problem ..... 79
2.4.4 Epsilon-nets and chaining ..... 85
2.4.5 $\triangleright$ Data science: Johnson-Lindenstrauss lemma and appli- cation to compressed sensing ..... 93
2.4.6 $\triangleright$ Data science: classification, empirical risk minimization and VC dimension ..... 102
3 Martingales and potentials ..... 122
3.1 Background ..... 122
3.1.1 Stopping times ..... 123
3.1.2 $\triangleright$ Markov chains: exponential tail of hitting times and some cover time bounds ..... 130
3.1.3 Martingales ..... 133
3.1.4 $\triangleright$ Percolation: critical regime on infinite $d$-regular tree ..... 144
3.2 Concentration for martingales and applications ..... 145
3.2.1 Azuma-Hoeffding inequality ..... 145
3.2.2 Method of bounded differences ..... 147
3.2.3 $\triangleright$ Random graphs: exposure martingale and application to the chromatic number in Erdös-Rényi model ..... 158
3.2.4 $\triangleright$ Random graphs: degree sequence of preferential attach- ment graphs ..... 163
3.2.5 $\triangleright$ Data science: stochastic bandits and the slicing method ..... 170
3.2.6 Coda: Talagrand's inequality ..... 178
3.3 Potential theory and electrical networks ..... 180
3.3.1 Martingales, the Dirichlet problem and Lyapounov functions ..... 182
3.3.2 Basic electrical network theory ..... 192
3.3.3 Bounding the effective resistance via variational principles ..... 202
3.3.4 $\triangleright$ Random walks: Pólya's theorem, two ways ..... 215
3.3.5 $\triangleright$ Randomized algorithms: Wilson's method for generating uniform spanning trees ..... 218
4 Coupling ..... 234
4.1 Background ..... 234
4.1.1 Basic definitions ..... 235
4.1.2 $\triangleright$ Random walks: harmonic functions on lattices and infi- nite d-regular trees ..... 237
4.1.3 Total variation distance and coupling inequality ..... 240
4.1.4 $\triangleright$ Random graphs: degree sequence in Erdős-Rényi model ..... 246
4.2 Stochastic domination ..... 249
4.2.1 Definitions ..... 250
4.2.2 $\triangleright$ Ising model: boundary conditions ..... 260
4.2.3 Correlation inequalities: FKG and Holley's inequalities ..... 263
4.2.4 $\triangleright$ Random graphs: Janson's inequality and application to the clique number in the Erdös-Rényi model ..... 269
4.2.5 $\triangleright$ Percolation: RSW theory and a proof of Harris' theorem ..... 272
4.3 Coupling of Markov chains and application to mixing ..... 281
4.3.1 Bounding the mixing time via coupling ..... 281
4.3.2 $\triangleright$ Random walks: mixing on cycles, hypercubes, and trees ..... 284
4.3.3 Path coupling ..... 293
4.3.4 $\triangleright$ Ising model: Glauber dynamics at high temperature ..... 297
4.4 Chen-Stein method ..... 300
4.4.1 Main bounds and examples ..... 301
4.4.2 Some motivation and proof ..... 308
4.4.3 $\triangleright$ Random graphs: clique number at the threshold in the Erdös-Rényi model ..... 318
5 Spectral methods ..... 327
5.1 Background ..... 327
5.1.1 Eigenvalues and their variational characterization ..... 328
5.1.2 Elements of spectral graph theory ..... 332
5.1.3 Perturbation results ..... 341
5.1.4 $\triangleright$ Data science: community recovery ..... 345
5.2 Spectral techniques for reversible Markov chains ..... 351
5.2.1 Spectral gap ..... 353
5.2.2 $\triangleright$ Random walks: a spectral look at cycles and hypercubes ..... 359
5.2.3 $\triangleright$ Markov chains: Varopoulos-Carne and diameter-based bounds on the mixing time ..... 365
5.2.4 $\triangleright$ Randomized algorithms: Markov chain Monte Carlo and a quantitative ergodic theorem ..... 370
5.2.5 Spectral radius ..... 375
5.3 Geometric bounds ..... 380
5.3.1 Cheeger's inequality ..... 385
5.3.2 $\triangleright$ Random walks: trees, cycles, and hypercubes revisited ..... 391
5.3.3 $\triangleright$ Random graphs: existence of an expander family and application to mixing ..... 394
5.3.4 $\triangleright$ Ising model: Glauber dynamics on complete graphs and expanders ..... 400
5.3.5 Congestion ratio ..... 405
6 Branching processes ..... 415
6.1 Background ..... 415
6.1.1 Basic definitions ..... 416
6.1.2 Extinction ..... 417
6.1.3 $\triangleright$ Percolation: Galton-Watson trees ..... 422
6.1.4 Multitype branching processes ..... 423
6.2 Random-walk representation ..... 429
6.2.1 Exploration process ..... 429
6.2.2 Duality principle ..... 432
6.2.3 Hitting-time theorem ..... 433
6.2.4 $\triangleright$ Percolation: critical exponents on the infinite b-ary tree ..... 438
6.3 Applications ..... 442
6.3.1 $\triangleright$ Probabilistic analysis of algorithms: binary search tree ..... 442
6.3.2 $\triangleright$ Data science: the reconstruction problem, the Kesten- Stigum bound and a phase transition in phylogenetics ..... 451
$6.4 \triangleright$ Finale: the phase transition of the Erdös-Rényi model ..... 466
6.4.1 Statement and proof sketch ..... 466
6.4.2 Bounding cluster size: domination by branching processes ..... 468
6.4.3 Concentration of cluster size: second moment bounds ..... 480
6.4.4 Critical case via martingales ..... 486
6.4.5 $\triangleright$ Encore: random walk on the Erdös-Rényi graph ..... 490
A Useful combinatorial formulas ..... 502
B Measure-theoretic foundations ..... 503
B. 1 Probability spaces ..... 503
B. 2 Random variables ..... 506
B. 3 Independence ..... 510
B. 4 Expectation ..... 513
B. 5 Fubini's theorem ..... 519
B. 6 Conditional expectation ..... 522
B. 7 Filtered spaces ..... 528

## Preface

This book arose from a set of lecture notes prepared for a one-semester topics course I taught at the University of Wisconsin-Madison in 2014, 2017, 2020 and 2023 which attracted a wide spectrum of students in mathematics, computer sciences, engineering, and statistics.

## What is it about?

The purpose of the book is to provide a graduate-level introduction to discrete probability. Topics covered are drawn primarily from stochastic processes on graphs: percolation, random graphs, Markov random fields, random walks on graphs, etc. No attempt is made at covering these broad areas in depth. Rather, the emphasis is on illustrating important techniques used to analyze such processes. Along the way, many standard results regarding discrete probability models are worked out.

The "modern" in the title refers to the (non-exclusive) focus on nonasymptotic methods and results, reflecting the impact of the theoretical computer science literature on the trajectory of this field. In particular several applications in randomized algorithms, probabilistic analysis of algorithms and theoretical machine learning are used throughout to motivate the techniques described (although, again, these areas are not covered exhaustively).

Of course the selection of topics is somewhat arbitrary and driven in part by personal interests. But the choice was guided by a desire to introduce techniques that are widely used across discrete probability and its applications. The material discussed here is developed in much greater depth in the following (incomplete list of) excellent textbooks and expository monographs, many of which influenced various sections of this book:

- Agarwal, Jiang, Kakade, Sun. Reinforcement learning: Theory and algorithms. [AJKS22]
- Aldous, Fill. Reversible Markov chains and random walks on graphs. [AF]
- Alon, Spencer. The Probabilistic Method. [AS11]
- Béla Bollobás. Random graphs. [Bol01]
- Boucheron, Lugosi, Massart. Concentration Inequalities: A Nonasymptotic Theory of Independence. [BLM13]
- Chung, Lu. Complex graphs and networks. [CL06]
- Durrett. Random Graph Dynamics. [Dur06]
- Frieze and Karoński. Introduction to random graphs. [FK16]
- Grimmett. Percolation. [Gri10b]
- Janson, Luczak, Rucinski. Random Graphs. [JLR11]
- Lattimore, Szepesvári. Bandit Algorithms. [LS20]
- Levin, Peres, Wilmer. Markov chains and mixing times. [LPW06]
- Lyons, Peres. Probability on trees and networks. [LP16]
- Mitzenmacher, Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. [MU05]
- Motwani, Raghavan. Randomized algorithms. [MR95]
- Rassoul-Agha, Seppäläinen. A course on large deviations with an introduction to Gibbs measures. [RAS15]
- S. Shalev-Shwartz and S. Ben-David. Understanding Machine Learning: From Theory to Algorithms. [SSBD14]
- van Handel. Probability in high dimension. [vH16]
- van der Hofstad. Random graphs and complex networks. Vol. 1. [vdH17]
- Vershynin. High-Dimensional Probability: An Introduction with Applications in Data Science. [Ver18]

In fact the book is meant as a first foray into the basic results and/or toolkits detailed in these more specialized references. My hope is that, by the end, the reader will have picked up sufficient fundamental background to learn advanced material on their own with some ease. I should add that I used many additional helpful sources; they are acknowledged in the "Bibliographic remarks" at the end of each chapter. It is impossible to cover everything. Some notable omissions include, e.g., graph limits [Lov12], influence [KS05], and group-theoretic methods [Dia88], among others.

Much of the material covered here (and more) can also be found in [HMRAR98], [Gri10a], and [Bre17] with a different emphasis and scope.

## Prerequisites

It is assumed throughout that the reader is fluent in undergraduate linear algebra, for example, at the level of [Axl15], and basic real analysis, for example, at the level of [Mor05].

In addition, it is recommended that the reader has taken at least one semester of graduate probability at the level of [Dur10]. I am also particularly fond of [Wil91], which heavily influenced the appendix where measure-theoretic background is reviewed. Some familiarity with countable Markov chain theory is necessary, as covered for instance in [Dur10, Chapter 6]. An advanced undergraduate or Masters' level treatment such as [Dur12], [Nor98], [GS20], [Law06] or [Bre20] will suffice however.

## Organization

The book is organized around five major "tools." The reader will have likely encountered those tools in prior probability courses. The goal here is to develop them further, specifically with their application to discrete random structures in mind, and to illustrate them in this setting on a variety of major, classical results and applications.

In the interest of keeping the book relatively self-contained and serving the widest spectrum of readers, each chapter begins with a "background" section which reviews the basic material on which the rest of the chapter builds. The remaining sections then proceed to expand on two or three important specializations of the tools. While the chapters are meant to be somewhat modular, results from previous chapters do occasionally make an appearance.

The techniques are illustrated throughout with simple examples first, and then with more substantial ones in separate sections marked with the symbol $\triangleright$. I have attempted to provide applications from many areas of discrete probability and theoretical computer science, although some techniques are better suited for certain types of models or questions. The examples and applications are important: many of the tools are quite straightforward (or even elementary), and it is only when seen in action that their full power can be appreciated. Moreover, the $\triangleright$ sections serve as an excuse to introduce the reader to classical results and important applicationsbeyond their reliance on specific tools.

Chapter 1 introduces some of the main models from probability on graphs that we come back to repeatedly throughout the book. It begins with a brief review of graph theory and Markov chain theory.

Chapter 2 starts out with the probabilistic method, including the first moment principle and second moment method, and then it moves on to concentration inequalities for sums of independent random variables, mostly sub-Gaussian and sub-exponential variables. It also discusses techniques to analyze the suprema of random processes.

Chapter 3 turns to martingales. The first main topic there is the Azuma-

Hoeffding inequality and the method of bounded differences with applications to random graphs and stochastic bandit problems. The second main topic is electrical network theory for random walks on graphs.

Chapter 4 introduces coupling. It covers stochastic domination and correlation inequalities as well as couplings of Markov chains with applications to mixing. It also discusses the Chen-Stein method for Poisson approximation.

Chapter 5 is concerned with spectral methods. A major topic there is the use of the spectral theorem and geometric bounds on the spectral gap to control the mixing time of a reversible Markov chain. The chapter also introduces spectral methods for community recovery in network analysis.

Chapter 6 ends the book with applications of branching processes. Among other applications, an introduction to the reconstruction problem on trees is provided. The final section gives a detailed analysis of the phase transition of the Erdös-Rényi graph, where techniques from all chapters of the book are brought to bear.

## Acknowledgments

The lecture notes on which this book is based were influenced by graduate courses of David Aldous, Steve Evans, Elchanan Mossel, Yuval Peres, Alistair Sinclair, and Luca Trevisan at UC Berkeley, where I learned much of this material. In particular scribe notes for some of these courses helped shape early iterations of this book.

I have also learned a lot over the years from my collaborators and mentors as well as my former and current Ph.D. students and postdocs. I am particularly grateful to Elchanan Mossel and Allan Sly for encouragements to finish this project and to the UW-Madison students who have taken the various iterations of the course that inspired the book for their invaluable feedback.

Warm thanks to everyone in the departments of mathematics at UCLA and UWMadison who have provided the stimulating environments that made this project possible. Beyond my current department, I am particularly indebted to my colleagues in the NSF-funded Institute for Foundations of Data Science (IFDS) who have significantly expanded my knowledge of applications of this material in machine learning and statistics.

Finally I thank my parents, my wife and my son for their love, patience and support.

## Notation

Throughout the book, we will use the following notation.

- The real numbers are denoted by $\mathbb{R}$, the nonnegative reals are denoted by $\mathbb{R}_{+}$, the integers are denoted by $\mathbb{Z}$, the nonnegative integers are denoted by $\mathbb{Z}_{+}$, the natural numbers (i.e., positive integers) are denoted by $\mathbb{N}$ and the rational numbers are denoted by $\mathbb{Q}$. We will also use the notation $\overline{\mathbb{Z}}_{+}:=$ $\{0,1, \ldots,+\infty\}$.
- For two reals $a, b \in \mathbb{R}$,

$$
a \wedge b:=\min \{a, b\}, \quad a \vee b:=\max \{a, b\},
$$

and

$$
a^{+}=0 \vee a, \quad a^{-}=0 \vee(-a) .
$$

- For a real $a,\lfloor a\rfloor$ is the largest integer that is smaller than or equal to $a$ and $\lceil a\rceil$ is the smallest integer that is larger than or equal to $a$.
- For $x \in \mathbb{R}$, the natural (i.e., base $e$ ) logarithm of $x$ is denoted by $\log x$. We also let $\exp (x)=e^{x}$.
- For a positive integer $n \in \mathbb{N}$, we let

$$
[n]:=\{1, \ldots, n\} .
$$

- Let $A$ be a set. The cardinality of $A$ is denoted by $|A|$. The power set of $A$, i.e., the collection of all of its subsets, is denoted by $2^{A}$.
- For two sets $A, B$, their cartesian product is denoted by $A \times B$.
- We will use the following notation for standard vectors: $\mathbf{0}$ is the all-zero vector, $\mathbf{1}$ is the all-one vector, and $\mathbf{e}_{i}$ is the standard basis vector with a one in coordinate $i$ and a zero elsewhere. In each case, the dimension is implicit, as well as whether it is a row or column vector.
- For a vector $\mathbf{u}=\left(u_{1}, \ldots, u_{n}\right) \in \mathbb{R}^{n}$ and real $p>0$, its $p$-norm (or $\ell^{p}$-norm) is

$$
\|\mathbf{u}\|_{p}:=\left(\sum_{i=1}^{n}\left|u_{i}\right|^{p}\right)^{1 / p}
$$

When $p=+\infty$, we have

$$
\|\mathbf{u}\|_{\infty}:=\max _{i}\left|u_{i}\right| .
$$

We also use the notation $\|\mathbf{u}\|_{0}$ to denote the number of nonzero coordinates of $\mathbf{u}$ (although it is not a norm; see Exercise 1.1). For two vectors $\mathbf{u}=$ $\left(u_{1}, \ldots, u_{n}\right), \mathbf{v}=\left(v_{1}, \ldots, v_{n}\right) \in \mathbb{R}^{n}$, their inner product is

$$
\langle\mathbf{u}, \mathbf{v}\rangle:=\sum_{i=1}^{n} u_{i} v_{i} .
$$

The same notations apply to row vectors.

- For a matrix $A$, we denote the entries of $A$ by $A(i, j), A_{i, j}$, or $A_{i j}$. The $i$-th row of $A$ is denoted by $A(i, \cdot)$ or $A_{i, .}$ The $j$-th column of $A$ is denoted by $A(\cdot, j)$ or $A_{\cdot, j}$. The transpose of $A$ is $A^{T}$.
- For a vector $\mathbf{z}=\left(z_{1}, \ldots, z_{d}\right)$, we let $\operatorname{diag}(\mathbf{z})$ be the diagonal matrix with diagonal entries $z_{1}, \ldots, z_{d}$.
- The binomial coefficients are defined as

$$
\binom{n}{k}=\frac{n!}{k!(n-k)!}
$$

binomial
coefficients
where $k, n \in \mathbb{N}$ with $k \leq n$ and $n!=1 \times 2 \times \cdots \times n$ is the factorial of $n$. Some standard approximations for $\binom{n}{k}$ and $n!$ are listed in Appendix A. See also Exercises 1.2, 1.3, and 1.4.

- We use the abbreviation a.s. for "almost surely," that is, with probability 1. We use "w.p." for "with probability."
- Convergence in probability is denoted as $\rightarrow_{\mathrm{p}}$. Convergence in distribution is denoted as $\xrightarrow{\text { d }}$.
- For a random variable $X$ and a probability distribution $\mu$, we write $X \sim \mu$ to indicate that $X$ has distribution $\mu$. We write $X \stackrel{\text { d }}{=} Y$ if the random variables $X$ and $Y$ have the same distribution.
- For an event $A$, the random variable $\mathbf{1}_{A}$ is the indicator of $A$, that is, it is 1 if $A$ occurs and 0 otherwise. We also use $1\{A\}$.
- For probability measures $\mu, \nu$ on a countable set $S$, their total variation distance is

$$
\|\mu-\nu\|_{\mathrm{TV}}:=\sup _{A \subseteq S}|\mu(A)-\nu(A)| .
$$

- For nonnegative functions $f(n), g(n)$ of $n \in \mathbb{Z}_{+}$we write $f(n)=O(g(n))$ if there exists a positive constant $C>0$ such that $f(n) \leq C g(n)$ for all $n$ large enough. Similarly, $f(n)=\Omega(g(n))$ means that $f(n) \geq c g(n)$ for some constant $c>0$ for all $n$ large enough. The notation $f(n)=\Theta(g(n))$ indicates that both $f(n)=O(g(n))$ and $f(n)=\Omega(g(n))$ hold. We also write $f(n)=o(g(n))$ or $g(n)=\omega(f(n))$ or $f(n) \ll g(n)$ or $g(n) \gg f(n)$ if $f(n) / g(n) \rightarrow 0$ as $n \rightarrow+\infty$. If $f(n) / g(n) \rightarrow 1$ we write $f(n) \sim g(n)$. The same notations are used for functions of a real variable $x$ as $x \rightarrow+\infty$.


## Chapter 1

## Introduction

In this chapter we describe a few discrete probability models to which we will come back repeatedly throughout the book. While there exists a vast array of well-studied random combinatorial structures (permutations, partitions, urn models, Boolean functions, polytopes, etc.), our focus is primarily on a limited number of graph-based processes, namely percolation, random graphs, the Ising model, and random walks on networks. We will not attempt to derive the theory of these models exhaustively here. Instead we will employ them to illustrate some essential techniques from discrete probability. Note that the toolkit developed in this book is meant to apply to other probabilistic models of interest as well, and in fact many more will be encountered along the way. After a brief review of graph basics and Markov chains theory in Section 1.1, we formally introduce our main models in Section 1.2. We also formulate various questions about these models that will be answered (at least partially) later on. We assume that the reader is familiar with the measure-theoretic foundations of probability. A refresher of all required concepts and results is provided in Appendix B.

### 1.1 Background

We start with a brief review of graph terminology and standard countable-space Markov chains results.


Figure 1.1: Petersen graph.

### 1.1. 1 Review of graph theory

Basic definitions An undirected graph (or graph for short) is a pair $G=(V, E)$ where $V$ is the set of vertices (or nodes or sites) and

$$
E \subseteq\{\{u, v\}: u, v \in V\},
$$

is the set of edges (or bonds). See Figure 1.1 for an example. We occasionally write $V(G)$ and $E(G)$ for the vertices and edges of the graph $G$. The set of vertices $V$ is either finite or countably infinite. Edges of the form $\{u\}$ are called self-loops. In general, we do not allow $E$ to be a multiset unless otherwise stated. But, when $E$ is a multiset, $G$ is called a multigraph.

A vertex $v \in V$ is incident with an edge $e \in E$ (or vice versa) if $v \in e$. The incident vertices of an edge are called endvertices. Two vertices $u, v \in V$ are adjacent (or neighbors), denoted by $u \sim v$, if $\{u, v\} \in E$. The set of adjacent vertices of $v$, denoted by $N(v)$, is called the neighborhood of $v$ and its size, that is, $\delta(v):=|N(v)|$, is the degree of $v$. A vertex $v$ with $\delta(v)=0$ is called isolated. A graph is called $d$-regular if all its degrees are $d$. A countable graph is locally finite if all its vertices have a finite degree.

Example 1.1.1 (Petersen graph). All vertices in the Petersen graph in Figure 1.1 have degree 3 , that is, it is a 3 -regular graph. In particular it has no isolated vertex.

A convenient (and mathematically useful) way to specify a graph is the following matrix representation. Assume the graph $G=(V, E)$ has $n=|V|$ vertices. Assume that the vertices are numbered $1, \ldots, n$. The adjacency matrix $A$ of $G$ is
the $n \times n$ symmetric matrix defined as

$$
A_{x, y}= \begin{cases}1 & \text { if }\{x, y\} \in E, \\ 0 & \text { otherwise } .\end{cases}
$$

Example 1.1.2 (Triangle). The adjacency matrix of a triangle, that is, a 3-vertex graph with all possible non-loop edges, is

$$
A=\left[\begin{array}{lll}
0 & 1 & 1 \\
1 & 0 & 1 \\
1 & 1 & 0
\end{array}\right]
$$

There exist other matrix representations. Here is one. Let $m=|E|$ and assume that the edges are labeled arbitrarily as $e_{1}, \ldots, e_{m}$. The incidence matrix of an undirected graph $G=(V, E)$ is the $n \times m$ matrix $B$ such that $B_{i, j}=1$ if vertex $i$ and edge $e_{j}$ are incident and 0 otherwise.

Subgraphs, paths, and cycles A subgraph of $G=(V, E)$ is a graph $G^{\prime}=$ ( $V^{\prime}, E^{\prime}$ ) with $V^{\prime} \subseteq V$ and $E^{\prime} \subseteq E$. Implied in this definition is the fact that the edges in $E^{\prime}$ are incident with $V^{\prime}$ only. The subgraph $G^{\prime}$ is said to be induced if

$$
E^{\prime}=\left\{\{x, y\}: x, y \in V^{\prime},\{x, y\} \in E\right\}
$$

that is, it contains exactly those edges of $G$ that are between vertices in $V^{\prime}$. In that case the notation $G^{\prime}:=G\left[V^{\prime}\right]$ is used. A subgraph is said to be spanning if $V^{\prime}=V$. A subgraph containing all possible non-loop edges between its vertices is called a clique (or complete subgraph). A clique with $k$ nodes is referred to as a $k$-clique.

Example 1.1.3 (Petersen graph (continued)). The Petersen graph contains no triangle, that is, 3 -clique, induced or not.

A walk in $G$ is a sequence of (not necessarily distinct) vertices $x_{0} \sim x_{1} \sim$ $\cdots \sim x_{k}$. Note the requirement that consecutive vertices of a walk are adjacent. The number $k \geq 0$ is the length of the walk. If the endvertices $x_{0}, x_{k}$ coincide, that is, $x_{0}=x_{k}$, we refer to the walk as closed. If the vertices of a walk are all distinct, we call it a path (or self-avoiding walk). If the vertices of a closed walk are all distinct except for the endvertices and its length is at least 3 , we call it a cycle. A path or cycle can be seen as a (not necessarily induced) subgraph of $G$. The length of the shortest path connecting two distinct vertices $u, v$ is the graph
distance between $u$ and $v$, denoted by $d_{G}(u, v)$. It can be checked that the graph distance is a metric (and that, in particular, it satisfies the triangle inequality; see Exercise 1.6). The minimum length of a cycle in a graph is its girth.

We write $u \leftrightarrow v$ if there is a path between $u$ and $v$. It can be checked that the binary relation $\leftrightarrow$ is an equivalence relation (i.e, it is reflexive, symmetric and transitive; see Exercise 1.6). Its equivalence classes are called connected components. A graph is connected if any two vertices are linked by a path, that is, if $u \leftrightarrow v$ for all $u, v \in V$. Or put differently, if there is only one connected component.

Example 1.1.4 (Petersen graph (continued)). The Petersen graph is connected.

Trees A forest is a graph with no cycle, or acyclic graph. A tree is a connected forest. Vertices of degree 1 are called leaves. A spanning tree of $G$ is a subgraph which is a tree and is also spanning. A tree is said to be rooted if it has a single distinguished vertex called the root.

Trees will play a key role and we collect several important facts about them (mostly without proofs). The following characterizations of trees will be useful. The proof is left as an exercise (see Exercise 1.8). We write $G+e$ (respectively $G-e$ ) to indicate the graph $G$ with edge $e$ added (respectively removed).

Theorem 1.1.5 (Trees: characterizations). The following are equivalent.
(i) The graph $T$ is a tree.
(ii) For any two vertices in $T$, there is a unique path between them.
(iii) The graph $T$ is connected, but $T-e$ is not for any edge e in $T$.
(iv) The graph $T$ is acyclic, but $T+\{x, y\}$ is not for any pair of non-adjacent vertices $x, y$.

Here are two important implications.
Corollary 1.1.6. If $G$ is connected, then it has at least one spanning tree.
Proof. Indeed, from Theorem 1.1.5, a graph is a tree if and only if it is minimally connected, in the sense that removing any of its edges disconnects it. So a spanning tree can be obtained by removing edges of $G$ that do not disconnect it until it is not possible anymore.

The following characterization is proved in Exercise 1.7.
Corollary 1.1.7. A connected graph with $n$ vertices is a tree if and only if it has $n-1$ edges.

And here is a related fact.
Corollary 1.1.8. Let $G$ be a graph with $n$ vertices. If an acyclic subgraph $H$ has $n$ vertices and $n-1$ edges, then it is a spanning tree of $G$.

Proof. If $H$ is not connected, then it has at least two connected components. Each of them is acyclic and therefore a tree. By applying Corollary 1.1.7 to the connected components and summing up, we see that the total number of edges in $H$ is $\leq n-2$, a contradiction. So $H$ is connected and therefore a spanning tree.

Finally, a classical formula:
Theorem 1.1.9 (Cayley's formula). There are $k^{k-2}$ trees on a set of $k$ labeled vertices.

We give a proof of Cayley's formula based on branching processes in Exercise 6.19.
Some standard graphs Here are a few more examples of finite graphs.

- Complete graph $K_{n}$ : This graph is made of $n$ vertices with all non-loop edges.
- Cycle graph $C_{n}$ (or $n$-cycle): The vertex set is $\{0,1, \ldots, n-1\}$ and two vertices $i \neq j$ are adjacent if and only if $|i-j|=1$ or $n-1$.
- Torus $\mathbb{L}_{n}^{d}$ : The vertex set is $\{0,1, \ldots, n-1\}^{d}$ and two vertices $\mathbf{x} \neq \mathbf{y}$ are adjacent if and only if there is a coordinate $i$ such that $\left|x_{i}-y_{i}\right|=1$ or $n-1$ and all other coordinates $j \neq i$ satisfy $x_{j}=y_{j}$.
- Hypercube $\mathbb{Z}_{2}^{n}$ (or $n$-dimensional hypercube): The vertex set is $\{0,1\}^{n}$ and two vertices $\mathbf{x} \neq \mathbf{y}$ are adjacent if and only if $\|\mathbf{x}-\mathbf{y}\|_{1}=1$.
- Rooted $b$-ary tree $\widehat{\mathbb{T}}_{b}^{\ell}$ : This graph is a tree with $\ell$ levels. The unique vertex on level 0 is called the root. For $j=1, \ldots, \ell-1$, level $j$ has $b^{j}$ vertices, each of which has exactly one neighbor on level $j-1$ (its parent) and $b$ neighbors on level $j+1$ (its children). The $b^{\ell}$ vertices on level $\ell$ are leaves.

Here are a few examples of infinite graphs, that is, a graph with a countably infinite number of vertices and edges.

Cayley's formula
complete graph
cycle graph

- Infinite d-regular tree $\mathbb{T}_{d}$ : This is an infinite tree where each vertex has exactly $d$ neighbors. The rooted version, that is, $\widehat{\mathbb{T}}_{b}^{\ell}$ with $\ell=+\infty$ levels, is denoted by $\widehat{\mathbb{T}}_{b}$.
- Lattice $\mathbb{L}^{d}$ : The vertex set is $\mathbb{Z}^{d}$ and two vertices $\boldsymbol{x} \neq \boldsymbol{y}$ are adjacent if and only if $\|\boldsymbol{x}-\boldsymbol{y}\|_{1}=1$.

A bipartite graph $G=(L \cup R, E)$ is a graph whose vertex set is composed of the union of two disjoint sets $L, R$ and whose edge set $E$ is a subset of $\{\{\ell, r\}$ : $\ell \in L, r \in R\}$. That is, there is no edge between vertices in $L$, and likewise for $R$.

Example 1.1.10 (Some bipartite graphs). The cycle graph $C_{2 n}$ is a bipartite graph. So is the complete bipartite graph $K_{n, m}$ with vertex set $\left\{\ell_{1}, \ldots, \ell_{n}\right\} \cup\left\{r_{1}, \ldots, r_{m}\right\}$ and edge set $\left\{\left\{\ell_{i}, r_{j}\right\}: i \in[n], j \in[m]\right\}$.

In a bipartite graph $G=(L \cup R, E)$, a perfect matching is a collection of edges $M \subseteq E$ such that each vertex is incident with exactly one edge in $M$.

An automorphism of a graph $G=(V, E)$ is a bijection $\phi$ of $V$ to itself that preserves the edges, that is, such that $\{x, y\} \in E$ if and only if $\{\phi(x), \phi(y)\} \in E$. A graph $G=(V, E)$ is vertex-transitive if for any $u, v \in V$ there is an automorphism mapping $u$ to $v$.

Example 1.1.11 (Petersen graph (continued)). For any $\ell \in \mathbb{Z}$, a ( $2 \pi \ell / 5$ )-rotation of the planar representation of the Petersen graph in Figure 1.1 corresponds to an automorphism.
Example 1.1.12 (Trees). The graph $\mathbb{T}_{d}$ is vertex-transitive. The graph $\widehat{\mathbb{T}}_{b}$ on the other hand has many automorphisms, but is not vertex-transitive.

Flows Let $G=(V, E)$ be a connected graph with two distinguished disjoint vertex sets, a source-set (or source for short) $A \subseteq V$ and a sink-set (or sink for short) $Z$. Let $\kappa: E \rightarrow \mathbb{R}_{+}$be a capacity function.

Definition 1.1.13 (Flow). A flow from source $A$ to $\operatorname{sink} Z$ is a function $f: V \times V \rightarrow{ }_{\text {flow }}$ $\mathbb{R}$ such that:
$F 1$ (Antisymmetry) $f(x, y)=-f(y, x), \forall x, y \in V$.
$F 2$ (Capacity constraint) $|f(x, y)| \leq \kappa(e), \forall e=\{x, y\} \in E$, and $f(x, y)=0$ otherwise.

F3 (Flow-conservation constraint)

$$
\sum_{y: y \sim x} f(x, y)=0, \quad \forall x \in V \backslash(A \cup Z) .
$$

For $U, W \subseteq V$, let $f(U, W):=\sum_{u \in U, w \in W} f(u, w)$. The strength of $f$ is $\|f\|:=$ $f\left(A, A^{c}\right)$.

One useful consequence of antisymmetry is that, for any $U \subseteq V$, we have $f(U, U)=0$ since each distinct pair $x \neq y \in U$ appears exactly twice in the sum, once in each ordering. Also if $W_{1}$ and $W_{2}$ are disjoint, then $f\left(U, W_{1} \cup W_{2}\right)=$ $f\left(U, W_{1}\right)+f\left(U, W_{2}\right)$. In particular, combining both observations, $f(U, W)=$ $f(U, W \backslash U)=-f(W \backslash U, U)$.

For $F \subseteq E$, let $\kappa(F):=\sum_{e \in F} \kappa(e)$. We call $F$ a cutset separating $A$ and $Z$ (or cutset for short) if all paths connecting $A$ and $Z$ include an edge in $F$. For such an $F$, let $A_{F}$ be the set of vertices not separated from $A$ by $F$, that is, vertices from which there is a path to $A$ not crossing an edge in $F$. Clearly $A \subseteq A_{F}$ but $A_{F} \cap Z=\emptyset$.

Lemma 1.1.14 (Max flow $\leq \min$ cut). For any flow $f$ and cutset $F$,

$$
\begin{equation*}
\|f\|=f\left(A_{F}, A_{F}^{c}\right) \leq \sum_{\{x, y\} \in F}|f(x, y)| \leq \kappa(F) . \tag{1.1.1}
\end{equation*}
$$

Proof. Since $F$ is a cutset, $\left(A_{F} \backslash A\right) \cap(A \cup Z)=\emptyset$. So, by (F3),

$$
\begin{aligned}
f\left(A, A^{c}\right)= & f\left(A, A^{c}\right)+\sum_{u \in A_{F} \backslash A} f(u, V) \\
= & f\left(A, A_{F} \backslash A\right)+f\left(A, A_{F}^{c}\right) \\
& +f\left(A_{F} \backslash A, A_{F}\right)+f\left(A_{F} \backslash A, A_{F}^{c}\right) \\
= & f\left(A, A_{F} \backslash A\right)+f\left(A, A_{F}^{c}\right) \\
& +f\left(A_{F} \backslash A, A\right)+f\left(A_{F} \backslash A, A_{F}^{c}\right) \\
= & f\left(A_{F}, A_{F}^{c}\right) \\
\leq & \sum_{\{x, y\} \in F}|f(x, y)|,
\end{aligned}
$$

where we used (F1) twice. The last line is justified by the fact that the edges between a vertex in $A_{F}$ and a vertex in $A_{F}^{c}$ have to be in $F$ by definition of $A_{F}$. That proves the first inequality in the claim. Condition (F2) implies the second one.

Remarkably, this bound is tight, in the following sense.
Theorem 1.1.15 (Max-flow min-cut theorem). Let $G$ be a finite connected graph with source $A$ and sink $Z$, and let $\kappa$ be a capacity function. Then the following holds

$$
\sup \{\|f\|: \text { flow } f\}=\min \{\kappa(F): \text { cutset } F\} .
$$

Proof. Note that, by compactness, the supremum on the left-hand side is achieved. Let $f$ be an optimal flow. The idea of the proof is to construct a "matching" cutset.

An augmentable path is a path $x_{0} \sim \cdots \sim x_{k}$ with $x_{0} \in A, x_{i} \notin A \cup Z$ for all $i \neq 0$ or $k$, and $f\left(x_{i-1}, x_{i}\right)<\kappa\left(\left\{x_{i-1}, x_{i}\right\}\right)$ for all $i \neq 0$. By default, each vertex in $A$ is an augmentable path. Moreover, by the optimality of $f$ there cannot be an augmentable path with $x_{k} \in Z$. Indeed, otherwise, we could "push more flow through that path" and increase the strength of $f$-a contradiction.

Let $B \subseteq V$ be the set of all final vertices in some augmentable path and let $F$ be the edge set between $B$ and $B^{c}:=V \backslash B$. Note that, again by contradiction, all vertices in $B$ can be reached from $A$ without crossing $F$ and that $f(x, y)=\kappa(e)$ for all $e=\{x, y\} \in F$ with $x \in B$ and $y \in B^{c}$. Furthermore $F$ is a cutset separating $A$ from $Z$ : trivially $A \subseteq B ; Z \subseteq B^{c}$ as argued above; and any path from $A$ to $Z$ must exit $B$ and enter $B^{c}$ through an edge in $F$. Thus $A_{F}=B$ and we have equality in (1.1.1). That concludes the proof.

Colorings, independent sets, and matchings A coloring of a graph $G=(V, E)$ is an assignment of colors to each vertex in $G$. In a coloring, two vertices may share the same color. A coloring is proper if for every edge $e$ in $G$ the endvertices of $e$ have distinct colors. The smallest number of colors in a proper coloring of a graph $G$ is called the chromatic number $\chi(G)$ of $G$.

An independent vertex set (or independent set for short) of $G=(V, E)$ is a subset of vertices $W \subseteq V$ such that all pairs of vertices in $W$ are non-adjacent. Likewise, two edges are independent if they are not incident with the same vertex. A matching is a set of pairwise independent edges. A matching $F$ is perfect if every vertex in $G$ is incident with an edge of $F$.

Edge-weighted graphs We refer to an edge-weighted graph $G=(V, E, w)$ as a network. Here $w: E \rightarrow \mathbb{R}_{+}$is a function that assigns positive real weights to the edges. Definitions can be generalized naturally. In particular, one defines the degree of a vertex $i$ as

$$
\delta(i)=\sum_{j: e=\{i, j\} \in E} w_{e} .
$$

The adjacency matrix $A$ of $G$ is the $n \times n$ symmetric matrix defined as

$$
A_{i, j}= \begin{cases}w_{e} & \text { if } e=\{i, j\} \in E, \\ 0 & \text { otherwise }\end{cases}
$$

where we denote the vertices $\{1, \ldots, n\}$.

Directed graphs A directed graph (or digraph for short) is a pair $G=(V, E)$ where $V$ is a set of vertices (or nodes or sites) and $E \subseteq V^{2}$ is a set of directed edges (or $\operatorname{arcs}$ ). A directed edge from $x$ to $y$ is typically denoted by $(x, y)$, or occasionally by $\langle x, y\rangle$. A directed path is a sequence of vertices $x_{0}, \ldots, x_{k}$, all distinct, with $\left(x_{i-1}, x_{i}\right) \in E$ for all $i=1, \ldots, k$. We write $u \rightarrow v$ if there is such a directed path with $x_{0}=u$ and $x_{k}=v$. We say that $u, v \in V$ communicate, denoted by $u \leftrightarrow v$, if $u \rightarrow v$ and $v \rightarrow u$. In particular, we always have $u \leftrightarrow u$ for every state $u$. The binary relation $\leftrightarrow$ relation is an equivalence relation (see Exercise 1.6). The equivalence classes of $\leftrightarrow$ are called the strongly connected components of $G$.

The following definition will prove useful.
Definition 1.1.16 (Oriented incidence matrix). Let $G=(V, E)$ be an undirected graph. Assume that the vertices of $G=(V, E)$ are numbered $1, \ldots,|V|$ and that the edges are labeled arbitrarily as $e_{1}, \ldots, e_{|E|}$. An orientation of $G$ is the choice of a direction $\vec{e}_{i}$ for each edge e $e_{i}$, turning it into a digraph $\vec{G}$. An oriented incidence matrix of $G$ is the incidence matrix of an orientation, that is, the matrix $B$ such that $B_{i j}=-1$ if edge $\vec{e}_{j}$ leaves vertex $i, B_{i j}=1$ if edge $\vec{e}_{j}$ enters vertex $i$, and 0 otherwise.

### 1.1.2 Review of Markov chain theory

Informally, a Markov chain (or Markov process) is a time-indexed stochastic process satisfying the property: conditioned on the present, the future is independent of the past. We restrict ourselves to the discrete-time, time-homogeneous, countable-space case, where such a process is characterized by its initial distribution and a transition matrix.

Construction of a Markov chain For our purposes, it will suffice to "define" a Markov chain through a particular construction. Let $V$ be a finite or countable space. Recall that a stochastic matrix on $V$ is a nonnegative matrix $P=$ $(P(i, j))_{i, j \in V}$ satisfying

$$
\sum_{j \in V} P(i, j)=1, \quad \forall i \in V .
$$

We think of $P(i, \cdot)$ as a probability distribution on $V$. In particular, for a set of states $A \subseteq V$, we let

$$
P(i, A)=\sum_{j \in A} P(i, j) .
$$

Let $\mu$ be a probability measure on $V$ and let $P$ be a stochastic matrix on $V$. One way to construct a Markov chain $\left(X_{t}\right)_{t \geq 0}$ on $V$ with transition matrix $P$ and initial distribution $\mu$ is the following:

- Pick $X_{0} \sim \mu$ and let $(Y(i, n))_{i \in V, n \geq 1}$ be a mutually independent array of random variables with $Y(i, n) \sim P(i, \cdot)$.
- Set inductively $X_{n}:=Y\left(X_{n-1}, n\right), n \geq 1$.

So in particular:

$$
\mathbb{P}\left[X_{0}=x_{0}, \ldots, X_{t}=x_{t}\right]=\mu\left(x_{0}\right) P\left(x_{0}, x_{1}\right) \cdots P\left(x_{t-1}, x_{t}\right)
$$

We use the notation $\mathbb{P}_{x}, \mathbb{E}_{x}$ for the probability distribution and expectation under the chain started at $x$. Similarly for $\mathbb{P}_{\mu}, \mathbb{E}_{\mu}$ where $\mu$ is a probability distribution.

Example 1.1.17 (Simple random walk on a graph). Let $G=(V, E)$ be a finite or infinite, locally finite graph. Simple random walk on $G$ is the Markov chain on $V$, started at an arbitrary vertex, which at each time picks a uniformly chosen neighbor of the current state. (Exercise 1.9 asks for the transition matrix.)

Markov property Let $\left(X_{t}\right)_{t \geq 0}$ be a Markov chain (or chain for short) with transition matrix $P$ and initial distribution $\mu$. Define the filtration $\left(\mathcal{F}_{t}\right)_{t \geq 0}$ with $\mathcal{F}_{t}=\sigma\left(X_{0}, \ldots, X_{t}\right)$ (see Appendix B). As mentioned above the defining property of Markov chains, known as the Markov property, is that: given the present, the future is independent of the past. In its simplest form, that can be interpreted as $\mathbb{P}\left[X_{t+1}=y \mid \mathcal{F}_{t}\right]=\mathbb{P}_{X_{t}}\left[X_{t+1}=y\right]=P\left(X_{t}, y\right)$. More formally:

Theorem 1.1.18 (Markov property). Let $f: V^{\infty} \rightarrow \mathbb{R}$ be bounded, measurable
simple
random
walk
on a graph

Markov property and let $F(x):=\mathbb{E}_{x}\left[f\left(\left(X_{t}\right)_{t \geq 0}\right)\right]$, then

$$
\mathbb{E}\left[f\left(\left(X_{s+t}\right)_{t \geq 0}\right) \mid \mathcal{F}_{s}\right]=F\left(X_{s}\right) \quad \text { a.s. }
$$

Remark 1.1.19. We will come back to the "strong" Markov property in Chapter 3.
We define $P^{t}(x, y):=\mathbb{P}_{x}\left[X_{t}=y\right]$. An important consequence of the Markov property (Theorem 1.1.18) is the following.

Theorem 1.1.20 (Chapman-Kolmogorov).

$$
P^{t}(x, z)=\sum_{y \in V} P^{s}(x, y) P^{t-s}(y, z), \quad s \in\{0,1, \ldots, t\} .
$$

Proof. This follows from the Markov property. Indeed note that $\mathbb{P}_{x}\left[X_{t}=z \mid \mathcal{F}_{s}\right]=$ $F\left(X_{s}\right)$ with $F(y):=\mathbb{P}_{y}\left[X_{t-s}=z\right]$ and take $\mathbb{E}_{x}$ on each side.

Example 1.1.21 (Random walk on $\mathbb{Z}$ ). Let $\left(X_{t}\right)$ be simple random walk on $\mathbb{Z}$ interpreted as a graph (i.e., $\mathbb{L}$ ) where $i \sim j$ if $|i-j|=1$.* Then $P(0, x)=1 / 2$ if $|x|=1$. And $P^{2}(0, x)=1 / 4$ if $|x|=2$ and $P^{2}(0,0)=1 / 2$.

As is conventional in Markov chain theory, we think of probability distributions over the state space as row vectors. We will typically denote them by Greek letters (e.g., $\mu, \pi$ ). If we write $\mu_{s}$ for the law of $X_{s}$ as a row vector, then

$$
\mu_{s}=\mu_{0} P^{s}
$$

where here $P^{s}$ is the matrix product of $P$ by itself $s$ times.

Stationarity The transition graph of a chain is the directed graph on $V$ whose edges are the transitions with strictly positive probability. A chain is irreducible if $V$ is the unique (strongly) connected component of its transition graph, that is, if all pairs of states have a directed path between them in the transition graph.

Example 1.1.22 (Simple random walk on a graph (continued)). Simple random walk on $G$ is irreducible if and only if $G$ is connected.

A stationary measure $\pi$ is a measure on $V$ such that

$$
\sum_{x \in V} \pi(x) P(x, y)=\pi(y), \quad \forall y \in V,
$$

or in matrix form $\pi=\pi P$. We say that $\pi$ is a stationary distribution if in addition $\pi$ is a probability measure.

Example 1.1.23 (Random walk on $\mathbb{Z}^{d}$ ). The all-one measure $\pi \equiv 1$ is stationary for simple random walk on $\mathbb{L}^{d}$.

Finite, irreducible chains always have a unique stationary distribution.
Theorem 1.1.24 (Existence and uniqueness: finite case). If $P$ is irreducible and has a finite state space, then:
(i) (Existence) it has a stationary distribution which, furthermore, is strictly positive;
(ii) (Uniqueness) the stationary distribution is unique.
*On $\mathbb{Z}$, simple random walk often refers to any nearest-neighbor random walk, whereas the example here is called simple symmetric random walk. We will not adopt this terminology here.

This result follows from Perron-Frobenius theory (a version of which is stated as Theorem 6.1.17). We give a self-contained proof.

Proof of Theorem 1.1.24 (i). We begin by proving existence. Denote by $n$ the number of states. Because $P$ is stochastic, we have by definition that $P \mathbf{1}=\mathbf{1}$, where $\mathbf{1}$ is the all-one vector. Put differently,

$$
(P-I) \mathbf{1}=\mathbf{0} .
$$

In particular, the columns of $P-I$ are linearly dependent, that is, the rank of $P-I$ is $<n$. That, in turn, implies that the rows of $P-I$ are linearly dependent since row rank and column rank are equal. Hence there exists a non-zero row vector $\mathbf{z} \in \mathbb{R}^{n}$ such that $\mathbf{z}(P-I)=\mathbf{0}$, or after rearranging,

$$
\begin{equation*}
\mathbf{z} P=\mathbf{z} . \tag{1.1.2}
\end{equation*}
$$

The rest of the proof is broken up into a series of lemmas. To take advantage of irreducibility, we first construct a positive stochastic matrix with $\mathbf{z}$ as a left eigenvector with eigenvalue 1 . We then show that all entries of $\mathbf{z}$ have the same sign. Finally, we normalize z.

Lemma 1.1.25 (Existence: Step 1). There exists a non-negative integer $h$ such that

$$
R=\frac{1}{h+1}\left[I+P+P^{2}+\cdots+P^{h}\right]
$$

is a stochastic matrix with strictly positive entries which satisfies

$$
\begin{equation*}
\mathbf{z} R=\mathbf{z} . \tag{1.1.3}
\end{equation*}
$$

Lemma 1.1.26 (Existence: Step 2). The entries of z are either all nonnegative or all nonpositive.
Lemma 1.1.27 (Existence: Step 3). Let

$$
\pi=\frac{\mathbf{z}}{\mathrm{z} \mathbf{1}} .
$$

Then $\pi$ is a strictly positive stationary distribution.
We denote the entries of $R$ and $P^{s}$ by $R_{x, y}$ and $P_{x, y}^{s}, x, y=1, \ldots, n$, respectively.
Proof of Lemma 1.1.25. By irreducibility (see Exercise 1.10), for any $x, y \in[n]$ there is $h_{x y}$ such that $P_{x, y}^{h_{x y}}>0$. Now define

$$
h=\max _{x, y \in[n]} h_{x y} .
$$

The matrix $P^{s}$, as a product of stochastic matrices, is a stochastic matrix for all $s$ (see Exercise 1.11). In particular, it has nonnegative entries. Hence, for each $x, y$,

$$
\begin{aligned}
R_{x, y} & =\frac{1}{h+1}\left[I_{x, y}+P_{x, y}+P_{x, y}^{2}+\cdots+P_{x, y}^{h}\right] \\
& \geq \frac{1}{h+1} P_{x, y}^{h_{x, y}}>0 .
\end{aligned}
$$

Moreover the matrix $R$, as a convex combination of stochastic matrices, is a stochastic matrix (see Exercise 1.11).

Since $\mathbf{z} P=\mathbf{z}$, it follows by induction that $\mathbf{z} P^{s}=\mathbf{z}$ for all $s$. Therefore,

$$
\begin{aligned}
\mathbf{z} R & =\frac{1}{h+1}\left[\mathbf{z} I+\mathbf{z} P+\mathbf{z} P^{2}+\cdots+\mathbf{z} P^{h}\right] \\
& =\frac{1}{h+1}[\mathbf{z}+\mathbf{z}+\mathbf{z}+\cdots+\mathbf{z}] \\
& =\mathbf{z} .
\end{aligned}
$$

That concludes the proof.
Proof of Lemma 1.1.26. We argue by contradiction. Suppose that two entries of $\mathbf{z}=\left(z_{1}, \ldots, z_{n}\right)$ have different signs, say $z_{i}>0$ while $z_{j}<0$. By the previous lemma,

$$
\begin{aligned}
\left|z_{y}\right| & =\left|\sum_{x} z_{x} R_{x, y}\right| \\
& =\left|\sum_{x: z_{x} \geq 0} z_{x} R_{x, y}+\sum_{x: z_{x}<0} z_{x} R_{x, y}\right| .
\end{aligned}
$$

Moreover, $R_{x, y}>0$ for all $x, y$. Therefore, the first term on the last line is strictly positive (since it is at least $z_{i} R_{i, y}>0$ ) while the second term is strictly negative (since it is at most $z_{j} R_{j, y}<0$ ). Hence, because of cancellations (see Exercise 1.13), the expression in the previous display is strictly smaller than the sum of the absolute values, that is,

$$
\left|z_{y}\right|<\sum_{x}\left|z_{x}\right| R_{x, y} .
$$

Since $R$ is also stochastic by the previous lemma, we deduce after summing over $y$ that

$$
\sum_{y}\left|z_{y}\right|<\sum_{y} \sum_{x}\left|z_{x}\right| R_{x, y}=\sum_{x}\left|z_{x}\right| \sum_{y} R_{x, y}=\sum_{x}\left|z_{x}\right|,
$$

a contradiction, thereby proving the claim.

Proof of Lemma 1.1.27. Now define $\pi$ entrywise by

$$
\pi_{x}=\frac{z_{x}}{\sum_{i} z_{i}}=\frac{\left|z_{x}\right|}{\sum_{i}\left|z_{i}\right|} \geq 0,
$$

where the second equality comes from the previous lemma. We also used the fact that $\mathbf{z} \neq \mathbf{0}$.

For all $y$, by definition of $\mathbf{z}$,

$$
\sum_{x} \pi_{x} P_{x, y}=\sum_{x} \frac{z_{x}}{\sum_{i} z_{i}} P_{x, y}=\frac{1}{\sum_{i} z_{i}} \sum_{x} z_{x} P_{x, y}=\frac{z_{y}}{\sum_{i} z_{i}}=\pi_{y} .
$$

The same holds with $P_{x, y}$ replaced by $R_{x, y}$ from (1.1.3). Since $R_{x, y}>0$ and $\mathbf{z} \neq \mathbf{0}$ it follows that $\pi_{y}>0$ for all $y$. That proves the claim.

That concludes the proof of the existence claim.
It remains to prove uniqueness. See Exercise 1.14 for an alternative proof based on the maximum principle (to which we come back in Theorem 3.3.9 and Exercise 3.12).

Proof of Theorem 1.1.24 (ii). Suppose there are two distinct stationary distributions $\pi_{1}$ and $\pi_{2}$ (which must be strictly positive). Since they are distinct and both sum to 1 , they are not a multiple of each other and therefore are linearly independent. Apply the Gram-Schmidt procedure:

$$
\mathbf{q}_{1}=\frac{\pi_{1}}{\left\|\pi_{1}\right\|_{2}} \quad \text { and } \quad \mathbf{q}_{2}=\frac{\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}}{\left\|\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}\right\|_{2}}
$$

Then

$$
\mathbf{q}_{1} P=\frac{\pi_{1}}{\left\|\pi_{1}\right\|_{2}} P=\frac{\pi_{1} P}{\left\|\pi_{1}\right\|_{2}}=\frac{\pi_{1}}{\left\|\pi_{1}\right\|_{2}}=\mathbf{q}_{1},
$$

and all entries of $\mathbf{q}_{1}$ are strictly positive.
Similarly,

$$
\begin{aligned}
\mathbf{q}_{2} P & =\frac{\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}}{\left\|\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}\right\|_{2}} P \\
& =\frac{\pi_{2} P-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1} P}{\left\|\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}\right\|_{2}} \\
& =\frac{\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}}{\left\|\pi_{2}-\left\langle\pi_{2}, \mathbf{q}_{1}\right\rangle \mathbf{q}_{1}\right\|_{2}} \\
& =\mathbf{q}_{2} .
\end{aligned}
$$

Since $\mathbf{z}:=\mathbf{q}_{2}$ satisfies (1.1.2), by Lemmas 1.1.25-1.1.27 there is a multiple of $\mathbf{q}_{2}$, say $\mathbf{q}_{2}^{\prime}=\alpha \mathbf{q}_{2}$ with $\alpha \neq 0$, such that $\mathbf{q}_{2}^{\prime} P=\mathbf{q}_{2}^{\prime}$ and all entries of $\mathbf{q}_{2}^{\prime}$ are strictly positive. By the Gram-Schmidt procedure,

$$
\left\langle\mathbf{q}_{1}, \mathbf{q}_{2}^{\prime}\right\rangle=\left\langle\mathbf{q}_{1}, \alpha \mathbf{q}_{2}\right\rangle=\alpha\left\langle\mathbf{q}_{1}, \mathbf{q}_{2}\right\rangle=0 .
$$

But this is a contradiction since both vectors are strictly positive. That concludes the proof of the uniqueness claim.

Reversibility A transition matrix $P$ is reversible with respect to a measure $\eta$ if

$$
\eta(x) P(x, y)=\eta(y) P(y, x)
$$

for all $x, y \in V$. These equations are known as detailed balance. Here is the key observation: by summing over $y$ and using the fact that $P$ is stochastic, such $a$ measure is necessarily stationary. (Exercise 1.12 explains the name.)

Example 1.1.28 (Random walk on $\mathbb{Z}^{d}$ (continued)). The measure $\eta \equiv 1$ is reversible for simple random walk on $\mathbb{L}^{d}$.

Example 1.1.29 (Simple random walk on a graph (continued)). Let ( $X_{t}$ ) be simple random walk on a connected graph $G=(V, E)$. Then $\left(X_{t}\right)$ is reversible with respect to $\eta(v):=\delta(v)$, where recall that $\delta(v)$ is the degree of $v$. Indeed, for all $\{u, v\} \in E$,

$$
\delta(u) P(u, v)=\delta(u) \frac{1}{\delta(u)}=1=\delta(v) \frac{1}{\delta(v)}=\delta(v) P(v, u) .
$$

(See Exercise 1.9 for the transition matrix of simple random walk on a graph.)

Example 1.1.30 (Metropolis chain). The Metropolis algorithm modifies an irreducible, symmetric (i.e., whose transition matrix is a symmetric matrix) chain $Q$ to produce a new chain $P$ with the same transition graph and a prescribed positive stationary distribution $\pi$. The idea is simple. For each pair $x \neq y$, either we multiply $Q(x, y)$ by $\pi(y) / \pi(x)$ and leave $Q(y, x)$ intact, or vice versa. Detailed balance immediately follows. To ensure that the new transition matrix remains stochastic, for each pair we make the choice that lowers the transition probabilities; then we add the lost probability to the diagonal (i.e., to the probability of staying put).

Formally, the definition of the new chain is

$$
P(x, y):= \begin{cases}Q(x, y)\left[\frac{\pi(y)}{\pi(x)} \wedge 1\right] & \text { if } x \neq y \\ 1-\sum_{z \neq x} Q(x, z)\left[\frac{\pi(z)}{\pi(x)} \wedge 1\right] & \text { otherwise }\end{cases}
$$

Metropolis
algorithm

Note that, by definition of $P$ and the fact that $Q$ is stochastic, we have $P(x, y) \leq$ $Q(x, y)$ for all $x \neq y$ so

$$
\sum_{y \neq x} P(x, y) \leq 1,
$$

and hence $P$ is well-defined as a transition matrix. We claim further that $P$ is reversible with respect to $\pi$. Suppose $x \neq y$ and assume without loss of generality that $\pi(x) \geq \pi(y)$. Then, by definition of $P$, we have

$$
\begin{aligned}
\pi(x) P(x, y) & =\pi(x) Q(x, y) \frac{\pi(y)}{\pi(x)} \\
& =Q(x, y) \pi(y) \\
& =Q(y, x) \pi(y) \\
& =P(y, x) \pi(y),
\end{aligned}
$$

where we used the symmetry of $Q$.

Convergence and mixing time A key property of Markov chains is that, under suitable assumptions, they converge to a stationary regime. We need one more definition before stating the theorem. A chain is said to be aperiodic if, for all $x \in V$, the greatest common divisor of $\left\{t: P^{t}(x, x)>0\right\}$ is 1 .

Example 1.1.31 (Lazy random walk on a graph). The lazy simple random walk on $G$ is the Markov chain such that, at each time, it stays put with probability $1 / 2$ or chooses a uniformly random neighbor of the current state otherwise. Such a chain is aperiodic.

Lemma 1.1.32 (Consequence of aperiodicity). If $P$ is aperiodic, irreducible and has a finite state space, then there is a positive integer $t_{0}$ such that for all $t \geq t_{0}$ the matrix $P^{t}$ has strictly positive entries.

We can now state the convergence theorem. For probability measures $\mu, \nu$ on $V$, their total variation distance is

$$
\begin{equation*}
\|\mu-\nu\|_{\mathrm{TV}}:=\sup _{A \subseteq V}|\mu(A)-\nu(A)| . \tag{1.1.4}
\end{equation*}
$$

Theorem 1.1.33 (Convergence theorem). Suppose $P$ is irreducible, aperiodic and has stationary distribution $\pi$. Then, for all $x$

$$
\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}} \rightarrow 0,
$$

as $t \rightarrow+\infty$.

We give a proof in the finite case in Example 4.3.3. In particular, the convergence theorem implies that for all $x, y$,

$$
P^{t}(x, y) \rightarrow \pi(y) .
$$

Without aperiodicity, it can be shown that we have the weaker claim

$$
\begin{equation*}
\frac{1}{t} \sum_{s=1}^{t} P^{s}(x, y) \rightarrow \pi(y), \tag{1.1.5}
\end{equation*}
$$

as $t \rightarrow+\infty$.
We will be interested in quantifying the speed of convergence in Theorem 1.1.33. For this purpose, we define

$$
\begin{equation*}
d(t):=\sup _{x \in V}\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}} \tag{1.1.6}
\end{equation*}
$$

Lemma 1.1.34 (Monotonicity of $d(t)$ ). The function $d(t)$ is non-increasing in $t$.
Proof. Note that, by definition of $P^{t+1}$,

$$
\begin{aligned}
d(t+1) & =\sup _{x \in V} \sup _{A \subseteq V}\left|P^{t+1}(x, A)-\pi(A)\right| \\
& =\sup _{x \in V} \sup _{A \subseteq V}\left|\sum_{z} P(x, z)\left(P^{t}(z, A)-\pi(A)\right)\right| \\
& \leq \sup _{x \in V} \sum_{z} P(x, z) \sup _{A \subseteq V}\left|P^{t}(z, A)-\pi(A)\right| \\
& \leq \sup _{z \in V} \sup _{A \subseteq V}\left|P^{t}(z, A)-\pi(A)\right| \\
& =d(t)
\end{aligned}
$$

where on the second and fourth line we used that $P$ is a stochastic matrix.
The following concept will play a key role in quantifying the "speed of convergence" to stationarity.

Definition 1.1.35 (Mixing time). For a fixed $\varepsilon>0$, the mixing time is defined as

$$
\mathrm{t}_{\text {mix }}(\varepsilon):=\inf \{t \geq 0: d(t) \leq \varepsilon\} .
$$

### 1.2 Some discrete probability models

With the necessary background covered, we are now in a position to define formally a few important discrete probability models that will be ubiquitous in this book. These are all graph-based processes. Many more interesting random discrete structures and other related probabilistic models will be encountered throughout (and defined where needed).

Percolation Percolation processes are meant to model the movement of a fluid through a porous medium. There are several types of percolation models. We focus here on bond percolation. In words, edges of a graph are "open" at random, indicating that fluid is passing through. We are interested in the "open clusters," that is, the regions reached by the fluid.

Definition 1.2.1 (Bond percolation). Let $G=(V, E)$ be a finite or infinite graph. The bond percolation process on $G$ with density $p \in[0,1]$, whose probability measure is denoted by $\mathbb{P}_{p}$, is defined as follows: each edge of $G$ is independently set to open with probability $p$, otherwise it is set to closed. Write $x \Leftrightarrow y$ if $x, y \in V$ are connected by a path all of whose edges are open. The open cluster of $x$ is

$$
\mathcal{C}_{x}:=\{y \in V: x \Leftrightarrow y\} .
$$

We will mostly consider bond percolation on the infinite graphs $\mathbb{L}^{d}$ or $\mathbb{T}_{d}$. The main question we will ask is: For which values of $p$ is there an infinite open cluster?

Random graphs Random graphs provide a natural framework to study complex networks. Different behaviors are observed depending on the modeling choices made. Perhaps the simplest and most studied is the Erdős-Rényi random graph model. We consider the version due to Gilbert. Here the edges are present independently with a fixed probability. Despite its simplicity, this model exhibits a rich set of phenomena that make it a prime example for the use of a variety of probabilistic techniques.

Definition 1.2.2 (Erdős-Rényi graph model). Let $n \in \mathbb{N}$ and $p \in[0,1]$. Set $V:=$ $[n]$. Under the Erdős-Rényi graph model on $n$ vertices with density $p$, a random graph $G=(V, E)$ is generated as follows: for each pair $x \neq y$ in $V$, the edge $\{x, y\}$ is in $E$ with probability $p$ independently of all other edges. We write $G \sim$ $\mathbb{G}_{n, p}$ and we denote the corresponding probability measure by $\mathbb{P}_{n, p}$.

Typical questions regarding the Erdős-Rényi graph model (and random graphs more generally) include: How are degrees distributed? Is G connected? What
is the (asymptotic) probability of observing a particular subgraph, for example, a triangle? What is the typical chromatic number?

As one alternative to the Erdôs-Rényi model, we will also encounter preferential attachment graphs. These are meant to model the growth of a network where new edges are more likely to be incident with vertices of high degree, a reasonable assumption in some applied settings. Such a process produces graphs with properties that differ from those of the Erdős-Rényi model; in particular they tend to have a "fatter" degree distribution tail. In the definition of preferential attachment graphs, we restrict ourselves to the tree case (see Exercise 1.15).

Definition 1.2.3 (Preferential attachment graph). The preferential attachment graph process produces a sequence of graphs $\left(G_{t}\right)_{t \geq 1}$ as follows. We start at time 1 with two vertices, denoted $v_{0}$ and $v_{1}$, connected by an edge. At time $t$, we add vertex $v_{t}$ with a single edge connecting it to an old vertex, which is picked proportionally to its degree. We write $\left(G_{t}\right)_{t \geq 1} \sim \mathrm{PA}_{1}$.

Markov random fields Another common class of graph-based processes involves the assignment of random "states" to the vertices of a fixed graph. The state distribution is typically specified through "interactions between neighboring vertices." Such models are widely studied in statistical physics and also have important applications in statistics. We focus on models with a Markovian (i.e., conditional independence) structure that makes them particularly amenable to rigorous analysis and computational methods. We start with Gibbs random fields, a broad class of such models.

Definition 1.2.4 (Gibbs random field). Let $S$ be a finite set and let $G=(V, E)$ be a finite graph. Denote by $\mathcal{K}$ the set of all cliques of $G$. A positive probability measure $\mu$ on $\mathcal{X}:=S^{V}$ is called a Gibbs random field if there exist clique potentials $\phi_{K}$ : $S^{K} \rightarrow \mathbb{R}, K \in \mathcal{K}$, such that

$$
\mu(\sigma)=\frac{1}{\mathcal{Z}} \exp \left(\sum_{K \in \mathcal{K}} \phi_{K}\left(\sigma_{K}\right)\right),
$$

where $\sigma_{K}$ is the vector $\sigma \in \mathcal{X}$ restricted to the vertices (i.e., coordinates) in $K$ and $\mathcal{Z}$ is a normalizing constant.

The following example introduces the primary Gibbs random field we will encounter.

Example 1.2.5 (Ising model). For $\beta>0$, the (ferromagnetic) Ising model with inverse temperature $\beta$ is the Gibbs random field with $S:=\{-1,+1\}, \phi_{\{i, j\}}\left(\sigma_{\{i, j\}}\right)=$
$\beta \sigma_{i} \sigma_{j}$ and $\phi_{K} \equiv 0$ if $|K| \neq 2$. The function $\mathcal{H}(\sigma):=-\sum_{\{i, j\} \in E} \sigma_{i} \sigma_{j}$ is known as the Hamiltonian. The normalizing constant $\mathcal{Z}:=\mathcal{Z}(\beta)$ is called the partition function. The states $\left(\sigma_{i}\right)_{i \in V}$ are referred to as spins.

Typical questions regarding Ising models include: How fast is correlation decaying down the graph? How well can one guess the state at an unobserved vertex? We will also consider certain Markov chains related to Ising models (see Definition 1.2.8).

Random walks on graphs and reversible Markov chains The last class of processes we focus on are random walks on graphs and their generalizations. Recall the following definition.
Definition 1.2.6 (Simple random walk on a graph). Let $G=(V, E)$ be a finite or countable, locally finite graph. Simple random walk on $G$ is the Markov chain on $V$, started at an arbitrary vertex, which at each time picks a uniformly chosen neighbor of the current state.

We generalize the definition by adding weights to the edges. In this context, we denote edge weights by $c(e)$ for "conductance" (see Section 3.3).

Definition 1.2.7 (Random walk on a network). Let $G=(V, E)$ be a finite or countably infinite graph. Let $c: E \rightarrow \mathbb{R}_{+}$be a positive edge weight function on $G$. Recall that we call $\mathcal{N}=(G, c)$ a network. We assume that for all $u \in V$

$$
c(u):=\sum_{e=\{u, v\} \in E} c(e)<+\infty .
$$

random
walk on a
network
Random walk on network $\mathcal{N}$ is the Markov chain on $V$, started at an arbitrary vertex, which at each time picks a neighbor of the current state proportionally to the weight of the corresponding edge. That is, the transition matrix is given by

$$
P(u, v)= \begin{cases}\frac{c(e)}{c(u)} & \text { if } e=\{u, v\} \in E, \\ 0 & \text { otherwise }\end{cases}
$$

By definition of $P$, it is immediate that this Markov chain is reversible with respect to the measure $\eta(u):=c(u)$. In fact, conversely, any countable reversible Markov chain can be seen as a random walk on a network by setting $c(e):=\pi(x) P(x, y)=$ $\pi(y) P(y, x)$ for all $x, y$ such that $P(x, y)>0$.

Typical questions include: How long does it take to visit all vertices at least once or a particular subset of vertices for the first time? How fast does the walk approach stationarity? How often does the walk return to its starting point?

We will also encounter a particular class of Markov chains related to Ising models, the Glauber dynamics.

Definition 1.2.8 (Glauber dynamics of the Ising model). Let $\mu_{\beta}$ be the Ising model with inverse temperature $\beta>0$ on a graph $G=(V, E)$. The (single-site) Glauber dynamics is the Markov chain on $\mathcal{X}:=\{-1,+1\}^{V}$ which at each time:

- selects a site $i \in V$ uniformly at random, and
- updates the spin at $i$ according to $\mu_{\beta}$ conditioned on agreeing with the current state at all sites in $V \backslash\{i\}$.

Specifically, for $\gamma \in\{-1,+1\}, i \in V$, and $\sigma \in \mathcal{X}$, let $\sigma^{i, \gamma}$ be the configuration $\sigma$ with the spin at $i$ being set to $\gamma$. Let $n=|V|$ and $S_{i}(\sigma):=\sum_{j \sim i} \sigma_{j}$. The nonzero entries of the transition matrix are

$$
Q_{\beta}\left(\sigma, \sigma^{i, \gamma}\right):=\frac{1}{n} \cdot \frac{e^{\gamma \beta S_{i}(\sigma)}}{e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}} .
$$

This chain is irreducible since we can flip each site one by one to go from any state to any other. It is straightforward to check that $Q_{\beta}\left(\sigma, \sigma^{i, \gamma}\right)$ is a stochastic matrix. The next theorem shows that $\mu_{\beta}$ is its stationary distribution.

Theorem 1.2.9. The Glauber dynamics is reversible with respect to $\mu_{\beta}$.
Proof. For all $\sigma \in \mathcal{X}$ and $i \in V$, let

$$
S_{\neq i}(\sigma):=\mathcal{H}\left(\sigma^{i,+}\right)+S_{i}(\sigma)=\mathcal{H}\left(\sigma^{i,-}\right)-S_{i}(\sigma) .
$$

We have

$$
\begin{aligned}
\mu_{\beta}\left(\sigma^{i,-}\right) Q_{\beta}\left(\sigma^{i,-}, \sigma^{i,+}\right) & =\frac{e^{-\beta S_{\neq i}(\sigma)} e^{-\beta S_{i}(\sigma)}}{\mathcal{Z}(\beta)} \cdot \frac{e^{\beta S_{i}(\sigma)}}{n\left[e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}\right]} \\
& =\frac{e^{-\beta S_{\neq i}(\sigma)}}{n \mathcal{Z}(\beta)\left[e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}\right]} \\
& =\frac{e^{-\beta S_{\neq i}(\sigma)} e^{\beta S_{i}(\sigma)}}{\mathcal{Z}(\beta)} \cdot \frac{e^{-\beta S_{i}(\sigma)}}{n\left[e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}\right]} \\
& =\mu_{\beta}\left(\sigma^{i,+}\right) Q_{\beta}\left(\sigma^{i,+}, \sigma^{i,-}\right) .
\end{aligned}
$$

That concludes the proof.

## Exercises

Exercise 1.1 (0-norm). Show that $\|\mathbf{u}\|_{0}$ does not define a norm.
Exercise 1.2 (A factorial bound: one way). Let $\ell$ be a positive integer.
(i) Use the bound $1+x \leq e^{x}$ to show that

$$
\frac{k+1}{k} \leq e^{1 / k},
$$

and

$$
\frac{k}{k+1} \leq e^{1 /(k+1)}
$$

for all positive integers $k$.
(ii) Use part (i) and the quantity

$$
\prod_{k=1}^{\ell-1} \frac{(k+1)^{k}}{k^{k}}
$$

to show that

$$
\ell!\geq \frac{\ell^{\ell}}{e^{\ell-1}}
$$

(iii) Use part (i) and the quantity

$$
\prod_{k=1}^{\ell-1} \frac{k^{k+1}}{(k+1)^{k+1}},
$$

to show that

$$
\ell!\leq \frac{\ell^{\ell+1}}{e^{\ell-1}}
$$

Exercise 1.3 (A factorial bound: another way). Let $\ell$ be a positive integer. Show that

$$
\frac{\ell^{\ell}}{e^{\ell-1}} \leq \ell!\leq \frac{\ell^{\ell+1}}{e^{\ell-1}}
$$

by considering the logarithm of $\ell$ !, interpreting the resulting quantity as a Riemann sum, and bounding above and below by an integral.

Exercise 1.4 (A binomial bound). Show that for integers $0<d \leq n$,

$$
\sum_{k=0}^{d}\binom{n}{k} \leq\left(\frac{e n}{d}\right)^{d}
$$

[Hint: Multiply the left-hand side of the inequality by $(d / n)^{d} \leq(d / n)^{k}$ and use the binomial theorem.]

Exercise 1.5 (Powers of the adjacency matrix). Let $A^{n}$ be the $n$-th matrix power of the adjacency matrix $A$ of a graph $G=(V, E)$. Prove that the $(i, j)$-th entry $a_{i j}^{n}$ is the number of walks of length exactly $n$ between vertices $i$ and $j$ in $G$. [Hint: Use induction on $n$.]

Exercise 1.6 (Paths). Let $u, v$ be vertices of a graph $G=(V, E)$.
(i) Show that the graph distance $d_{G}(u, v)$ is a metric.
(ii) Show that the binary relation $u \leftrightarrow v$ is an equivalence relation.
(iii) Prove (ii) in the directed case.

Exercise 1.7 (Trees: number of edges). Prove that a connected graph with $n$ vertices is a tree if and only if it has $n-1$ edges. [Hint: Proceed by induction. Then use Corollary 1.1.6.]

Exercise 1.8 (Trees: characterizations). Prove Theorem 1.1.5.
Exercise 1.9 (Simple random walk on a graph). Let $\left(X_{t}\right)_{t \geq 0}$ be simple random walk on a finite graph $G=(V, E)$. Suppose the vertex set is $V=[n]$. Write down an expression for the transition matrix of $\left(X_{t}\right)$.

Exercise 1.10 (Communication lemma). Let $\left(X_{t}\right)$ be a finite Markov chain. Show that, if $x \rightarrow y$, then there is an integer $r \geq 1$ such that

$$
\mathbb{P}\left[X_{r}=y \mid X_{0}=x\right]=\left(P^{r}\right)_{x, y}>0
$$

Exercise 1.11 (Stochastic matrices from stochastic matrices). Let

$$
P^{(1)}, P^{(2)}, \ldots, P^{(r)} \in \mathbb{R}^{n \times n}
$$

be stochastic matrices.
(i) Show that $P^{(1)} P^{(2)}$ is a stochastic matrix. That is, a product of stochastic matrices is a stochastic matrix.
(ii) Show that for any $\alpha_{1}, \ldots, \alpha_{r} \in[0,1]$ with $\sum_{i=1}^{r} \alpha_{i}=1$,

$$
\sum_{i=1}^{r} \alpha_{i} P^{(i)}
$$

is stochastic. That is, a convex combination of stochastic matrices is a stochastic matrix.

Exercise 1.12 (Reversing time). Let $\left(X_{t}\right)$ be a finite Markov chain with transition matrix $P$. Assume $P$ is reversible with respect to a probability distribution $\pi$. Assume that the initial distribution is $\pi$. Show that for any sequence of states $z_{0}, \ldots, z_{s}$, the reversed sequence has the same probability, that is,

$$
\mathbb{P}\left[X_{s}=z_{0}, \ldots, X_{0}=z_{s}\right]=\mathbb{P}\left[X_{s}=z_{s}, \ldots, X_{0}=z_{0}\right]
$$

Exercise 1.13 (A strict inequality). Let $a, b \in \mathbb{R}$ with $a<0$ and $b>0$. Show that

$$
|a+b|<|a|+|b|
$$

[Hint: Consider the cases $a+b \geq 0$ and $a+b<0$ separately.]
Exercise 1.14 (Uniqueness: maximum principle). Let $P=\left(P_{i, j}\right)_{i, j=1}^{n} \in \mathbb{R}^{n}$ be a transition matrix.
(i) Let $\alpha_{1}, \ldots, \alpha_{m}>0$ such that $\sum_{i=1}^{m} \alpha_{i}=1$. Let $\mathbf{x}=\left(x_{1}, \ldots, x_{m}\right) \in \mathbb{R}^{n}$. Show that

$$
\sum_{i=1}^{m} \alpha_{i} x_{i} \leq \max _{i} x_{i}
$$

and that equality holds if and only if $x_{1}=x_{2}=\cdots=x_{m}$.
(ii) Let $\mathbf{0} \neq \mathbf{y} \in \mathbb{R}^{n}$ be a right eigenvector of $P$ with eigenvalue 1 , that is, $P \mathbf{y}=\mathbf{y}$. Assume that $\mathbf{y}$ is not a constant vector, that is, there is $i \neq j$ such that $y_{i} \neq y_{j}$. Let $k$ be such that $y_{k}=\max _{i \in[n]} y_{i}$. Show that for any $\ell$ such that $P_{k, \ell}>0$ we necessarily have $y_{\ell}=y_{k}$. [Hint: Use that $y_{k}=\sum_{i=1}^{n} P_{k, i} y_{i}$ and apply (i).]
(iii) Assume that $P$ is irreducible. Let $\mathbf{0} \neq \mathbf{y} \in \mathbb{R}^{n}$ again be a right eigenvector of $P$ with eigenvalue 1 . Use (ii) to show that $\mathbf{y}$ is necessarily a constant vector.
(iv) Assume again that $P$ is irreducible. Use (iii) to conclude that the dimension of the null space of $P^{T}-I$ is exactly 1. [Hint: Use the Rank-Nullity Theorem.]

Exercise 1.15 (Preferential attachment trees). Let $\left(G_{t}\right)_{t \geq 1} \sim \mathrm{PA}_{1}$ as in Definition 1.2.3. Show that $G_{t}$ is a tree with $t+1$ vertices for all $t \geq 1$.

Exercise 1.16 (Warm-up: a little calculus). Prove the following inequalities which we will encounter throughout. [Hint: Basic calculus should do.]
(i) Show that $1-x \leq e^{-x}$ for all $x \in \mathbb{R}$.
(ii) Show that $1-x \geq e^{-x-x^{2}}$ for $x \in[0,1 / 2]$.
(iii) Show that $\log (1+x) \leq x-x^{2} / 4$ for $x \in[0,1]$.
(iv) Show that $\log x \leq x-1$ for all $x \in \mathbb{R}_{+}$.
(v) Show that $\cos x \leq e^{-x^{2} / 2}$ for $x \in[0, \pi / 2)$. [Hint: Consider the function $h(x)=\log \left(e^{x^{2} / 2} \cos x\right)$ and recall that the derivative of $\tan x$ is $1+\tan ^{2} x$.]

## Bibliographic Remarks

Section 1.1 For an introduction to graphs see for example [Die10] or [Bol98]. Four different proofs of Cayley's formula are detailed in the delightful [AZ18]. Markov chain theory is covered in details in [Dur10, Chapter 6]. For a more gentle introduction, see for example [Dur12, Chapter 1], [Nor98, Chapter 1], or [Res92, Chapter 2].

Section 1.2 For book-length treatments of percolation theory see [BR06a, Gri10b]. The version of the Erdős-Rényi random graph model we consider here is due to Gilbert [Gil59]. For much deeper accounts of the theory of random graphs and related processes, see for example [Bol01, Dur06, JLR11, vdH17, FK16]. Two standard references on finite Markov chains and mixing times are [AF, LPW06].

## Chapter 2

## Moments and tails

In this chapter we look at the moments of a random variable. Specifically we demonstrate that moments capture useful information about the tail of a random variable while often being simpler to compute or at least bound. Several wellknown inequalities quantify this intuition. Although they are straightforward to derive, such inequalities are surprisingly powerful. Through a range of applications, we illustrate the utility of controlling the tail of a random variable, typically by allowing one to dismiss certain "bad events" as rare. We begin in Section 2.1 by recalling the classical Markov and Chebyshev's inequalities. Then we discuss three of the most fundamental tools in discrete probability and probabilistic combinatorics. In Sections 2.2 and 2.3, we derive the complementary first and second moment methods, and give several standard applications, especially to phase transitions in random graphs and percolation. In Section 2.4 we develop the ChernoffCramér method, which relies on the moment-generating function and is the building block for a large class of tail bounds. Two key applications in data science are briefly introduced: sparse recovery and empirical risk minimization.

### 2.1 Background

We start with a few basic definitions and standard inequalities. See Appendix B for a refresher on random variables and their expectation.

[^0]
### 2.1.1 Definitions

Moments As a quick reminder, let $X$ be a random variable with $\mathbb{E}|X|^{k}<+\infty$ for some non-negative integer $k$. In that case we write $X \in L^{k}$. Recall that the quantities $\mathbb{E}\left[X^{k}\right]$ and $\mathbb{E}\left[(X-\mathbb{E} X)^{k}\right]$, which are well-defined when $X \in L^{k}$, are called respectively the $k$-th moment and $k$-th central moment of $X$. The first moment and the second central moment are known as the mean and variance, the square root of which is the standard deviation. A random variable is said to be centered if its mean is 0 . Recall that for a non-negative random variable $X$, the $k$-th moment can be expressed as

$$
\begin{equation*}
\mathbb{E}\left[X^{k}\right]=\int_{0}^{+\infty} k x^{k-1} \mathbb{P}[X>x] \mathrm{d} x \tag{2.1.1}
\end{equation*}
$$

The moment-generating function (or exponential moment) of $X$ is the function

$$
M_{X}(s):=\mathbb{E}\left[e^{s X}\right],
$$

defined for all $s \in \mathbb{R}$ where it is finite, which includes at least $s=0$. If $M_{X}(s)$ is defined on $\left(-s_{0}, s_{0}\right)$ for some $s_{0}>0$ then $X$ has finite moments of all orders, for any $k \in \mathbb{Z}$,

$$
\begin{equation*}
\frac{\mathrm{d}^{k}}{\mathrm{~d} s} M_{X}(s)=\mathbb{E}\left[X^{k} e^{s X}\right] \tag{2.1.2}
\end{equation*}
$$

and the following expansion holds

$$
M_{X}(s)=\sum_{k \geq 0} \frac{s^{k}}{k!} \mathbb{E}\left[X^{k}\right], \quad|s|<s_{0}
$$

The moment-generating function plays nicely with sums of independent random variables. Specifically, if $X_{1}$ and $X_{2}$ are independent random variables with $M_{X_{1}}$ and $M_{X_{2}}$ defined over a joint interval $\left(-s_{0}, s_{0}\right)$, then for $s$ in that interval

$$
\begin{align*}
M_{X_{1}+X_{2}}(s) & =\mathbb{E}\left[e^{s\left(X_{1}+X_{2}\right)}\right] \\
& =\mathbb{E}\left[e^{s X_{1}} e^{s X_{2}}\right] \\
& =\mathbb{E}\left[e^{s X_{1}}\right] \mathbb{E}\left[e^{s X_{2}}\right] \\
& =M_{X_{1}}(s) M_{X_{2}}(s), \tag{2.1.3}
\end{align*}
$$

where we used independence in the third equality.
One more piece of notation: if $A$ is an event and $X \in L^{1}$, then we use the shorthand

$$
\mathbb{E}[X ; A]=\mathbb{E}\left[X \mathbf{1}_{A}\right] .
$$

Tails We refer to a probability of the form $\mathbb{P}[X \geq x]$ as an upper tail (or right tail) probability. Typically $x$ is (much) greater than the mean or median of $X$. Similarly we refer to $\mathbb{P}[X \leq x]$ as a lower tail (or left tail) probability. Our general goal in this chapter is to bound tail probabilities using moments and moment-generating functions.

Tail bounds arise naturally in many contexts, as events of interest can often be framed in terms of a random variable being unusually large or small. Such probabilities are often hard to compute directly however. As we will see in this chapter, moments offer an effective means to control tail probabilities for two main reasons: (i) moments contain information about the tails of a random variable, as (2.1.1) below makes explicit for instance; and (ii) they are typically easier to compute-or, at least, to approximate.

As we will see, tail bounds are also useful to study the maximum of a collection of random variables.

### 2.1.2 Basic inequalities

Markov's inequality Our first bound on the tail of a random variable is Markov's inequality. In words, for a non-negative random variable: the heavier the tail, the larger the expectation. This simple inequality is in fact a key ingredient in more sophisticated tail bounds as we will see.
Theorem 2.1.1 (Markov's inequality). Let $X$ be a non-negative random variable. Then, for all $b>0$,

$$
\begin{equation*}
\mathbb{P}[X \geq b] \leq \frac{\mathbb{E} X}{b} \tag{2.1.4}
\end{equation*}
$$

Proof.

$$
\mathbb{E} X \geq \mathbb{E}[X ; X \geq b] \geq \mathbb{E}[b ; X \geq b]=b \mathbb{P}[X \geq b]
$$

See Figure 2.1 for a proof by picture. Note that this inequality is non-trivial only when $b>\mathbb{E} X$.

Chebyshev's inequality An application of Markov's inequality (Theorem 2.1.1) to $|X-\mathbb{E} X|^{2}$ gives a classical tail bound featuring the second moment of a random variable.
Theorem 2.1.2 (Chebyshev's inequality). Let $X$ be a random variable with $\mathbb{E} X^{2}<$ $+\infty$. Then, for all $\beta>0$,

Chebyshev's
inequality

$$
\begin{equation*}
\mathbb{P}[|X-\mathbb{E} X|>\beta] \leq \frac{\operatorname{Var}[X]}{\beta^{2}} \tag{2.1.5}
\end{equation*}
$$



Figure 2.1: Proof of Markov's inequality: taking expectations of the two functions depicted above yields the inequality.

Proof. This follows immediately by applying (2.1.4) to $|X-\mathbb{E} X|^{2}$ with $b=\beta^{2}$.

Of course this bound is non-trivial only when $\beta$ is larger than the standard deviation. Results of this type that quantify the probability of deviating from the mean are referred to as concentration inequalities. Chebyshev's inequality is perhaps the simplest instance-we will derive many more. To bound the variance, the following standard formula is sometimes useful

$$
\begin{equation*}
\operatorname{Var}\left[\sum_{i=1}^{n} X_{i}\right]=\sum_{i=1}^{n} \operatorname{Var}\left[X_{i}\right]+2 \sum_{i<j} \operatorname{Cov}\left[X_{i}, X_{j}\right], \tag{2.1.6}
\end{equation*}
$$

where recall that the covariance of $X_{i}$ and $X_{j}$ is

$$
\operatorname{Cov}\left[X_{i}, X_{j}\right]:=\mathbb{E}\left[X_{i} X_{j}\right]-\mathbb{E}\left[X_{i}\right] \mathbb{E}\left[X_{j}\right]
$$

When $X_{i}$ and $X_{j}$ are independent, then $\operatorname{Cov}\left[X_{i}, X_{j}\right]=0$.
Example 2.1.3. Let $X$ be a Gaussian random variable with mean $\mu$ and variance $\sigma^{2}$, that is, whose density is

$$
f_{X}(x)=\frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left(-\frac{(x-\mu)^{2}}{2 \sigma^{2}}\right), \quad x \in \mathbb{R}
$$



Figure 2.2: Comparison of Markov's and Chebyshev's inequalities: the squared deviation from the mean (solid) gives a better approximation of the indicator function (dotted) close to the mean than the absolute deviation (dashed).

We write $X \sim \mathrm{~N}\left(\mu, \sigma^{2}\right)$. A direct computation shows that $\mathbb{E}|X-\mu|=\sigma \sqrt{\frac{2}{\pi}}$. Hence Markov's inequality gives

$$
\mathbb{P}[|X-\mu| \geq b] \leq \frac{\mathbb{E}|X-\mu|}{b}=\sqrt{\frac{2}{\pi}} \cdot \frac{\sigma}{b},
$$

while Chebyshev's inequality (Theorem 2.1.2) gives

$$
\mathbb{P}[|X-\mu| \geq b] \leq\left(\frac{\sigma}{b}\right)^{2}
$$

Hence, for $b$ large enough, Chebyshev's inequality produces a stronger bound. See Figure 2.2 for some insight.

Example 2.1.4 (Coupon collector's problem). Let $\left(X_{t}\right)_{t \in \mathbb{N}}$ be i.i.d. uniform random variables over $[n]$, that is, that are equally likely to take any value in $[n]$. Let $T_{n, i}$ be the first time that $i$ elements of $[n]$ have been picked, that is,

$$
T_{n, i}=\inf \left\{t \geq 1:\left|\left\{X_{1}, \ldots, X_{t}\right\}\right|=i\right\}
$$

with $T_{n, 0}:=0$. We prove that the time it takes to pick all elements at least onceor "collect each coupon"-has the following tail. For any $\varepsilon>0$, we have as $n \rightarrow+\infty$ :

## Claim 2.1.5.

$$
\mathbb{P}\left[\left|T_{n, n}-n \sum_{j=1}^{n} j^{-1}\right| \geq \varepsilon n \log n\right] \rightarrow 0
$$

To prove this claim we note that the time elapsed between $T_{n, i-1}$ and $T_{n, i}$, which we denote by $\tau_{n, i}:=T_{n, i}-T_{n, i-1}$, is geometric with success probability $1-\frac{i-1}{n}$. And all $\tau_{n, i} \mathrm{~s}$ are independent. Recall that a geometric random variable $Z$ with success probability $p$ has probability mass function $\mathbb{P}[Z=z]=(1-p)^{z-1} p$ for $z \in \mathbb{N}$ and has mean $1 / p$ and variance $(1-p) / p^{2}$. So, the expectation and variance of $T_{n, n}=\sum_{i=1}^{n} \tau_{n, i}$ are

$$
\begin{equation*}
\mathbb{E}\left[T_{n, n}\right]=\sum_{i=1}^{n}\left(1-\frac{i-1}{n}\right)^{-1}=n \sum_{j=1}^{n} j^{-1}=\Theta(n \log n), \tag{2.1.7}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Var}\left[T_{n, n}\right] \leq \sum_{i=1}^{n}\left(1-\frac{i-1}{n}\right)^{-2}=n^{2} \sum_{j=1}^{n} j^{-2} \leq n^{2} \sum_{j=1}^{+\infty} j^{-2}=\Theta\left(n^{2}\right) . \tag{2.1.8}
\end{equation*}
$$

So by Chebyshev's inequality

$$
\begin{aligned}
\mathbb{P}\left[\left|T_{n, n}-n \sum_{j=1}^{n} j^{-1}\right| \geq \varepsilon n \log n\right] & \leq \frac{\operatorname{Var}\left[T_{n, n}\right]}{(\varepsilon n \log n)^{2}} \\
& \leq \frac{n^{2} \sum_{j=1}^{+\infty} j^{-2}}{(\varepsilon n \log n)^{2}} \\
& \rightarrow 0,
\end{aligned}
$$

by (2.1.7) and (2.1.8).
A classical implication of Chebyshev's inequality is (a version of) the law of large numbers. Recall that a sequence of random variables $\left(X_{n}\right)_{n \geq 1}$ converges in probability to a random variable $X$, denoted by $X_{n} \rightarrow_{\mathrm{p}} X$, if for all $\varepsilon>0$

$$
\lim _{n \rightarrow+\infty} \mathbb{P}\left[\left|X_{n}-X\right| \geq \varepsilon\right] \rightarrow 0
$$

Theorem 2.1.6 ( $L^{2}$ weak law of large numbers). Let $X_{1}, X_{2}, \ldots$ be uncorrelated random variables, that is, $\mathbb{E}\left[X_{i} X_{j}\right]=\mathbb{E}\left[X_{i}\right] \mathbb{E}\left[X_{j}\right]$ for $i \neq j$, with $\mathbb{E}\left[X_{i}\right]=\mu<$ $+\infty$ and $\sup _{i} \operatorname{Var}\left[X_{i}\right]<+\infty$. Then

$$
\frac{1}{n} \sum_{k \leq n} X_{k} \rightarrow_{\mathrm{p}} \mu .
$$

See Exercise 2.5 for a proof. When the $X_{k}$ s are i.i.d. and integrable (but not necessarily square integrable), convergence is almost sure. That result, the strong law of large numbers, also follows from Chebyshev's inequality (and other ideas), but we will not prove it here.

### 2.2 First moment method

In this section, we develop some techniques based on the first moment. Recall that the expectation of a random variable has an elementary, yet handy, property: linearity. That is, if random variables $X_{1}, \ldots, X_{k}$ defined on a joint probability space have finite first moments, then

$$
\begin{equation*}
\mathbb{E}\left[X_{1}+\cdots+X_{k}\right]=\mathbb{E}\left[X_{1}\right]+\cdots+\mathbb{E}\left[X_{k}\right], \tag{2.2.1}
\end{equation*}
$$

without any further assumption. In particular linearity holds whether or not the $X_{i} \mathrm{~s}$ are independent.

### 2.2.1 The probabilistic method

A key technique of probabilistic combinatorics is the so-called probabilistic method. The idea is that one can establish the existence of an object satisfying a certain property-without having to construct one explicitly. Instead one argues that a randomly chosen object exhibits the given property with positive probability. The following "obvious" observation, sometimes referred to as the first moment principle, plays a key role in this context.
Theorem 2.2.1 (First moment principle). Let $X$ be a random variable with finite expectation. Then, for any $\mu \in \mathbb{R}$,

$$
\mathbb{E} X \leq \mu \Longrightarrow \mathbb{P}[X \leq \mu]>0
$$

Proof. We argue by contradiction. Assume $\mathbb{E} X \leq \mu$ and $\mathbb{P}[X \leq \mu]=0$. Write $\{X \leq \mu\}=\bigcap_{n \geq 1}\{X<\mu+1 / n\}$. That implies by monotonicity (see Lemma B.2.6) that, for any $\varepsilon \in(0,1)$, it holds that $\mathbb{P}[X<\mu+1 / n]<\varepsilon$ for $n$ large enough. Hence, because we assume that $\mathbb{P}[X \leq \mu]=0$,

$$
\begin{aligned}
\mu & \geq \mathbb{E} X \\
& =\mathbb{E}[X ; X<\mu+1 / n]+\mathbb{E}[X ; X \geq \mu+1 / n] \\
& \geq \mu \mathbb{P}[X<\mu+1 / n]+(\mu+1 / n)(1-\mathbb{P}[X<\mu+1 / n]) \\
& =\mu+n^{-1}(1-\mathbb{P}[X<\mu+1 / n]) \\
& >\mu+n^{-1}(1-\varepsilon) \\
& >\mu,
\end{aligned}
$$

a contradiction.
The power of this principle is easier to appreciate on an example.
Example 2.2.2 (Balancing vectors). Let $\boldsymbol{v}_{1}, \ldots, \boldsymbol{v}_{n}$ be arbitrary unit vectors in $\mathbb{R}^{n}$. How small can we make the 2 -norm of the linear combination

$$
x_{1} \boldsymbol{v}_{1}+\cdots+x_{n} \boldsymbol{v}_{n}
$$

by appropriately choosing $x_{1}, \ldots, x_{n} \in\{-1,+1\}$ ? We claim that it can be as small as $\sqrt{n}$, for any collection of $\boldsymbol{v}_{i} \mathrm{~s}$. At first sight, this may appear to be a complicated geometry problem. But the proof is trivial once one thinks of choosing the $x_{i} \mathrm{~s}$ at random. Let $X_{1}, \ldots, X_{n}$ be independent random variables uniformly distributed in $\{-1,+1\}$. Then, since $\mathbb{E}\left[X_{i} X_{j}\right]=\mathbb{E}\left[X_{i}\right] \mathbb{E}\left[X_{j}\right]=0$ for all $i \neq j$ but $\mathbb{E}\left[X_{i}^{2}\right]=1$ for all $i$,

$$
\begin{aligned}
\mathbb{E}\left\|X_{1} \boldsymbol{v}_{1}+\cdots+X_{n} \boldsymbol{v}_{n}\right\|_{2}^{2} & =\mathbb{E}\left[\sum_{i, j} X_{i} X_{j}\left\langle\boldsymbol{v}_{i}, \boldsymbol{v}_{j}\right\rangle\right] \\
& =\sum_{i, j} \mathbb{E}\left[X_{i} X_{j}\left\langle\boldsymbol{v}_{i}, \boldsymbol{v}_{j}\right\rangle\right] \\
& =\sum_{i, j}\left\langle\boldsymbol{v}_{i}, \boldsymbol{v}_{j}\right\rangle \mathbb{E}\left[X_{i} X_{j}\right] \\
& =\sum_{i}\left\|\boldsymbol{v}_{i}\right\|_{2}^{2} \\
& =n,
\end{aligned}
$$

where we used the linearity of expectation on the second line. Hence the random variable $Z=\left\|X_{1} \boldsymbol{v}_{1}+\cdots+X_{n} \boldsymbol{v}_{n}\right\|^{2}$ has expectation $\mathbb{E} Z=n$ and must take a value $\leq n$ with positive probability by the first moment principle (Theorem 2.2.1). In other words, there must be a choice of $X_{i} \mathrm{~s}$ such that $Z \leq n$. That proves the claim.

Here is a slightly more subtle example of the probabilistic method, where one has to modify the original random choice.

Example 2.2.3 (Independent sets). For $d \in \mathbb{N}$, let $G=(V, E)$ be a $d$-regular graph with $n$ vertices. Such a graph necessarily has $m=n d / 2$ edges. Our goal is derive a lower bound on the size, $\alpha(G)$, of the largest independent set in $G$. Recall that an independent set is a set of vertices in a graph, no two of which are adjacent. Again, at first sight, this may seem like a rather complicated graph-theoretic problem. But an appropriate random choice gives a non-trivial bound. Specifically:

## Claim 2.2.4.

$$
\alpha(G) \geq \frac{n}{2 d} .
$$

Proof. The proof proceeds in two steps:

1. We first prove the existence of a subset $S$ of vertices with relatively few edges.
2. We remove vertices from $S$ to obtain an independent set.

Step 1. Let $0<p<1$ to be chosen below. To form the set $S$, pick each vertex in $V$ independently with probability $p$. Letting $X$ be the number of vertices in $S$, we have by the linearity of expectation that

$$
\mathbb{E} X=\mathbb{E}\left[\sum_{v \in V} \mathbf{1}_{v \in S}\right]=n p,
$$

where we used that $\mathbb{E}\left[\mathbf{1}_{v \in S}\right]=p$. Letting $Y$ be the number of edges between vertices in $S$, we have by the linearity of expectation that

$$
\mathbb{E} Y=\mathbb{E}\left[\sum_{\{i, j\} \in E} \mathbf{1}_{i \in S} \mathbf{1}_{j \in S}\right]=\frac{n d}{2} p^{2},
$$

where we also used that $\mathbb{E}\left[\mathbf{1}_{i \in S} \mathbf{1}_{j \in S}\right]=p^{2}$ by independence. Hence, subtracting,

$$
\mathbb{E}[X-Y]=n p-\frac{n d}{2} p^{2},
$$

which, as a function of $p$, is maximized at $p=1 / d$ where it takes the value $n /(2 d)$. As a result, by the first moment principle applied to $X-Y$, there must exist a set $S$ of vertices in $G$ such that

$$
\begin{equation*}
|S|-|\{\{i, j\} \in E: i, j \in S\}| \geq \frac{n}{2 d} . \tag{2.2.2}
\end{equation*}
$$

Step 2. For each edge $e$ connecting two vertices in $S$, remove one of the endvertices of $e$. By construction, the remaining set of vertices: (i) forms an independent set; and (ii) has a size larger or equal than the left hand side of (2.2.2). That inequality implies the claim.

Note that a graph $G$ made of $n /(d+1)$ cliques of size $d+1$ (with no edge between the cliques) has $\alpha(G)=n /(d+1)$, showing that our bound is tight up to a constant. This is known as a Turán graph.

Remark 2.2.5. The previous result can be strengthened to

$$
\alpha(G) \geq \sum_{v \in V} \frac{1}{\delta(v)+1}
$$

for a general graph $G=(V, E)$, where $\delta(v)$ is the degree of $v$. This bound is achieved for Turán graphs. See, for example, [AS11, The probabilistic lens: Turán's theorem].

The previous example also illustrates the important indicator trick, that is, writing a random variable as a sum of indicators, which is naturally used in combination with the linearity of expectation.

### 2.2.2 Boole's inequality

One implication of the first moment principle (Theorem 2.2.1) is that: if a nonnegative, integer-valued random variable $X$ has expectation strictly smaller than 1 , then its value is 0 with positive probability. The following application of Markov's inequality (Theorem 2.1.1) adds a quantitative twist: if that same $X$ has a "small" expectation, then its value is 0 with "large" probability.

Theorem 2.2.6 (First moment method). If $X$ is a non-negative, integer-valued random variable, then

$$
\begin{equation*}
\mathbb{P}[X>0] \leq \mathbb{E} X \tag{2.2.3}
\end{equation*}
$$

Proof. Take $b=1$ in Markov's inequality.
This simple fact is typically used in the following manner: one wants to show that a certain "bad event" does not occur with probability approaching 1 ; the random variable $X$ then counts the number of such "bad events." In that case, $X$ is a sum of indicators and Theorem 2.2.6 reduces simply to the standard union bound, also known as Boole's inequality. We record one useful version of this setting in the next corollary.

Corollary 2.2.7. Let $B_{n}=A_{n, 1} \cup \cdots \cup A_{n, m_{n}}$, where $A_{n, 1}, \ldots, A_{n, m_{n}}$ is a collection of events for each $n$. Then, letting

$$
\mu_{n}:=\sum_{i=1}^{m_{n}} \mathbb{P}\left[A_{n, i}\right]
$$

we have

$$
\mathbb{P}\left[B_{n}\right] \leq \mu_{n}
$$

In particular, if $\mu_{n} \rightarrow 0$ as $n \rightarrow+\infty$, then $\mathbb{P}\left[B_{n}\right] \rightarrow 0$.

Proof. Take $X:=X_{n}=\sum_{i=1}^{m_{n}} \mathbf{1}_{A_{n, i}}$ in Theorem 2.2.6.
A useful generalization of the union bound is given in Exercise 2.2. We will refer to applications of Theorem 2.2.6 as the first moment method.

Example 2.2.8 (Random $k$-SAT threshold). For $r \in \mathbb{R}_{+}$, let $\Phi_{n, r}:\{0,1\}^{n} \rightarrow$ $\{0,1\}$ be a random $k$-CNF formula on $n$ Boolean variables $z_{1}, \ldots, z_{n}$ with $\lceil r n\rceil$
moment
method clauses. That is, $\Phi_{n, r}$ is an AND of $\lceil r n\rceil$ ORs, each obtained by picking independently $k$ literals uniformly at random (with replacement). Recall that a literal is a variable $z_{i}$ or its negation $\bar{z}_{i}$. The formula $\Phi_{n, r}$ is said to be satisfiable if there exists an assignment $z=\left(z_{1}, \ldots, z_{n}\right)$ such that $\Phi_{n, r}(z)=1$. Clearly the higher the value of $r$, the less likely it is for $\Phi_{n, r}$ to be satisfiable. In fact it is natural to conjecture that a sharp transition takes place, that is, that there exists an $r_{k}^{*} \in \mathbb{R}_{+}$ (depending on $k$ but not on $n$ ) such that

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left[\Phi_{n, r} \text { is satisfiable }\right]=\left\{\begin{array}{l}
0, \text { if } r>r_{k}^{*}  \tag{2.2.4}\\
1, \text { if } r<r_{k}^{*}
\end{array}\right.
$$

Studying such threshold phenomena is a major theme of modern discrete probability. Using the first moment method (Theorem 2.2.6), we give an upper bound on the threshold. Formally:

Claim 2.2.9.

$$
r>2^{k} \log 2 \Longrightarrow \limsup _{n \rightarrow \infty} \mathbb{P}\left[\Phi_{n, r} \text { is satisfiable }\right]=0
$$

Proof. How to start the proof should be obvious: let $X_{n}$ be the number of satisfying assignments of $\Phi_{n, r}$. Applying the first moment method, since

$$
\mathbb{P}\left[\Phi_{n, r} \text { is satisfiable }\right]=\mathbb{P}\left[X_{n}>0\right],
$$

it suffices to show that $\mathbb{E} X_{n} \rightarrow 0$. To compute $\mathbb{E} X_{n}$, we use the indicator trick

$$
X_{n}=\sum_{z \in\{0,1\}^{n}} \mathbf{1}_{\left\{z \text { satisfies } \Phi_{n, r}\right\}} .
$$

There are $2^{n}$ possible assignments. Each fixed assignment satisfies the random choice of clauses $\Phi_{n, r}$ with probability $\left(1-2^{-k}\right)^{\lceil r n\rceil}$. Indeed note that the $r n$ clauses are picked independently and each clause literal picked is satisfied with probability $1 / 2$. Therefore, by the assumption on $r$, for $\varepsilon>0$ small enough and $n$
large enough

$$
\begin{aligned}
\mathbb{E} X_{n} & =2^{n}\left(1-2^{-k}\right)^{[r n\rceil} \\
& \leq 2^{n}\left(1-2^{-k}\right)^{\left(2^{k} \log 2\right)(1+\varepsilon) n} \\
& \leq 2^{n} e^{-(\log 2)(1+\varepsilon) n} \\
& =2^{-\varepsilon n} \\
& \rightarrow 0,
\end{aligned}
$$

where we used that $(1-1 / \ell)^{\ell} \leq e^{-1}$ for all $\ell \in \mathbb{N}$ (see Exercise 1.16). Theorem 2.2.6 implies the claim.

Remark 2.2.10. Bounds in the other direction are also known. For instance, for $k \geq 3$, it has been shown that if $r<2^{k} \log 2-k$

$$
\liminf _{n \rightarrow \infty} \mathbb{P}\left[\Phi_{n, r} \text { is satisfiable }\right]=1
$$

See [ANP05]. For the $k=2$ case, it is known that (2.2.4) in fact holds with $r_{2}^{*}=1$ [CR92]. A breakthrough of [DSS22] also establishes (2.2.4) for large $k$; the threshold $r_{k}^{*}$ is characterized as the root of a certain equation coming from statistical physics.

### 2.2.3 $\triangleright$ Random permutations: longest increasing subsequence

In this section, we bound the expected length of a longest increasing subsequence in a random permutation. Let $\sigma_{n}=\left(\sigma_{n}(1), \ldots, \sigma_{n}(n)\right)$ be a uniformly random permutation of $[n]:=\{1, \ldots, n\}$ (i.e., a bijection of $[n]$ to itself chosen uniformly at random among all such mappings) and let $L_{n}$ be the length of a longest increasing subsequence of $\sigma_{n}$ (i.e., a sequence of indices $i_{1}<\cdots<i_{k}$ such that $\left.\sigma_{n}\left(i_{1}\right)<\cdots<\sigma_{n}\left(i_{k}\right)\right)$.

## Claim 2.2.11.

$$
\mathbb{E} L_{n}=\Theta(\sqrt{n})
$$

Proof. We first prove that

$$
\begin{equation*}
\limsup _{n \rightarrow \infty} \frac{\mathbb{E} L_{n}}{\sqrt{n}} \leq e \tag{2.2.5}
\end{equation*}
$$

which implies half of the claim. Bounding the expectation of $L_{n}$ is not straightforward as it is the expectation of a maximum. A natural way to proceed is to find a value $\ell$ for which $\mathbb{P}\left[L_{n} \geq \ell\right]$ is "small." More formally, we bound the expectation as follows

$$
\begin{equation*}
\mathbb{E} L_{n} \leq \ell \mathbb{P}\left[L_{n}<\ell\right]+n \mathbb{P}\left[L_{n} \geq \ell\right] \leq \ell+n \mathbb{P}\left[L_{n} \geq \ell\right], \tag{2.2.6}
\end{equation*}
$$

for an $\ell$ chosen below. To bound the probability on the right-hand side, we appeal to the first moment method (Theorem 2.2.6) by letting $X_{n}$ be the number of increasing subsequences of length $\ell$. We also use the indicator trick, that is, we think of $X_{n}$ as a sum of indicators over subsequences (not necessarily increasing) of length $\ell$.

There are $\binom{n}{\ell}$ such subsequences, each of which is increasing with probability $1 / \ell$ !. Note that these subsequences are not independent. Nevertheless, by the linearity of expectation and the first moment method,

$$
\mathbb{P}\left[L_{n} \geq \ell\right]=\mathbb{P}\left[X_{n}>0\right] \leq \mathbb{E} X_{n}=\frac{1}{\ell!}\binom{n}{\ell} \leq \frac{n^{\ell}}{(\ell!)^{2}} \leq \frac{n^{\ell}}{e^{2}[\ell / e]^{2 \ell}} \leq\left(\frac{e \sqrt{n}}{\ell}\right)^{2 \ell}
$$

where we used a standard bound on factorials recalled in Appendix A. Note that, in order for this bound to go to 0 , we need $\ell>e \sqrt{n}$. Then (2.2.5) follows by taking $\ell=(1+\delta) e \sqrt{n}$ in (2.2.6), for an arbitrarily small $\delta>0$.

For the other half of the claim, we show that

$$
\frac{\mathbb{E} L_{n}}{\sqrt{n}} \geq 1
$$

This part does not rely on the first moment method (and may be skipped). We seek a lower bound on the expected length of a longest increasing subsequence. The proof uses the following two ideas. First observe that there is a natural symmetry between the lengths of the longest increasing and decreasing subsequencesthey are identically distributed. Moreover if a permutation has a "short" longest increasing subsequence, then intuitively it must have a "long" decreasing subsequence, and vice versa. Combining these two observations gives a lower bound on the expectation of $L_{n}$. Formally, let $D_{n}$ be the length of a longest decreasing subsequence. By symmetry and the arithmetic mean-geometric mean inequality, note that

$$
\mathbb{E} L_{n}=\mathbb{E}\left[\frac{L_{n}+D_{n}}{2}\right] \geq \mathbb{E} \sqrt{L_{n} D_{n}}
$$

We show that $L_{n} D_{n} \geq n$, which proves the claim. Let $L_{n}^{(k)}$ be the length of a longest increasing subsequence ending at position $k$, and similarly for $D_{n}^{(k)}$. It suffices to show that the pairs $\left(L_{n}^{(k)}, D_{n}^{(k)}\right), 1 \leq k \leq n$, are distinct. Indeed, noting that $L_{n}^{(k)} \leq L_{n}$ and $D_{n}^{(k)} \leq D_{n}$, the number of pairs in $\left[L_{n}\right] \times\left[D_{n}\right]$ is at most $L_{n} D_{n}$ which must then be at least $n$.

Let $1 \leq j<k \leq n$. If $\sigma_{n}(k)>\sigma_{n}(j)$ then we see that $L_{n}^{(k)}>L_{n}^{(j)}$ by appending $\sigma_{n}(k)$ to the subsequence ending at position $j$ achieving $L_{n}^{(j)}$. If the opposite holds, then we have instead $D_{n}^{(k)}>D_{n}^{(j)}$. Either way, $\left(L_{n}^{(j)}, D_{n}^{(j)}\right)$ and $\left(L_{n}^{(k)}, D_{n}^{(k)}\right)$ must be distinct. This clever combinatorial argument is known as the Erdös-Szekeres Theorem. That concludes the proof of the second claim.

Remark 2.2.12. It has been shown that in fact

$$
\mathbb{E} L_{n}=2 \sqrt{n}+c n^{1 / 6}+o\left(n^{1 / 6}\right),
$$

as $n \rightarrow+\infty$, where $c=-1.77$... [BDJ99].

### 2.2.4 $\triangleright$ Percolation: existence of a non-trivial critical value on $\mathbb{Z}^{2}$

In this section we use the first moment method (Theorem 2.2.6) to prove the existence of a non-trivial threshold in bond percolation on the two-dimensional lattice. We begin with some background.

Critical value in bond percolation Consider bond percolation (Definition 1.2.1) on the two-dimensional lattice $\mathbb{L}^{2}$ (see Section 1.1.1) with density $p$. Let $\mathbb{P}_{p}$ denote the corresponding probability measure. Recall that paths are "self-avoiding" by definition (see Section 1.1.1). We say that a path is open if all edges in the induced subgraph are open. Writing $x \Leftrightarrow y$ if $x, y \in \mathbb{L}^{2}$ are connected by an open path, recall that the open cluster of $x$ is

$$
\mathcal{C}_{x}:=\left\{y \in \mathbb{Z}^{2}: x \Leftrightarrow y\right\} .
$$

The percolation function is defined as

$$
\theta(p):=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right],
$$

that is, $\theta(p)$ is the probability that the origin is connected by open paths to infinitely many vertices. It is intuitively clear that the function $\theta(p)$ is non-decreasing. Indeed consider the following alternative representation of the percolation process: to each edge $e$, assign a uniform $[0,1]$ random variable $U_{e}$ and declare the edge open if $U_{e} \leq p$. Using the same $U_{e} \mathrm{~s}$ for densities $p_{1}<p_{2}$, it follows immediately from the monotonicity of the construction that $\theta\left(p_{1}\right) \leq \theta\left(p_{2}\right)$. (We will have much more to say about this type of "coupling" argument in Chapter 4.) Moreover note that $\theta(0)=0$ and $\theta(1)=1$. The critical value is defined as
percolation
function

$$
p_{c}\left(\mathbb{L}^{2}\right):=\sup \{p \geq 0: \theta(p)=0\},
$$

the point at which the probability that the origin is contained in an infinite open cluster becomes positive. Note that by a union bound over all vertices, when $\theta(p)=0$, we have that $\mathbb{P}_{p}\left[\exists x,\left|\mathcal{C}_{x}\right|=+\infty\right]=0$. Conversely, because $\left\{\exists x,\left|\mathcal{C}_{x}\right|=\right.$ $+\infty\}$ is a tail event (see Definition B.3.9) for any enumeration of the edges, by Kolmogorov's 0 -1 law (Theorem B.3.11) it holds that $\mathbb{P}_{p}\left[\exists x,\left|\mathcal{C}_{x}\right|=+\infty\right]=1$ when $\theta(p)>0$.

Using the first moment method we show that the critical value is non-trivial, that is, it is strictly between 0 and 1 . This is a different example of a threshold phenomenon.

## Claim 2.2.13.

$$
p_{c}\left(\mathbb{L}^{2}\right) \in(0,1) .
$$

Proof. We first show that, for any $p<1 / 3, \theta(p)=0$. In order to apply the first moment method, roughly speaking, we need to reduce the problem to counting the number of instances of an appropriately chosen substructure. The key observation is the following:

An infinite $\mathcal{C}_{0}$ contains an open path starting at 0 of infinite length and, as a result, of all lengths.

Hence, we let $X_{n}$ be the number of open paths of length $n$ starting at 0 . Then, by monotonicity,

$$
\begin{equation*}
\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right] \leq \mathbb{P}_{p}\left[\cap_{n}\left\{X_{n}>0\right\}\right]=\lim _{n} \mathbb{P}_{p}\left[X_{n}>0\right] \leq \limsup _{n} \mathbb{E}_{p}\left[X_{n}\right], \tag{2.2.7}
\end{equation*}
$$

where the last inequality follows from Theorem 2.2.6. We bound the number of paths of length $n$ (each of which is open with probability $p^{n}$ ) by noting that they cannot backtrack. That gives 4 choices at the first step, and at most 3 choices at each subsequent step. Hence, we get the following bound

$$
\mathbb{E}_{p} X_{n} \leq 4\left(3^{n-1}\right) p^{n} .
$$

The right-hand side goes to 0 for all $p<1 / 3$. When combined with (2.2.7), that proves half of the claim:

$$
p_{c}\left(\mathbb{L}^{2}\right)>0 .
$$

For the other direction, we show that $\theta(p)>0$ for $p$ close enough to 1 . This time, we count "dual cycles." This type of proof is known as a contour argument, or Peierls' argument, and is based on the following construction. Consider the dual lattice $\widetilde{\mathbb{L}}^{2}$ whose vertices are $\mathbb{Z}^{2}+(1 / 2,1 / 2)$ and whose edges connect vertices $u, v$ with $\|u-v\|_{1}=1$. See Figure 2.3. Note that each edge in the primal lattice $\mathbb{L}^{2}$ has a unique corresponding edge in the dual lattice which crosses it perpendicularly. We make the same assignment, open or closed, for corresponding primal and dual edges. The following graph-theoretic lemma, whose proof is sketched below, forms the basis of contour arguments. Recall that cycles are "self-avoiding" by definition (see Section 1.1.1). We say that a cycle is closed if all edges in the induced subgraph are closed, that is, are not open.


Figure 2.3: Primal (solid) and dual (dotted) lattices.

Lemma 2.2.14 (Contour lemma). If $\left|\mathcal{C}_{0}\right|<+\infty$, then there is a closed cycle contour lemma around the origin in the dual lattice $\mathbb{L}^{2}$.
To prove that $\theta(p)>0$ for $p$ close enough to 1 , the idea is to apply the first moment method to $Z_{n}$ equal to the number of closed dual cycles of length $n$ surrounding the origin. We bound from above the number of dual cycles of length $n$ around the origin by the number of choices for the starting edge across the upper $y$-axis and for each $n-1$ subsequent non-backtracking choices. Namely,

$$
\begin{aligned}
\mathbb{P}\left[\left|\mathcal{C}_{0}\right|<+\infty\right] & \leq \mathbb{P}\left[\exists n \geq 4, Z_{n}>0\right] \\
& \leq \sum_{n \geq 4} \mathbb{P}\left[Z_{n}>0\right] \\
& \leq \sum_{n \geq 4} \mathbb{E} Z_{n} \\
& \leq \sum_{n \geq 4} \frac{n}{2} 3^{n-1}(1-p)^{n} \\
& =\frac{3^{3}(1-p)^{4}}{2} \sum_{m \geq 1}(m+3)(3(1-p))^{m-1} \\
& =\frac{3^{3}(1-p)^{4}}{2}\left(\frac{1}{(1-3(1-p))^{2}}+3 \frac{1}{1-3(1-p)}\right)
\end{aligned}
$$

when $p>2 / 3$, where the first term in parentheses on the last line comes from differentiating with respect to $q$ the geometric series $\sum_{m \geq 0} q^{m}$ and setting $q:=1-$ $p$. The expression on the last line can be made smaller than 1 if we let $p$ approach 1. We have shown that $\theta(p)>0$ for $p$ close enough to 1 , and that concludes the proof. (Exercise 2.3 sketches a proof that $\theta(p)>0$ for all $p>2 / 3$.)

It is straightforward to extend the claim to $\mathbb{L}^{d}$. (Exercise 2.4 asks for the details.)

Proof of the contour lemma We conclude this section by sketching the proof of the contour lemma, which relies on topological arguments.

Proof of Lemma 2.2.14. Assume $\left|\mathcal{C}_{0}\right|<+\infty$. Imagine identifying each vertex in $\mathbb{L}^{2}$ with a square of side 1 centered around it so that the sides line up with dual edges. Paint green the squares of vertices in $\mathcal{C}_{0}$. Paint red the squares of vertices in $\mathcal{C}_{0}^{c}$ which share a side with a green square. Leave the other squares white. Let $u_{0}$ be a highest vertex in $\mathcal{C}_{0}$ along the $y$-axis and let $v_{0}$ and $v_{1}$ be the dual vertices corresponding to the upper left and right corners respectively of the square of $u_{0}$. Because $u_{0}$ is highest, it must be that the square above it is red. Walk along the dual edge $\left\{v_{0}, v_{1}\right\}$ separating the squares of $u_{0}$ and $u_{0}+(0,1)$ from $v_{0}$ to $v_{1}$. Notice that this edge satisfies what we call the red-green property: as you traverse it from $v_{0}$ to $v_{1}$, a red square sits on your left and a green square is on your right. Proceed further by iteratively walking along an incident dual edge with the following rule. Choose an edge satisfying the red-green property, with the edges to your left, straight ahead, and to your right in decreasing order of priority. Stop when a previously visited dual vertex is reached. The claim is that this procedure constructs the desired cycle. Let $v_{0}, v_{1}, v_{2}, \ldots$ be the dual vertices visited. By construction $\left\{v_{i-1}, v_{i}\right\}$ is a dual edge for all $i$.

- A dual cycle is produced. We first argue that this procedure cannot get stuck. Let $\left\{v_{i-1}, v_{i}\right\}$ be the edge just crossed and assume that it has the red-green property. If there is a green square to the left ahead, then the edge to the left, which has highest priority, has the red-green property. If the left square ahead is not green, but the right one is, then the left square must in fact be red by construction (i.e., it cannot be white). In that case, the edge straight ahead has the red-green property. Finally, if neither square ahead is green, then the right square must in fact be red because the square behind to the right is green by assumption. That implies that the edge to the right has the red-green property. Hence we have shown that the procedure does not get stuck. Moreover, because by assumption the number of green squares
is finite, this procedure must eventually terminate when a previously visited dual vertex is reached, forming a cycle (of length at least 4).
- The origin lies within the cycle. The inside of a cycle in the plane is welldefined by the Jordan curve theorem. So the dual cycle produced above has its adjacent green squares either on the inside (negative orientation) or on the outside (positive orientation). In the former case the origin must lie inside the cycle as otherwise the vertices corresponding to the green squares on the inside would not be in $\mathcal{C}_{0}$, as they could not be connected to the origin with open paths.
So it remains to consider the latter case, where through a similar reasoning the origin must lie outside the cycle. Let $v_{j}$ be the repeated dual vertex. Assume first that $v_{j} \neq v_{0}$ and let $v_{j-1}$ and $v_{j+1}$ be the dual vertices preceding and following $v_{j}$ during the first visit to $v_{j}$. Let $v_{k}$ be the dual vertex preceding $v_{j}$ on the second visit. After traversing the edge from $v_{j-1}$ to $v_{j}$, $v_{k}$ cannot be to the left or to the right because in those cases the red-green properties of the two corresponding edges (i.e., $\left\{v_{j-1}, v_{j}\right\}$ and $\left\{v_{k}, v_{j}\right\}$ ) are not compatible. So $v_{k}$ is straight ahead and, by the priority rules, $v_{j+1}$ must be to the left upon entering $v_{j}$ the first time. But in that case, for the origin to lie outside the cycle as we are assuming and for the cycle to avoid the path $v_{0}, \ldots, v_{j-1}$, we must traverse the cycle with a negative orientation, that is, the green squares adjacent to the cycle must be on the inside, a contradiction.
So, finally, assume $v_{0}$ is the repeated vertex. If the cycle is traversed with a positive orientation and the origin is on the outside, it must be that the cycle crosses the $y$-axis at least once above $u_{0}+(0,1)$, again a contradiction.
Hence we have shown that the origin is inside the cycle.


## That concludes the proof.

Remark 2.2.15. It turns out that $p_{c}\left(\mathbb{L}^{2}\right)=1 / 2$. We will prove $p_{c}\left(\mathbb{L}^{2}\right) \geq 1 / 2$, known as Harris' Theorem, in Section 4.2.5. The other direction is due to Kesten [Kes80].

### 2.3 Second moment method

The first moment method (Theorem 2.2.6) gives an upper bound on the probability that a non-negative, integer-valued random variable is positive-which is nontrivial provided its expectation is small enough. In this section we seek a lower bound on that probability. We first note that a large expectation does not suffice in general. Say $X_{n}$ is $n^{2}$ with probability $1 / n$, and 0 otherwise. Then $\mathbb{E} X_{n}=n \rightarrow+\infty$, yet


Figure 2.4: Second moment method: if the standard deviation $\sigma_{X}$ of $X$ is less than its expectation $\mu_{X}$, then the probability that $X$ is 0 is bounded away from 1 .
$\mathbb{P}\left[X_{n}>0\right] \rightarrow 0$. That is, although the expectation diverges, the probability that $X_{n}$ is positive can be arbitrarily small.

So we turn to the second moment. Intuitively the basis for the so-called second moment method is that, if the expectation of $X_{n}$ is large and its variance is relatively small, then we can bound the probability that $X_{n}$ is close to 0 . As we will see in applications, the first and second moment methods often work hand in hand.

### 2.3.1 Paley-Zygmund inequality

As an immediate corollary of Chebyshev's inequality (Theorem 2.1.2), we get a first version of the second moment method: if the standard deviation of $X$ is less than its expectation, then the probability that $X$ is 0 is bounded away from 1 . See Figure 2.4. Formally, let $X$ be a nonnegative random variable (not identically zero). Then

$$
\begin{equation*}
\mathbb{P}[X>0] \geq 1-\frac{\operatorname{Var}[X]}{(\mathbb{E} X)^{2}} \tag{2.3.1}
\end{equation*}
$$

Indeed, by (2.1.5),

$$
\mathbb{P}[X=0] \leq \mathbb{P}[|X-\mathbb{E} X| \geq \mathbb{E} X] \leq \frac{\operatorname{Var}[X]}{(\mathbb{E} X)^{2}}
$$

The following tail bound, a simple application of Cauchy-Schwarz (Theorem B.4.8), leads to an improved version of this inequality.

Theorem 2.3.1 (Paley-Zygmund inequality). Let $X$ be a nonnegative random variable. For all $0<\theta<1$,

Paley-Zygmund inequality

$$
\begin{equation*}
\mathbb{P}[X \geq \theta \mathbb{E} X] \geq(1-\theta)^{2} \frac{(\mathbb{E} X)^{2}}{\mathbb{E}\left[X^{2}\right]} \tag{2.3.2}
\end{equation*}
$$

Proof. We have

$$
\begin{aligned}
\mathbb{E} X & =\mathbb{E}\left[X \mathbf{1}_{\{X<\theta \mathbb{E} X\}}\right]+\mathbb{E}\left[X \mathbf{1}_{\{X \geq \theta \mathbb{E} X\}}\right] \\
& \leq \theta \mathbb{E} X+\sqrt{\mathbb{E}\left[X^{2}\right] \mathbb{P}[X \geq \theta \mathbb{E} X]},
\end{aligned}
$$

where we used Cauchy-Schwarz. Rearranging gives the result.
As an immediate application:
Theorem 2.3.2 (Second moment method). Let $X$ be a nonnegative random variable (not identically zero). Then

$$
\begin{equation*}
\mathbb{P}[X>0] \geq \frac{(\mathbb{E} X)^{2}}{\mathbb{E}\left[X^{2}\right]} \tag{2.3.3}
\end{equation*}
$$

Proof. Take $\theta \downarrow 0$ in (2.3.2).
Since

$$
\frac{(\mathbb{E} X)^{2}}{\mathbb{E}\left[X^{2}\right]}=1-\frac{\operatorname{Var}[X]}{(\mathbb{E} X)^{2}+\operatorname{Var}[X]},
$$

we see that (2.3.3) is stronger than (2.3.1).
We typically apply the second moment method to a sequence of random variables $\left(X_{n}\right)$. The previous theorem gives a uniform lower bound on the probability that $\left\{X_{n}>0\right\}$ when $\mathbb{E}\left[X_{n}^{2}\right] \leq C \mathbb{E}\left[X_{n}\right]^{2}$ for some $C>0$. Just like the first moment method, the second moment method is often applied to a sum of indicators (but see Section 2.3.3 for a weighted case). We record in the next corollary a convenient version of the method.

Corollary 2.3.3. Let $B_{n}=A_{n, 1} \cup \cdots \cup A_{n, m_{n}}$, where $A_{n, 1}, \ldots, A_{n, m_{n}}$ is a collection of events for each $n$. Write $i \stackrel{n}{\sim} j$ if $i \neq j$ and $A_{n, i}$ and $A_{n, j}$ are not independent. Then, letting

$$
\mu_{n}:=\sum_{i=1}^{m_{n}} \mathbb{P}\left[A_{n, i}\right], \quad \gamma_{n}:=\sum_{i \sim j} \mathbb{P}\left[A_{n, i} \cap A_{n, j}\right],
$$

where the second sum is over ordered pairs, we have $\lim _{n} \mathbb{P}\left[B_{n}\right]>0$ whenever $\mu_{n} \rightarrow+\infty$ and $\gamma_{n} \leq C \mu_{n}^{2}$ for some $C>0$. If moreover $\gamma_{n}=o\left(\mu_{n}^{2}\right)$ then $\lim _{n} \mathbb{P}\left[B_{n}\right]=1$.

Proof. We apply the second moment method to $X_{n}:=\sum_{i=1}^{m_{n}} \mathbf{1}_{A_{n, i}}$ so that $B_{n}=$ $\left\{X_{n}>0\right\}$. Note that

$$
\operatorname{Var}\left[X_{n}\right]=\sum_{i} \operatorname{Var}\left[\mathbf{1}_{A_{n, i}}\right]+\sum_{i \neq j} \operatorname{Cov}\left[\mathbf{1}_{A_{n, i}}, \mathbf{1}_{A_{n, j}}\right],
$$

where

$$
\operatorname{Var}\left[\mathbf{1}_{A_{n, i}}\right]=\mathbb{E}\left[\left(\mathbf{1}_{A_{n, i}}\right)^{2}\right]-\left(\mathbb{E}\left[\mathbf{1}_{A_{n, i}}\right]\right)^{2} \leq \mathbb{P}\left[A_{n, i}\right],
$$

and, if $A_{n, i}$ and $A_{n, j}$ are independent,

$$
\operatorname{Cov}\left[\mathbf{1}_{A_{n, i}}, \mathbf{1}_{A_{n, j}}\right]=0,
$$

whereas, if $i \stackrel{n}{\sim} j$,

$$
\operatorname{Cov}\left[\mathbf{1}_{A_{n, i}}, \mathbf{1}_{A_{n, j}}\right]=\mathbb{E}\left[\mathbf{1}_{A_{n, i}} \mathbf{1}_{A_{n, j}}\right]-\mathbb{E}\left[\mathbf{1}_{A_{n, i}}\right] \mathbb{E}\left[\mathbf{1}_{A_{n, j}}\right] \leq \mathbb{P}\left[A_{n, i} \cap A_{n, j}\right]
$$

Hence

$$
\frac{\operatorname{Var}\left[X_{n}\right]}{\left(\mathbb{E} X_{n}\right)^{2}} \leq \frac{\mu_{n}+\gamma_{n}}{\mu_{n}^{2}}=\frac{1}{\mu_{n}}+\frac{\gamma_{n}}{\mu_{n}^{2}}
$$

Noting

$$
\frac{\left(\mathbb{E} X_{n}\right)^{2}}{\mathbb{E}\left[X_{n}^{2}\right]}=\frac{\left(\mathbb{E} X_{n}\right)^{2}}{\left(\mathbb{E} X_{n}\right)^{2}+\operatorname{Var}\left[X_{n}\right]}=\frac{1}{1+\operatorname{Var}\left[X_{n}\right] /\left(\mathbb{E} X_{n}\right)^{2}},
$$

and applying Theorem 2.3.2 gives the result.

### 2.3.2 $\triangleright$ Random graphs: subgraph containment and connectivity in the Erdốs-Rényi model

Threshold phenomena are also common in random graphs. We consider here the Erdős-Rényi random graph model (Definition 1.2.2). In this context a threshold function for a graph property $P$ is a function $r(n)$ such that

$$
\lim _{n} \mathbb{P}_{n, p_{n}}\left[G_{n} \text { has property } P\right]= \begin{cases}0, & \text { if } p_{n} \ll r(n) \\ 1, & \text { if } p_{n} \gg r(n)\end{cases}
$$

threshold
function
where $G_{n} \sim \mathbb{G}_{n, p_{n}}$ is a random graph with $n$ vertices and density $p_{n}$. In this section, we illustrate this type of phenomenon on two properties: the containment of small subgraphs and connectivity.

## Subgraph containment

We first consider the clique number, then we turn to more general subgraphs.

Cliques Let $\omega(G)$ be the clique number of a graph $G$, that is, the size of its largest clique.

Claim 2.3.4. The property $\omega\left(G_{n}\right) \geq 4$ has threshold function $n^{-2 / 3}$.
Proof. Let $X_{n}$ be the number of 4-cliques in the random graph $G_{n} \sim \mathbb{G}_{n, p_{n}}$. Then, noting that there are $\binom{4}{2}=6$ edges in a 4 -clique,

$$
\mathbb{E}_{n, p_{n}}\left[X_{n}\right]=\binom{n}{4} p_{n}^{6}=\Theta\left(n^{4} p_{n}^{6}\right)
$$

which goes to 0 when $p_{n} \ll n^{-2 / 3}$. Hence the first moment method (Theorem 2.2.6) gives one direction: $\mathbb{P}_{n, p_{n}}\left[\omega\left(G_{n}\right) \geq 4\right] \rightarrow 0$ in that case.

For the other direction, we apply the second moment method for sums of indicators, that is, Corollary 2.3.3. We use the notation from that corollary. For an enumeration $S_{1}, \ldots, S_{m_{n}}$ of the 4 -tuples of vertices in $G_{n}$, let $A_{n, 1}, \ldots, A_{n, m_{n}}$ be the events that the corresponding 4-clique is present. By the calculation above we have $\mu_{n}=\Theta\left(n^{4} p_{n}^{6}\right)$ which goes to $+\infty$ when $p_{n} \gg n^{-2 / 3}$. Also $\mu_{n}^{2}=\Theta\left(n^{8} p_{n}^{12}\right)$ so it suffices to show that $\gamma_{n}=o\left(n^{8} p_{n}^{12}\right)$. Note that two 4 -cliques with disjoint edge sets (but possibly sharing one vertex) are independent (i.e., their presence or absence is independent). Suppose $S_{i}$ and $S_{j}$ share 3 vertices. Then $i \stackrel{n}{\sim} j$ and

$$
\mathbb{P}_{n, p_{n}}\left[A_{n, i} \mid A_{n, j}\right]=p_{n}^{3}
$$

as the event $A_{n, j}$ implies that all edges between three of the vertices in $S_{i}$ are already present, and there are 3 edges between the remaining vertex and the rest of $S_{i}$. Similarly, if $\left|S_{i} \cap S_{j}\right|=2$, we have again $i \stackrel{n}{\sim} j$ and this time $\mathbb{P}_{n, p_{n}}\left[A_{n, i} \mid A_{n, j}\right]=$ $p_{n}^{5}$. Putting these together, we get by the definition of the conditional probability
(see Appendix B) and the fact that $\mathbb{P}_{n, p_{n}}\left[A_{n, j}\right]=p_{n}^{6}$

$$
\begin{aligned}
\gamma_{n} & =\sum_{i \sim j} \mathbb{P}\left[A_{n, i} \cap A_{n, j}\right] \\
& =\sum_{i \sim j} \mathbb{P}_{n, p_{n}}\left[A_{n, j}\right] \mathbb{P}_{n, p_{n}}\left[A_{n, i} \mid A_{n, j}\right] \\
& =\sum_{j} \mathbb{P}_{n, p_{n}}\left[A_{n, j}\right] \sum_{i: i \sim j} \mathbb{P}_{n, p_{n}}\left[A_{n, i} \mid A_{n, j}\right] \\
& =\binom{n}{4} p_{n}^{6}\left[\binom{4}{3}(n-4) p_{n}^{3}+\binom{4}{2}\binom{n-4}{2} p_{n}^{5}\right] \\
& =O\left(n^{5} p_{n}^{9}\right)+O\left(n^{6} p_{n}^{11}\right) \\
& =O\left(\frac{n^{8} p_{n}^{12}}{n^{3} p_{n}^{3}}\right)+O\left(\frac{n^{8} p_{n}^{12}}{n^{2} p_{n}}\right) \\
& =o\left(n^{8} p_{n}^{12}\right) \\
& =o\left(\mu_{n}^{2}\right),
\end{aligned}
$$

where we used that $p_{n} \gg n^{-2 / 3}$ (so that for example $n^{3} p_{n}^{3} \gg 1$ ). Corollary 2.3.3 gives the result: $\mathbb{P}_{n, p_{n}}\left[\cup_{i} A_{n, i}\right] \rightarrow 1$ when $p_{n} \gg n^{-2 / 3}$.

Roughly speaking, the first and second moments suffice to pinpoint the threshold in this case because the indicators in $X_{n}$ are "mostly" pairwise independent and, as a result, the sum is "concentrated around its mean."

General subgraphs The methods of Claim 2.3.4 can be applied to more general subgraphs. However the situation is somewhat more complicated than it is for cliques. For a graph $H_{0}$, let $v_{H_{0}}$ and $e_{H_{0}}$ be the number of vertices and edges of $H_{0}$ respectively. Let $X_{n}$ be the number of (not necessarily induced) copies of $H_{0}$ in $G_{n} \sim \mathbb{G}_{n, p_{n}}$. By the first moment method,

$$
\mathbb{P}\left[X_{n}>0\right] \leq \mathbb{E}\left[X_{n}\right]=\Theta\left(n^{v_{H_{0}}} p_{n}^{e_{H_{0}}}\right) \rightarrow 0,
$$

when $p_{n} \ll n^{-v_{H_{0}} / e_{H_{0}}}$. The constant factor, which does not play a role in the asymptotics, accounts in particular for the number of automorphisms of $H_{0}$. Indeed note that a fixed set of $v_{H_{0}}$ vertices can contain several distinct copies of $H_{0}$, depending on its structure (and unlike the clique case).

From the proof of Claim 2.3.4, one might guess that the threshold function is $n^{-v_{H_{0}} / e_{H_{0}}}$. That is not the case in general. To see what can go wrong, consider the graph $H_{0}$ in Figure 2.5 whose edge density is $\frac{e_{H_{0}}}{v_{H_{0}}}=\frac{6}{5}$. When $p_{n} \gg n^{-5 / 6}$,


Figure 2.5: Graph $H_{0}$ and subgraph $H$.
the expected number of copies of $H_{0}$ in $G_{n}$ tends to $+\infty$. But observe that the subgraph $H$ of $H_{0}$ has the higher density $5 / 4$ and, hence, when $n^{-5 / 6} \ll p_{n} \ll$ $n^{-4 / 5}$ the expected number of copies of $H$ tends to 0 . By the first moment method, the probability that a copy of $H_{0}$-and therefore $H$-is present in that regime is asymptotically negligible despite its diverging expectation. This leads to the following definition

$$
r_{H_{0}}:=\max \left\{\frac{e_{H}}{v_{H}}: \text { subgraphs } H \subseteq H_{0}, e_{H}>0\right\} .
$$

Assume $H_{0}$ has at least one edge.
Claim 2.3.5. "Having a copy of $H_{0}$ " has threshold $n^{-1 / r_{H_{0}}}$.
Proof. We proceed as in Claim 2.3.4. Let $H_{0}^{*}$ be a subgraph of $H_{0}$ achieving $r_{H_{0}}$. When $p_{n} \ll n^{-1 / r_{H_{0}}}$, the probability that a copy of $H_{0}^{*}$ is in $G_{n}$ tends to 0 by the argument above. Therefore the same conclusion holds for $H_{0}$ itself.

Assume $p_{n} \gg n^{-1 / r_{H_{0}}}$. Let $S_{1}, \ldots, S_{m_{n}}$ be an enumeration of the copies (as subgraphs) of $H_{0}$ in a complete graph on the vertices of $G_{n}$. Let $A_{n, i}$ be the event that $S_{i} \subseteq G_{n}$. Using again the notation of Corollary 2.3.3,

$$
\mu_{n}=\Theta\left(n^{v_{H_{0}}} p_{n}^{e_{H_{0}}}\right)=\Omega\left(\Phi_{H_{0}}(n)\right),
$$

where

$$
\Phi_{H_{0}}(n):=\min \left\{n^{v_{H}} p_{n}^{e_{H}}: \text { subgraphs } H \subseteq H_{0}, e_{H}>0\right\}
$$

Note that $\Phi_{H_{0}}(n) \rightarrow+\infty$ when $p_{n} \gg n^{-1 / r_{H_{0}}}$ by definition of $r_{H_{0}}$. The events $A_{n, i}$ and $A_{n, j}$ are independent if $S_{i}$ and $S_{j}$ share no edge. Otherwise we write $i \stackrel{n}{\sim} j$. Note that there are $\Theta\left(n^{v_{H}} n^{2\left(v_{H_{0}}-v_{H}\right)}\right)$ pairs $S_{i}, S_{j}$ whose intersection is isomorphic to $H$. The probability that both $S_{i}$ and $S_{j}$ of such a pair are present in $G_{n}$ is $\Theta\left(p_{n}^{e_{H}} p_{n}^{2\left(e_{H_{0}}-e_{H}\right)}\right)$. Hence

$$
\begin{aligned}
\gamma_{n} & =\sum_{i \sim j} \mathbb{P}\left[A_{n, i} \cap A_{n, j}\right] \\
& =\sum_{H \subseteq H_{0}, e_{H}>0} \Theta\left(n^{2 v_{H_{0}}-v_{H}} p_{n}^{2 e_{H_{0}}-e_{H}}\right) \\
& \leq \frac{\Theta\left(\mu_{n}^{2}\right)}{\Theta\left(\Phi_{H_{0}}(n)\right)} \\
& =o\left(\mu_{n}^{2}\right),
\end{aligned}
$$

where we used that $\Phi_{H_{0}}(n) \rightarrow+\infty$. The result follows from Corollary 2.3.3.
Going back to the example of Figure 2.5, the proof above confirms that when $n^{-5 / 6} \ll p_{n} \ll n^{-4 / 5}$ the second moment method fails for $H_{0}$ since $\Phi_{H_{0}}(n) \rightarrow 0$. In that regime, although there is in expectation a large number of copies of $H_{0}$, those copies are highly correlated as they are produced from a small (vanishing in expectation) number of copies of $H$-producing a large variance that helps to explain the failure of the second moment method.

## Connectivity threshold

Next we use the second moment method to show that the threshold function for connectivity in the Erdős-Rényi random graph model is $\frac{\log n}{n}$. In fact we prove this result by deriving the threshold function for the presence of isolated vertices. The connection between the two is obvious in one direction. Isolated vertices imply a disconnected graph. What is less obvious is that it also works the other way in the following sense: the two thresholds actually coincide.

Isolated vertices We begin with isolated vertices.
Claim 2.3.6. "Not having an isolated vertex" has threshold function $\frac{\log n}{n}$.
Proof. Let $X_{n}$ be the number of isolated vertices in the random graph $G_{n} \sim \mathbb{G}_{n, p_{n}}$. Using $1-x \leq e^{-x}$ for all $x \in \mathbb{R}$ (see Exercise 1.16),

$$
\begin{equation*}
\mathbb{E}_{n, p_{n}}\left[X_{n}\right]=n\left(1-p_{n}\right)^{n-1} \leq e^{\log n-(n-1) p_{n}} \rightarrow 0 \tag{2.3.4}
\end{equation*}
$$

when $p_{n} \gg \frac{\log n}{n}$. So the first moment method gives one direction: $\mathbb{P}_{n, p_{n}}\left[X_{n}>\right.$ $0] \rightarrow 0$.

For the other direction, we use the second moment method. Let $A_{n, j}$ be the event that vertex $j$ is isolated. By the computation above, using $1-x \geq e^{-x-x^{2}}$ for $x \in[0,1 / 2]$ (see Exercise 1.16 again),

$$
\begin{equation*}
\mu_{n}=\sum_{i} \mathbb{P}_{n, p_{n}}\left[A_{n, i}\right]=n\left(1-p_{n}\right)^{n-1} \geq e^{\log n-n p_{n}-n p_{n}^{2}}, \tag{2.3.5}
\end{equation*}
$$

which goes to $+\infty$ when $p_{n} \ll \frac{\log n}{n}$. Note that $A_{n, i}$ and $A_{n, j}$ are not independent for all $i \neq j$ (because the absence of an edge between $i$ and $j$ is part of both events) and

$$
\mathbb{P}_{n, p_{n}}\left[A_{n, i} \cap A_{n, j}\right]=\left(1-p_{n}\right)^{2(n-2)+1}
$$

so that

$$
\gamma_{n}=\sum_{i \neq j} \mathbb{P}_{n, p_{n}}\left[A_{n, i} \cap A_{n, j}\right]=n(n-1)\left(1-p_{n}\right)^{2 n-3} .
$$

Because $\gamma_{n}$ is not $o\left(\mu_{n}^{2}\right)$, we cannot apply Corollary 2.3.3. Instead we use Theorem 2.3.2 directly. We have

$$
\begin{align*}
\frac{\mathbb{E}_{n, p_{n}}\left[X_{n}^{2}\right]}{\mathbb{E}_{n, p_{n}}\left[X_{n}\right]^{2}} & =\frac{\mu_{n}+\gamma_{n}}{\mu_{n}^{2}} \\
& \leq \frac{n\left(1-p_{n}\right)^{n-1}+n^{2}\left(1-p_{n}\right)^{2 n-3}}{n^{2}\left(1-p_{n}\right)^{2 n-2}} \\
& \leq \frac{1}{n\left(1-p_{n}\right)^{n-1}}+\frac{1}{1-p_{n}} \tag{2.3.6}
\end{align*}
$$

which is $1+o(1)$ when $p_{n} \ll \frac{\log n}{n}$ by (2.3.5). The second moment method implies that $\mathbb{P}_{n, p_{n}}\left[X_{n}>0\right] \rightarrow 1$ in that case.

Connectivity We use Claim 2.3.6 to study the threshold for connectivity.
Claim 2.3.7. Connectivity has threshold function $\frac{\log n}{n}$.
Proof. We start with the easy direction. If $p_{n} \ll \frac{\log n}{n}$, Claim 2.3.6 implies that the graph has at least one isolated vertex-and therefore is necessarily disconnectedwith probability going to 1 as $n \rightarrow+\infty$.

Assume now that $p_{n} \gg \frac{\log n}{n}$. Let $\mathcal{D}_{n}$ be the event that $G_{n}$ is disconnected. To bound $\mathbb{P}_{n, p_{n}}\left[\mathcal{D}_{n}\right]$, we let $Y_{k}$ be the number of subsets of $k$ vertices that are
disconnected from all other vertices in the graph for $k \in\{1, \ldots, n / 2\}$. Then, by the first moment method,

$$
\mathbb{P}_{n, p_{n}}\left[\mathcal{D}_{n}\right]=\mathbb{P}_{n, p_{n}}\left[\sum_{k=1}^{n / 2} Y_{k}>0\right] \leq \sum_{k=1}^{n / 2} \mathbb{E}_{n, p_{n}}\left[Y_{k}\right] .
$$

The expectation of $Y_{k}$ is straightforward to bound. Using that $k \leq n / 2$ and $\binom{n}{k} \leq$ $n^{k}$,

$$
\mathbb{E}_{n, p_{n}}\left[Y_{k}\right]=\binom{n}{k}\left(1-p_{n}\right)^{k(n-k)} \leq\left(n\left(1-p_{n}\right)^{n / 2}\right)^{k} .
$$

The expression in parentheses is $o(1)$ when $p_{n} \gg \frac{\log n}{n}$ by a calculation similar to (2.3.4). Summing over $k$,

$$
\mathbb{P}_{n, p_{n}}\left[\mathcal{D}_{n}\right] \leq \sum_{k=1}^{+\infty}\left(n\left(1-p_{n}\right)^{n / 2}\right)^{k}=O\left(n\left(1-p_{n}\right)^{n / 2}\right)=o(1),
$$

where we used that the geometric series (started at $k=1$ ) is dominated asymptotically by its first term. So the probability of being disconnected goes to 0 when $p_{n} \gg \frac{\log n}{n}$ and we have proved the claim.

A closer look We have shown that connectivity and the absence of isolated vertices have the same threshold function. In fact, in a sense, isolated vertices are the "last obstacle" to connectivity. A slight modification of the proof above leads to the following more precise result. For $k \in\{1, \ldots, n / 2\}$, let $Z_{k}$ be the number of connected components of size $k$ in $G_{n}$. In particular, $Z_{1}$ is the number of isolated vertices. We consider the "critical window" $p_{n}=\frac{c_{n}}{n}$ where $c_{n}:=\log n+s$ for some fixed $s \in \mathbb{R}$. We show that, in that regime, asymptotically the graph is typically composed of a large connected component together with some isolated vertices. Formally, we prove the following claim which says that with probability close to one: either the graph is connected or there are some isolated vertices together with a (necessarily unique) connected component of size greater than $n / 2$.

## Claim 2.3.8.

$$
\mathbb{P}_{n, p_{n}}\left[Z_{1}>0\right] \geq \frac{1}{1+e^{s}}+o(1) \quad \text { and } \quad \mathbb{P}_{n, p_{n}}\left[\sum_{k=2}^{n / 2} Z_{k}>0\right]=o(1) .
$$

The limit of $\mathbb{P}_{n, p_{n}}\left[Z_{1}>0\right]$ can be computed explicitly using the method of moments. See Exercise 2.19.

Proof of Claim 2.3.8. We first consider isolated vertices. From (2.3.5), (2.3.6) and the second moment method,

$$
\mathbb{P}_{n, p_{n}}\left[Z_{1}>0\right] \geq\left(e^{-\log n+n p_{n}+n p_{n}^{2}}+\frac{1}{1-p_{n}}\right)^{-1}=\frac{1}{1+e^{s}}+o(1)
$$

as $n \rightarrow+\infty$ by our choice of $p_{n}$.
To bound the number of components of size $k>1$, we note first that the random variable $Y_{k}$ used in the previous claim (which imposes no condition on the edges between the vertices in the subsets of size $k$ ) is too loose to provide a suitable bound. Instead, to bound the probability that a subset of $k$ vertices forms a connected component, we observe that a connected component is characterized by two properties: it is disconnected from the rest of the graph; and it contains a spanning tree. Formally, for $k=2, \ldots, n / 2$, we let $Z_{k}^{\prime}$ be the number of (not necessarily induced) maximal trees of size $k$ or, put differently, the number of spanning trees of connected components of size $k$. Then, by the first moment method, the probability that a connected component of size $>1$ is present in $G_{n}$ is bounded by

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\sum_{k=2}^{n / 2} Z_{k}>0\right] \leq \mathbb{P}_{n, p_{n}}\left[\sum_{k=2}^{n / 2} Z_{k}^{\prime}>0\right] \leq \sum_{k=2}^{n / 2} \mathbb{E}_{n, p_{n}}\left[Z_{k}^{\prime}\right] . \tag{2.3.7}
\end{equation*}
$$

To bound the expectation of $Z_{k}^{\prime}$, we use Cayley's formula which states that there are $k^{k-2}$ trees on a set of $k$ labeled vertices. Recall further that a tree on $k$ vertices has $k-1$ edges (see Exercise 1.7). Hence,

$$
\mathbb{E}_{n, p_{n}}\left[Z_{k}^{\prime}\right]=\underbrace{\binom{n}{k} k^{k-2}}_{\text {(a) }} \underbrace{p_{n}^{k-1}}_{\text {(b) }} \underbrace{\left(1-p_{n}\right)^{k(n-k)}}_{\text {(c) }},
$$

where (a) is the number of trees of size $k$ (as subgraphs) in a complete graph of size $n$, (b) is the probability that such a tree is present in the graph, and (c) is the probability that this tree is disconnected from every other vertex in the graph. Using that $k!\geq(k / e)^{k}$ (see Appendix A) and $1-x \leq e^{-x}$ for all $x \in \mathbb{R}$ (see

Exercise 1.16),

$$
\begin{aligned}
\mathbb{E}_{n, p_{n}}\left[Z_{k}^{\prime}\right] & \leq \frac{n^{k}}{k!} k^{k-2} p_{n}^{k-1}\left(1-p_{n}\right)^{k(n-k)} \\
& \leq \frac{n^{k} e^{k}}{k^{k}} k^{k} n p_{n}^{k} e^{-p_{n} k(n-k)} \\
& =n\left(e c_{n} e^{-\left(1-\frac{k}{n}\right) c_{n}}\right)^{k} \\
& =n\left(e(\log n+s) e^{-\left(1-\frac{k}{n}\right)(\log n+s)}\right)^{k} .
\end{aligned}
$$

For $k \leq n / 2$, the expression in parentheses is $o(1)$. In fact, for $2 \leq k \leq n / 2$, $\mathbb{E}_{n, p_{n}}\left[Z_{k}^{\prime}\right]=o(1)$. Furthermore, summing over $k>2$,

$$
\sum_{k=3}^{n / 2} \mathbb{E}_{n, p_{n}}\left[Z_{k}^{\prime}\right] \leq \sum_{k=3}^{+\infty} n\left(e(\log n+s) e^{-\frac{1}{2}(\log n+s)}\right)^{k}=O\left(n^{-1 / 2} \log ^{3} n\right)=o(1)
$$

Plugging this back into (2.3.7) gives the second claim in the statement.

### 2.3.3 $\triangleright$ Percolation: critical value on trees and branching number

Consider bond percolation (see Definition 1.2 .1 ) on the infinite $d$-regular tree $\mathbb{T}_{d}$. Root the tree arbitrarily at a vertex 0 and let $\mathcal{C}_{0}$ be the open cluster of the root. In this section we illustrate the use of the first and second moment methods on the identification of the critical value

$$
p_{c}\left(\mathbb{T}_{d}\right)=\sup \{p \in[0,1]: \theta(p)=0\},
$$

where recall that the percolation function is $\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right]$. We then consider general trees, introduce the branching number, and present a weighted version of the second moment method.

Regular tree Our main result for $\mathbb{T}_{d}$ is the following.
Claim 2.3.9.

$$
p_{c}\left(\mathbb{T}_{d}\right)=\frac{1}{d-1} .
$$

Proof. Let $\partial_{n}$ be the $n$-th level of $\mathbb{T}_{d}$, that is, the set of vertices at graph distance $n$ from 0 . Let $X_{n}$ be the number of vertices in $\partial_{n} \cap \mathcal{C}_{0}$. In order for the open cluster of


Figure 2.6: Most recent common ancestor of $x$ and $y$.
the root to be infinite, there must be at least one vertex on the $n$-th level connected to the root by an open path. By the first moment method (Theorem 2.2.6),

$$
\begin{equation*}
\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right] \leq \mathbb{P}_{p}\left[X_{n}>0\right] \leq \mathbb{E}_{p} X_{n}=d(d-1)^{n-1} p^{n} \rightarrow 0, \tag{2.3.8}
\end{equation*}
$$

as $n \rightarrow+\infty$, for any $p<\frac{1}{d-1}$. Here we used that there is a unique path between 0 and any vertex in the tree to deduce that $\mathbb{P}_{p}\left[x \in \mathcal{C}_{0}\right]=p^{n}$ for $x \in \partial_{n}$. Equation (2.3.8) implies half of the claim: $p_{c}\left(\mathbb{T}_{d}\right) \geq \frac{1}{d-1}$.

The second moment method gives a lower bound on $\mathbb{P}_{p}\left[X_{n}>0\right]$. To simplify the notation, it is convenient to introduce the "branching ratio" $b:=d-1$. We say that $x$ is a descendant of $z$ if the path between 0 and $x$ goes through $z$. Each $z \neq 0$ has $d-1$ descendant subtrees, that is, subtrees of $\mathbb{T}_{d}$ rooted at $z$ made of all descendants of $z$. Let $x \wedge y$ be the most recent common ancestor of $x$ and $y$, that is, the furthest vertex from 0 that lies on both the path from 0 to $x$ and the path from 0 to $y$; see Figure 2.6. Letting

$$
\mu_{n}:=\mathbb{E}_{p}\left[X_{n}\right]=\mathbb{E}_{p}\left[\sum_{x \in \partial_{n}} \mathbf{1}_{\left\{x \in \mathcal{C}_{0}\right\}}\right]=(b+1) b^{n-1} p^{n},
$$

we have

$$
\begin{aligned}
\mathbb{E}_{p}\left[X_{n}^{2}\right] & =\mathbb{E}_{p}\left[\left(\sum_{x \in \partial_{n}} \mathbf{1}_{\left\{x \in \mathcal{C}_{0}\right\}}\right)^{2}\right] \\
& =\sum_{x, y \in \partial_{n}} \mathbb{P}_{p}\left[x, y \in \mathcal{C}_{0}\right] \\
& =\sum_{x \in \partial_{n}} \mathbb{P}_{p}\left[x \in \mathcal{C}_{0}\right]+\sum_{m=0}^{n-1} \sum_{x, y \in \partial_{n}} \mathbf{1}_{\left\{x \wedge y \in \partial_{m}\right\}} p^{m} p^{2(n-m)} \\
& =\mu_{n}+(b+1) b^{n-1} \sum_{m=0}^{n-1}(b-1) b^{(n-m)-1} p^{2 n-m} \\
& \leq \mu_{n}+(b+1)(b-1) b^{2 n-2} p^{2 n} \sum_{m=0}^{+\infty}(b p)^{-m} \\
& =\mu_{n}+\mu_{n}^{2} \cdot \frac{b-1}{b+1} \cdot \frac{1}{1-(b p)^{-1}},
\end{aligned}
$$

where, on the fourth line, we used that all vertices on the $n$-th level are equivalent and that, for a fixed $x$, the set $\left\{y: x \wedge y \in \partial_{m}\right\}$ is composed of those vertices in $\partial_{n}$ that are descendants of $x \wedge y$ but not in the descendant subtree of $x \wedge y$ containing $x$. When $p>\frac{1}{d-1}=\frac{1}{b}$, dividing by $\left(\mathbb{E}_{p} X_{n}\right)^{2}=\mu_{n}^{2} \rightarrow+\infty$, we get

$$
\begin{align*}
\frac{\mathbb{E}_{p}\left[X_{n}^{2}\right]}{\left(\mathbb{E}_{p} X_{n}\right)^{2}} & \leq \frac{1}{\mu_{n}}+\frac{b-1}{b+1} \cdot \frac{1}{1-(b p)^{-1}}  \tag{2.3.9}\\
& \leq 1+\frac{b-1}{b+1} \cdot \frac{1}{1-(b p)^{-1}} \\
& =C_{b, p} .
\end{align*}
$$

By the second moment method (Theorem 2.3.2) and monotonicity,

$$
\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right]=\mathbb{P}_{p}\left[\forall n, X_{n}>0\right]=\lim _{n} \mathbb{P}_{p}\left[X_{n}>0\right] \geq C_{b, p}^{-1}>0
$$

which concludes the proof. (Note that the version of the second moment method in Equation (2.3.1) does not work here. Subtract 1 in (2.3.9) and take $p$ close to $1 / b$.)

The argument above relies crucially on the fact that, in a tree, any two vertices are connected by a unique path. For instance, approximating $\mathbb{P}_{p}\left[x \in \mathcal{C}_{0}\right]$ is much harder on a lattice. Note furthermore that, intuitively, the reason why the first moment captures the critical threshold exactly in this case is that bond percolation on $\mathbb{T}_{d}$ is a "branching process" (defined formally and studied at length in Chapter 6),
where $X_{n}$ represents the "population size at generation $n$." The qualitative behavior of a branching process is governed by its expectation: when the mean number of children $b p$ exceeds 1 , the process grows exponentially on average and "explodes" with positive probability (see Theorem 6.1.6). We will come back to this point of view in Section 6.2.4 where branching processes are used to give a more refined analysis of bond percolation on $\mathbb{T}_{d}$.

General trees Let $\mathscr{T}$ be a locally finite tree (i.e., all its degrees are finite) with root 0 . For an edge $e$, let $v_{e}$ be the endvertex of $e$ furthest from the root. We denote by $|e|$ the graph distance between 0 and $v_{e}$. Generalizing a previous definition from Section 1.1.1 to infinite, locally finite graphs, a cutset separating 0 and $+\infty$ is a finite set of edges $\Pi$ such that all infinite paths (which, recall, are self-avoiding by definition) starting at 0 go through $\Pi$. (For our purposes, it will suffice to assume that cutsets are finite by default.) For a cutset $\Pi$, we let $\Pi_{\mathrm{v}}:=\left\{v_{e}: e \in \Pi\right\}$. Repeating the argument in (2.3.8), for any cutset $\Pi$,

$$
\begin{align*}
\theta(p) & =\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right] \\
& \leq \mathbb{P}_{p}\left[\mathcal{C}_{0} \cap \Pi_{\mathrm{v}} \neq \emptyset\right] \\
& \leq \sum_{u \in \Pi_{\mathrm{v}}} \mathbb{P}_{p}\left[u \in \mathcal{C}_{0}\right] \\
& =\sum_{e \in \Pi} p^{|e|} . \tag{2.3.10}
\end{align*}
$$

This bound naturally leads to the following definition.
Definition 2.3.10 (Branching number). The branching number of $\mathscr{T}$ is given by

$$
\begin{equation*}
\operatorname{br}(\mathscr{T})=\sup \left\{\lambda \geq 1: \inf _{\text {cutset } \Pi} \sum_{e \in \Pi} \lambda^{-|e|}>0\right\} . \tag{2.3.11}
\end{equation*}
$$

branching
number

Using the max-flow min-cut theorem (Theorem 1.1.15), the branching number can also be characterized in terms of a "flow to $+\infty$." We will not do this here. (But see Theorem 3.3.30.)

Equation (2.3.10) implies that $p_{c}(\mathscr{T}) \geq \frac{1}{\operatorname{br}(\mathscr{T})}$. Remarkably, this bound is tight. The proof is based on a "weighted" second moment method argument.

Claim 2.3.11. For any rooted, locally finite tree $\mathscr{T}$,

$$
p_{c}(\mathscr{T})=\frac{1}{\operatorname{br}(\mathscr{T})} .
$$

Proof. Suppose $p<\frac{1}{\operatorname{br}(\mathscr{I})}$. Then $p^{-1}>\operatorname{br}(\mathscr{T})$ and the sum in (2.3.10) can be made arbitrarily small by definition of the branching number, that is, $\theta(p)=0$. Hence we have shown that $p_{c}(\mathscr{T}) \geq \frac{1}{\operatorname{br}(\mathscr{T})}$.

To argue in the other direction, let $p>\frac{1}{\operatorname{br}(\mathscr{F})}, p^{-1}<\lambda<\operatorname{br}(\mathscr{T})$, and $\varepsilon>0$ such that

$$
\begin{equation*}
\sum_{e \in \Pi} \lambda^{-|e|} \geq \varepsilon \tag{2.3.12}
\end{equation*}
$$

for all cutsets $\Pi$. The existence of such an $\varepsilon$ is guaranteed by the definition of the branching number. As in the proof of Claim 2.3.9, we use that $\theta(p)$ is the limit as $n \rightarrow+\infty$ of the probability that $\mathcal{C}_{0}$ reaches the $n$-th level (i.e., the vertices at graph distance $n$ from the root 0 , which is necessarily a finite set in a locally finite tree). However, this time, we use a weighted count on the $n$-th level. Let $\mathscr{T}_{n}$ be the first $n$ levels of $\mathscr{T}$ and, as before, let $\partial_{n}$ be the vertices on the $n$-th level. For a probability measure $\nu_{n}$ on $\partial_{n}$, we define the weighted count

$$
X_{n}=\sum_{z \in \partial_{n}} \mathbf{1}_{\left\{z \in \mathcal{C}_{0}\right\}} \frac{\nu_{n}(z)}{\mathbb{P}_{p}\left[z \in \mathcal{C}_{0}\right]} .
$$

The purpose of the denominator is normalization, that is,

$$
\mathbb{E}_{p} X_{n}=\sum_{z \in \partial_{n}} \nu_{n}(z)=1
$$

Observe that, while $\nu_{n}(z)$ may be 0 for some $z$ s (but not all), we still have that $X_{n}>0, \forall n$ implies $\left\{\left|\mathcal{C}_{0}\right|=+\infty\right\}$, which is what we need to apply the second moment method.

Because of (2.3.12), a natural choice of $\nu_{n}$ follows from the max-flow mincut theorem (Theorem 1.1.15) applied to $\mathscr{T}_{n}$ with source $0, \operatorname{sink} \partial_{n}$ and capacity constraint $|\phi(x, y)| \leq \kappa(e):=\varepsilon^{-1} \lambda^{-|e|}$ for all edges $e=\{x, y\}$. Indeed, for all cutsets $\Pi$ in $\mathscr{T}_{n}$ separating 0 and $\partial_{n}$, we have $\sum_{e \in \Pi} \kappa(e)=\sum_{e \in \Pi} \varepsilon^{-1} \lambda^{-|e|} \geq 1$ by (2.3.12). That then guarantees by Theorem 1.1.15 the existence of a unit flow $\phi$ from 0 to $\partial_{n}$ satisfying the capacity constraints. Define $\nu_{n}(z)$ to be the flow entering $z \in \partial_{n}$ under $\phi$. In particular, because $\phi$ is a unit flow, $\nu_{n}$ defines a probability measure. It remains to bound the second moment of $X_{n}$ under this
choice. We have

$$
\begin{aligned}
\mathbb{E}_{p} X_{n}^{2} & =\mathbb{E}_{p}\left[\left(\sum_{z \in \partial_{n}} \mathbf{1}_{\left\{z \in \mathcal{C}_{0}\right\}} \frac{\nu_{n}(z)}{\mathbb{P}_{p}\left[z \in \mathcal{C}_{0}\right]}\right)^{2}\right] \\
& =\sum_{x, y \in \partial_{n}} \nu_{n}(x) \nu_{n}(y) \frac{\mathbb{P}_{p}\left[x, y \in \mathcal{C}_{0}\right]}{\mathbb{P}_{p}\left[x \in \mathcal{C}_{0}\right] \mathbb{P}_{p}\left[y \in \mathcal{\mathcal { C } _ { 0 } ]}\right.} \\
& =\sum_{m=0}^{n} \sum_{x, y \in \partial_{n}} \mathbf{1}_{\left\{x \wedge y \in \partial_{m}\right\}} \nu_{n}(x) \nu_{n}(y) \frac{p^{m} p^{2(n-m)}}{p^{2 n}} \\
& =\sum_{m=0}^{n} p^{-m} \sum_{z \in \partial_{m}}\left(\sum_{x, y \in \partial_{n}} \mathbf{1}_{\{x \wedge y=z\}} \nu_{n}(x) \nu_{n}(y)\right) .
\end{aligned}
$$

In the expression in parentheses, for each $x$ descendant of $z$, the sum over $y$ is at most $\nu_{n}(x) \nu_{n}(z)$ by the definition of a flow; then the sum over those $x$ s gives at most $\nu_{n}(z)^{2}$. So

$$
\begin{aligned}
\mathbb{E}_{p} X_{n}^{2} & \leq \sum_{m=0}^{n} p^{-m} \sum_{z \in \partial_{m}} \nu_{n}(z)^{2} \\
& \leq \sum_{m=0}^{n} p^{-m} \sum_{z \in \partial_{m}}\left(\varepsilon^{-1} \lambda^{-m}\right) \nu_{n}(z) \\
& \leq \varepsilon^{-1} \sum_{m=0}^{+\infty}(p \lambda)^{-m} \\
& =\frac{\varepsilon^{-1}}{1-(p \lambda)^{-1}}=: C_{\varepsilon, \lambda, p}<+\infty
\end{aligned}
$$

where the second line follows from the capacity constraint, and we used $p \lambda>1$ on the last line. From the second moment method (recalling that $\mathbb{E}_{p} X_{n}=1$ ),

$$
\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right] \geq \mathbb{P}_{p}\left[\forall n, X_{n}>0\right]=\lim _{n} \mathbb{P}_{p}\left[X_{n}>0\right] \geq C_{\varepsilon, \lambda, p}^{-1}>0,
$$

where the second equality follows from the construction of $\nu_{n}$. It follows that

$$
\theta(p) \geq C_{\varepsilon, \lambda, p}^{-1}>0
$$

and $p_{c}(\mathscr{T}) \leq \frac{1}{\operatorname{br}(\mathscr{F})}$. That concludes the proof.
Note that Claims 2.3.9 and 2.3.11 imply that $\operatorname{br}\left(\mathbb{T}_{d}\right)=d-1$. The next example is more striking and insightful.

Example 2.3.12 (The 3-1 tree). The $3-1$ tree $\widehat{\mathscr{T}}_{3-1}$ is an infinite rooted tree. We give a planar description. The root $\rho$ (level 0 ) is at the top. It has two children below it (level 1). Then on level $n$, for $n \geq 1$, the first $2^{n-1}$ vertices starting from the left have exactly 1 child and the next $2^{n-1}$ vertices have exactly 3 children. In particular level $n$ has $2^{n}$ vertices, which we denote by $u_{n, 1}, \ldots, u_{n, 2^{n}}$. For vertex $u_{n, j}$ we refer to $j / 2^{n}$ as its relative position (on level $n$ ). So vertices have 1 or 3 children according to whether their relative position is $\leq 1 / 2$ or $>1 / 2$.

Because the level size is growing at rate 2 , it is tempting to conjecture that the branching number is 2 -but that turns out to be way off.

Claim 2.3.13. $\operatorname{br}\left(\widehat{\mathscr{T}}_{3-1}\right)=1$.
What makes this tree entirely different from the infinite 2 -ary tree, despite having the same level growth, is that each infinite path from the root in $\widehat{\mathscr{T}}_{3-1}$ eventually "stops branching," with the sole exception of the rightmost path which we refer to as the main path. Indeed, let $\Gamma=v_{0} \sim v_{1} \sim v_{2} \sim \cdots$ with $v_{0}=\rho$ be an infinite path distinct from the main path. Let $x_{i}$ be the relative position of $v_{i}, i \geq 1$. Let $v_{k}$ be the first vertex of $\Gamma$ not on the main path. It lies on the $k$-th level.

Lemma 2.3.14. Letv be a vertex that is not on the main path with relative position $x$ and assume that $0 \leq x \leq \alpha<1$. Let $w$ be a child of $v$ and denote by $y$ its relative position. Then

$$
y \leq \begin{cases}\frac{1}{2} x & \text { if } x \leq 1 / 2 \\ x-\frac{1}{2}(1-\alpha) & \text { otherwise }\end{cases}
$$

Proof. Assume without loss of generality that $v=u_{n, j}$ for some $n$ and $j<2^{n}$. If $j \leq 2^{n-1}$, then by construction $v$ has exactly one child with relative position

$$
y=\frac{j}{2^{n+1}}=\frac{1}{2} x .
$$

That proves the first claim.
If $j>2^{n-1}$, then all vertices to the right of $v$ have 3 children, all of whom are to the right of the children of $v$. Hence the children of $v$ have relative position at most

$$
y \leq \frac{2^{n+1}-3\left(2^{n}-j\right)}{2^{n+1}}=\frac{3 j-2^{n}}{2^{n+1}}=\frac{3}{2} x-\frac{1}{2} .
$$

Subtracting $x$ and using $x \leq \alpha$ gives the second claim.
We now apply Lemma 2.3.14 to $v_{k}$ as defined above and its descendants on $\Gamma$ with $\alpha=1-1 / 2^{k}$. We get that the relative position decreases from $v_{k}$ by $1 / 2^{k+1}$ on
each level until it falls below $1 / 2$ at which point it gets cut in half at each level. Once this last regime is reached, each vertex on $\Gamma$ from then on has exactly one child-that is, there is no more branching.

We are now ready to prove the claim.
Proof of Claim 2.3.13. Take any $\lambda>1$. From the definition of the branching number (Definition 2.3.10), it suffices to find a sequence of cutsets $\left(\Pi_{n}\right)_{n}$ such that

$$
\sum_{e \in \Pi_{n}} \lambda^{-|e|} \rightarrow 0,
$$

as $n \rightarrow+\infty$. What does not work is to choose $\Pi_{n}:=\Lambda_{n}$ to be the edges between level $n-1$ and level $n$, since we then have

$$
\sum_{e \in \Lambda_{n}} \lambda^{-|e|}=2^{n} \lambda^{-n},
$$

which diverges whenever $\lambda<2$. Instead we construct a new cutset $\Phi_{n}$ based on $\Lambda_{n}$ as follows. We divide up $\Lambda_{n}$ into the disjoint union $\Lambda_{n}^{-} \cup \Lambda_{n}^{+}$, where $\Lambda_{n}^{-}$are the edges whose endvertex on level $n$ has relative position $\leq 1 / 2$ and $\Lambda_{n}^{+}$are the rest of the edges. Start with $\Phi_{n}:=\emptyset$.
Step 1. For each edge $e$ in $\Lambda_{n}^{-}$, letting $v$ be the endvertex of $e$ on level $n$, add to $\Phi_{n}$ the edge $\left\{v^{\prime}, v^{\prime \prime}\right\}$ where $v^{\prime}$ and $v^{\prime \prime}$ are the unique descendants of $v$ on level $m_{n}-1$ and $m_{n}$ respectively. The value of $m_{n} \geq n$ is chosen so that

$$
\begin{equation*}
2^{n-1} \lambda^{-m_{n}} \leq \frac{1}{2 n} . \tag{2.3.13}
\end{equation*}
$$

Any infinite path from the root going through one of the edges in $\Lambda_{n}^{-}$has to go through the edge that replaced it in $\Phi_{n}$ since there is no branching below that point by Lemma 2.3.14.
Step 2. We also add to $\Phi_{n}$ the edge $\left\{w^{\prime}, w^{\prime \prime}\right\}$ on the main path where $w^{\prime}=$ $u_{\ell_{n}-1,2^{\ell_{n}-1}}$ is on level $\ell_{n}-1$ and $w^{\prime \prime}=u_{\ell_{n}, 2^{\ell_{n}}}$ is on level $\ell_{n}$. We mean for the value of $\ell_{n}$ to be such that any infinite path going through an edge in $\Lambda_{n}^{+}$has to go through $\left\{w^{\prime}, w^{\prime \prime}\right\}$ first. That is, we need all vertices of level $n$ with relative position $>1 / 2$ to be a descendant of $w^{\prime \prime}$. The number of descendants of $w^{\prime \prime}$ on level $J>\ell_{n}$ is $3^{J-\ell_{n}}$ until the last $J$ such that it is $\leq 2^{J-1}$, which we denote by $J^{*}$. A quick calculation gives

$$
J^{*}=\left\lfloor\frac{\ell_{n} \log 3-\log 2}{\log 3-\log 2}\right\rfloor .
$$

After level $J^{*}$, the leftmost descendant of $w^{\prime \prime}$ has relative position $\leq 1 / 2$ by Lemma 2.3.14. Therefore we need $n>J^{*}$ to ensure that $w^{\prime \prime}$ has the desired property. Taking

$$
\begin{equation*}
\ell_{n}=\left\lfloor\frac{\log 3 / 2}{\log 3} n\right\rfloor, \tag{2.3.14}
\end{equation*}
$$

will do for $n$ large enough.
Finishing up. By construction, $\Phi_{n}$ is a cutset for all $n \geq n_{0}$. Moreover

$$
\sum_{e \in \Phi_{n}} \lambda^{-|e|}=2^{n-1} \lambda^{-m_{n}}+\lambda^{-\ell_{n}}<\frac{1}{n},
$$

for $n$ large enough, where we used (2.3.13) and (2.3.14). Taking $n \rightarrow+\infty$ gives the claim.

As a consequence of Claims 2.3.11 and 2.3.13, $\left|\mathcal{C}_{\rho}\right|<+\infty$ almost surely for all $p<1$ on $\widehat{\mathscr{T}}_{3}$.

### 2.4 Chernoff-Cramér method

Chebyshev's inequality (Theorem 2.1.2) gives a bound on the concentration around its mean of a square integrable random variable. It is, in general, best possible. Indeed take $X$ to be $\mu+b \sigma$ or $\mu-b \sigma$ with probability $\left(2 b^{2}\right)^{-1}$ each, and $\mu$ otherwise. Then $\mathbb{E} X=\mu, \operatorname{Var} X=\sigma^{2}$, and for $\beta=b \sigma$,

$$
\mathbb{P}[|X-\mathbb{E} X| \geq \beta]=\mathbb{P}[|X-\mathbb{E} X|=\beta]=\frac{1}{b^{2}}=\frac{\operatorname{Var} X}{\beta^{2}}
$$

However, in many cases, much stronger bounds can be derived. For instance, if $X \sim \mathrm{~N}(0,1)$, by the following lemma

$$
\begin{equation*}
\mathbb{P}[|X-\mathbb{E} X| \geq \beta] \sim \sqrt{\frac{2}{\pi}} \beta^{-1} \exp \left(-\beta^{2} / 2\right) \ll \frac{1}{\beta^{2}}, \tag{2.4.1}
\end{equation*}
$$

as $\beta \rightarrow+\infty$. Indeed:
Lemma 2.4.1. For $x>0$,

$$
\left(x^{-1}-x^{-3}\right) e^{-x^{2} / 2} \leq \int_{x}^{+\infty} e^{-y^{2} / 2} \mathrm{~d} y \leq x^{-1} e^{-x^{2} / 2}
$$

Proof. By the change of variable $y=x+z$ and using $e^{-z^{2} / 2} \leq 1$

$$
\int_{x}^{+\infty} e^{-y^{2} / 2} \mathrm{~d} y \leq e^{-x^{2} / 2} \int_{0}^{+\infty} e^{-x z} \mathrm{~d} z=e^{-x^{2} / 2} x^{-1}
$$

For the other direction, by differentiation

$$
\int_{x}^{+\infty}\left(1-3 y^{-4}\right) e^{-y^{2} / 2} \mathrm{~d} y=\left(x^{-1}-x^{-3}\right) e^{-x^{2} / 2}
$$

In this section we discuss the Chernoff-Cramér method, which produces exponential tail bounds, provided the moment-generating function (see Section 2.1.1) is finite in a neighborhood of 0 .

### 2.4.1 Tail bounds via the moment-generating function

Under a finite variance, squaring within Markov's inequality (Theorem 2.1.1) produces Chebyshev's inequality (Theorem 2.1.2). This "boosting" can be pushed further when stronger integrability conditions hold.

Chernoff-Cramér We refer to (2.4.2) in the next lemma as the Chernoff-Cramér bound.

Lemma 2.4.2 (Chernoff-Cramér bound). Assume $X$ is a random variable such that $M_{X}(s)<+\infty$ for $s \in\left(-s_{0}, s_{0}\right)$ for some $s_{0}>0$. For any $\beta>0$ and $s>0$,

$$
\begin{equation*}
\mathbb{P}[X \geq \beta] \leq \exp \left[-\left\{s \beta-\Psi_{X}(s)\right\}\right], \tag{2.4.2}
\end{equation*}
$$

where

$$
\Psi_{X}(s):=\log M_{X}(s),
$$

is the cumulant-generating function of $X$.
Proof. Exponentiating within Markov's inequality gives for $s>0$

$$
\mathbb{P}[X \geq \beta]=\mathbb{P}\left[e^{s X} \geq e^{s \beta}\right] \leq \frac{M_{X}(s)}{e^{s \beta}}=\exp \left[-\left\{s \beta-\Psi_{X}(s)\right\}\right] .
$$

Returning to the Gaussian case, let $X \sim N(0, \nu)$ where $\nu>0$ is the variance and note that

$$
\begin{aligned}
M_{X}(s) & =\int_{-\infty}^{+\infty} e^{s x} \frac{1}{\sqrt{2 \pi \nu}} e^{-\frac{x^{2}}{2 \nu}} \mathrm{~d} x \\
& =\int_{-\infty}^{+\infty} e^{\frac{s^{2} \nu}{2}} \frac{1}{\sqrt{2 \pi \nu}} e^{-\frac{(x-s \nu)^{2}}{2 \nu}} \mathrm{~d} x \\
& =\exp \left(\frac{s^{2} \nu}{2}\right) .
\end{aligned}
$$

By straightforward calculus, the optimal choice of $s$ in (2.4.2) gives the exponent

$$
\begin{equation*}
\sup _{s>0}\left(s \beta-s^{2} \nu / 2\right)=\frac{\beta^{2}}{2 \nu}, \tag{2.4.3}
\end{equation*}
$$

achieved at $s_{\beta}=\beta / \nu$. For $\beta>0$, this leads to the bound

$$
\begin{equation*}
\mathbb{P}[X \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right) \tag{2.4.4}
\end{equation*}
$$

which is much sharper than Chebyshev's inequality for large $\beta$-compare to (2.4.1).
As another toy example, we consider simple random walk on $\mathbb{Z}$.
Lemma 2.4.3 (Chernoff bound for simple random walk on $\mathbb{Z}$ ). Let $Z_{1}, \ldots, Z_{n}$ be independent Rademacher variables, that is, they are $\{-1,1\}$-valued random variables with $\mathbb{P}\left[Z_{i}=1\right]=\mathbb{P}\left[Z_{i}=-1\right]=1 / 2$. Let $S_{n}=\sum_{i \leq n} Z_{i}$. Then, for any $\beta>0$,

$$
\begin{equation*}
\mathbb{P}\left[S_{n} \geq \beta\right] \leq e^{-\beta^{2} / 2 n} \tag{2.4.5}
\end{equation*}
$$

Proof. The moment-generating function of $Z_{1}$ can be bounded as follows

$$
\begin{equation*}
M_{Z_{1}}(s)=\frac{e^{s}+e^{-s}}{2}=\sum_{j \geq 0} \frac{s^{2 j}}{(2 j)!} \leq \sum_{j \geq 0} \frac{\left(s^{2} / 2\right)^{j}}{j!}=e^{s^{2} / 2} . \tag{2.4.6}
\end{equation*}
$$

Using independence and taking $s=\beta / n$ in the Chernoff-Cramér bound (2.4.2), we get

$$
\begin{aligned}
\mathbb{P}\left[S_{n} \geq \beta\right] & \leq \exp \left(-s \beta+n \Psi_{Z_{1}}(s)\right) \\
& \leq \exp \left(-s \beta+n s^{2} / 2\right) \\
& =e^{-\beta^{2} / 2 n}
\end{aligned}
$$

which concludes the proof.

Observe the similarity between (2.4.5) and the Gaussian bound (2.4.4), if one takes $\nu$ to be the variance of $S_{n}$, that is,

$$
\nu=\operatorname{Var}\left[S_{n}\right]=n \operatorname{Var}\left[Z_{1}\right]=n \mathbb{E}\left[Z_{1}^{2}\right]=n,
$$

where we used that $Z_{1}$ is centered. The central limit theorem says that simple random walk is well approximated by a Gaussian in the "bulk" of the distribution; the bound above extends the approximation in the "large deviation" regime. The bounding technique used in the proof of Lemma 2.4 .3 will be substantially extended in Section 2.4.2.

Example 2.4.4 (Set balancing). Let $\boldsymbol{v}_{1}, \ldots, \boldsymbol{v}_{m}$ be arbitrary non-zero vectors in $\{0,1\}^{n}$. Think of $\boldsymbol{v}_{i}=\left(v_{i, 1}, \ldots, v_{i, n}\right)$ as representing a subset of $[n]=\{1, \ldots, n\}$ : $v_{i, j}=1$ indicates that $j$ is in subset $i$. Suppose we want to partition $[n]$ into two groups such that the subsets corresponding to the $\boldsymbol{v}_{i}$ s are as balanced as possible, that is, are as close as possible to having the same number of elements from each group. More formally, we seek a vector $\boldsymbol{x}=\left(x_{1}, \ldots, x_{n}\right) \in\{-1,+1\}^{n}$ such that $B^{*}=\max _{i=1, \ldots, m}\left|\boldsymbol{x} \cdot \boldsymbol{v}_{i}\right|$ is as small as possible.

A simple random choice does well: select each $x_{i}$ independently, uniformly at random in $\{-1,+1\}$. Fix $\varepsilon>0$. We claim that

$$
\begin{equation*}
\mathbb{P}\left[B^{*} \geq \sqrt{2 n\left(\log m+\log \left(2 \varepsilon^{-1}\right)\right)}\right] \leq \varepsilon . \tag{2.4.7}
\end{equation*}
$$

Indeed, by (2.4.5) (considering only the non-zero entries of $\boldsymbol{v}_{i}$ ),

$$
\begin{aligned}
\mathbb{P}\left[\left|\boldsymbol{x} \cdot \boldsymbol{v}_{i}\right|\right. & \left.\geq \sqrt{2 n\left(\log m+\log \left(2 \varepsilon^{-1}\right)\right)}\right] \\
& \leq 2 \exp \left(-\frac{2 n\left(\log m+\log \left(2 \varepsilon^{-1}\right)\right)}{2\left\|\boldsymbol{v}_{i}\right\|_{1}}\right) \\
& \leq \frac{\varepsilon}{m}
\end{aligned}
$$

where we used that $\left\|\boldsymbol{v}_{i}\right\|_{1} \leq n$. Taking a union bound over the $m$ vectors gives the result. In (2.4.7), the $\sqrt{n}$ term on the right-hand side of the inequality is to be expected since it is the standard deviation of $\left|\boldsymbol{x} \cdot \boldsymbol{v}_{i}\right|$ in the worst case. The power of the exponential tail bound (2.4.5) appears in the logarithmic terms, which would have been replaced with something much larger if one had used Chebyshev's inequality instead.

The Chernoff-Cramér bound is particularly useful for sums of independent random variables as the moment-generating function then factorizes; see (2.1.3). Let

$$
\Psi_{X}^{*}(\beta)=\sup _{s \in \mathbb{R}_{+}}\left(s \beta-\Psi_{X}(s)\right),
$$

be the Fenchel-Legendre dual of the cumulant-generating function of $X$.

Theorem 2.4.5 (Chernoff-Cramér method). Let $S_{n}=\sum_{i \leq n} X_{i}$, where the $X_{i} s$ are i.i.d. random variables. Assume $M_{X_{1}}(s)<+\infty$ on $s \in\left(-s_{0}, s_{0}\right)$ for some $s_{0}>0$. For any $\beta>0$,

$$
\begin{equation*}
\mathbb{P}\left[S_{n} \geq \beta\right] \leq \exp \left(-n \Psi_{X_{1}}^{*}\left(\frac{\beta}{n}\right)\right) \tag{2.4.8}
\end{equation*}
$$

In particular, in the large deviations regime, that is, when $\beta=b n$ for some $b>0$, we have

$$
\begin{equation*}
-\underset{n}{\limsup } \frac{1}{n} \log \mathbb{P}\left[S_{n} \geq b n\right] \geq \Psi_{X_{1}}^{*}(b) . \tag{2.4.9}
\end{equation*}
$$

Proof. Observe that, by taking a logarithm in (2.1.3), it holds that

$$
\Psi_{S_{n}}^{*}(\beta)=\sup _{s>0}\left(s \beta-n \Psi_{X_{1}}(s)\right)=\sup _{s>0} n\left(s\left(\frac{\beta}{n}\right)-\Psi_{X_{1}}(s)\right)=n \Psi_{X_{1}}^{*}\left(\frac{\beta}{n}\right) .
$$

Now optimize over $s$ in (2.4.2).
We use the Chernoff-Cramér method to derive a few standard bounds.

Poisson variables We start with the Poisson case. Let $Z \sim \operatorname{Poi}(\lambda)$ be Poisson with mean $\lambda$, where recall that

$$
\mathbb{P}[Z=k]=e^{-\lambda} \frac{\lambda^{k}}{k!}, \quad k \in \mathbb{Z}_{+} .
$$

Letting $X=Z-\lambda$,

$$
\begin{aligned}
\Psi_{X}(s) & =\log \left(\sum_{\ell \geq 0} e^{-\lambda} \frac{\lambda^{\ell}}{\ell!} e^{s(\ell-\lambda)}\right) \\
& =\log \left(e^{-(1+s) \lambda} \sum_{\ell \geq 0} \frac{\left(e^{s} \lambda\right)^{\ell}}{\ell!}\right) \\
& =\log \left(e^{-(1+s) \lambda} e^{e^{s} \lambda}\right) \\
& =\lambda\left(e^{s}-s-1\right)
\end{aligned}
$$

so that straightforward calculus gives for $\beta>0$

$$
\begin{aligned}
\Psi_{X}^{*}(\beta) & =\sup _{s>0}\left(s \beta-\lambda\left(e^{s}-s-1\right)\right) \\
& =\lambda\left[\left(1+\frac{\beta}{\lambda}\right) \log \left(1+\frac{\beta}{\lambda}\right)-\frac{\beta}{\lambda}\right] \\
& =\lambda h\left(\frac{\beta}{\lambda}\right),
\end{aligned}
$$

achieved at $s_{\beta}=\log \left(1+\frac{\beta}{\lambda}\right)$, where $h$ is defined as the expression in square brackets above. Plugging $\Psi_{X}^{*}(\beta)$ into Theorem 2.4.5 leads for $\beta>0$ to the bound

$$
\begin{equation*}
\mathbb{P}[Z \geq \lambda+\beta] \leq \exp \left(-\lambda h\left(\frac{\beta}{\lambda}\right)\right) \tag{2.4.10}
\end{equation*}
$$

A similar calculation for $-(Z-\lambda)$ gives for $\beta<0$

$$
\begin{equation*}
\mathbb{P}[Z \leq \lambda+\beta] \leq \exp \left(-\lambda h\left(\frac{\beta}{\lambda}\right)\right) \tag{2.4.11}
\end{equation*}
$$

If $S_{n}$ is a sum of $n$ i.i.d. $\operatorname{Poi}(\lambda)$ variables, then by (2.4.9) for $a>\lambda$

$$
\begin{align*}
-\limsup _{n} \frac{1}{n} \log \mathbb{P}\left[S_{n} \geq a n\right] & \geq \lambda h\left(\frac{a-\lambda}{\lambda}\right) \\
& =a \log \left(\frac{a}{\lambda}\right)-a+\lambda \\
& =: I_{\lambda}^{\mathrm{Poi}}(a), \tag{2.4.12}
\end{align*}
$$

and similarly for $a<\lambda$

$$
\begin{equation*}
-\underset{n}{\limsup } \frac{1}{n} \log \mathbb{P}\left[S_{n} \leq a n\right] \geq I_{\lambda}^{\mathrm{Poi}}(a) . \tag{2.4.13}
\end{equation*}
$$

In fact, these bounds follow immediately from (2.4.10) and (2.4.11) by noting that $S_{n} \sim \operatorname{Poi}(n \lambda)$ (see, e.g., Exercise 6.7).

Binomial variables and Chernoff bounds Let $Z \sim \operatorname{Bin}(n, p)$ be a binomial random variable with parameters $n$ and $p$. Recall that $Z$ is a sum of i.i.d. indicators $Y_{1}, \ldots, Y_{n}$ equal to 1 with probability $p$. The $Y_{i}$ s are also known as Bernoulli random variables or Bernoulli trials, and their law is denote by $\operatorname{Ber}(p)$. We also refer to $p$ as the success probability. Letting $X_{i}=Y_{i}-p$ and $S_{n}=Z-n p$,

$$
\Psi_{X_{1}}(s)=\log \left(p e^{s}+(1-p)\right)-p s
$$

For $b \in(0,1-p)$, letting $a=b+p$, direct calculation gives

$$
\begin{align*}
\Psi_{X_{1}}^{*}(b) & =\sup _{s>0}\left(s b-\left(\log \left[p e^{s}+(1-p)\right]-p s\right)\right) \\
& =(1-a) \log \frac{1-a}{1-p}+a \log \frac{a}{p}=: D(a \| p), \tag{2.4.14}
\end{align*}
$$

achieved at $s_{b}=\log \frac{(1-p) a}{p(1-a)}$. The function $D(a \| p)$ in (2.4.14) is the so-called Kullback-Leibler divergence or relative entropy between two Bernoulli variables
with parameters $a$ and $p$ respectively. By (2.4.8) for $\beta>0$

$$
\mathbb{P}[Z \geq n p+\beta] \leq \exp (-n D(p+\beta / n \| p)) .
$$

Applying the same argument to $Z^{\prime}=n-Z$ gives a bound in the other direction.
Remark 2.4.6. In the large deviations regime, it can be shown that the previous bound is tight in the sense that

$$
-\frac{1}{n} \log \mathbb{P}[Z \geq n p+b n] \rightarrow D(p+b \| p)=: I_{n, p}^{\mathrm{Bin}}(b),
$$

as $n \rightarrow+\infty$. The theory of large deviations provides general results of this type. See for example [Dur10, Section 2.6]. Upper bounds will be enough for our purposes.

The following related bounds, proved in Exercise 2.7, are often useful.
Theorem 2.4.7 (Chernoff bounds for Poisson trials). Let $Y_{1}, \ldots, Y_{n}$ be independent $\{0,1\}$-valued random variables with $\mathbb{P}\left[Y_{i}=1\right]=p_{i}$ and $\mu=\sum_{i} p_{i}$. These are called Poisson trials. Let $Z=\sum_{i} Y_{i}$. Then:
(a) For any $\delta>0$,

$$
\mathbb{P}[Z \geq(1+\delta) \mu] \leq\left(\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}}\right)^{\mu}
$$

(b) For any $0<\delta \leq 1$,

$$
\mathbb{P}[Z \geq(1+\delta) \mu] \leq e^{-\mu \delta^{2} / 3}
$$

(ii) Below the mean
(a) For any $0<\delta<1$,

$$
\mathbb{P}[Z \leq(1-\delta) \mu] \leq\left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\mu}
$$

(b) For any $0<\delta<1$,

$$
\mathbb{P}[Z \leq(1-\delta) \mu] \leq e^{-\mu \delta^{2} / 2}
$$

### 2.4.2 Sub-Gaussian and sub-exponential random variables

The bounds in Section 2.4.1 were obtained by computing the moment-generating function explicitly (possibly with some approximations). This is not always possible. In this section, we give some important examples of tail bounds derived from the Chernoff-Cramér method for broad classes of random variables under natural conditions on their distributions.

## Sub-Gaussian random variables

We begin with sub-Gaussian random variables which, as the name suggests, have a moment-generating function that is bounded by that of a Gaussian.

General case Here is our key definition.
Definition 2.4.8 (Sub-Gaussian random variables). We say that a random variable $X$ with mean $\mu$ is sub-Gaussian with variance factor $\nu$ if

$$
\Psi_{X-\mu}(s) \leq \frac{s^{2} \nu}{2}, \quad \forall s \in \mathbb{R}
$$

sub-Gaussian
variable
for some $\nu>0$. We use the notation $X \in \operatorname{sG}(\nu)$.
Note that the right-hand side in (2.4.15) is the cumulant-generating function of a $\mathrm{N}(0, \nu)$. By the Chernoff-Cramér method and (2.4.3) it follows immediately that

$$
\begin{equation*}
\mathbb{P}[X-\mu \leq-\beta] \vee \mathbb{P}[X-\mu \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right) \tag{2.4.16}
\end{equation*}
$$

where we used that $X \in \operatorname{s\mathcal {G}}(\nu)$ implies $-X \in \mathrm{~s} \mathcal{G}(\nu)$. As a quick example, note that this is the approach we took in Lemma 2.4.3, that is, we showed that a uniform random variable in $\{-1,1\}$ (i.e., a Rademacher variable) is sub-Gaussian with variance factor 1 .

When considering (weighted) sums of independent sub-Gaussian random variables, we get the following.

Theorem 2.4.9 (General Hoeffding inequality). Suppose $X_{1}, \ldots, X_{n}$ are independent random variables where, for each $i, X_{i} \in \operatorname{s\mathcal {G}}\left(\nu_{i}\right)$ with $0<\nu_{i}<+\infty$. For $w_{1}, \ldots, w_{n} \in \mathbb{R}$, let $S_{n}=\sum_{i \leq n} w_{i} X_{i}$. Then

$$
S_{n} \in \mathrm{~s} \mathcal{G}\left(\sum_{i=1}^{n} w_{i}^{2} \nu_{i}\right) .
$$

In particular, for all $\beta>0$,

$$
\mathbb{P}\left[S_{n}-\mathbb{E} S_{n} \geq \beta\right] \leq \exp \left(-\frac{\beta^{2}}{2 \sum_{i=1}^{n} w_{i}^{2} \nu_{i}}\right)
$$

Proof. Assume the $X_{i} \mathrm{~s}$ are centered. By independence and (2.1.3),

$$
\Psi_{S_{n}}(s)=\sum_{i \leq n} \Psi_{w_{i} X_{i}}(s)=\sum_{i \leq n} \Psi_{X_{i}}\left(s w_{i}\right) \leq \sum_{i \leq n} \frac{\left(s w_{i}\right)^{2} \nu_{i}}{2}=\frac{s^{2} \sum_{i \leq n} w_{i}^{2} \nu_{i}}{2}
$$

Bounded random variables For bounded random variables, the previous inequality reduces to a standard bound.

Theorem 2.4.10 (Hoeffding's inequality for bounded variables). Let $X_{1}, \ldots, X_{n}$ be independent random variables where, for each $i, X_{i}$ takes values in $\left[a_{i}, b_{i}\right]$ with $-\infty<a_{i} \leq b_{i}<+\infty$. Let $S_{n}=\sum_{i \leq n} X_{i}$. For all $\beta>0$,

$$
\mathbb{P}\left[S_{n}-\mathbb{E} S_{n} \geq \beta\right] \leq \exp \left(-\frac{2 \beta^{2}}{\sum_{i \leq n}\left(b_{i}-a_{i}\right)^{2}}\right)
$$

By Theorem 2.4.9, it suffices to show that $X_{i} \in \mathrm{sG}\left(\nu_{i}\right)$ with $\nu_{i}=\frac{1}{4}\left(b_{i}-a_{i}\right)^{2}$. We first give a quick proof of a weaker version that uses a trick called symmetrization. Suppose the $X_{i}$ s are centered and satisfy $\left|X_{i}\right| \leq c_{i}$ for some $c_{i}>0$. Let $X_{i}^{\prime}$ be an independent copy of $X_{i}$ and let $Z_{i}$ be an independent uniform random variable in $\{-1,1\}$. For any $s$,

$$
\begin{aligned}
\mathbb{E}\left[e^{s X_{i}}\right] & =\mathbb{E}\left[e^{s \mathbb{E}\left[X_{i}-X_{i}^{\prime} \mid X_{i}\right]}\right] \\
& \leq \mathbb{E}\left[\mathbb{E}\left[e^{s\left(X_{i}-X_{i}^{\prime}\right)} \mid X_{i}\right]\right] \\
& =\mathbb{E}\left[e^{s\left(X_{i}-X_{i}^{\prime}\right)}\right],
\end{aligned}
$$

since $\left|X_{i}\right|,\left|X_{i}^{\prime}\right| \leq c_{i}$. Putting everything together, we arrive at

$$
\mathbb{E}\left[e^{s X_{i}}\right] \leq e^{4 c_{i}^{2} s^{2} / 2}
$$

That is, $X_{i}$ is sub-Gaussian with variance factor $4 c_{i}^{2}$. By Theorem 2.4.9, $S_{n}$ is sub-Gaussian with variance factor $\sum_{i \leq n} 4 c_{i}^{2}$ and

$$
\mathbb{P}\left[S_{n} \geq t\right] \leq \exp \left(-\frac{t^{2}}{8 \sum_{i \leq n} c_{i}^{2}}\right)
$$

Proof of Theorem 2.4.10. As pointed out above, it suffices to show that $X_{i}$ is subGaussian with variance factor $\frac{1}{4}\left(b_{i}-a_{i}\right)^{2}$. This is the content of Hoeffding's lemma below (which we will use again in Chapter 3). First an observation:
Lemma 2.4.11 (Variance of bounded random variables). For any random variable $Z$ taking values in $[a, b]$ with $-\infty<a \leq b<+\infty$, we have

$$
\operatorname{Var}[Z] \leq \frac{1}{4}(b-a)^{2}
$$

Proof. Indeed

$$
\left|Z-\frac{a+b}{2}\right| \leq \frac{b-a}{2},
$$

and

$$
\operatorname{Var}[Z]=\operatorname{Var}\left[Z-\frac{a+b}{2}\right] \leq \mathbb{E}\left[\left(Z-\frac{a+b}{2}\right)^{2}\right] \leq\left(\frac{b-a}{2}\right)^{2}
$$

Lemma 2.4.12 (Hoeffding's lemma). Let $X$ be a random variable taking values in $[a, b]$ for $-\infty<a \leq b<+\infty$. Then $X \in \mathrm{~s} \mathcal{G}\left(\frac{1}{4}(b-a)^{2}\right)$.
Proof. Note first that $X-\mathbb{E} X \in[a-\mathbb{E} X, b-\mathbb{E} X]$ and $\frac{1}{4}((b-\mathbb{E} X)-(a-\mathbb{E} X))^{2}=$ $\frac{1}{4}(b-a)^{2}$. So without loss of generality we assume that $\mathbb{E} X=0$. Because $X$ is bounded, $M_{X}(s)$ is finite for all $s \in \mathbb{R}$. Hence, by (2.1.2),

$$
\Psi_{X}(0)=\log M_{X}(0)=0, \quad \Psi_{X}^{\prime}(0)=\frac{M_{X}^{\prime}(0)}{M_{X}(0)}=\mathbb{E} X=0,
$$

and by a Taylor expansion

$$
\Psi_{X}(s)=\Psi_{X}(0)+s \Psi_{X}^{\prime}(0)+\frac{s^{2}}{2} \Psi_{X}^{\prime \prime}\left(s^{*}\right)=\frac{s^{2}}{2} \Psi_{X}^{\prime \prime}\left(s^{*}\right),
$$

Hoeffding's
lemma
for some $s^{*} \in[0, s]$. Therefore it suffices to show that for all $s$

$$
\begin{equation*}
\Psi_{X}^{\prime \prime}(s) \leq \frac{1}{4}(b-a)^{2} . \tag{2.4.17}
\end{equation*}
$$

Note that

$$
\begin{aligned}
\Psi_{X}^{\prime \prime}(s) & =\frac{M_{X}^{\prime \prime}(s)}{M_{X}(s)}-\left(\frac{M_{X}^{\prime}(s)}{M_{X}(s)}\right)^{2} \\
& =\frac{1}{M_{X}(s)} \mathbb{E}\left[X^{2} e^{s X}\right]-\left(\frac{1}{M_{X}(s)} \mathbb{E}\left[X e^{s X}\right]\right)^{2} \\
& =\mathbb{E}\left[X^{2} \frac{e^{s X}}{M_{X}(s)}\right]-\left(\mathbb{E}\left[X \frac{e^{s X}}{M_{X}(s)}\right]\right)^{2} .
\end{aligned}
$$

The trick to conclude is to notice that $\frac{e^{s x}}{M_{X}^{(s)}}$ defines a density on $[a, b]$ with respect to the law of $X$. The variance under this density-the last line above-is less than $\frac{1}{4}(b-a)^{2}$ by Lemma 2.4.11. This establishes (2.4.17) and concludes the proof.

Remark 2.4.13. The change of measure above is known as tilting and is a standard trick in large deviations theory. See for example [Dur10, Section 2.6].

Since we have shown that $X_{i}$ is sub-Gaussian with variance factor $\frac{1}{4}\left(b_{i}-a_{i}\right)^{2}$, Theorem 2.4.10 follows from Theorem 2.4.9.

## Sub-exponential random variables

Unfortunately, not every random variable of interest is sub-Gaussian. A simple example is the square of a Gaussian variable. Indeed, suppose $X \sim N(0,1)$. Then $W=X^{2}$ is $\chi^{2}$-distributed and its moment-generating function can be computed explicitly. Using the change of variable $u=x \sqrt{1-2 s}$, for $s<1 / 2$,

$$
\begin{align*}
M_{W}(s) & =\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{+\infty} e^{s x^{2}} e^{-x^{2} / 2} \mathrm{~d} x \\
& =\frac{1}{\sqrt{1-2 s}} \times \frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{+\infty} e^{-u^{2} / 2} \mathrm{~d} u \\
& =\frac{1}{(1-2 s)^{1 / 2}} . \tag{2.4.18}
\end{align*}
$$

When $s \geq 1 / 2$, however, we have $M_{W}(s)=+\infty$. In particular, $W$ cannot be sub-Gaussian for any variance factor $\nu>0$. (Note that centering $W$ produces an additional factor of $e^{-s}$ in the moment-generating function which does not prevent
it from diverging.) Further confirming this observation, arguing as in (2.4.1), the upper tail of $W$ decays as

$$
\begin{aligned}
\mathbb{P}[W \geq \beta] & =\mathbb{P}[X \geq \sqrt{\beta}] \\
& \sim \sqrt{\frac{1}{2 \pi}}[\sqrt{\beta}]^{-1} \exp \left(-[\sqrt{\beta}]^{2} / 2\right) \\
& \sim \sqrt{\frac{1}{2 \pi \beta}} \exp (-\beta / 2),
\end{aligned}
$$

as $\beta \rightarrow+\infty$. That is, it decays exponentially with $\beta$, but slower than the Gaussian tail.

General case We now define a broad class of distributions which have such exponential tail decay.

Definition 2.4.14 (Sub-exponential random variable). We say that a random variable $X$ with mean $\mu$ is sub-exponential with parameters $(\nu, \alpha)$ if

$$
\Psi_{X-\mu}(s) \leq \frac{s^{2} \nu}{2}, \quad \forall|s| \leq \frac{1}{\alpha}
$$

sub-exponential
variable
for some $\nu, \alpha>0$. We write $X \in \operatorname{s\mathcal {E}}(\nu, \alpha)$.*
Observe that the key difference between (2.4.15) and (2.4.19) is the interval of $s$ over which it holds. As we will see below, the parameter $\alpha$ dictates the exponential decay rate of the tail. The specific form of the bound in (2.4.19) is natural once one notices that, as $|s| \rightarrow 0$, a centered random variable with variance $\nu$ (and a finite moment-generating function) should roughly satisfy

$$
\log \mathbb{E}\left[e^{s X}\right] \approx \log \left\{1+s \mathbb{E}[X]+\frac{s^{2}}{2} \mathbb{E}\left[X^{2}\right]\right\} \approx \log \left\{1+\frac{s^{2} \nu}{2}\right\} \approx \frac{s^{2} \nu}{2} .
$$

[^1]Returning to the $\chi^{2}$ distribution, note that from (2.4.18) we have for $|s| \leq 1 / 4$

$$
\begin{aligned}
\Psi_{W-1}(s) & =-s-\frac{1}{2} \log (1-2 s) \\
& =-s-\frac{1}{2}\left[-\sum_{i=1}^{+\infty} \frac{(2 s)^{i}}{i}\right] \\
& =\frac{s^{2}}{2}\left[4 \sum_{i=2}^{+\infty} \frac{(2 s)^{i-2}}{i}\right] \\
& \leq \frac{s^{2}}{2}\left[2 \sum_{i=2}^{+\infty}|1 / 2|^{i-2}\right] \\
& \leq \frac{s^{2}}{2} \times 4 .
\end{aligned}
$$

Hence $W \in \operatorname{s\mathcal {E}}(4,4)$.
Using the Chernoff-Cramér bound (Lemma 2.4.2), we obtain the following tail bound for sub-exponential variables.

Theorem 2.4.15 (Sub-exponential tail bound). Suppose the random variable $X$ with mean $\mu$ is sub-exponential with parameters $(\nu, \alpha)$. Then for all $\beta \in \mathbb{R}_{+}$

$$
\mathbb{P}[X-\mu \geq \beta] \leq \begin{cases}\exp \left(-\frac{\beta^{2}}{2 \nu}\right), & \text { if } 0 \leq \beta \leq \nu / \alpha,  \tag{2.4.20}\\ \exp \left(-\frac{\beta}{2 \alpha}\right), & \text { if } \beta>\nu / \alpha .\end{cases}
$$

In words, the tail decays exponentially fast at large deviations but behaves as in the sub-Gaussian case for smaller deviations. We will see below that this double-tail allows to extrapolate naturally between different regimes. First we prove the claim.

Proof of Theorem 2.4.15. We start by applying the Chernoff-Cramér bound. For any $\beta>0$ and $|s| \leq 1 / \alpha$

$$
\mathbb{P}[X-\mu \geq \beta] \leq \exp \left(-s \beta+\Psi_{X}(s)\right) \leq \exp \left(-s \beta+s^{2} \nu / 2\right) .
$$

At this point, the proof diverges from the sub-Gaussian case because the optimal choice of $s$ depends on $\beta$ because of the additional constraint $|s| \leq 1 / \alpha$. When $s^{*}=\beta / \nu$ satisfies $s^{*} \leq 1 / \alpha$, the quadratic function of $s$ in the exponent is minimized at $s^{*}$, giving the bound

$$
\mathbb{P}[X \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right)
$$

for $0 \leq \beta \leq \nu / \alpha$.

On the other hand, when $\beta>\nu / \alpha$, the exponent is strictly decreasing over the interval $s \leq 1 / \alpha$. Hence the optimal choice is $s^{*}=1 / \alpha$, which produces the bound

$$
\begin{aligned}
\mathbb{P}[X \geq \beta] & \leq \exp \left(-\frac{\beta}{\alpha}+\frac{\nu}{2 \alpha^{2}}\right) \\
& <\exp \left(-\frac{\beta}{\alpha}+\frac{\beta}{2 \alpha}\right) \\
& =\exp \left(-\frac{\beta}{2 \alpha}\right),
\end{aligned}
$$

where we used that $\nu<\beta \alpha$ on the second line.
For (weighted) sums of independent sub-exponential random variables, we get the following.

Theorem 2.4.16 (General Bernstein inequality). Suppose $X_{1}, \ldots, X_{n}$ are independent random variables where, for each $i, X_{i} \in \operatorname{s\mathcal {E}}\left(\nu_{i}, \alpha_{i}\right)$ with $0<\nu_{i}, \alpha_{i}<$ $+\infty$. For $w_{1}, \ldots, w_{n} \in \mathbb{R}$, let $S_{n}=\sum_{i \leq n} w_{i} X_{i}$. Then

$$
S_{n} \in \operatorname{s\mathcal {E}}\left(\sum_{i=1}^{n} w_{i}^{2} \nu_{i}, \max _{i}\left|w_{i}\right| \alpha_{i}\right) .
$$

In particular, for all $\beta>0$,

$$
\mathbb{P}\left[S_{n}-\mathbb{E} S_{n} \geq \beta\right] \leq \begin{cases}\exp \left(-\frac{\beta^{2}}{2 \sum_{i=1}^{n} w_{i}^{2} \nu_{i}}\right), & \text { if } 0 \leq \beta \leq \frac{\sum_{i=1}^{n} w_{i}^{2} \nu_{i}}{\max x_{i}\left|w_{i}\right| \alpha_{i}}, \\ \exp \left(-\frac{\beta}{2 \max _{i}\left|w_{i}\right| \alpha_{i}}\right), & \text { if } \beta>\frac{\sum_{i=1}^{n} w_{i}^{2} \nu_{i}}{\max _{i}\left|w_{i}\right| \alpha_{i}} .\end{cases}
$$

Proof. By independence and (2.1.3),

$$
\Psi_{S_{n}}(s)=\sum_{i \leq n} \Psi_{w_{i} X_{i}}(s)=\sum_{i \leq n} \Psi_{X_{i}}\left(s w_{i}\right) \leq \sum_{i \leq n} \frac{\left(s w_{i}\right)^{2} \nu_{i}}{2}=\frac{s^{2} \sum_{i \leq n} w_{i}^{2} \nu_{i}}{2},
$$

provided $\left|s w_{i}\right| \leq 1 / \alpha_{i}$ for all $i$, that is,

$$
|s| \leq \frac{1}{\max _{i}\left|w_{i}\right| \alpha_{i}}
$$

Bounded random variables: revisited We apply the previous result to bounded random variables.

Theorem 2.4.17 (Bernstein's inequality for bounded variables). Let $X_{1}, \ldots, X_{n}$ be independent random variables where, for each $i, X_{i}$ has mean $\mu_{i}$, variance $\sigma_{i}^{2}$ and satisfies $\left|X_{i}-\mu_{i}\right| \leq c$ for some $0<c<+\infty$. Let $S_{n}=\sum_{i \leq n} X_{i}$. For all $\beta>0$,

$$
\mathbb{P}\left[S_{n}-\mathbb{E} S_{n} \geq \beta\right] \leq \begin{cases}\exp \left(-\frac{\beta^{2}}{4 \sum_{i=1}^{n} \sigma_{i}^{2}}\right), & \text { if } 0 \leq \beta \leq \frac{\sum_{i=1}^{n} \sigma_{i}^{2}}{c}, \\ \exp \left(-\frac{\beta}{4 c}\right), & \text { if } \beta>\frac{\sum_{i=1}^{n} \sigma_{i}^{2}}{c} .\end{cases}
$$

Proof. We claim that $X_{i} \in \operatorname{se}\left(2 \sigma_{i}^{2}, 2 c\right)$. To establish the claim, we derive a bound on all moments of $X_{i}$. Note that for all integers $k \geq 2$

$$
\mathbb{E}\left|X_{i}-\mu_{i}\right|^{k} \leq c^{k-2} \mathbb{E}\left|X_{i}-\mu_{i}\right|^{2}=c^{k-2} \sigma_{i}^{2} .
$$

Hence, first applying the dominated convergence theorem (Proposition B.4.14) to establish the limit, we have for $|s| \leq \frac{1}{2 c}$

$$
\begin{aligned}
\mathbb{E}\left[e^{s\left(X_{i}-\mu_{i}\right)}\right] & =\sum_{k=0}^{+\infty} \frac{s^{k}}{k!} \mathbb{E}\left[\left(X_{i}-\mu_{i}\right)^{k}\right] \\
& \leq 1+s \mathbb{E}\left[\left(X_{i}-\mu_{i}\right)\right]+\sum_{k=2}^{+\infty} \frac{s^{k}}{k!} c^{k-2} \sigma_{i}^{2} \\
& \leq 1+\frac{s^{2} \sigma_{i}^{2}}{2}+\frac{s^{2} \sigma_{i}^{2}}{3!} \sum_{k=3}^{+\infty}(c s)^{k-2} \\
& =1+\frac{s^{2} \sigma_{i}^{2}}{2}\left\{1+\frac{1}{3} \frac{c s}{1-c s}\right\} \\
& \leq 1+\frac{s^{2} \sigma_{i}^{2}}{2}\left\{1+\frac{1}{3} \frac{1 / 2}{1-1 / 2}\right\} \\
& \leq 1+\frac{s^{2}}{2} 2 \sigma_{i}^{2} \\
& \leq \exp \left(\frac{s^{2}}{2} 2 \sigma_{i}^{2}\right) .
\end{aligned}
$$

Using the general Bernstein inequality (Theorem 2.4.16) gives the result.
It may seem counter-intuitive to derive a tail bound based on the sub-exponential property of bounded random variables when we have already done so using their
sub-Gaussian behavior. After all, the latter is on the surface a strengthening of the former. However, note that we have obtained a better bound in Theorem 2.4.17 than we did in Theorem 2.4.10-when $\beta$ is not too large. That improvement stems from the use of the (actual) variance for moderate deviations. This is easier to appreciate on an example.

Example 2.4.18 (Erdős-Rényi: maximum degree). Let $G_{n}=\left(V_{n}, E_{n}\right) \sim \mathbb{G}_{n, p_{n}}$ be a random graph with $n$ vertices and density $p_{n}$ under the Erdős-Rényi model (Definition 1.2.2). Recall that two vertices $u, v \in V_{n}$ are adjacent if $\{u, v\} \in E_{n}$ and that the set of adjacent vertices of $v$, denoted by $N(v)$, is called the neighborhood of $v$. The degree of $v$ is the size of its neighborhood, that is, $\delta(v)=|N(v)|$. Here we study the maximum degree of $G_{n}$

$$
D_{n}=\max _{v \in V_{n}} \delta(v) .
$$

We focus on the regime $n p_{n}=\omega(\log n)$. Note that, for any vertex $v \in V_{n}$, its degree is $\operatorname{Bin}\left(n-1, p_{n}\right)$ by independence of the edges. In particular its expected degree is $(n-1) p_{n}$. To prove a high-probability upper bound on the maximum $D_{n}$, we need to control the deviation of the degree of each vertex from its expectation. Observe that the degrees are not independent. Instead we apply a union bound over all vertices, after using a tail bound.

Claim 2.4.19. For any $\varepsilon>0$, as $n \rightarrow+\infty$,

$$
\mathbb{P}\left[\left|D_{n}-(n-1) p_{n}\right| \geq 2 \sqrt{(1+\varepsilon) n p_{n} \log n}\right] \rightarrow 0
$$

Proof. For a fixed vertex $v$, think of $\delta(v)=S_{n-1} \sim \operatorname{Bin}\left(n-1, p_{n}\right)$ as a sum of $n-1$ independent $\{0,1\}$-valued random variables, one for each possible edge. That is, $S_{n-1}=\sum_{i=1}^{n-1} X_{i}$ where the $X_{i}$ s are bounded random variables. The mean of $X_{i}$ is $p_{n}$ and its variance is $p_{n}\left(1-p_{n}\right)$. So in Bernstein's inequality (Theorem 2.4.17), we can take $\mu_{i}:=p_{n}, \sigma_{i}^{2}:=p_{n}\left(1-p_{n}\right)$ and $c:=1$ for all $i$. We get

$$
\mathbb{P}\left[S_{n-1} \geq(n-1) p_{n}+\beta\right] \leq \begin{cases}\exp \left(-\frac{\beta^{2}}{4 \nu}\right), & \text { if } 0 \leq \beta \leq \nu, \\ \exp \left(-\frac{\beta}{4}\right), & \text { if } \beta>\nu\end{cases}
$$

where $\nu=(n-1) p_{n}\left(1-p_{n}\right)=\omega(\log n)$ by assumption. We choose $\beta$ to be the smallest value that will produce a tail probability less than $n^{-1-\varepsilon}$ for $\varepsilon>0$, that is,

$$
\beta=\sqrt{4(n-1) p_{n}\left(1-p_{n}\right)} \times \sqrt{(1+\varepsilon) \log n}=o(\nu),
$$

which falls in the lower regime of the tail bound. In particular, $\beta=o\left(n p_{n}\right)$ (i.e., the deviation is much smaller than the expectation). Finally by a union bound over $v \in V_{n}$

$$
\mathbb{P}\left[D_{n} \geq(n-1) p_{n}+\sqrt{4(1+\varepsilon) p_{n}\left(1-p_{n}\right)(n-1) \log n}\right] \leq n \times \frac{1}{n^{1+\varepsilon}} \rightarrow 0 .
$$

The same holds in the other direction. That proves the claim.
Had we used Hoeffding's inequality (Theorem 2.4.10) in the proof of Claim 2.4.19 we would have had to take $\beta=\sqrt{(1+\varepsilon) n \log n}$. That would have produced a much weaker bound when $p_{n}=o(1)$. Indeed the advantage of Bernstein's inequality is that it makes explicit use of the variance, which when $p_{n}=o(1)$ is much smaller than the worst case for bounded variables.

### 2.4.3 $\triangleright$ Probabilistic analysis of algorithms: knapsack problem

In a knapsack problem, we have $n$ items. Item $i$ has weight $W_{i} \in[0,1]$ and value $V_{i} \in[0,1]$. Given a weight bound $\mathcal{W}$, we want to pack as valuable a collection of items in the knapsack under the constraint that the total weight is less or equal than $\mathcal{W}$. Formally we seek a solution to the optimization problem

$$
\begin{equation*}
Z^{*}=\max \left\{\sum_{j=1}^{n} x_{j} V_{j}: x_{1}, \ldots, x_{n} \in[0,1], \sum_{j=1}^{n} x_{j} W_{j} \leq \mathcal{W}\right\} . \tag{2.4.21}
\end{equation*}
$$

This is the fractional knapsack problem, where we allow a fraction of an item to be added to the knapsack.

It is used as a computationally tractable relaxation of the 0-1 knapsack problem, which also includes the combinatorial constraint $x_{j} \in\{0,1\}, \forall j$. Indeed, it turns out that the optimization problem (2.4.21) is solved exactly by a simple greedy solution (see Exercise 2.8 for a formal proof of correctness): let $\pi$ be a permutation of $\{1, \ldots, n\}$ that puts the items in decreasing order of value per unit weight

$$
\frac{V_{\pi(1)}}{W_{\pi(1)}} \geq \frac{V_{\pi(2)}}{W_{\pi(2)}} \geq \cdots \geq \frac{V_{\pi(n)}}{W_{\pi(n)}}
$$

add the items in that order until the first time the weight constraints is violated; include whatever fraction of that last item that will fit. This greedy algorithm has a natural geometric interpretation, depicted in Figure 2.7, that will be useful. We associate item $j$ to a point $\left(W_{j}, V_{j}\right) \in[0,1]^{2}$ and keep only those items falling on or above a line with slope $\theta$ chosen to satisfy the total weight constraint. Specifically, let

$$
\Delta_{\theta}=\left\{j \in[n]: V_{j}>\theta W_{j}\right\},
$$



Figure 2.7: Vizualization of the greedy algorithm.

$$
\Lambda_{\theta}=\left\{j \in[n]: V_{j}=\theta W_{j}\right\},
$$

and

$$
\Theta^{*}=\inf \left\{\theta \geq 0: W_{\Delta_{\theta}}<\mathcal{W}\right\} .
$$

where, for a subset of items $J \subset[n], W_{J}=\sum_{j \in J} W_{j}$ (and $V_{J}$ is similarly defined).

We consider a stochastic version of the fractional knapsack problem where the weights and values are i.i.d. random variables picked uniformly at random in $[0,1]$. Characterizing $Z^{*}$ (e.g., its moments or distribution) is not straightforward. Here we show that $Z^{*}$ is highly concentrated around a natural quantity. Observe that, under our probabilistic model, almost surely $\left|\Lambda_{\theta}\right| \in\{0,1\}$ for any $\theta \geq 0$. Hence, there are two cases. Either $\Theta^{*}=0$, in which case all items fit in the knapsack so $Z^{*}=\sum_{j=1}^{n} V_{j}$. Or $\Theta^{*}>0$, in which case $\left|\Lambda_{\Theta^{*}}\right|=1$ and

$$
\begin{equation*}
Z^{*}=V_{\Delta_{\Theta^{*}}}+\frac{\mathcal{W}-W_{\Delta_{\Theta^{*}}}}{W_{\Lambda_{\Theta^{*}}}} V_{\Lambda_{\Theta^{*}}} . \tag{2.4.22}
\end{equation*}
$$

One interesting regime is $\mathcal{W}=\tau n$ for some constant $\tau>0$. Clearly, $\tau>1$ is trivial. In fact, because

$$
\mathbb{E}\left[\sum_{j=1}^{n} W_{j}\right]=n \mathbb{E}\left[W_{1}\right]=\frac{1}{2} n,
$$

we assume that $\tau \leq 1 / 2$. To further simplify the calculations, we restrict ourselves to the case $\tau \in(1 / 6,1 / 2)$. (See Exercise 2.8 for the remaining case.) In this regime, we show that $Z^{*}$ grows linearly with $n$ and give a bound on its deviation.

Although $Z^{*}$ is technically a sum of random variables, the choice of $\Theta^{*}$ correlates them and we cannot apply our concentration bounds directly. Instead we show that $\Theta^{*}$ itself can be controlled well. It is natural to conjecture that $\Theta^{*}$ is approximately equal to a solution $\theta_{\tau}$ of the expected constraint equation $\mathbb{E}\left[W_{\Delta_{\theta_{\tau}}}\right]=\mathcal{W}$, that is,

$$
\begin{equation*}
n \bar{w}_{\theta_{\tau}}=n \tau, \tag{2.4.23}
\end{equation*}
$$

where $\bar{w}_{\theta}$ is defined through

$$
\begin{aligned}
\mathbb{E}\left[W_{\Delta_{\theta}}\right] & =\mathbb{E}\left[\sum_{j \in \Delta_{\theta}} W_{j}\right] \\
& =\mathbb{E}\left[\sum_{j=1}^{n} \mathbf{1}\left\{V_{j}>\theta W_{j}\right\} W_{j}\right] \\
& =n \mathbb{E}\left[\mathbf{1}\left\{V_{1}>\theta W_{1}\right\} W_{1}\right] \\
& =n \bar{w}_{\theta} .
\end{aligned}
$$

Similarly, we define

$$
\bar{v}_{\theta}:=\mathbb{E}\left[\mathbf{1}\left\{V_{1}>\theta W_{1}\right\} V_{1}\right] .
$$

We see directly from the definitions that both $\bar{w}_{\theta}$ and $\bar{v}_{\theta}$ are monotone as functions of $\theta$.

Our main claim is the following.
Claim 2.4.20. There is a constant $c>0$ such that for any $\delta>0$

$$
\mathbb{P}\left[\left|Z^{*}-n \bar{v}_{\theta_{\tau}}\right| \geq \sqrt{c n \log \delta^{-1}}\right] \leq \delta,
$$

for all $n$ large enough.
Proof. Because all weights and values are in [0, 1], it follows from (2.4.22) that

$$
\begin{equation*}
V_{\Delta_{\Theta^{*}}} \leq Z^{*} \leq V_{\Delta_{\Theta^{*}}}+1, \tag{2.4.24}
\end{equation*}
$$

and it will suffice to work with $V_{\Delta_{\Theta^{*}}}$. The idea of the proof is to show that $\Theta^{*}$ is close to $\theta_{\tau}$ by establishing that $W_{\Delta_{\theta}}$ is highly likely to be less than $\tau n$ when $\theta>\theta_{\tau}$, while the opposite holds when $\theta<\theta_{\tau}$. For this, we view $W_{\Delta_{\theta}}$ as a sum of independent bounded random variables and use Hoeffding's inequality (Theorem 2.4.10).

Controlling $\Theta^{*}$. First, it will be useful to compute $\bar{w}_{\theta}$ and $\theta_{\tau}$ analytically. By definition,

$$
\begin{align*}
\bar{w}_{\theta} & =\mathbb{E}\left[\mathbf{1}\left\{V_{1}>\theta W_{1}\right\} W_{1}\right] \\
& =\int_{0}^{1} \int_{0}^{1} \mathbf{1}\{y>\theta x\} x \mathrm{~d} y \mathrm{~d} x \\
& =\int_{0}^{1 \wedge 1 / \theta}(1-\theta x) x \mathrm{~d} x \\
& = \begin{cases}\frac{1}{2}-\frac{1}{3} \theta & \text { if } \theta \leq 1, \\
\frac{1}{6 \theta^{2}} & \text { otherwise. }\end{cases} \tag{2.4.25}
\end{align*}
$$

Plugging back into (2.4.23), we get the unique solution

$$
\theta_{\tau}:=3\left(\frac{1}{2}-\tau\right) \in(0,1)
$$

for the range $\tau \in(1 / 6,1 / 2)$.
Now observe that, for each fixed $\theta$, the quantity

$$
W_{\Delta_{\theta}}=\sum_{j=1}^{n} \mathbf{1}\left\{V_{j}>\theta W_{j}\right\} W_{j},
$$

is a sum of independent random variables taking values in $[0,1]$. Hence, for any $\beta>0$, Hoeffding's inequality gives

$$
\mathbb{P}\left[W_{\Delta_{\theta}}-n \bar{w}_{\theta} \geq \beta\right] \leq \exp \left(-\frac{2 \beta^{2}}{n}\right) .
$$

Using this inequality with $\theta=\theta_{\tau}+\frac{C}{\sqrt{n}}$ (with $n$ large enough that $\theta<1$ ) and $\beta=(C / 3) \sqrt{n}$ gives

$$
\mathbb{P}\left[W_{\Delta_{\theta_{\tau}+\frac{C}{\sqrt{n}}}}-n\left(\frac{1}{2}-\frac{1}{3} \theta_{\tau}-\frac{C / 3}{\sqrt{n}}\right) \geq(C / 3) \sqrt{n}\right] \leq \exp \left(-2(C / 3)^{2}\right),
$$

where we used (2.4.25). After rearranging and using that $n\left(\frac{1}{2}-\frac{1}{3} \theta_{\tau}\right)=n \tau$ by (2.4.23) and (2.4.25), this gives

$$
\mathbb{P}\left[\Theta^{*} \geq \theta_{\tau}+\frac{C}{\sqrt{n}}\right]=\mathbb{P}\left[W_{\Delta_{\theta_{\tau}+\frac{C}{\sqrt{n}}}} \geq n \tau\right] \leq \exp \left(-2(C / 3)^{2}\right) .
$$

Applying the same argument to $-W_{\Delta_{\theta}}$ with $\theta=\theta_{\tau}-\frac{C}{\sqrt{n}}$ and combining with the previous inequality gives

$$
\begin{equation*}
\mathbb{P}\left[\left|\Theta^{*}-\theta_{\tau}\right|>\frac{C}{\sqrt{n}}\right] \leq 2 \exp \left(-2(C / 3)^{2}\right), \tag{2.4.26}
\end{equation*}
$$

assuming $n$ is large enough.
Controlling $Z^{*}$. We conclude by applying Hoeffding's inequality to $V_{\Delta_{\theta}}$. Arguing as above with the same $\theta$ 's and $\beta$ (but the roles of the two cases reversed), we obtain

$$
\begin{equation*}
\mathbb{P}\left[V_{\Delta_{\theta_{T}-\frac{C}{\sqrt{n}}}}-n \bar{v}_{\theta_{\tau}-\frac{C}{\sqrt{n}}} \geq(C / 3) \sqrt{n}\right] \leq \exp \left(-2(C / 3)^{2}\right) \tag{2.4.27}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{P}\left[V_{\Delta_{\theta_{\tau}+\frac{C}{\sqrt{n}}}}-n \bar{v}_{\theta_{\tau}+\frac{C}{\sqrt{n}}} \leq-(C / 3) \sqrt{n}\right] \leq \exp \left(-2(C / 3)^{2}\right) . \tag{2.4.28}
\end{equation*}
$$

Again, it will be useful to compute $\bar{v}_{\theta}$ analytically. By definition,

$$
\begin{aligned}
\bar{v}_{\theta} & =\mathbb{E}\left[\mathbf{1}\left\{V_{1}>\theta W_{1}\right\} V_{1}\right] \\
& =\int_{0}^{1} \int_{0}^{1} \mathbf{1}\{y>\theta x\} y \mathrm{~d} x \mathrm{~d} y \\
& =\int_{0}^{1 \wedge \theta} \frac{y^{2}}{\theta} \mathrm{~d} y+\int_{1 \wedge \theta}^{1} y \mathrm{~d} y \\
& = \begin{cases}\frac{1}{2}-\frac{1}{6} \theta^{2} & \text { if } \theta \leq 1, \\
\frac{1}{3 \theta} & \text { otherwise. }\end{cases}
\end{aligned}
$$

Assuming $n$ is large enough (recall that $\theta_{\tau}<1$ ), we get

$$
\bar{v}_{\theta_{\tau}}-\bar{v}_{\theta_{\tau}+\frac{C}{\sqrt{n}}}=\frac{1}{6}\left(2 \frac{C}{\sqrt{n}} \theta_{\tau}+\frac{C^{2}}{n}\right) \leq \frac{C}{\sqrt{n}} .
$$

A quick check reveals that, similarly, $\bar{v}_{\theta_{\tau}-\frac{C}{\sqrt{n}}}-\bar{v}_{\theta_{\tau}} \leq \frac{C}{\sqrt{n}}$. Plugging back into (2.4.27) and (2.4.28) gives

$$
\begin{equation*}
\mathbb{P}\left[V_{\Delta_{\theta_{\tau}-\frac{C}{\sqrt{n}}}} \geq n \bar{v}_{\theta_{\tau}}+2 C \sqrt{n}\right] \leq \exp \left(-2(C / 3)^{2}\right) \tag{2.4.29}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{P}\left[V_{\Delta_{\theta_{\tau}+\frac{C}{\sqrt{n}}}} \leq n \bar{v}_{\theta_{\tau}}-2 C \sqrt{n}\right] \leq \exp \left(-2(C / 3)^{2}\right) . \tag{2.4.30}
\end{equation*}
$$

Observe that the following monotonicity property holds almost surely

$$
\begin{equation*}
\theta_{0} \leq \theta_{1} \leq \theta_{2} \Longrightarrow V_{\Delta_{\theta_{0}}} \geq V_{\Delta_{\theta_{1}}} \geq V_{\Delta_{\theta_{2}}} \tag{2.4.31}
\end{equation*}
$$

Combining (2.4.24), (2.4.26), (2.4.29), (2.4.30) and (2.4.31), we obtain

$$
\mathbb{P}\left[\left|Z^{*}-n \bar{v}_{\theta_{\tau}}\right|>2 C \sqrt{n}\right] \leq 4 \exp \left(-2(C / 3)^{2}\right),
$$

for $n$ large enough. Choosing $C$ appropriately gives the claim.
A similar bound is proved for the 0-1 knapsack problem in Exercise 2.9.

### 2.4.4 Epsilon-nets and chaining

Suppose we are interested in bounding the expectation or tail of the supremum of a stochastic process

$$
\sup _{t \in \mathscr{T}} X_{t}
$$

where $\mathscr{T}$ is an arbitrary index set and the $X_{t} \mathrm{~s}$ are real-valued random variables. To avoid measurability issues, we assume throughout that $\mathscr{T}$ is countable. ${ }^{\dagger}$ Note that $t$ does not in general need to be a "time" index.

So far we have developed tools that can handle cases where $\mathscr{T}$ is finite. When the supremum is over an infinite index set, however, new ideas are required. One way to proceed is to apply a tail inequality to a sufficiently dense finite subset of the index set and then extend the resulting bound by a Lipschitz continuity argument. We present this type of approach in this section, as well as a multi-scale version known as chaining.

First we summarize one important special case that will be useful below: $\mathscr{T}$ is finite and $X_{t}$ is sub-Gaussian.

Theorem 2.4.21 (Maximal inequalities: sub-Gaussian case). Let $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ be a stochastic process with finite index set $\mathscr{T}$. Assume that there is $\nu>0$ such that, for all $t, X_{t} \in \operatorname{sG}(\nu)$ and $\mathbb{E}\left[X_{t}\right]=0$. Then

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] \leq \sqrt{2 \nu \log |\mathscr{T}|},
$$

and, for all $\beta>0$,

$$
\mathbb{P}\left[\sup _{t \in \mathscr{T}} X_{t} \geq \sqrt{2 \nu \log |\mathscr{T}|}+\beta\right] \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right) .
$$

Proof. For the expectation, we apply a variation on the Chernoff-Cramér method (Section 2.4). Naively, we could bound the supremum $\sup _{t \in \mathscr{T}} X_{t}$ by the sum $\sum_{t \in \mathscr{T}}\left|X_{t}\right|$, but that would lead to a bound growing linearly with the cardinality $|\mathscr{T}|$. Instead we first take an exponential, which tends to amplify the largest term and produces a much stronger bound. Specifically, by Jensen's inequality (Theorem B.4.15), for any $s>0$

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right]=\frac{1}{s} \mathbb{E}\left[\sup _{t \in \mathscr{T}} s X_{t}\right] \leq \frac{1}{s} \log \mathbb{E}\left[\exp \left(\sup _{t \in \mathscr{T}} s X_{t}\right)\right] .
$$

[^2]Since $e^{a \vee b} \leq e^{a}+e^{b}$ by the non-negativity of the exponential, we can bound

$$
\begin{aligned}
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] & \leq \frac{1}{s} \log \left[\sum_{t \in \mathscr{T}} \mathbb{E}\left[\exp \left(s X_{t}\right)\right]\right] \\
& =\frac{1}{s} \log \left[\sum_{t \in \mathscr{T}} M_{X_{t}}(s)\right] \\
& \leq \frac{1}{s} \log \left[|\mathscr{T}| e^{\frac{s^{2} \nu}{2}}\right] \\
& =\frac{\log |\mathscr{T}|}{s}+\frac{s \nu}{2} .
\end{aligned}
$$

The optimal choice of $s$ (i.e., leading to the least upper bound) is when the two terms in the sum above are equal, that is, $s=\sqrt{2 \nu^{-1} \log |\mathscr{T}|}$, which gives finally

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] \leq \sqrt{2 \nu \log |\mathscr{T}|},
$$

as claimed.
For the tail inequality, we use a union bound and (2.4.16)

$$
\begin{aligned}
\mathbb{P}\left[\sup _{t \in \mathscr{T}} X_{t} \geq \sqrt{2 \nu \log |\mathscr{T}|}+\beta\right] & \leq \sum_{t \in \mathscr{T}} \mathbb{P}\left[X_{t} \geq \sqrt{2 \nu \log |\mathscr{T}|}+\beta\right] \\
& \leq|\mathscr{T}| \exp \left(-\frac{(\sqrt{2 \nu \log |\mathscr{T}|}+\beta)^{2}}{2 \nu}\right) \\
& \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right)
\end{aligned}
$$

as claimed, where we used that $\beta>0$ on the last line.

## Epsilon-nets and covering numbers

Moving on to infinite index sets, we first define the notion of an $\varepsilon$-net. This notion requires that a pseudometric $\rho$ (i.e., $\rho: \mathscr{T} \times \mathscr{T} \rightarrow \mathbb{R}_{+}$is symmetric and satisfies the triangle inequality) be defined over $\mathscr{T}$.

Definition 2.4.22 ( $\varepsilon$-net). Let $\mathscr{T}$ be a subset of a pseudometric space $(M, \rho)$ and let $\varepsilon>0$. The collection of points $N \subseteq M$ is called an $\varepsilon$-net of $\mathscr{T}$ if

$$
\mathscr{T} \subseteq \bigcup_{t \in N} B_{\rho}(t, \varepsilon),
$$

where $B_{\rho}(t, \varepsilon)=\{s \in \mathscr{T}: \rho(s, t) \leq \varepsilon\}$, that is, each element of $\mathscr{T}$ is within distance $\varepsilon$ of an element in $N$. The smallest cardinality of an $\varepsilon$-net of $\mathscr{T}$ is called the covering number

$$
\mathcal{N}(\mathscr{T}, \rho, \varepsilon)=\inf \{|N|: N \text { is an } \varepsilon \text {-net of } \mathscr{T}\} .
$$

A natural way to construct an $\varepsilon$-net is the following algorithm. Start with $N=\emptyset$ and successively add a point from $\mathscr{T}$ to $N$ at distance at least $\varepsilon$ from all other previous points until it is not possible to do so anymore. Provided $\mathscr{T}$ is compact, this procedure will terminate after a finite number of steps. This leads to the following dual perspective.

Definition 2.4.23 ( $\varepsilon$-packing). Let $\mathscr{T}$ be a subset of a pseudometric space $(M, \rho)$ and let $\varepsilon>0$. The collection of points $N \subseteq \mathscr{T}$ is called an $\varepsilon$-packing of $\mathscr{T}$ if

$$
t \notin B_{\rho}\left(t^{\prime}, \varepsilon\right), \quad \forall t \neq t^{\prime} \in N
$$

that is, every pair of elements of $N$ is at distance strictly greater than $\varepsilon$. The largest cardinality of an $\varepsilon$-packing of $\mathscr{T}$ is called the packing number

$$
\mathcal{P}(\mathscr{T}, \rho, \varepsilon)=\sup \{|N|: N \text { is an } \varepsilon \text {-packing of } \mathscr{T}\} .
$$

Lemma 2.4.24 (Covering and packing numbers). For any $\mathscr{T} \subseteq M$ and all $\varepsilon>0$,

$$
\mathcal{N}(\mathscr{T}, \rho, \varepsilon) \leq \mathcal{P}(\mathscr{T}, \rho, \varepsilon)
$$

Proof. Observe that a maximal $\varepsilon$-packing $N$ is an $\varepsilon$-net. Indeed, by maximality, any element of $\mathscr{T} \backslash N$ is at distance at most $\varepsilon$ from an element of $N$.

Example 2.4.25 (Sphere in $\mathbb{R}^{k}$ ). We let $\mathbb{B}^{k}(x, \varepsilon)$ be the ball of radius $\varepsilon$ around $x \in \mathbb{R}^{k}$ with the Euclidean metric. We let $\mathbb{S}^{k-1}$ be the sphere of radius 1 centered around the origin $\mathbf{0}$, that is, the surface of $\mathbb{B}^{k}(\mathbf{0}, 1)$. Let $0<\varepsilon<1$.

Claim 2.4.26. For $S:=\mathbb{S}^{k-1}$,

$$
\mathcal{N}(S, \rho, \varepsilon) \leq\left(\frac{3}{\varepsilon}\right)^{k}
$$

Proof. Let $N$ be any maximal $\varepsilon$-packing of $S$. We show that $|N| \leq(3 / \varepsilon)^{k}$, which implies the claim by Lemma 2.4.24. The balls of radius $\varepsilon / 2$ around points in $N$, $\left\{\mathbb{B}^{k}\left(x_{i}, \varepsilon / 2\right): x_{i} \in N\right\}$, satisfy two properties:

1. They are pairwise disjoint: if $z \in \mathbb{B}^{k}\left(x_{i}, \varepsilon / 2\right) \cap \mathbb{B}^{k}\left(x_{j}, \varepsilon / 2\right)$, then $\| x_{i}$ $x_{j}\left\|_{2} \leq\right\| x_{i}-z\left\|_{2}+\right\| x_{j}-z \|_{2} \leq \varepsilon$, a contradiction.
2. They are included in the ball of radius $3 / 2$ around the origin: if $z \in \mathbb{B}^{k}\left(x_{i}, \varepsilon / 2\right)$, then $\|z\|_{2} \leq\left\|z-x_{i}\right\|_{2}+\left\|x_{i}\right\| \leq \varepsilon / 2+1 \leq 3 / 2$.

The volume of a ball of radius $\varepsilon / 2$ is $\frac{\pi^{k / 2}(\varepsilon / 2)^{k}}{\Gamma(k / 2+1)}$ and that of a ball of radius $3 / 2$ is $\frac{\pi^{k / 2}(3 / 2)^{k}}{\Gamma(k / 2+1)}$. Dividing one by the other proves the claim.

This bound will be useful later.
The basic approach to use an $\varepsilon$-net for controlling the supremum of a stochastic process is the following. We say that a stochastic process $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ is Lipschitz for pseudometric $\rho$ on $\mathscr{T}$ if there is a random variable $0<K<+\infty$ such that

$$
\left|X_{t}-X_{s}\right| \leq K \rho(s, t), \quad \forall s, t \in \mathscr{T}
$$

If in addition $X_{t}$ is sub-Gaussian for all $t$, then we can bound the expectation or tail probability of the supremum of $\left\{X_{t}\right\}_{t \in \mathscr{T}}$-if we can bound the expectation or tail probability of the (random) Lipschitz constant $K$ itself. To see this, let $N \subseteq \mathscr{T}$ be an $\varepsilon$-net of $\mathscr{T}$ and, for each $t \in \mathscr{T}$, let $\pi(t)$ be the closest element of $N$ to $t$. We will refer to $\pi$ as the projection map of $N$. We then have the inequality

$$
\begin{equation*}
\sup _{t \in \mathscr{T}} X_{t} \leq \sup _{t \in \mathscr{T}}\left(X_{t}-X_{\pi(t)}\right)+\sup _{t \in \mathscr{T}} X_{\pi(t)} \leq K \varepsilon+\sup _{s \in N} X_{s} \tag{2.4.32}
\end{equation*}
$$

where we can use Theorem 2.4.21 to bound the last term. ${ }^{\ddagger}$ We give an example of this type of argument next (although we do not apply the above bound directly). Another example (where (2.4.32) is used this time) can be found in Section 2.4.5.

Example 2.4.27 (Spectral norm of a random matrix). For an $m \times n$ matrix $A \in$ $\mathbb{R}^{m \times n}$, the spectral norm (or induced 2-norm, or 2-norm for short) is defined as

$$
\begin{equation*}
\|A\|_{2}:=\sup _{\mathbf{x} \in \mathbb{R}^{n} \backslash\{0\}} \frac{\|A \mathbf{x}\|_{2}}{\|\mathbf{x}\|_{2}}=\sup _{\mathbf{x} \in \mathbb{S}^{n-1}}\|A \mathbf{x}\|_{2}=\sup _{\substack{\mathbf{x} \in \mathbb{S}^{n-1} \\ \mathbf{y} \in \mathbb{S}^{m-1}}}\langle A \mathbf{x}, \mathbf{y}\rangle \tag{2.4.33}
\end{equation*}
$$

where $\mathbb{S}^{n-1}$ is the sphere of Euclidean radius 1 around the origin in $\mathbb{R}^{n}$. The rightmost expression, which is central to our developments, is justified in Exercise 5.4.

We will be interested in the case where $A$ is a random matrix with independent entries. One key observation is that the quantity $\langle A \mathbf{x}, \mathbf{y}\rangle$ can then be seen as a linear combination of independent random variables

$$
\langle A \mathbf{x}, \mathbf{y}\rangle=\sum_{i, j} x_{j} y_{i} A_{i j}
$$

[^3]Hence we will be able to apply our previous tail bounds. However, we also need to deal with the supremum.

Theorem 2.4.28 (Upper tail of the spectral norm). Let $A \in \mathbb{R}^{m \times n}$ be a random matrix whose entries are centered, independent and sub-Gaussian with variance factor $\nu$. Then there exists a constant $0<C<+\infty$ such that, for all $t>0$,

$$
\|A\|_{2} \leq C \sqrt{\nu}(\sqrt{m}+\sqrt{n}+t)
$$

with probability at least $1-e^{-t^{2}}$.
Without the independence assumption, the norm can be much larger in general (see Exercise 2.15).

Proof. Fix $\varepsilon=1 / 4$. By Claim 2.4.26, there is an $\varepsilon$-net $N \subseteq \mathbb{S}^{n-1}$ (respectively $M \subseteq \mathbb{S}^{m-1}$ ) of $\mathbb{S}^{n-1}$ (respectively $\mathbb{S}^{m-1}$ ) with $|N| \leq 12^{n}$ (respectively $|M| \leq$ $12^{m}$ ). We proceed in two steps:

1. We first apply the general Hoeffding inequality (Theorem 2.4.9) to control the deviations of the supremum in (2.4.33) restricted to $N$ and $M$.
2. We then extend the bound to the full supremum by Lipschitz continuity.

Formally, the result follows from the following two lemmas.
Lemma 2.4.29. Let $N$ and $M$ be as above. There is a constant $C$ large enough (not depending on $n, m$ ) such that, for all $t>0$,

$$
\mathbb{P}\left[\max _{\substack{\mathbf{x} \in N \\ \mathbf{y} \in M}}\langle A \mathbf{x}, \mathbf{y}\rangle \geq \frac{1}{2} C \sqrt{\nu}(\sqrt{m}+\sqrt{n}+t)\right] \leq e^{-t^{2}} .
$$

Lemma 2.4.30. For any $\varepsilon$-nets $N \subseteq \mathbb{S}^{n-1}$ and $M \subseteq \mathbb{S}^{m-1}$ of $\mathbb{S}^{n-1}$ and $\mathbb{S}^{m-1}$ respectively, the following inequalities hold

$$
\sup _{\substack{\mathbf{x} \in N \\ \mathbf{y} \in M}}\langle A \mathbf{x}, \mathbf{y}\rangle \leq\|A\|_{2} \leq \frac{1}{1-2 \varepsilon} \sup _{\substack{\mathbf{x} \in N \\ \mathbf{y} \in M}}\langle A \mathbf{x}, \mathbf{y}\rangle .
$$

Proof of Lemma 2.4.29. Recall that

$$
\langle A \mathbf{x}, \mathbf{y}\rangle=\sum_{i, j} x_{j} y_{i} A_{i j}
$$

is a linear combination of independent random variables. By the general Hoeffding inequality, $\langle A \mathbf{x}, \mathbf{y}\rangle$ is sub-Gaussian with variance factor

$$
\sum_{i, j}\left(x_{i} y_{j}\right)^{2} \nu=\|\mathbf{x}\|_{2}^{2}\|\mathbf{y}\|_{2}^{2} \nu=\nu,
$$

for all $\mathbf{x} \in N$ and $\mathbf{y} \in M$. In particular, for all $\beta>0$,

$$
\mathbb{P}[\langle A \mathbf{x}, \mathbf{y}\rangle \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2 \nu}\right)
$$

Hence, by a union bound over $N$ and $M$,

$$
\begin{aligned}
& \mathbb{P}\left[\max _{\substack{\mathbf{x} \in N \\
\mathbf{y} \in M}}\langle A \mathbf{x}, \mathbf{y}\rangle \geq \frac{1}{2} C \sqrt{\nu}(\sqrt{m}+\sqrt{n}+t)\right] \\
& \quad \leq \sum_{\substack{\mathbf{x} \in N \\
\mathbf{y} \in M}} \mathbb{P}\left[\langle A \mathbf{x}, \mathbf{y}\rangle \geq \frac{1}{2} C \sqrt{\nu}(\sqrt{m}+\sqrt{n}+t)\right] \\
& \quad \leq|N||M| \exp \left(-\frac{1}{2 \nu}\left\{\frac{1}{2} C \sqrt{\nu}(\sqrt{m}+\sqrt{n}+t)\right\}^{2}\right) \\
& \quad \leq 12^{n+m} \exp \left(-\frac{C^{2}}{8}\left\{m+n+t^{2}\right\}\right) \\
& \quad \leq e^{-t^{2}}
\end{aligned}
$$

for $C^{2} / 8=\log 12 \geq 1$, where in the third inequality we ignored all cross-products since they are nonnegative.

Proof of Lemma 2.4.30. The first inequality is immediate by definition of the spectral norm. For the second inequality, we will use the following observation

$$
\begin{equation*}
\langle A \mathbf{x}, \mathbf{y}\rangle-\left\langle A \mathbf{x}_{0}, \mathbf{y}_{0}\right\rangle=\left\langle A \mathbf{x}, \mathbf{y}-\mathbf{y}_{0}\right\rangle+\left\langle A\left(\mathbf{x}-\mathbf{x}_{0}\right), \mathbf{y}_{0}\right\rangle . \tag{2.4.34}
\end{equation*}
$$

Fix $\mathbf{x} \in \mathbb{S}^{n-1}$ and $\mathbf{y} \in \mathbb{S}^{m-1}$ such that $\langle A \mathbf{x}, \mathbf{y}\rangle=\|A\|_{2}$ (which exist by compactness), and let $\mathbf{x}_{0} \in N$ and $\mathbf{y}_{0} \in M$ such that

$$
\left\|\mathbf{x}-\mathbf{x}_{0}\right\|_{2} \leq \varepsilon \quad \text { and } \quad\left\|\mathbf{y}-\mathbf{y}_{0}\right\|_{2} \leq \varepsilon .
$$

Then (2.4.34), Cauchy-Schwarz and the definition of the spectral norm imply

$$
\|A\|_{2}-\left\langle A \mathbf{x}_{0}, \mathbf{y}_{0}\right\rangle \leq\|A\|_{2}\|\mathbf{x}\|_{2}\left\|\mathbf{y}-\mathbf{y}_{0}\right\|_{2}+\|A\|_{2}\left\|\mathbf{x}-\mathbf{x}_{0}\right\|_{2}\left\|\mathbf{y}_{0}\right\|_{2} \leq 2 \varepsilon\|A\|_{2} .
$$

Rearranging gives the claim.
Putting the two lemmas together concludes the proof of Theorem 2.4.28.
We will give an application of this bound in Section 5.1.4.

## Chaining method

We go back to the inequality

$$
\begin{equation*}
\sup _{t \in \mathscr{T}} X_{t} \leq \sup _{t \in \mathscr{T}}\left(X_{t}-X_{\pi(t)}\right)+\sup _{t \in \mathscr{T}} X_{\pi(t)} . \tag{2.4.35}
\end{equation*}
$$

Previously we controlled the first term on the right-hand side with a random Lipschitz constant and the second term with a maximal inequality for finite sets. Now we consider cases where we may not have a good almost sure bound on the Lipschitz constant, but where we can control increments uniformly in the following probabilistic sense. We say that a stochastic process $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ has sub-Gaussian increments on $(\mathscr{T}, \rho)$ if there exists a deterministic constant $0<\mathcal{K}<+\infty$ such that

$$
X_{t}-X_{s} \in \operatorname{sG}\left(\mathcal{K}^{2} \rho(s, t)^{2}\right), \quad \forall s, t \in \mathscr{T} .
$$

Even with this assumption, in (2.4.35) the first term on the right-hand side remains a supremum over an infinite set. To control it, the chaining method repeats the argument above at progressively smaller scales, leading to the following inequality. The diameter of $\mathscr{T}$, denoted by $\operatorname{diam}(\mathscr{T})$, is defined as

$$
\operatorname{diam}(\mathscr{T})=\sup \{\rho(s, t): s, t, \in \mathscr{T}\} .
$$

Theorem 2.4.31 (Discrete Dudley inequality). Let $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ be a zero-mean stochastic process with sub-Gaussian increments on $(\mathscr{T}, \rho)$ and assume $\operatorname{diam}(\mathscr{T}) \leq 1$. Then

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] \leq C \sum_{k=0}^{+\infty} 2^{-k} \sqrt{\log \mathcal{N}\left(\mathscr{T}, \rho, 2^{-k}\right)} .
$$

for some constant $0 \leq C<+\infty$.
Proof. Recall that we assume that $\mathscr{T}$ is countable. Let $\mathscr{T}_{j} \subseteq \mathscr{T}, j \geq 1$, be a sequence of finite sets such that $\mathscr{T}_{j} \uparrow \mathscr{T}$. By monotone convergence (Proposition B.4.14),

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right]=\sup _{j \geq 1} \mathbb{E}\left[\sup _{t \in \mathscr{T}_{j}} X_{t}\right]
$$

Moreover, $\mathcal{N}\left(\mathscr{T}_{j}, \rho, \varepsilon\right) \leq \mathcal{N}(\mathscr{T}, \rho, \varepsilon)$ for any $\varepsilon>0$ since $\mathscr{T}_{j} \subseteq \mathscr{T}$. Hence it suffices to handle the case $|\mathscr{T}|<+\infty$.
$\varepsilon$-nets at all scales. For each $k \geq 0$, let $N_{k}$ be an $2^{-k}$-net of $\mathscr{T}$ with $\left|N_{k}\right|=$ $\mathcal{N}\left(\mathscr{T}, \rho, 2^{-k}\right)$ and projection map $\pi_{k}$. Because $\operatorname{diam}(\mathscr{T}) \leq 1$, we can set $N_{0}=$ $\left\{t_{0}\right\}$ where $t_{0} \in \mathscr{T}$ can be taken arbitrarily. Moreover, because $\mathscr{T}$ is finite, there
is $1 \leq \kappa<+\infty$ such that we can take $N_{k}=\mathscr{T}$ for all $k \geq \kappa^{\S}$. In particular, $\pi_{\kappa}(t)=t$ for all $t \in \mathscr{T}$. By a telescoping argument,

$$
X_{t}=X_{t_{0}}+\sum_{k=0}^{\kappa-1}\left(X_{\pi_{k+1}(t)}-X_{\pi_{k}(t)}\right) .
$$

Taking a supremum and then an expectation gives

$$
\begin{equation*}
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] \leq \sum_{k=0}^{\kappa-1} \mathbb{E}\left[\sup _{t \in \mathscr{T}}\left(X_{\pi_{k+1}(t)}-X_{\pi_{k}(t)}\right)\right], \tag{2.4.36}
\end{equation*}
$$

where we used $\mathbb{E}\left[X_{t_{0}}\right]=0$.
Sub-Gaussian bound. We use the maximal inequality (Theorem 2.4.21) to bound the expectation in (2.4.36). For each $k$, the number of distinct elements in the supremum is at most

$$
\begin{aligned}
\left|\left\{\left(\pi_{k}(t), \pi_{k+1}(t)\right): t \in \mathscr{T}\right\}\right| & \leq\left|N_{k} \times N_{k+1}\right| \\
& =\left|N_{k}\right| \times\left|N_{k+1}\right| \\
& \leq \mathcal{N}\left(\mathscr{T}, \rho, 2^{-k-1}\right)^{2} .
\end{aligned}
$$

For any $t \in \mathscr{T}$, by the triangle inequality,

$$
\rho\left(\pi_{k}(t), \pi_{k+1}(t)\right) \leq \rho\left(\pi_{k}(t), t\right)+\rho\left(t, \pi_{k+1}(t)\right) \leq 2^{-k}+2^{-k-1} \leq 2^{-k+1}
$$

so that

$$
X_{\pi_{k+1}(t)}-X_{\pi_{k}(t)} \in \mathrm{s} \mathcal{G}\left(\mathcal{K}^{2} 2^{-2 k+2}\right),
$$

for some $0<\mathcal{K}<+\infty$ by the sub-Gaussian increments assumption. We can therefore apply Theorem 2.4.21 to get

$$
\begin{aligned}
\mathbb{E}\left[\sup _{t \in \mathscr{T}}\left(X_{\pi_{k+1}(t)}-X_{\pi_{k}(t)}\right)\right] & \leq \sqrt{2 \mathcal{K}^{2} 2^{-2 k+2} \log \left(\mathcal{N}\left(\mathscr{T}, \rho, 2^{-k-1}\right)^{2}\right)} \\
& \leq C 2^{-k-1} \sqrt{\log \mathcal{N}\left(\mathscr{T}, \rho, 2^{-k-1}\right)}
\end{aligned}
$$

for some constant $0 \leq C<+\infty$ (depending on $\mathcal{K}$ ).
To finish the argument, we plug back into (2.4.36),

$$
\mathbb{E}\left[\sup _{t \in \mathscr{T}} X_{t}\right] \leq \sum_{k=0}^{\kappa-1} C 2^{-k-1} \sqrt{\log \mathcal{N}\left(\mathscr{T}, \rho, 2^{-k-1}\right)},
$$

which implies the claim.

[^4]Using a similar argument, one can derive a tail inequality.
Theorem 2.4.32 (Chaining tail inequality). Let $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ be a zero-mean stochastic process with sub-Gaussian increments on $(\mathscr{T}, \rho)$ and assume that $\operatorname{diam}(\mathscr{T}) \leq$ 1. Then, for all $t_{0} \in \mathscr{T}$ and $\beta>0$,

$$
\mathbb{P}\left[\sup _{t \in \mathscr{T}}\left(X_{t}-X_{t_{0}}\right) \geq C \sum_{k=0}^{+\infty} 2^{-k} \sqrt{\log \mathcal{N}\left(\mathscr{T}, \rho, 2^{-k}\right)}+\beta\right] \leq C \exp \left(-\frac{\beta^{2}}{C}\right),
$$

for some constant $0 \leq C<+\infty$.
We give an application of the discrete Dudley inequality in Section 2.4.6.

### 2.4.5 $\triangleright$ Data science: Johnson-Lindenstrauss lemma and application to compressed sensing

In this section we discuss an application of the Chernoff-Cramér method (Section 2.4.1) to dimension reduction in data science. We use once again an $\varepsilon$-net argument (Section 2.4.4).

## Johnson-Lindenstrauss lemma

The Johnson-Lindenstrauss lemma states roughly that, for any collection of points in a high-dimensional Euclidean space, one can find an embedding of much lower dimension that roughly preserves the metric relationships of the points, that is, their distances. Remarkably, no structure is assumed on the original points and the result is independent of the input dimension. The method of proof simply involves performing a random projection.
Lemma 2.4.33 (Johnson-Lindenstrauss lemma). For any set of points $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)}$ in $\mathbb{R}^{n}$ and $\theta \in(0,1)$, there exists a mapping $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{d}$ with $d=\Theta\left(\theta^{-2} \log m\right)$ such that the following holds: for all $i, j$

$$
\begin{equation*}
(1-\theta)\left\|\mathbf{x}^{(i)}-\mathbf{x}^{(j)}\right\|_{2} \leq\left\|f\left(\mathbf{x}^{(i)}\right)-f\left(\mathbf{x}^{(j)}\right)\right\|_{2} \leq(1+\theta)\left\|\mathbf{x}^{(i)}-\mathbf{x}^{(j)}\right\|_{2} . \tag{2.4.37}
\end{equation*}
$$

We use the probabilistic method: we derive a "distributional" version of the result that, in turn, implies Lemma 2.4 .33 by showing that a mapping with the desired properties exists with positive probability. Before stating this claim formally, we define the explicit random linear mapping we will employ. Let $A$ be a $d \times n$ matrix whose entries are independent $\mathrm{N}(0,1)$. Note that, for any fixed $\mathbf{z} \in \mathbb{R}^{n}$,

$$
\begin{equation*}
\mathbb{E}\|A \mathbf{z}\|_{2}^{2}=\mathbb{E}\left[\sum_{i=1}^{d}\left(\sum_{j=1}^{n} A_{i j} z_{j}\right)^{2}\right]=d \operatorname{Var}\left[\sum_{j=1}^{n} A_{1 j} z_{j}\right]=d\|\mathbf{z}\|_{2}^{2}, \tag{2.4.38}
\end{equation*}
$$

where we used the independence of the $A_{i j} \mathrm{~s}$ (and, in particular, of the rows of $A$ ) and the fact that

$$
\begin{equation*}
\mathbb{E}\left[\sum_{j=1}^{n} A_{i j} z_{j}\right]=0 . \tag{2.4.39}
\end{equation*}
$$

Hence the normalized mapping

$$
L=\frac{1}{\sqrt{d}} A
$$

preserves the squared Euclidean norm "on average," that is, $\mathbb{E}\|L \mathbf{z}\|_{2}^{2}=\|\mathbf{z}\|_{2}^{2}$. We use the Chernoff-Cramér method to prove a high-probability result.

Lemma 2.4.34. Fix $\delta, \theta \in(0,1)$. Then the random linear mapping $L$ above with $d=\Theta\left(\theta^{-2} \log \delta^{-1}\right)$ is such that for any $\mathbf{z} \in \mathbb{R}^{n}$ with $\|\mathbf{z}\|_{2}=1$

$$
\begin{equation*}
\mathbb{P}\left[\left|\|L \mathbf{z}\|_{2}-1\right| \geq \theta\right] \leq \delta . \tag{2.4.40}
\end{equation*}
$$

Before proving Lemma 2.4.34, we argue that it implies the Johnson-Lindenstrauss lemma (Lemma 2.4.33). Simply take $\delta=1 /\left(2\binom{m}{2}\right)$, apply the previous lemma to each normalized pairwise difference $\mathbf{z}=\left(\mathbf{x}^{(i)}-\mathbf{x}^{(j)}\right) /\left\|\mathbf{x}^{(i)}-\boldsymbol{x}^{(j)}\right\|_{2}$, and use a union bound over all $\binom{m}{2}$ such pairs. The probability that any of the inequalities (2.4.37) is not satisfied by the linear mapping $f(\mathbf{z})=L \mathbf{z}$ is then at most $1 / 2$. Hence a mapping with the desired properties exists for $d=\Theta\left(\theta^{-2} \log m\right)$.

Proof of Lemma 2.4.34. We prove one direction. Specifically, we establish

$$
\begin{equation*}
\mathbb{P}\left[\|L \mathbf{z}\|_{2} \geq 1+\theta\right] \leq \exp \left(-\frac{3}{4} d \theta^{2}\right) \tag{2.4.41}
\end{equation*}
$$

Note that the right-hand side is $\leq \delta$ for $d=\Theta\left(\theta^{-2} \log \delta^{-1}\right)$. An inequality in the other direction can be proved similarly by working with $-W$ (where $W$ is defined below).

Recall that a sum of independent Gaussians is Gaussian (just compute the convolution and complete the squares). So

$$
(A \mathbf{z})_{k} \sim \mathrm{~N}\left(0,\|\mathbf{z}\|_{2}^{2}\right)=\mathrm{N}(0,1), \quad \forall k
$$

where we argued as in (2.4.38) to compute the variance. Hence

$$
W=\|A \mathbf{z}\|_{2}^{2}=\sum_{k=1}^{d}(A \mathbf{z})_{k}^{2}
$$

is a sum of squares of independent Gaussians, that is, $\chi^{2}$-distributed random variables. By (2.4.18) and independence,

$$
M_{W}(s)=\frac{1}{(1-2 s)^{d / 2}}
$$

Applying the Chernoff-Cramér bound (2.4.2) with $s=\frac{1}{2}(1-d / \beta)$ gives

$$
\mathbb{P}[W \geq \beta] \leq \frac{M_{W}(s)}{e^{s \beta}}=\frac{1}{e^{s \beta}(1-2 s)^{d / 2}}=e^{(d-\beta) / 2}\left(\frac{\beta}{d}\right)^{d / 2}
$$

(Alternatively, we could have used the general Bernstein inequality (Theorem 2.4.16).) Finally, take $\beta=d(1+\theta)^{2}$. Rearranging we get

$$
\begin{aligned}
\mathbb{P}\left[\|L \mathbf{z}\|_{2} \geq 1+\theta\right] & =\mathbb{P}\left[\|A \mathbf{z}\|_{2}^{2} \geq d(1+\theta)^{2}\right] \\
& =\mathbb{P}[W \geq \beta] \\
& \leq e^{d\left[1-(1+\theta)^{2}\right] / 2}\left[(1+\theta)^{2}\right]^{d / 2} \\
& =\exp \left(-d\left(\theta+\theta^{2} / 2-\log (1+\theta)\right)\right) \\
& \leq \exp \left(-\frac{3}{4} d \theta^{2}\right)
\end{aligned}
$$

where we used that $\log (1+x) \leq x-x^{2} / 4$ on $[0,1]$ (see Exercise 1.16).
Remark 2.4.35. The Johnson-Lindenstrauss lemma is essentially optimal [Alo03, Section 9]: any set of $n$ points with all pairwise distances in $[1-\theta, 1+\theta]$ requires at least $\Omega\left(\log n /\left(\theta^{2} \log \theta^{-1}\right)\right)$ dimensions. Note however that it relies crucially on the use of the Euclidean norm [BC03].

To give some further geometric insights into the proof, we make a series of observations:

1. The $d$ rows of $\frac{1}{\sqrt{n}} A$ are "on average" orthonormal. Indeed, note that for $i \neq j$

$$
\mathbb{E}\left[\frac{1}{n} \sum_{k=1}^{n} A_{i k} A_{j k}\right]=\mathbb{E}\left[A_{i 1}\right] \mathbb{E}\left[A_{j 1}\right]=0,
$$

by independence and

$$
\mathbb{E}\left[\frac{1}{n} \sum_{k=1}^{n} A_{i k}^{2}\right]=\mathbb{E}\left[A_{i 1}^{2}\right]=1
$$

since the $A_{i k} \mathrm{~s}$ have mean 0 and variance 1 . When $n$ is large, those two quantities are concentrated around their mean. Fix a unit vector $\mathbf{z}$. Then $\frac{1}{\sqrt{n}} A \mathbf{z}$ corresponds approximately to an orthogonal projection of $\mathbf{z}$ onto a uniformly chosen random subspace of dimension $d$.
2. Now observe that projecting $\mathbf{z}$ on a uniform random subspace of dimension $d$ can be done in the following way: first apply a uniformly chosen random rotation to $\mathbf{z}$; and then project the resulting vector on the first $d$ dimensions. In other words, $\frac{1}{\sqrt{n}}\|A \mathbf{z}\|_{2}$ is approximately distributed as the norm of the first $d$ components of a uniform unit vector in $\mathbb{R}^{n}$. To analyze this quantity, note that a vector in $\mathbb{R}^{n}$ whose components are independent $\mathrm{N}(0,1)$, when divided by its norm, produces a uniform vector in $\mathbb{R}^{n}$. When $d$ is large, the norm of the first $d$ components of that vector is therefore a ratio whose numerator is concentrated around $\sqrt{d}$ and whose denominator is concentrated around $\sqrt{n}$ (by calculations similar to those in the first point above).
3. Hence $\|L \mathbf{z}\|_{2}=\sqrt{\frac{n}{d}} \times \frac{1}{\sqrt{n}}\|A \mathbf{z}\|_{2}$ should be concentrated around 1 .

The Johnson-Lindenstrauss lemma makes it possible to solve certain computational problems (e.g., finding the nearest point to a query) more efficiently by working in a smaller dimension. We discuss a different application of the "random projection method" next.

## Compressed sensing

In the compressed sensing problem, one seeks to recover a signal $\mathbf{x} \in \mathbb{R}^{n}$ from a small number of linear measurements $(L \mathbf{x})_{i}, i=1, \ldots, d$. In complete generality, one needs $n$ such measurements to recover any unknown $\mathbf{x} \in \mathbb{R}^{n}$ as the sensing matrix $L$ must be invertible (or, more precisely, injective). However, by imposing extra structure on the signal and choosing the sensing matrix appropriately, much better results can be obtained. Compressed sensing relies on sparsity.

Definition 2.4.36 (Sparse vectors). We say that a vector $\mathbf{z} \in \mathbb{R}^{n}$ is $k$-sparse if it has at most $k$ non-zero entries. We let $\mathscr{S}_{k}^{n}$ be the set of $k$-sparse vectors in $\mathbb{R}^{n}$.

Note that $\mathscr{S}_{k}^{n}$ is a union of $\binom{n}{k}$ linear subspaces, one for each support of the nonzero entries.

To solve the compressed sensing problem over $k$-sparse vectors, it suffices to find a sensing matrix $L$ satisfying that all subsets of $2 k$ columns are linearly independent. Indeed, if $\mathbf{x}, \mathbf{x}^{\prime} \in \mathscr{S}_{k}^{n}$, then $\mathbf{x}-\mathbf{x}^{\prime}$ has at most $2 k$ nonzero entries. Hence, in order to have $L\left(\mathbf{x}-\mathbf{x}^{\prime}\right)=0$, it must be that $\mathbf{x}-\mathbf{x}^{\prime}=0$ under the previous condition on $L$. That implies the required injectivity. The implication goes in the other
direction as well. Observe for instance that the matrix used in the proof of the Johnson-Lindenstrauss lemma satisfies this property as long as $d \geq 2 k$ : because of the continuous density of its entries, the probability that $2 k$ of its columns are linearly dependent is 0 when $d \geq 2 k$. For practical applications, however, other requirements must be met, in particular, computational efficiency and robustness. We describe such an approach.

The following definition will play a key role. Roughly speaking, a restricted isometry preserves enough of the metric structure of $\mathscr{S}_{k}^{n}$ to be invertible on its image.

Definition 2.4.37 (Restricted isometry property). A $d \times n$ linear mapping $L$ satisfies the ( $k, \theta$ )-restricted isometry property (RIP) if for all $\mathbf{z} \in \mathscr{S}_{k}^{n}$

$$
\begin{equation*}
(1-\theta)\|\mathbf{z}\|_{2} \leq\|L \mathbf{z}\|_{2} \leq(1+\theta)\|\mathbf{z}\|_{2} . \tag{2.4.42}
\end{equation*}
$$

We say that $L$ is $(k, \theta)-R I P$.
Given a $(k, \theta)$-RIP matrix $L$, can we recover $\mathbf{z} \in \mathscr{S}_{k}^{n}$ from $L \mathbf{z}$ ? And how small can $d$ be? The next two claims answer these questions.

Lemma 2.4.38 (Sensing matrix). Let A be a $d \times n$ matrix whose entries are i.i.d. $N(0,1)$ and let $L=\frac{1}{\sqrt{d}}$ A. There is a constant $0<C<+\infty$ such that if $d \geq C k \log n$ then $L$ is $(10 k, 1 / 3)$-RIP with probability at least $1-1 / n$.
Lemma 2.4.39 (Sparse signal recovery). Let L be ( $10 k, 1 / 3$ )-RIP. Then for any $\mathrm{x} \in \mathscr{S}_{k}^{n}$, the unique solution to the following minimization problem

$$
\begin{equation*}
\min _{\mathbf{z} \in \mathbb{R}^{n}}\|\mathbf{z}\|_{1} \quad \text { subject to } \quad L \mathbf{z}=L \mathbf{x} \tag{2.4.43}
\end{equation*}
$$

is $\mathbf{z}^{*}=\mathbf{x}$.
It may seem that a more natural approach, compared to (2.4.43), would be to instead minimize the number of non-zero entries in $\mathbf{z}$, that is, $\|\mathbf{z}\|_{0}$. However the advantage of the $\ell^{1}$ norm is that the problem can then be formulated as a linear program, that is, the minimization of a linear objective subject to linear inequalities (see Exercise 2.13). This permits much faster computation of the solution using standard techniques-while still leading to a sparse solution. See Figure 2.8 for some insights into to why $\ell^{1}$ indeed promotes sparsity.

Putting the two lemmas together shows that:
Claim 2.4.40. Let $L$ be as above with $d=\Theta(k \log n)$ as required by Lemma 2.4.38. With probability $1-o(1)$, any $\mathbf{x} \in \mathscr{S}_{k}^{n}$ can be recovered from the input $L \mathbf{x}$ by solving (2.4.43).
Note that $d$ can in general be much smaller than $n$ and not far from the $2 k$ bound we derived above.


Figure 2.8: Because $\ell^{1}$ balls (squares) have corners, minimizing the $\ell^{1}$ norm over a linear subspace (line) tends to produce sparse solutions.
$\varepsilon$-net argument We start with the proof of Lemma 2.4.38. The claim does not follow immediately from the (distributional) Johnson-Lindenstrauss lemma (i.e., Lemma 2.4.34). Indeed that lemma implies that a (normalized) matrix with i.i.d. standard Gaussian entries is an approximate isometry on a finite set of points. Here we need a linear mapping that is an approximate isometry for all vectors in $\mathscr{S}_{k}^{n}$, an uncountable space.

For a subset of indices $J \subseteq[n]$ and a vector $\mathbf{y} \in \mathbb{R}^{n}$, we let $\mathbf{y}_{J}$ be the vector $\mathbf{y}$ restricted to the entries in $J$, that is, the subvector $\left(y_{j}\right)_{j \in J}$. Fix a subset of indices $I \subseteq[n]$ of size $10 k$. We need the RIP condition (Definition 2.4.37) to hold for all $\mathbf{z} \in \mathbb{R}^{n}$ with non-zero entries in $I$ (and all such $I$ ). The way to achieve this is to use an $\varepsilon$-net argument, as described in Section 2.4.4. Indeed, notice that, for $\mathbf{z} \neq \mathbf{0}$, the function $\|L \mathbf{z}\|_{2} /\|\mathbf{z}\|_{2}$ :

1. does not depend on the norm of $\mathbf{z}$, so that we can restrict ourselves to the compact set $\partial B_{I}:=\left\{\mathbf{z}: \mathbf{z}_{[n] \backslash I}=\mathbf{0},\|\mathbf{z}\|_{2}=1\right\}$; and
2. is continuous on $\partial B_{I}$, so that it suffices to construct a fine enough covering of $\partial B_{I}$ by a finite collection of balls (i.e., an $\varepsilon$-net) and apply Lemma 2.4.34 to the centers of those balls.

Proof of Lemma 2.4.38. Let $I \subseteq[n]$ be a subset of indices of size $k^{\prime}:=10 k$. There are $\binom{n}{k^{\prime}} \leq n^{k^{\prime}}=\exp \left(k^{\prime} \log n\right)$ such subsets and we denote their collection by $\mathcal{I}\left(k^{\prime}, n\right)$. We let $N_{I}$ be an $\varepsilon$-net of $\partial B_{I}$. By Claim 2.4.26, we can choose one in $\partial B_{I}$ of size at most $(3 / \varepsilon)^{k^{\prime}}$. We take

$$
\varepsilon=\frac{1}{C^{\prime} \sqrt{6 n \log n}}
$$

for a constant $C^{\prime}$ that will be determined below. The reason for this choice will become clear when we set $C^{\prime}$. The union of all $\varepsilon$-nets has size

$$
\left|\cup_{I \in \mathcal{I}\left(k^{\prime}, n\right)} N_{I}\right| \leq n^{k^{\prime}}\left(\frac{3}{\varepsilon}\right)^{k^{\prime}} \leq \exp \left(C^{\prime \prime} k^{\prime} \log n\right),
$$

for some $C^{\prime \prime}>0$. Our goal is to show that

$$
\begin{equation*}
\sup _{\mathbf{z} \in \cup_{I \in \mathcal{I}\left(k^{\prime}, n\right)} \partial B_{I}}\left|\|L \mathbf{z}\|_{2}-1\right| \leq \frac{1}{3} \tag{2.4.44}
\end{equation*}
$$

We seek to apply the inequality (2.4.32).
Applying the "distributional" Johnson-Lindenstrauss lemma to the $\varepsilon$-nets: The first step is to control the supremum in (2.4.44)—restricted to the $\varepsilon$-nets. Lemma 2.4.34 is exactly what we need for this. Take $\theta=1 / 6, \delta=1 /\left(2 n\left|\cup_{I} N_{I}\right|\right)$, and

$$
d=\Theta\left(\theta^{-2} \log \left(2 n\left|\cup_{I} N_{I}\right|\right)\right)=\Theta\left(k^{\prime} \log n\right),
$$

as required by the lemma. Then, by a union bound over the $N_{I} \mathrm{~s}$, with probability $1-1 /(2 n)$ we have

$$
\begin{equation*}
\sup _{\mathbf{z} \in \cup_{I} N_{I}}\left|\|L \mathbf{z}\|_{2}-1\right| \leq \frac{1}{6} . \tag{2.4.45}
\end{equation*}
$$

Lipschitz continuity: The next step is to establish Lipschitz continuity of $\left\|\|L z\|_{2}-\right.$ $1 \mid$. For vectors $\mathbf{y}, \mathbf{z} \in \mathbb{R}^{n}$, by repeated applications of the triangle inequality, we have

$$
\left|\|L \mathbf{z}\|_{2}-1\right|-\left|\|L \mathbf{y}\|_{2}-1\left\|\leq\left|\|L \mathbf{z}\|_{2}-\|L \mathbf{y}\|_{2}\right| \leq\right\| L(\mathbf{z}-\mathbf{y}) \|_{2} .\right.
$$

To bound the rightmost expression, we let $A_{*}$ be the largest entry of $A$ in absolute value and note that

$$
\begin{aligned}
\|L(\mathbf{z}-\mathbf{y})\|_{2}^{2} & =\sum_{i=1}^{d}\left(\sum_{j=1}^{n} L_{i j}\left(z_{j}-y_{j}\right)\right)^{2} \\
& \leq \sum_{i=1}^{d}\left(\sum_{j=1}^{n} L_{i j}^{2}\right)\left(\sum_{j=1}^{n}\left(z_{j}-y_{j}\right)^{2}\right) \\
& \leq d n\left(\frac{1}{\sqrt{d}} A_{*}\right)^{2}\|\mathbf{z}-\mathbf{y}\|_{2}^{2} \\
& \leq n A_{*}^{2}\|\mathbf{z}-\mathbf{y}\|_{2}^{2}
\end{aligned}
$$

where we used Cauchy-Schwarz (Theorem B.4.8) on the second line. Taking a square root, we see that the (random) Lipschitz constant of $\left|\|L \mathbf{z}\|_{2}-1\right|$ (with respect to the Euclidean metric) is at most $K:=\sqrt{n} A_{*}$.

Controlling the Lipschitz constant: So it remains to control $A_{*}$. For this we use the Chernoff-Cramér bound for Gaussians (see (2.4.4)) which implies by a union bound over the entries of $A$ that

$$
\begin{aligned}
\mathbb{P}\left[A_{*} \geq C^{\prime} \sqrt{\log n}\right] & \leq \mathbb{P}\left[\exists i, j,\left|A_{i j}\right| \geq C^{\prime} \sqrt{\log n}\right] \\
& \leq n^{2} \exp \left(-\frac{\left(C^{\prime} \sqrt{\log n}\right)^{2}}{2}\right) \\
& \leq \frac{1}{2 n}
\end{aligned}
$$

for a $C^{\prime}>0$ large enough. Hence with probability $1-1 /(2 n)$, we have $A_{*}<$ $C^{\prime} \sqrt{\log n}$ and

$$
\begin{equation*}
K \varepsilon \leq \frac{1}{6}, \tag{2.4.46}
\end{equation*}
$$

by the choice of $\varepsilon$ made previously.
Putting everything together: We apply (2.4.32). Combining (2.4.45) and (2.4.46), with probability $1-1 / n$, the claim (2.4.44) holds. That concludes the proof.
$\ell^{1}$ minimization Finally we prove Lemma 2.4.39 (which can be skipped).
Proof of Lemma 2.4.39. Let $\mathbf{z}^{*}$ be a solution to (2.4.43) and note that such a solution exists because $\mathbf{z}=\mathbf{x}$ satisfies the constraint. Without loss of generality assume that only the first $k$ entries of $\mathbf{x}$ are nonzero, that is, $\mathbf{x}_{[n] \backslash k]}=\mathbf{0}$. Moreover order the remaining entries of $\mathbf{x}$ so that the residual $\mathbf{r}=\mathbf{z}^{*}-\mathbf{x}$ has its entries $\mathbf{r}_{[n] \backslash k]}$ in nonincreasing order in absolute value. Our goal is to show that $\|\mathbf{r}\|_{2}=0$.

In order to leverage the RIP condition, we break up the vector $\mathbf{r}$ into $9 k$-long subvectors. Let

$$
I_{0}=[k], \quad I_{i}=\{(9(i-1)+1) k+1, \ldots,(9 i+1) k\}, \forall i \geq 1,
$$

and $\bar{I}_{i}=\bigcup_{j>i} I_{j}$. We will also need $I_{01}=I_{0} \cup I_{1}$ and $\bar{I}_{01}=\bar{I}_{1}$.
We first use the optimality of $\mathbf{z}^{*}$. Note that $\mathbf{x}_{\bar{I}_{0}}=0$ implies that

$$
\left\|\mathbf{z}^{*}\right\|_{1}=\left\|\mathbf{z}_{I_{0}}^{*}\right\|_{1}+\left\|\mathbf{z}_{\bar{I}_{0}}^{*}\right\|_{1}=\left\|\mathbf{z}_{I_{0}}^{*}\right\|_{1}+\left\|\mathbf{r}_{\bar{I}_{0}}\right\|_{1},
$$

and

$$
\|\mathbf{x}\|_{1}=\left\|\mathbf{x}_{I_{0}}\right\|_{1} \leq\left\|\mathbf{z}_{I_{0}}^{*}\right\|_{1}+\left\|\mathbf{r}_{I_{0}}\right\|_{1}
$$

by the triangle inequality. Since $\left\|\mathbf{z}^{*}\right\|_{1} \leq\|\mathbf{x}\|_{1}$ by optimality (and the fact that $\mathbf{x}$ satisfies the constraint), we then have

$$
\begin{equation*}
\left\|\mathbf{r}_{\bar{I}_{0}}\right\|_{1} \leq\left\|\mathbf{r}_{I_{0}}\right\|_{1} . \tag{2.4.47}
\end{equation*}
$$

On the other hand, the RIP condition gives a similar inequality in the other direction. Indeed notice that $L \mathbf{r}=\mathbf{0}$ by the constraint in (2.4.43) or, put differently, $L \mathbf{r}_{I_{01}}=-\sum_{i \geq 2} L \mathbf{r}_{I_{i}}$. Then, by the RIP condition and the triangle inequality, we have that

$$
\begin{equation*}
\frac{2}{3}\left\|\mathbf{r}_{I_{01}}\right\|_{2} \leq\left\|L \mathbf{r}_{I_{01}}\right\|_{2} \leq \sum_{i \geq 2}\left\|L \mathbf{r}_{I_{i}}\right\|_{2} \leq \frac{4}{3} \sum_{i \geq 2}\left\|\mathbf{r}_{I_{i}}\right\|_{2}, \tag{2.4.48}
\end{equation*}
$$

where we used the fact that by construction $\mathbf{r}_{I_{01}}$ is $10 k$-sparse and each $\mathbf{r}_{I_{i}}$ is $9 k$ sparse.

We note that by the ordering of the entries of $\mathbf{x}$

$$
\begin{equation*}
\left\|\mathbf{r}_{I_{i+1}}\right\|_{2}^{2} \leq 9 k\left(\frac{\left\|\mathbf{r}_{I_{i}}\right\|_{1}}{9 k}\right)^{2}=\frac{\left\|\mathbf{r}_{I_{i}}\right\|_{1}^{2}}{9 k} \tag{2.4.49}
\end{equation*}
$$

where we bounded $\mathbf{r}_{I_{i+1}}$ entrywise by the expression in parenthesis. Combining (2.4.47) and (2.4.49), and using that $\left\|\mathbf{r}_{I_{0}}\right\|_{1} \leq \sqrt{k}\left\|\mathbf{r}_{I_{0}}\right\|_{2}$ by Cauchy-Schwarz, we have

$$
\sum_{i \geq 2}\left\|\mathbf{r}_{I_{i}}\right\|_{2} \leq \sum_{j \geq 1} \frac{\left\|\mathbf{r}_{I_{j}}\right\|_{1}}{\sqrt{9 k}}=\frac{\left\|\mathbf{r}_{\bar{I}_{0}}\right\|_{1}}{3 \sqrt{k}} \leq \frac{\left\|\mathbf{r}_{I_{0}}\right\|_{1}}{3 \sqrt{k}} \leq \frac{\left\|\mathbf{r}_{I_{0}}\right\|_{2}}{3} \leq \frac{\left\|\mathbf{r}_{I_{01}}\right\|_{2}}{3} .
$$

Plugging this back into (2.4.48) gives

$$
\left\|\mathbf{r}_{I_{01}}\right\|_{2} \leq 2 \sum_{i \geq 2}\left\|\mathbf{r}_{I_{i}}\right\|_{2} \leq \frac{2}{3}\left\|\mathbf{r}_{I_{01}}\right\|_{2}
$$

which implies $\mathbf{r}_{I_{01}}=\mathbf{0}$. In particular $\mathbf{r}_{I_{0}}=\mathbf{0}$ and, by (2.4.47), $\mathbf{r}_{\bar{I}_{0}}=\mathbf{0}$ as well. We have shown that $\mathbf{r}=\mathbf{0}$. Or, in other words, $\mathbf{z}^{*}=\mathbf{x}$.

Remark 2.4.41. Lemma 2.4.39 can be extended to noisy measurements using a modification of (2.4.43). This provides some robustness to noise which is important in applications. See [CRT06b].

### 2.4.6 $\triangleright$ Data science: classification, empirical risk minimization and VC dimension

In the binary classification problem, one is given samples $\mathcal{S}_{n}=\left\{\left(X_{i}, C\left(X_{i}\right)\right)\right\}_{i=1}^{n}$ where $X_{i} \in \mathbb{R}^{d}$ is a feature vector and $C\left(X_{i}\right) \in\{0,1\}$ is a label. The feature vectors are assumed to be independent samples from an unknown probability measure $\mu$ and $C: \mathbb{R}^{d} \rightarrow\{0,1\}$ is a measurable Boolean function. For instance, the feature vector might be an image (encoded as a vector) and the label might indicate "cat" (label 0) or "dog" (label 1). Our goal is learn the function (or concept) $C$ from the samples.

More precisely, we seek to construct a hypothesis $h: \mathbb{R}^{d} \rightarrow\{0,1\}$ that is a good approximation to $C$ in the sense that it predicts the label well on a new sample (from the same distribution). Formally, we want $h$ to have small true risk (or generalization error),

$$
R(h)=\mathbb{P}[h(X) \neq C(X)]
$$

where $X \sim \mu$. Because we only have access to the distribution $\mu$ through the samples, it is natural to estimate the true risk of the hypothesis $h$ using the samples as

$$
R_{n}(h)=\frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\left\{h\left(X_{i}\right) \neq C\left(X_{i}\right)\right\},
$$

which is called the empirical risk. Indeed observe that $\mathbb{E} R_{n}(h)=R(h)$ and, by the law of large numbers, $R_{n}(h) \rightarrow R(h)$ almost surely as $n \rightarrow+\infty$. Ignoring computational considerations, one can then formally define an empirical risk minimizer

$$
h^{*} \in \operatorname{ERM}_{\mathcal{H}}\left(\mathcal{S}_{n}\right)=\left\{h \in \mathcal{H}: R_{n}(h) \leq R_{n}\left(h^{\prime}\right), \forall h^{\prime} \in \mathcal{H}\right\},
$$

where $\mathcal{H}$, the hypothesis class, is a given collection of Boolean functions over $\mathbb{R}^{d}$. (We assume that $h^{*}$ can be defined as a measurable function of the samples.)

Overfitting Why restrict the hypothesis class? It turns out that minimizing the empirical risk over all Boolean functions makes it impossible to achieve an arbitrarily small risk. Intuitively considering too rich a class of functions, that is, functions that too intricately follow the data, leads to overfitting: the learned hypothesis will fit the sampled data, but it may not generalize well to unseen examples. A learner $\mathcal{A}$ is a map from samples to measurable Boolean functions over $\mathbb{R}^{d}$, that is, for any $n$ and any $\mathcal{S}_{n} \in\left(\mathbb{R}^{d} \times\{0,1\}\right)^{n}$, the learner outputs a function $\mathcal{A}\left(\cdot, \mathcal{S}_{n}\right): \mathbb{R}^{d} \rightarrow\{0,1\}$. The following theorem shows that any learner has fundamental limitations if all concepts are possible.

Theorem 2.4.42 (No Free Lunch). For any learner $\mathcal{A}$ and any finite $\mathcal{X} \subseteq \mathbb{R}^{d}$ of even size $|\mathcal{X}|=: 2 m>4$, there exist a concept $C: \mathcal{X} \rightarrow\{0,1\}$ and a distribution $\mu$ over $\mathcal{X}$ such that

$$
\begin{equation*}
\mathbb{P}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right) \geq 1 / 8\right] \geq 1 / 8 \tag{2.4.50}
\end{equation*}
$$

where $\mathcal{S}_{m}=\left\{\left(X_{i}, C\left(X_{i}\right)\right)\right\}_{i=1}^{m}$ with independent $X_{i} \sim \mu$.
The gist of the proof is intuitive. In essence, if the target concept is arbitrary and we only get to see half of the possible instances, then we have learned nothing about the other half and cannot expect low generalization error.

Proof of Theorem 2.4.42. We let $\mu$ be uniform over $\mathcal{X}$. To prove the existence of a concept satisfying (2.4.50), we use the probabilistic method (Section 2.2.1) and pick $C$ at random. For each $x \in \mathcal{X}$, we set $C(x):=Y_{x}$ where the $Y_{x} \mathrm{~s}$ are i.i.d. uniform in $\{0,1\}$.

We first bound $\mathbb{E}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right)\right]$, where the expectation runs over both random labels $\left\{Y_{x}\right\}_{x \in \mathcal{X}}$ and the samples $\mathcal{S}_{m}=\left\{\left(X_{i}, C\left(X_{i}\right)\right)\right\}_{i=1}^{m}$. For an additional independent sample $X \sim \mu$, we will need the event that the learner, given samples $\mathcal{S}_{m}$, makes an incorrect prediction on $X$

$$
\left.B=\left\{\mathcal{A}\left(X, \mathcal{S}_{m}\right)\right) \neq Y_{X}\right\}
$$

and the event that $X$ is observed in the samples $\mathcal{S}_{m}$

$$
O=\left\{X \in\left\{X_{1}, \ldots, X_{m}\right\}\right\} .
$$

By the tower property (Lemma B.6.16),

$$
\begin{aligned}
\mathbb{E}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right)\right] & =\mathbb{P}[B] \\
& =\mathbb{E}\left[\mathbb{P}\left[B \mid \mathcal{S}_{m}\right]\right] \\
& =\mathbb{E}\left[\mathbb{P}\left[B \mid O, \mathcal{S}_{m}\right] \mathbb{P}\left[O \mid \mathcal{S}_{m}\right]+\mathbb{P}\left[B \mid O^{c}, \mathcal{S}_{m}\right] \mathbb{P}\left[O^{c} \mid \mathcal{S}_{m}\right]\right] \\
& \geq \mathbb{E}\left[\mathbb{P}\left[B \mid O^{c}, \mathcal{S}_{m}\right] \mathbb{P}\left[O^{c} \mid \mathcal{S}_{m}\right]\right] \\
& \geq \frac{1}{2} \times \frac{1}{2},
\end{aligned}
$$

where we used that:

- $\mathbb{P}\left[O^{c} \mid \mathcal{S}_{m}\right] \geq 1 / 2$ because $|\mathcal{X}|=2 m$ and $\mu$ is uniform, and;
- $\mathbb{P}\left[B \mid O^{c}, \mathcal{S}_{m}\right]=1 / 2$ because for any $x \notin\left\{X_{1}, \ldots, X_{m}\right\}$ the prediction $\mathcal{A}\left(x, \mathcal{S}_{m}\right) \in\{0,1\}$ is independent of $Y_{x}$ and the latter is uniform.
Conditioning over the concept, we have proved that

$$
\mathbb{E}\left[\mathbb{E}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right) \mid\left\{Y_{x}\right\}_{x \in \mathcal{X}}\right]\right] \geq \frac{1}{4}
$$

Hence, by the first moment principle (Theorem 2.2.1),

$$
\mathbb{P}\left[\mathbb{E}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right) \mid\left\{Y_{x}\right\}_{x \in \mathcal{X}}\right] \geq 1 / 4\right]>0,
$$

where the probability is taken over $\left\{Y_{x}\right\}_{x \in \mathcal{X}}$. That is, there exists a choice $\left\{y_{x}\right\}_{x \in \mathcal{X}} \in$ $\{0,1\}^{\mathcal{X}}$ such that

$$
\begin{equation*}
\mathbb{E}\left[R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right) \mid\left\{Y_{x}=y_{x}\right\}_{x \in \mathcal{X}}\right] \geq 1 / 4 \tag{2.4.51}
\end{equation*}
$$

Finally, to prove (2.4.50), we use a variation on Markov's inequality (Theorem 2.1.1) for $[0,1]$-valued random variables. If $Z \in[0,1]$ is a random variable with $\mathbb{E}[Z]=\mu$ and $\alpha \in[0,1]$, then

$$
\mathbb{E}[Z] \leq \alpha \times \mathbb{P}[Z<\alpha]+1 \times \mathbb{P}[Z \geq \alpha] \leq \mathbb{P}[Z \geq \alpha]+\alpha
$$

Taking $\alpha=\mu / 2$ gives

$$
\mathbb{P}[Z \geq \mu / 2] \geq \mu / 2 .
$$

Going back to (2.4.51), we obtain

$$
\mathbb{P}\left[\left.R\left(\mathcal{A}\left(\cdot, \mathcal{S}_{m}\right)\right) \geq \frac{1}{8} \right\rvert\,\left\{Y_{x}=y_{x}\right\}_{x \in \mathcal{X}}\right] \geq \frac{1}{8}
$$

establishing the claim.

The way out is to "limit the complexity" of the hypotheses. For instance, we could restrict ourselves to half-spaces

$$
\mathcal{H}_{\mathrm{H}}=\left\{h(x)=\mathbf{1}\left\{x^{T} u \geq \alpha\right\}: u \in \mathbb{R}^{d}, \alpha \in \mathbb{R}\right\}
$$

or axis-aligned boxes

$$
\mathcal{H}_{\mathrm{B}}=\left\{h(x)=\mathbf{1}\left\{x_{i} \in\left[\alpha_{i}, \beta_{i}\right], \forall i\right\}:-\infty \leq \alpha_{i} \leq \beta_{i} \leq \infty, \forall i\right\} .
$$

In order for the empirical risk minimizer $h^{*}$ to have a generalization error close to the best achievable error, we need the empirical risk of the learned hypothesis $R_{n}\left(h^{*}\right)$ to be close to its expectation $R\left(h^{*}\right)$, which is guaranteed by the law of large numbers for sufficiently large $n$. But that is not enough, we also need that same property to hold for all hypotheses in $\mathcal{H}$ simultaneously. Otherwise we could be fooled by a poorly performing hypothesis with unusually good empirical risk on the samples. The hypothesis class is typically infinite and, therefore, controlling empirical risk deviations from their expectations uniformly over $\mathcal{H}$ is not straightforward.

Uniform deviations Our goal in this section is to show how to bound

$$
\begin{equation*}
\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{R_{n}(h)-R(h)\right\}\right]=\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{\frac{1}{n} \sum_{i=1}^{n} \ell\left(h, X_{i}\right)-\mathbb{E}[\ell(h, X)]\right\}\right] \tag{2.4.52}
\end{equation*}
$$

in terms of a measure of complexity of the class $\mathcal{H}$, where we defined the loss $\ell(h, x)=\mathbf{1}\{h(x) \neq C(x)\}$ to simplify the notation. We assume that $\mathcal{H}$ is countable. (Observe for instance that, for $\mathcal{H}_{\mathrm{H}}$ and $\mathcal{H}_{\mathrm{B}}$, nothing is lost by assuming that the parameters defining the hypotheses are rational-valued.)

Controlling deviations uniformly over $\mathcal{H}$ as in (2.4.52) allows one to provide guarantees on the empirical risk minimizer. Indeed, for any $h^{\prime} \in \mathcal{H}$,

$$
\begin{aligned}
R\left(h^{*}\right) & =R_{n}\left(h^{*}\right)+\left\{R\left(h^{*}\right)-R_{n}\left(h^{*}\right)\right\} \\
& \leq R_{n}\left(h^{*}\right)+\sup _{h \in \mathcal{H}}\left\{R(h)-R_{n}(h)\right\} \\
& \leq R_{n}\left(h^{\prime}\right)+\sup _{h \in \mathcal{H}}\left\{R(h)-R_{n}(h)\right\} \\
& =R\left(h^{\prime}\right)+\left\{R_{n}\left(h^{\prime}\right)-R\left(h^{\prime}\right)\right\}+\sup _{h \in \mathcal{H}}\left\{R(h)-R_{n}(h)\right\} \\
& \leq R\left(h^{\prime}\right)+\sup _{h \in \mathcal{H}}\left\{R_{n}(h)-R(h)\right\}+\sup _{h \in \mathcal{H}}\left\{R(h)-R_{n}(h)\right\},
\end{aligned}
$$

where, on the third line, we used the definition of the empirical risk minimizer. Taking an infimum over $h^{\prime}$, then an expectation over the samples, and rearranging gives

$$
\begin{align*}
& \mathbb{E}\left[R\left(h^{*}\right)\right]-\inf _{h^{\prime} \in \mathcal{H}} R\left(h^{\prime}\right) \\
& \quad \leq \mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{R_{n}(h)-R(h)\right\}\right]+\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{R(h)-R_{n}(h)\right\}\right] . \tag{2.4.53}
\end{align*}
$$

This inequality allows us to relate two quantities of interest: the expected true risk of the empirical risk minimizer (i.e., $\mathbb{E}\left[R\left(h^{*}\right)\right]$, where recall that $h^{*}$ is defined over the samples) and the best possible true risk (i.e., $\inf _{h^{\prime} \in \mathcal{H}} R\left(h^{\prime}\right)$ ). The first term on the right-hand side is (2.4.52) and the second one can be bounded in a similar fashion as we argue below. Observe that the suprema are inside the expectations and that the random variables $R_{n}(h)-R(h)$ are highly correlated. Indeed, two similar hypotheses will produce similar predictions. The correlation is ultimately what allows us to tackle infinite classes $\mathcal{H}$ - as we saw in Section 2.4.4.

Indeed, to bound (2.4.52), we use the methods of Section 2.4.4. As a first step, we apply the symmetrization trick, which we introduced in Section 2.4.2 to give a proof of Hoeffding's lemma (Lemma 2.4.12). Let $\left(\varepsilon_{i}\right)_{i=1}^{n}$ be i.i.d. uniform random variables in $\{-1,+1\}$ (i.e., Rademacher variables) and let $\left(X_{i}^{\prime}\right)_{i=1}^{n}$ be an independent copy of $\left(X_{i}\right)_{i=1}^{n}$. Then

$$
\begin{aligned}
& \mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{R_{n}(h)-R(h)\right\}\right] \\
& =\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{\frac{1}{n} \sum_{i=1}^{n} \ell\left(h, X_{i}\right)-\mathbb{E}[\ell(h, X)]\right\}\right] \\
& =\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{\frac{1}{n} \sum_{i=1}^{n}\left[\ell\left(h, X_{i}\right)-\mathbb{E}\left[\ell\left(h, X_{i}^{\prime}\right) \mid\left(X_{j}\right)_{j=1}^{n}\right]\right]\right\}\right] \\
& =\mathbb{E}\left[\sup _{h \in \mathcal{H}} \mathbb{E}\left[\left.\frac{1}{n} \sum_{i=1}^{n}\left[\ell\left(h, X_{i}\right)-\ell\left(h, X_{i}^{\prime}\right)\right] \right\rvert\,\left(X_{j}\right)_{j=1}^{n}\right]\right] \\
& \leq \mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{\frac{1}{n} \sum_{i=1}^{n}\left[\ell\left(h, X_{i}\right)-\ell\left(h, X_{i}^{\prime}\right)\right]\right\}\right],
\end{aligned}
$$

where on the fourth line we used "taking it out what is known" (Lemma B.6.13) and on the fifth line we used $\sup _{h} \mathbb{E} Y_{h} \leq \mathbb{E}\left[\sup _{h} Y_{h}\right]$ and the tower property. Next we note that $\ell\left(h, X_{i}\right)-\ell\left(h, X_{i}^{\prime}\right)$ is symmetric and independent of $\varepsilon_{i}$ (which is also
symmetric) to deduce that the last line above is

$$
\begin{aligned}
& =\mathbb{E}\left[\sup _{h \in \mathcal{H}}\left\{\frac{1}{n} \sum_{i=1}^{n} \varepsilon_{i}\left[\ell\left(h, X_{i}\right)-\ell\left(h, X_{i}^{\prime}\right)\right]\right\}\right] \\
& \leq \mathbb{E}\left[\sup _{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \varepsilon_{i} \ell\left(h, X_{i}\right)+\sup _{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n}\left(-\varepsilon_{i}\right) \ell\left(h, X_{i}^{\prime}\right)\right] \\
& =2 \mathbb{E}\left[\sup _{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \varepsilon_{i} \ell\left(h, X_{i}\right)\right] .
\end{aligned}
$$

The exact same argument also applies to the second term on the right-hand side of (2.4.53), so

$$
\begin{equation*}
\mathbb{E}\left[R\left(h^{*}\right)\right]-\inf _{h^{\prime} \in \mathcal{H}} R\left(h^{\prime}\right) \leq 4 \mathbb{E}\left[\sup _{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \varepsilon_{i} \ell\left(h, X_{i}\right)\right] . \tag{2.4.54}
\end{equation*}
$$

Changing the normalization, we define the process

$$
\begin{equation*}
Z_{n}(h)=\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \varepsilon_{i} \ell\left(h, X_{i}\right), \quad h \in \mathcal{H} . \tag{2.4.55}
\end{equation*}
$$

Our task reduces to upper bounding

$$
\begin{equation*}
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] . \tag{2.4.56}
\end{equation*}
$$

Note that we will not compute the best possible true risk (which in general could be "bad," i.e., large)—only how close the empirical risk minimizer gets to it.

VC dimension We make two observations about $Z_{n}(h)$.

1. It is centered. Also, as a weighted sum of independent random variables in $[-1,1]$, it is sub-Gaussian with variance factor 1 by the general Hoeffding inequality (Theorem 2.4.9) and Hoeffding's lemma (Lemma 2.4.12).
2. It depends only on the values of the hypothesis $h$ at a finite number of points, $X_{1}, \ldots, X_{n}$. Hence, while the supremum in (2.4.56) is over a potentially infinite class of functions $\mathcal{H}$, it is in effect a supremum over at most $2^{n}$ functions, that is, all the possible restrictions of the $h \mathrm{~s}$ to $\left(X_{i}\right)_{i=1}^{n}$.

A naive application of the maximal inequality in Lemma 2.4.21, together with the two observations above, gives

$$
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] \leq \sqrt{2 \log 2^{n}}=\sqrt{2 n \log 2} .
$$

Unfortunately, plugging this back into (2.4.54) gives an upper bound which fails to converge to 0 as $n \rightarrow+\infty$.

To obtain a better bound, we show that in general the number of distinct restrictions of $\mathcal{H}$ to $n$ points can grow much slower than $2^{n}$.
Definition 2.4.43 (Shattering). Let $\Lambda=\left\{\ell_{1}, \ldots, \ell_{n}\right\} \subseteq \mathbb{R}^{d}$ be a finite set and let $\mathcal{H}$ be a class of Boolean functions on $\mathbb{R}^{d}$. The restriction of $\mathcal{H}$ to $\Lambda$ is

$$
\mathcal{H}_{\Lambda}=\left\{\left(h\left(\ell_{1}\right), \ldots, h\left(\ell_{n}\right)\right): h \in \mathcal{H}\right\} .
$$

We say that $\Lambda$ is shattered by $\mathcal{H}$ if $\left|\mathcal{H}_{\Lambda}\right|=2^{|\Lambda|}$, that is, if all Boolean functions over $\Lambda$ can be obtained by restricting a function in $\mathcal{H}$ to the points in $\Lambda$.
Definition 2.4.44 (VC dimension). Let $\mathcal{H}$ be a class of Boolean functions on $\mathbb{R}^{d}$. The VC dimension of $\mathcal{H}$, denoted $\mathrm{vc}(\mathcal{H})$, is the maximum cardinality of a set shattered by $\mathcal{H}$.
We prove the following combinatorial lemma at the end of this section.
Lemma 2.4.45 (Sauer's lemma). Let $\mathcal{H}$ be a class of Boolean functions on $\mathbb{R}^{d}$. For any finite set $\Lambda=\left\{\ell_{1}, \ldots, \ell_{n}\right\} \subseteq \mathbb{R}^{d}$,

$$
\left|\mathcal{H}_{\Lambda}\right| \leq\left(\frac{e n}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})}
$$

That is, the number of distinct restrictions of $\mathcal{H}$ to any $n$ points grows at most as $\propto n^{\mathrm{vc}(\mathcal{H})}$.

Returning to $\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right]$, we get the following inequality.
Lemma 2.4.46. There exists a constant $0<C<+\infty$ such that, for any countable class of measurable Boolean functions $\mathcal{H}$ over $\mathbb{R}^{d}$,

$$
\begin{equation*}
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] \leq C \sqrt{\operatorname{vc}(\mathcal{H}) \log n} . \tag{2.4.57}
\end{equation*}
$$

Proof. Recall that $Z_{n}(h) \in \mathrm{s} \mathcal{G}(1)$. Since the supremum over $\mathcal{H}$, when seen as restricted to $\left\{X_{1}, \ldots, X_{n}\right\}$, is in fact a supremum over at $\operatorname{most}\left(\frac{e n}{\operatorname{vc}(\mathcal{H})}\right)^{\operatorname{vc}(\mathcal{H})}$ functions by Sauer's lemma (Lemma 2.4.45), we have by Lemma 2.4.21

$$
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] \leq \sqrt{2 \log \left[\left(\frac{e n}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})}\right]} .
$$

That proves the claim.
Returning to (2.4.54), the previous lemma finally implies

$$
\mathbb{E}\left[R\left(h^{*}\right)\right]-\inf _{h^{\prime} \in \mathcal{H}} R\left(h^{\prime}\right) \leq 4 C \sqrt{\frac{\mathrm{vc}(\mathcal{H}) \log n}{n}} .
$$

For hypothesis classes with finite VC dimension, the bound goes to 0 as $n \rightarrow+\infty$.
We give some examples.
Example 2.4.47 (VC dimension of half-spaces). Consider the class of half-spaces.
Claim 2.4.48.

$$
\operatorname{vc}\left(\mathcal{H}_{\mathrm{H}}\right)=d+1 .
$$

We only prove the case $d=1$, where $\mathcal{H}_{\mathrm{H}}$ reduces to half-lines $(-\infty, \gamma]$ or $[\gamma,+\infty)$. Clearly any set $\Lambda=\left\{\ell_{1}, \ell_{2}\right\} \subseteq \mathbb{R}$ with elements is shattered by $\mathcal{H}_{\mathrm{H}}$. On the other hand, for any $\Lambda=\left\{\ell_{1}, \ell_{2}, \ell_{3}\right\}$ with $\ell_{1}<\ell_{2}<\ell_{3}$, any half-line containing $\ell_{1}$ and $\ell_{3}$ necessarily includes $\ell_{2}$ as well. Hence no set of size 3 is shattered by $\mathcal{H}_{\mathrm{H}}$.

Example 2.4.49 (VC dimension of boxes). Consider the class of axis-aligned boxes.

## Claim 2.4.50.

$$
\operatorname{vc}\left(\mathcal{H}_{\mathrm{B}}\right)=2 d .
$$

We only prove the case $d=2$, where $\mathcal{H}_{\mathrm{B}}$ reduces to rectangles. The four-point set $\Lambda=\{(-1,0),(1,0),(0,-1),(0,1)\}$ is shattered by $\mathcal{H}_{\mathrm{B}}$. Indeed, the rectangle $[-1,1] \times[-1,1]$ contains $\Lambda$, with each side of the rectangle containing one of the points. Moving any side inward by $\varepsilon<1$ removes the corresponding point from the rectangle without affecting the other ones. Hence, any subset of $\Lambda$ can be obtained by this procedure.

On the other hand, let $\Lambda=\left\{\ell_{1}, \ldots, \ell_{5}\right\} \subseteq \mathbb{R}^{2}$ be any set of five distinct points. If the points all lie on the same axis-aligned line, then an argument similar to the half-line case in Claim 2.4.48 shows that $\Lambda$ is not shattered. Otherwise consider the axis-aligned rectangle with smallest area containing $\Lambda$. For each side of the rectangle, choose one point of $\Lambda$ that lies on it. These necessarily exist (otherwise the rectangle could be made even smaller) and denote them by $x_{\mathrm{N}}$ for the highest, $x_{\mathrm{E}}$ for the rightmost, $x_{\mathrm{S}}$ for the lowest, and $x_{\mathrm{W}}$ for the leftmost. Note that they may not be distinct, but in any case at least one point in $\Lambda$, say $\ell_{5}$ without loss of generality, is not in the list. Now observe that any axis-aligned rectangle containing $x_{\mathrm{N}}, x_{\mathrm{E}}, x_{\mathrm{S}}, x_{\mathrm{W}}$ must also contain $\ell_{5}$ since its coordinates are sandwiched between the bounds defined by those points. Hence no set of size 5 is shattered. That proves the claim.

These two examples also provide insights into Sauer's lemma. Consider the case of rectangles for instance. Over a collection of $n$ sample points, a rectangle defines the same $\{0,1\}$-labeling as the minimal-area rectangle containing the same points. Because each side of a minimal-area rectangle must touch at least one point in the sample, there are at most $n^{4}$ such rectangles, and hence there are at most $n^{4} \ll 2^{n}$ restrictions of $\mathcal{H}_{\mathrm{B}}$ to these sample points.

Application of chaining It turns out that the $\sqrt{\log n}$ factor in (2.4.57) is not optimal. We use chaining (Section 2.4.4) to improve the bound.

We claim that the process $\left\{Z_{n}(h)\right\}_{h \in \mathcal{H}}$ has sub-Gaussian increments under an appropriately defined pseudometric. Indeed, conditioning on $\left(X_{i}\right)_{i=1}^{n}$, by the general Hoeffding inequality (Theorem 2.4.9) and Hoeffding's lemma (Lemma 2.4.12), we have that the increment (as a function of the $\varepsilon_{i}$ s which have variance factor 1)

$$
Z_{n}(g)-Z_{n}(h)=\sum_{i=1}^{n} \varepsilon_{i} \frac{\ell\left(g, X_{i}\right)-\ell\left(h, X_{i}\right)}{\sqrt{n}},
$$

is sub-Gaussian with variance factor

$$
\sum_{i=1}^{n}\left(\frac{\ell\left(g, X_{i}\right)-\ell\left(h, X_{i}\right)}{\sqrt{n}}\right)^{2} \times 1=\frac{1}{n} \sum_{i=1}^{n}\left[\ell\left(g, X_{i}\right)-\ell\left(h, X_{i}\right)\right]^{2} .
$$

Define the pseudometric

$$
\rho_{n}(g, h)=\left[\frac{1}{n} \sum_{i=1}^{n}\left[\ell\left(g, X_{i}\right)-\ell\left(h, X_{i}\right)\right]^{2}\right]^{1 / 2}=\left[\frac{1}{n} \sum_{i=1}^{n}\left[g\left(X_{i}\right)-h\left(X_{i}\right)\right]^{2}\right]^{1 / 2},
$$

where we used that $\ell(h, x)=\mathbf{1}\{h(x) \neq C(x)\}$ by definition. It satisfies the triangle inequality since it can be expressed as a Euclidean norm. In fact, it will be useful to recast it in a more general setting. For a probability measure $\eta$ over $\mathbb{R}^{d}$, define

$$
\|g-h\|_{L^{2}(\eta)}^{2}=\int_{\mathbb{R}^{d}}(f(x)-g(x))^{2} \mathrm{~d} \eta(x) .
$$

Let $\mu_{n}$ be the empirical measure

$$
\begin{equation*}
\mu_{n}=\mu_{\left(X_{i}\right)_{i=1}^{n}}:=\frac{1}{n} \sum_{i=1}^{n} \delta_{X_{i}}, \tag{2.4.58}
\end{equation*}
$$

empirical
measure
where $\delta_{x}$ is the probability measure that puts mass 1 on $x$. Then, we can re-write

$$
\rho_{n}(g, h)=\|g-h\|_{L^{2}\left(\mu_{n}\right)} .
$$

Hence we have shown that, conditioned on the samples, the process $\left\{Z_{n}(h)\right\}_{h \in \mathcal{H}}$ has sub-Gaussian increments with respect to $\|\cdot\|_{L^{2}\left(\mu_{n}\right)}$. Note that the pseudometric here is random as it depends on the samples. Though, by the law of large numbers, $\|g-h\|_{L^{2}\left(\mu_{n}\right)}$ approaches its expectation, $\|g-h\|_{L^{2}(\mu)}$, as $n \rightarrow+\infty$.

Applying the discrete Dudley inequality (Theorem 2.4.31), we obtain the following bound.

Lemma 2.4.51. There exists a constant $0<C<+\infty$ such that, for any countable class of measurable Boolean functions $\mathcal{H}$ over $\mathbb{R}^{d}$,

$$
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] \leq C \mathbb{E}\left[\sum_{k=0}^{+\infty} 2^{-k} \sqrt{\log \mathcal{N}\left(\mathcal{H},\|\cdot\|_{L^{2}\left(\mu_{n}\right)}, 2^{-k}\right)}\right],
$$

where $\mu_{n}$ is the empirical measure over the samples $\left(X_{i}\right)_{i=1}^{n}$.
Proof. Because $\mathcal{H}$ comprises only Boolean functions, it follows that under the pseudometric $\|\cdot\|_{L^{2}\left(\mu_{n}\right)}$ the diameter is bounded by 1 . We apply the discrete Dudley inequality conditioned on $\left(X_{i}\right)_{i=1}^{n}$. Then we take an expectation over the samples.

Our use of the symmetrization trick is more intuitive than it may have appeared at first. The central limit theorem indicates that the fluctuations of centered averages such as

$$
\left(R_{n}(g)-R(g)\right)-\left(R_{n}(h)-R(h)\right)
$$

tend cancel out and that, in the limit, the variance alone characterizes the overall behavior. The $\varepsilon_{i}$ s in some sense explicitly capture the canceling part of this phenomenon while $\rho_{n}$ captures the scale of the resulting global fluctuations in the increments.

Our final task is to bound the covering numbers $\mathcal{N}\left(\mathcal{H},\|\cdot\|_{L^{2}\left(\mu_{n}\right)}, 2^{-k}\right)$.
Theorem 2.4.52 (Covering numbers via VC dimension). There exists a constant $0<C<+\infty$ such that, for any class of measurable Boolean functions $\mathcal{H}$ over $\mathbb{R}^{d}$, any probability measure $\eta$ over $\mathbb{R}^{d}$ and any $\varepsilon \in(0,1)$,

$$
\mathcal{N}\left(\mathcal{H},\|\cdot\|_{L^{2}(\eta)}, \varepsilon\right) \leq\left(\frac{2}{\varepsilon}\right)^{C \operatorname{vc}(\mathcal{H})} .
$$

Before proving Theorem 2.4.52, we derive its implications for uniform deviations. Compare the following bound to Lemma 2.4.46.

Lemma 2.4.53. There exists a constant $0<C<+\infty$ such that, for any countable class of measurable Boolean functions $\mathcal{H}$ over $\mathbb{R}^{d}$,

$$
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] \leq C \sqrt{\mathrm{vc}(\mathcal{H})} .
$$

Proof. By Lemma 2.4.51 and Theorem 2.4.52,

$$
\begin{aligned}
\mathbb{E}\left[\sup _{h \in \mathcal{H}} Z_{n}(h)\right] & \leq C \mathbb{E}\left[\sum_{k=0}^{+\infty} 2^{-k} \sqrt{\log \mathcal{N}\left(\mathcal{H},\|\cdot\|_{L^{2}\left(\mu_{n}\right)}, 2^{-k}\right)}\right] \\
& \leq C \mathbb{E}\left[\sum_{k=0}^{+\infty} 2^{-k} \sqrt{\log \left(\frac{2}{2^{-k}}\right)^{C^{\prime} \operatorname{vc}(\mathcal{H})}}\right] \\
& =C \sqrt{\operatorname{vc}(\mathcal{H})} \mathbb{E}\left[\sum_{k=0}^{+\infty} 2^{-k} \sqrt{k+1} \sqrt{C^{\prime} \log 2}\right] \\
& \leq C^{\prime \prime} \sqrt{\operatorname{vc}(\mathcal{H})}
\end{aligned}
$$

for some $0<C^{\prime \prime}<+\infty$.
It remains to prove Theorem 2.4.52.
Proof of Theorem 2.4.52. Let $\mathcal{G}=\left\{g_{1}, \ldots, g_{N}\right\} \subseteq \mathcal{H}$ be a maximal $\varepsilon$-packing of $\mathcal{H}$ with $N \geq \mathcal{N}\left(\mathcal{H},\|\cdot\|_{L^{2}(\eta)}, \varepsilon\right)$, which exists by Lemma 2.4.24. We use the probabilistic method (Section 2.2) and Hoeffding's inequality for bounded variables (Theorem 2.4.10) to show that there exists a small number of points $\left\{x_{1}, \ldots, x_{m}\right\}$ such that $\mathcal{G}$ is still a good packing when $\mathcal{H}$ is restricted to the $x_{i}$ s. Then we use Sauer's lemma (Lemma 2.4.45) to conclude.

1. Restriction. By construction, the collection $\mathcal{G}$ satisfies

$$
\left\|g_{i}-g_{j}\right\|_{L^{2}(\eta)}>\varepsilon, \quad \forall i \neq j .
$$

For an integer $m$ that we will choose as small as possible below, let $\mathbb{X}=$ $\left\{X_{1}, \ldots, X_{m}\right\}$ be i.i.d. samples from $\eta$ and let $\mu_{\mathbb{X}}$ be the corresponding empirical measure (as defined in (2.4.58)). Observe that, for any $i \neq j$,

$$
\mathbb{E}\left[\left\|g_{i}-g_{j}\right\|_{L^{2}\left(\mu_{\mathrm{x}}\right)}^{2}\right]=\mathbb{E}\left[\frac{1}{m} \sum_{k=1}^{m}\left[g_{i}\left(X_{k}\right)-g_{j}\left(X_{k}\right)\right]^{2}\right]=\left\|g_{i}-g_{j}\right\|_{L^{2}(\eta)}^{2} .
$$

Moreover $\left[g_{i}\left(X_{k}\right)-g_{j}\left(X_{k}\right)\right]^{2} \in[0,1]$. Hence, by Hoeffding's inequality there exists a constant $0<C<+\infty$ and an $m \leq C \varepsilon^{-4} \log N$ such that

$$
\begin{aligned}
& \mathbb{P}\left[\left\|g_{i}-g_{j}\right\|_{L^{2}(\eta)}^{2}-\left\|g_{i}-g_{j}\right\|_{L^{2}(\mu \mathbb{X})}^{2} \geq \frac{3 \varepsilon^{2}}{4}\right] \\
& =\mathbb{P}\left[m\left\|g_{i}-g_{j}\right\|_{L^{2}(\eta)}^{2}-\sum_{k=1}^{m}\left[g_{i}\left(X_{k}\right)-g_{j}\left(X_{k}\right)\right]^{2} \geq m \frac{3 \varepsilon^{2}}{4}\right] \\
& \leq \exp \left(-\frac{2\left(m \cdot 3 \varepsilon^{2} / 4\right)^{2}}{m}\right) \\
& =\exp \left(-\frac{9}{8} m \varepsilon^{4}\right) \\
& <\frac{1}{N^{2}} .
\end{aligned}
$$

That implies that, for this choice of $m$,

$$
\mathbb{P}\left[\left\|g_{i}-g_{j}\right\|_{L^{2}\left(\mu_{\mathbb{X}}\right)}>\frac{\varepsilon}{2}, \forall i \neq j\right]>0
$$

where the probability is over the samples and we used the assumption on the collection $\mathcal{G}$. Therefore, there must be a set $\mathcal{X}=\left\{x_{1}, \ldots, x_{m}\right\} \subseteq \mathbb{R}^{d}$ such that

$$
\begin{equation*}
\left\|g_{i}-g_{j}\right\|_{L^{2}\left(\mu_{\mathcal{X}}\right)}>\frac{\varepsilon}{2}, \forall i \neq j . \tag{2.4.59}
\end{equation*}
$$

2. VC bound. In particular, by (2.4.59), the functions in $\mathcal{G}$ restricted to $\mathcal{X}$ are distinct. By Sauer's lemma (Lemma 2.4.45),

$$
\begin{equation*}
N=\left|\mathcal{G X X}_{\mathcal{X}}\right| \leq\left|\mathcal{H}_{\mathcal{X}}\right| \leq\left(\frac{e m}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})} \leq\left(\frac{e C \varepsilon^{-4} \log N}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})} \tag{2.4.60}
\end{equation*}
$$

Using that $\frac{1}{2 D} \log N=\log N^{1 / 2 D} \leq N^{1 / 2 D}$ where $D=\operatorname{vc}(\mathcal{H})$, we get

$$
\begin{equation*}
\left(\frac{e C \varepsilon^{-4} \log N}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})} \leq\left(C^{\prime} \varepsilon^{-4}\right)^{\mathrm{vc}(\mathcal{H})} N^{1 / 2} \tag{2.4.61}
\end{equation*}
$$

where $C^{\prime}=2 e C$. Plugging (2.4.61) back into (2.4.60) and rearranging gives

$$
N \leq\left(C^{\prime} \varepsilon^{-4}\right)^{2 \operatorname{vc}(\mathcal{H})}
$$

That concludes the proof.

Proof of Sauer's lemma Recall from Appendix A (see also Exercise 1.4) that for integers $0<d \leq n$,

$$
\begin{equation*}
\sum_{k=0}^{d}\binom{n}{k} \leq\left(\frac{e n}{d}\right)^{d} \tag{2.4.62}
\end{equation*}
$$

Sauer's lemma (Lemma 2.4.45) follows from the following claim.
Lemma 2.4.54 (Pajor). Let $\mathcal{H}$ be a class of Boolean functions on $\mathbb{R}^{d}$ and let $\Lambda={ }^{\text {Pajor's lemma }}$ $\left\{\ell_{1}, \ldots, \ell_{n}\right\} \subseteq \mathbb{R}^{d}$ be any finite subset. Then

$$
\left|\mathcal{H}_{\Lambda}\right| \leq \mid\{S \subseteq \Lambda: S \text { is shattered by } \mathcal{H}\} \mid,
$$

where the right-hand side includes the empty set.
Going back to Sauer's lemma, by Lemma 2.4.54 we have the upper bound

$$
\left|\mathcal{H}_{\Lambda}\right| \leq \mid\{S \subseteq \Lambda: S \text { is shattered by } \mathcal{H}\} \mid .
$$

By definition of the VC-dimension (Definition 2.4.44), the subsets $S \subseteq \Lambda$ that are shattered by $\mathcal{H}$ have size at most $\operatorname{vc}(\mathcal{H})$. So the right-hand side is bounded above by the total number of subsets of size at most $d=\mathrm{vc}(\mathcal{H})$ of a set of size $n$. By (2.4.62), this gives

$$
\left|\mathcal{H}_{\Lambda}\right| \leq\left(\frac{e n}{\operatorname{vc}(\mathcal{H})}\right)^{\mathrm{vc}(\mathcal{H})}
$$

which establishes Sauer's lemma.
So it remain to prove Lemma 2.4.54.
Proof of Lemma 2.4.54. We prove the claim by induction on the size $n$ of $\Lambda$. The result is trivial for $n=1$. Assume the result is true for any $\mathcal{H}$ and any subset of size $n-1$. To apply induction, for $\iota=0,1$ we let

$$
\mathcal{H}^{\iota}=\left\{h \in \mathcal{H}: h\left(\ell_{n}\right)=\iota\right\},
$$

and we set

$$
\Lambda^{\prime}=\left\{\ell_{1}, \ldots, \ell_{n-1}\right\} .
$$

It will be convenient to introduce the following notation

$$
\mathcal{S}(\Lambda ; \mathcal{H})=\mid\{S \subseteq \Lambda: S \text { is shattered by } \mathcal{H}\} \mid .
$$

Because $\left|\mathcal{H}_{\Lambda}\right|=\left|\mathcal{H}_{\Lambda^{\prime}}^{0}\right|+\left|\mathcal{H}_{\Lambda^{\prime}}^{1}\right|$ and the induction hypothesis implies $\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{\iota}\right) \geq$ $\left|\mathcal{H}_{\Lambda^{\prime}}^{\iota}\right|$ for $\iota=0,1$, it suffices to show that

$$
\begin{equation*}
\mathcal{S}(\Lambda ; \mathcal{H}) \geq \mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{0}\right)+\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{1}\right) \tag{2.4.63}
\end{equation*}
$$

There are two types of sets that contribute to the right-hand side.

- One but not both. Let $S \subseteq \Lambda^{\prime}$ be a set that contributes to one of $\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{0}\right)$ or $\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{1}\right)$ but not both. Then $S$ is a subset of the larger set $\Lambda$ and it is certainly shattered by the larger collection $\mathcal{H}$. Hence it also contributes to the left-hand side of (2.4.63).
- Both. Let $S \subseteq \Lambda^{\prime}$ be a set that contributes to both $\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{0}\right)$ and $\mathcal{S}\left(\Lambda^{\prime} ; \mathcal{H}^{1}\right)$. Hence it contributes two to the right-hand side of (2.4.63). As in the previous point, it is also included in $\mathcal{S}(\Lambda ; \mathcal{H})$, but it only contributes one to the lefthand side of (2.4.63). It turns out that there is another set that contributes one to the left-hand side but zero to the right-hand side: the subset $S \cup\left\{\ell_{n}\right\}$. Indeed, by definition of $\mathcal{H}^{\iota}$, the subset $S \cup\left\{\ell_{n}\right\}$ cannot be shattered by it since all functions in it take the same value on $\ell_{n}$. On the other hand, any Boolean function $h$ on $S \cup\left\{\ell_{n}\right\}$ with $h\left(\ell_{n}\right)=\iota$ is realized in $\mathcal{H}^{\iota}$ since $S$ itself is shattered by $\mathcal{H}^{\iota}$.

That concludes the proof.

## Exercises

Exercise 2.1 (Moments of nonnegative random variables). Prove (B.5.1). [Hint: Use Fubini's Theorem to compute the integral.]

Exercise 2.2 (Bonferroni inequalities). Let $A_{1}, \ldots, A_{n}$ be events and $B_{n}:=\cup_{i} A_{i}$. Define

$$
S^{(r)}:=\sum_{1 \leq i_{1}<\cdots<i_{r} \leq n} \mathbb{P}\left[A_{i_{1}} \cap \cdots \cap A_{i_{r}}\right],
$$

and

$$
X_{n}:=\sum_{i=1}^{n} \mathbf{1}_{A_{i}} .
$$

(i) Let $x_{0} \leq x_{1} \leq \cdots \leq x_{s} \geq x_{s+1} \geq \cdots \geq x_{m}$ be a unimodal sequence of nonnegative reals such that $\sum_{j=0}^{m}(-1)^{j} x_{j}=0$. Show that $\sum_{j=0}^{\ell}(-1)^{j} x_{j}$ is $\geq 0$ for even $\ell$ and $\leq 0$ for odd $\ell$.
(ii) Show that, for all $r$,

$$
\sum_{1 \leq i_{1}<\cdots<i_{r} \leq n} \mathbf{1}_{A_{i_{1}}} \mathbf{1}_{A_{i_{2}}} \cdots \mathbf{1}_{A_{i_{r}}}=\binom{X_{n}}{r} .
$$

(iii) Use (i) and (ii) to show that when $\ell \in[n]$ is odd

$$
\mathbb{P}\left[B_{n}\right] \leq \sum_{r=1}^{\ell}(-1)^{r-1} S^{(r)}
$$

and when $\ell \in[n]$ is even

$$
\mathbb{P}\left[B_{n}\right] \geq \sum_{r=1}^{\ell}(-1)^{r-1} S^{(r)}
$$

These inequalities are called Bonferroni inequalities. The case $\ell=1$ is Boole's inequality.

Exercise 2.3 (Percolation on $\mathbb{Z}^{2}$ : a better bound). Let $E_{1}$ be the event that all edges are open in $[-N, N]^{2}$ and $E_{2}$ be the event that there is no closed self-avoiding dual cycle surrounding $[-N, N]^{2}$. By looking at $E_{1} \cap E_{2}$, show that $\theta(p)>0$ for $p>2 / 3$.

Exercise 2.4 (Percolation on $\mathbb{Z}^{d}$ : existence of critical threshold). Consider bond percolation on $\mathbb{L}^{d}$.
(i) Show that $p_{c}\left(\mathbb{L}^{d}\right)>0$. [Hint: Count self-avoiding paths.]
(ii) Show that $p_{c}\left(\mathbb{L}^{d}\right)<1$. [Hint: Use the result for $\mathbb{L}^{2}$.]

Exercise 2.5 (Sums of uncorrelated variables). Centered random variables $X_{1}, X_{2}, \ldots$ are uncorrelated if

$$
\mathbb{E}\left[X_{r} X_{s}\right]=0, \quad \forall r \neq s
$$

(i) Assume further that $\operatorname{Var}\left[X_{r}\right] \leq C<+\infty$ for all $r$. Show that

$$
\mathbb{P}\left[\frac{1}{n} \sum_{r \leq n} X_{r} \geq \beta\right] \leq \frac{C^{2}}{\beta^{2} n} .
$$

(ii) Use (i) to prove Theorem 2.1.6.

Exercise 2.6 (Pairwise independence: lack of concentration). Let $\boldsymbol{U}=\left(U_{1}, \ldots, U_{\ell}\right)$ be uniformly distributed over $\{0,1\}^{\ell}$. Let $n=2^{\ell}-1$. For all $\boldsymbol{v} \in\{0,1\}^{\ell} \backslash \mathbf{0}$, define

$$
X_{\boldsymbol{v}}=\langle\boldsymbol{U}, \boldsymbol{v}\rangle \quad \bmod 2 .
$$

(i) Show that the random variables $X_{\boldsymbol{v}}, \boldsymbol{v} \in\{0,1\}^{\ell} \backslash \mathbf{0}$, are uniformly distributed in $\{0,1\}$ and pairwise independent.
(ii) Show that for any event $A$ measurable with respect to $\sigma\left(X_{\boldsymbol{v}}, \boldsymbol{v} \in\{0,1\}^{\ell} \backslash \mathbf{0}\right)$, $\mathbb{P}[A]$ is either 0 or $\geq 1 /(n+1)$.

Exercise 2.5 shows that pairwise independence implies "polynomial concentration" of the average of square-integrable $X_{v}$ s. On the other hand, the current exercise suggests that in general pairwise independence cannot imply "exponential concentration."

Exercise 2.7 (Chernoff bound for Poisson trials). Using the Chernoff-Cramér method, prove part (i) of Theorem 2.4.7. Show that part (ii) follows from part (i).

Exercise 2.8 (Stochastic knapsack: some details). Consider the stochastic fractional knapsack problem in Section 2.4.3.
(i) Prove that the greedy algorithm described there gives an optimal solution to problem (2.4.21).
(ii) Prove Claim 2.4.20 for $\tau \in(0,1 / 6)$.

Exercise 2.9 (Stochastic knapsack: 0-1 version). Consider the stochastic fractional knapsack problem in Section 2.4.3.
(i) Adapt the greedy algorithm for the 0-1 knapsack problem and show that it is not optimal in general. [Hint: Construct a counter-example with two items.]
(ii) Prove Claim 2.4.20 for the greedy solution of (i).

Exercise 2.10 (A proof of Pólya's theorem). Let $\left(S_{t}\right)$ be simple random walk on $\mathbb{L}^{d}$ started at the origin 0 .
(i) For $d=1$, use Stirling's formula (see Appendix A) to show that $\mathbb{P}\left[S_{2 n}=\right.$ $0]=\Theta\left(n^{-1 / 2}\right)$.
(ii) For $j=1, \ldots, d$, let $N_{t}^{(j)}$ be the number of steps in the $j$-th coordinate by time $t$. Show that

$$
\mathbb{P}\left[N_{n}^{(j)} \in\left[\frac{n}{2 d}, \frac{3 n}{2 d}\right], \forall j\right] \geq 1-\exp \left(-\kappa_{d} n\right)
$$

for some constant $\kappa_{d}>0$.
(iii) Use (i) and (ii) to show that, for any $d \geq 3, \mathbb{P}\left[S_{2 n}=0\right]=O\left(n^{-d / 2}\right)$.

Exercise 2.11 (Maximum degree). Let $G_{n}=\left(V_{n}, E_{n}\right) \sim \mathbb{G}_{n, p_{n}}$ be an ErdősRényi graph with $n$ vertices and density $p_{n}$. Suppose $n p_{n}=C \log n$ for some $C>0$. Let $D_{n}$ be the maximum degree of $G_{n}$. Use Bernstein's inequality to show that for any $\varepsilon>0$

$$
\mathbb{P}\left[D_{n} \geq(n-1) p_{n}+\max \{C, 4(1+\varepsilon)\} \log n\right] \rightarrow 0,
$$

as $n \rightarrow+\infty$.
Exercise 2.12 (RIP vs. orthogonality). Show that a ( $k, 0$ )-RIP matrix with $k \geq 2$ is orthogonal, that is, its columns are orthonormal.

Exercise 2.13 (Compressed sensing: linear programming formulation). Formulate (2.4.43) as a linear program, that is, the minimization of a linear objective subject to linear inequalities.

Exercise 2.14 (Compressed sensing: almost sparse case). By adapting the proof of Lemma 2.4.39, show the following "almost sparse" version. Let $L$ be $(10 k, 1 / 3)$ RIP. Then, for any $\mathbf{x} \in \mathbb{R}^{n}$, the solution to (2.4.43) satisfies

$$
\left\|\mathbf{z}^{*}-\mathbf{x}\right\|_{2}=O(\eta(\mathbf{x}) / \sqrt{k}),
$$

where $\eta(\mathbf{x}):=\min _{\mathbf{x}^{\prime} \in \mathscr{S}_{k}^{n}}\left\|\mathbf{x}-\mathbf{x}^{\prime}\right\|_{1}$.

Exercise 2.15 (Spectral norm without independence). Give an example of a random matrix $A \in \mathbb{R}^{n \times n}$ whose entries are bounded, but not independent, such that the spectral norm is $\Omega(n)$ with high probability.

Exercise 2.16 (Spectral norm: symmetric matrix). Let $A \in \mathbb{R}^{n \times n}$ be a symmetric random matrix. We assume that entries on and above the diagonal $A_{i, j}, i \leq j$, are centered, independent and sub-Gaussian with variance factor $\nu$. Each entry below the diagonal is equal to the corresponding entry above it. Prove an analogue of Theorem 2.4.28 for $A$. [Hint: Mimic the proof of Theorem 2.4.28.]

Exercise 2.17 (Chaining tail inequality). Prove Theorem 2.4.32.
Exercise 2.18 (Poisson convergence: method of moments). Let $A_{1}, \ldots, A_{n}$ be events and $A:=\cup_{i} A_{i}$. Define

$$
S^{(r)}:=\sum_{1 \leq i_{1}<\cdots<i_{r} \leq n} \mathbb{P}\left[A_{i_{1}} \cap \cdots \cap A_{i_{r}}\right],
$$

and

$$
X_{n}:=\sum_{i=1}^{n} A_{i} .
$$

Assume that there is $\mu>0$ such that, for all $r$,

$$
S^{(r)} \rightarrow \frac{\mu^{r}}{r!},
$$

as $n \rightarrow+\infty$. Use Exercise 2.2 and a Taylor expansion of $e^{-\mu}$ to show that

$$
\mathbb{P}\left[X_{n}=0\right] \rightarrow e^{-\mu} .
$$

In fact, $X_{n} \xrightarrow{\mathrm{~d}} \operatorname{Poi}(\mu)$ (no need to prove this). This is a special case of the method of moments.

Exercise 2.19 (Connectivity: critical window). Using Exercise 2.18 show that, when $p_{n}=\frac{\log n+s}{n}$, the probability that an Erdős-Rényi graph $G_{n} \sim \mathbb{G}_{n, p_{n}}$ contains no isolated vertex converges to $e^{-e^{-s}}$.

## Bibliographic Remarks

Section 2.1 For more on moment-generating functions, see [Bil12, Section 21].

Section 2.2 The examples in Section 2.2.1 are taken from [AS11, Sections 2.4, 3.2]. A fascinating account of the longest increasing subsequence problem is given in [Rom15], from which the material in Section 2.2.3 is taken. The contour lemma, Lemma 2.2.14, is attributed to Whitney [Whi32] and is usually proved "by picture" [Gri10a, Figure 3.1]. A formal proof of the lemma can be found in [Kes82, Appendix A]. For much more on percolation, see [Gri10b]. A gentler introduction is provided in [Ste].

Section 2.3 The presentation in Section 2.3.2 follows [AS11, Section 4.4] and [JLR11, Section 3.1]. The result for general subgraphs is due to Bollobás [Bol81]. A special case (including cliques) was proved by Erdős and Rényi [ER60]. For variants of the small subgraph containment problem involving copies that are induced, disjoint, isolated, etc., see for example [JLR11, Chapter 3]. For corresponding results for larger subgraphs, such as cycles or matchings, see for example [Bol01]. The connectivity threshold in Section 2.3.2 is also due to the same authors [ER59]. The presentation here follows [vdH17, Section 5.2]. For more on the method of moments, see for example [Dur10, Section 3.3.5] or [JLR11, Section 6.1]. Claim 2.3.11 is due to R. Lyons [Lyo90].

Section 2.4 The use of the moment-generating function to derive tail bounds for sums of independent random variables was pioneered by Cramér [Cra38], Bernstein [Ber46], and Chernoff [Che52]. For much more on concentration inequalities, see for example [BLM13]. The basics of large deviations theory are covered in [Dur10, Section 2.6]. See also [RAS15] and [DZ10]. Section 2.4.2 is based partly on [Ver18] and [Lug, Section 3.2]. Section 2.4.3 is based on [FR98, Section 5.3]. Very insightful, and much deeper, treatment of the material in Section 2.4.4 can be found in [Ver18, vH16]. The presentation in Section 2.4.5 is inspired by [Har, Lectures 6 and 8] and [Tao]. The Johnson-Lindenstrauss lemma was first proved by Johnson and Lindenstrauss using non-probabilistic arguments [JL84]. The idea of using random projections to simplify the proof was introduced by Frankl and Maehara [FM88] and the proof presented here based on Gaussian projections is due to Indyk and Motwani [IM98]. See [Ach03] for an overview of the various proofs known. For more on the random projection method, see [Vem04]. For algorithmic applications of the Johnson-Lindenstrauss lemma, see for example [Har, Lecture 7]. Compressed sensing emerged in the works of Donoho [Don06]
and Candès, Romberg and Tao [CRT06a, CRT06b]. The restricted isometry property was introduced by Candès and Tao [CT05]. Lemma 2.4.39 is due to Candés, Romberg and Tao [CRT06b]. The proof of Lemma 2.4.38 presented here is due to Baraniuk et al. [BDDW08]. A survey of compressed sensing can be found in [CW08]. A thorough mathematical introduction to compressed sensing can be found in [FR13]. The material in Section 2.4.2 can be found in [BLM13, Chapter 2]. Hoeffding's lemma and inequality are due to Hoeffding [Hoe63]. Section 2.4.6 borrows from [Ver18, vH16, SSBD14, Haz16]. The proof of Sauer's lemma follows [Ver18, Section 8.3.3]. For a proof of Claim 2.4.48 in general dimension $d$, see for example [SSBD14, Section 9.1.3].

## Chapter 3

## Martingales and potentials

In this chapter we turn to martingales, which play a central role in probability theory. We illustrate their use in a number of applications to the analysis of discrete stochastic processes. After some background on stopping times and a brief review of basic martingale properties and results in Section 3.1, we develop two major directions. In Section 3.2, we show how martingales can be used to derive a substantial generalization of our previous concentration inequalities-from the sums of independent random variables we focused on in Chapter 2 to nonlinear functions with Lipschitz properties. In particular, we give several applications of the method of bounded differences to random graphs. We also discuss bandit problems in machine learning. In the second thread in Section 3.3, we give an introduction to potential theory and electrical network theory for Markov chains. This toolkit in particular provides bounds on hitting times for random walks on networks, with important implications in the study of recurrence among other applications. We also introduce Wilson's remarkable method for generating uniform spanning trees.

### 3.1 Background

We begin with a quick review of stopping times and martingales. Along the way, we prove a few useful results. In particular, we derive some bounds on hitting times and cover times of Markov chains.

Throughout, $\left(\Omega, \mathcal{F},\left(\mathcal{F}_{t}\right)_{t \in \mathbb{Z}_{+}}, \mathbb{P}\right)$ is a filtered space. See Appendix B for a formal definition. Recall that, intuitively, the $\sigma$-algebra $\mathcal{F}_{t}$ in the filtration $\left(\mathcal{F}_{t}\right)_{t}$ represents " 'the information known at time $t$." All time indices are discrete (in $\mathbb{Z}_{+}$

Version: December 20, 2023
Modern Discrete Probability: An Essential Toolkit
Copyright © 2023 Sébastien Roch
unless stated otherwise). We will also use the notation $\overline{\mathbb{Z}}_{+}:=\{0,1, \ldots,+\infty\}$ to allow time $+\infty$.

### 3.1.1 Stopping times

Definitions Roughly speaking, a stopping time is a random time whose value is determined by a rule not depending on the future. Formally:

Definition 3.1.1 (Stopping time). A random variable $\tau: \Omega \rightarrow \overline{\mathbb{Z}}_{+}$is called $a$ stopping time if

$$
\{\tau \leq t\} \in \mathcal{F}_{t}, \forall t \in \overline{\mathbb{Z}}_{+}
$$

or, equivalently,

$$
\{\tau=t\} \in \mathcal{F}_{t}, \forall t \in \overline{\mathbb{Z}}_{+} .
$$

To see the equivalence above, note that $\{\tau=t\}=\{\tau \leq t\} \backslash\{\tau \leq t-1\}$, and $\{\tau \leq t\}=\cup_{i \leq t}\{\tau=i\}$.

Example 3.1.2 (Hitting time). Let $\left(A_{t}\right)_{t \in \mathbb{Z}_{+}}$, with values in $(E, \mathcal{E})$, be adapted and let $B \in \mathcal{E}$. Then

$$
\tau=\inf \left\{t \geq 0: A_{t} \in B\right\},
$$

is a stopping time known as a hitting time. In contrast, the last visit to a set is typically not a stopping time.

Let $\tau$ be a stopping time. Denote by $\mathcal{F}_{\tau}$ the set of all events $F$ such that, $\forall t \in \overline{\mathbb{Z}}_{+}$, $F \cap\{\tau=t\} \in \mathcal{F}_{t}$. Intuitively, the $\sigma$-algebra $\mathcal{F}_{\tau}$ captures the information up to time $\tau$. The following lemmas help clarify the definition of $\mathcal{F}_{\tau}$.

Lemma 3.1.3. $\mathcal{F}_{\tau}=\mathcal{F}_{s}$ if $\tau:=s, \mathcal{F}_{\tau}=\mathcal{F}_{\infty}=\sigma\left(\cup_{t} \mathcal{F}_{t}\right)$ if $\tau:=+\infty$ and $\mathcal{F}_{\tau} \subseteq \mathcal{F}_{\infty}$ for any stopping time $\tau$.

Proof. In the first case, note that $F \cap\{\tau=t\}$ is empty if $t \neq s$ and is $F$ if $t=s$. So if $F \in \mathcal{F}_{\tau}$ then $F=F \cap\{\tau=s\} \in \mathcal{F}_{s}$ by definition of $\mathcal{F}_{\tau}$, and if $F \in \mathcal{F}_{s}$ then $F=F \cap\{\tau=t\} \in \mathcal{F}_{t}$ for all $t$ by definition of $\tau$. So we have proved both inclusions. This works also for $t=+\infty$. For the third claim note that, for any $F \in \mathcal{F}_{\tau}$,

$$
F=\cup_{t \in \overline{\mathbb{Z}}_{+}} F \cap\{\tau=t\} \in \mathcal{F}_{\infty},
$$

again by definition of $\mathcal{F}_{\tau}$.
Lemma 3.1.4. If $\left(X_{t}\right)$ is adapted and $\tau$ is a stopping time then $X_{\tau} \in \mathcal{F}_{\tau}$ (where we assume that $X_{\infty} \in \mathcal{F}_{\infty}$, e.g., by setting $X_{\infty}:=\liminf X_{n}$ ).

Proof. For $B \in \mathcal{E}$,

$$
\left\{X_{\tau} \in B\right\} \cap\{\tau=t\}=\left\{X_{t} \in B\right\} \cap\{\tau=t\} \in \mathcal{F}_{t}
$$ by definition of $\tau$. That shows $X_{\tau}$ is measurable with respect to $\mathcal{F}_{\tau}$ as claimed.

Lemma 3.1.5. If $\sigma, \tau$ are stopping times then $\mathcal{F}_{\sigma \wedge \tau} \subseteq \mathcal{F}_{\tau}$.
Proof. Let $F \in \mathcal{F}_{\sigma \wedge \tau}$. Note that

$$
F \cap\{\tau=t\}=\cup_{s \leq t}[(F \cap\{\sigma \wedge \tau=s\}) \cap\{\tau=t\}] \in \mathcal{F}_{t} .
$$

Indeed, the expression in parentheses is in $\mathcal{F}_{s} \subseteq \mathcal{F}_{t}$ by definition of $\mathcal{F}_{\sigma \wedge \tau}$ and $\{\tau=t\} \in \mathcal{F}_{t}$.

Let $\left(X_{t}\right)$ be a Markov chain on a countable space $V$. The following two examples of stopping times will play an important role.

Definition 3.1.6 (First visit and return). The first visit time and first return time to $x \in V$ are

$$
\tau_{x}:=\inf \left\{t \geq 0: X_{t}=x\right\} \quad \text { and } \quad \tau_{x}^{+}:=\inf \left\{t \geq 1: X_{t}=x\right\} .
$$

Similarly, $\tau_{B}$ and $\tau_{B}^{+}$are the first visit time and first return time to $B \subseteq V$.
Definition 3.1.7 (Cover time). Assume $V$ is finite. The cover time of $\left(X_{t}\right)$ is the first time that all states have been visited, that is,

$$
\tau_{\text {cov }}:=\inf \left\{t \geq 0:\left\{X_{0}, \ldots, X_{t}\right\}=V\right\}
$$

Strong Markov property Let $\left(X_{t}\right)$ be a Markov chain with transition matrix $P$ and initial distribution $\mu$. Let $\mathcal{F}_{t}=\sigma\left(X_{0}, \ldots, X_{t}\right)$. Recall that the Markov property (Theorem 1.1.18) says that, given the present, the future is independent of the past. The Markov property naturally extends to stopping times. Let $\tau$ be a stopping time with $\mathbb{P}[\tau<+\infty]>0$. In its simplest form we have:

$$
\mathbb{P}\left[X_{\tau+1}=y \mid \mathcal{F}_{\tau}\right]=\mathbb{P}_{X_{\tau}}\left[X_{\tau+1}=y\right]=P\left(X_{\tau}, y\right) .
$$

In words, the chain "starts fresh" at a stopping time with the state at that time as starting point. More generally:

Theorem 3.1.8 (Strong Markov property). Let $f_{t}: V^{\infty} \rightarrow \mathbb{R}$ be a sequence of measurable functions, uniformly bounded in $t$ and let $F_{t}(x):=\mathbb{E}_{x}\left[f_{t}\left(\left(X_{s}\right)_{s \geq 0}\right)\right]$. On $\{\tau<+\infty\}$,

$$
\mathbb{E}\left[f_{\tau}\left(\left(X_{\tau+t}\right)_{t \geq 0}\right) \mid \mathcal{F}_{\tau}\right]=F_{\tau}\left(X_{\tau}\right) .
$$

Throughout, when we say that two random variables $Y, Z$ are equal on an event $B$, we mean formally that $Y \mathbf{1}_{B}=Z \mathbf{1}_{B}$ almost surely.

Proof of Theorem 3.1.8. We use that

$$
\mathbb{E}\left[f_{\tau}\left(\left(X_{\tau+t}\right)_{t \geq 0}\right) \mid \mathcal{F}_{\tau}\right] \mathbf{1}_{\tau<+\infty}=\mathbb{E}\left[f_{\tau}\left(\left(X_{\tau+t}\right)_{t \geq 0}\right) \mathbf{1}_{\tau<+\infty} \mid \mathcal{F}_{\tau}\right] .
$$

Let $A \in \mathcal{F}_{\tau}$. Summing over the possible values of $\tau$, using the tower property (Lemma B.6.16) and then the Markov property

$$
\begin{aligned}
& \mathbb{E}\left[f_{\tau}\left(\left(X_{\tau+t}\right)_{t \geq 0}\right) \mathbf{1}_{\tau<+\infty} ; A\right] \\
& =\mathbb{E}\left[f_{\tau}\left(\left(X_{\tau+t}\right)_{t \geq 0}\right) ; A \cap\{\tau<+\infty\}\right] \\
& =\sum_{s \geq 0} \mathbb{E}\left[f_{s}\left(\left(X_{s+t}\right)_{t \geq 0}\right) ; A \cap\{\tau=s\}\right] \\
& =\sum_{s \geq 0} \mathbb{E}\left[\mathbb{E}\left[f_{s}\left(\left(X_{s+t}\right)_{t \geq 0}\right) ; A \cap\{\tau=s\} \mid \mathcal{F}_{s}\right]\right] \\
& =\sum_{s \geq 0} \mathbb{E}\left[\mathbf{1}_{A \cap\{\tau=s\}} \mathbb{E}\left[f_{s}\left(\left(X_{s+t}\right)_{t \geq 0}\right) \mid \mathcal{F}_{s}\right]\right] \\
& =\sum_{s \geq 0} \mathbb{E}\left[\mathbf{1}_{A \cap\{\tau=s\}} F_{s}\left(X_{s}\right)\right] \\
& =\sum_{s \geq 0} \mathbb{E}\left[F_{s}\left(X_{s}\right) ; A \cap\{\tau=s\}\right] \\
& =\mathbb{E}\left[F_{\tau}\left(X_{\tau}\right) ; A \cap\{\tau<+\infty\}\right] \\
& =\mathbb{E}\left[F_{\tau}\left(X_{\tau}\right) \mathbf{1}_{\tau<+\infty} ; A\right],
\end{aligned}
$$

where, on the fifth line, we used that $A \cap\{\tau=s\} \in \mathcal{F}_{s}$ by definition of $\mathcal{F}_{\tau}$ and taking out what is known (Lemma B.6.13). The definition of the conditional expectation (Theorem B.6.1) concludes the proof.

The following typical application of the strong Markov property (Theorem 3.1.8) is useful.

Theorem 3.1.9 (Reflection principle). Let $X_{1}, X_{2}, \ldots$ be i.i.d. with a distribution symmetric about 0 and let $S_{t}=\sum_{i \leq t} X_{i}$. Then, for $b>0$,

$$
\mathbb{P}\left[\sup _{i \leq t} S_{i} \geq b\right] \leq 2 \mathbb{P}\left[S_{t} \geq b\right]
$$

Proof. Let $\tau:=\inf \left\{i \leq t: S_{i} \geq b\right\}$. By the strong Markov property, on $\{\tau<$ $t\}, S_{t}-S_{\tau}$ is independent of $\mathcal{F}_{\tau}$ and is symmetric about 0 . In particular, it has
probability at least $1 / 2$ of being greater or equal to 0 by the first moment principle (Theorem 2.2.1), an event which implies that $S_{t}$ is greater than or equal to $b$. Hence

$$
\mathbb{P}\left[S_{t} \geq b\right] \geq \mathbb{P}[\tau=t]+\frac{1}{2} \mathbb{P}[\tau<t] \geq \frac{1}{2} \mathbb{P}[\tau \leq t]
$$

(Exercise 3.1 asks for a more formal proof.)
In the case of simple random walk on $\mathbb{Z}$, we get a stronger statement.
Theorem 3.1.10 (Reflection principle: simple random walk). Let $\left(S_{t}\right)$ be simple random walk on $\mathbb{Z}$ started at 0 . Then, $\forall a, b, t>0$,

$$
\mathbb{P}\left[S_{t}=b+a\right]=\mathbb{P}\left[S_{t}=b-a, \sup _{i \leq t} S_{i} \geq b\right]
$$

and

$$
\mathbb{P}\left[\sup _{i \leq t} S_{i} \geq b\right]=\mathbb{P}\left[S_{t}=b\right]+2 \mathbb{P}\left[S_{t}>b\right]
$$

Proof. For the first claim, reflect the sub-path after the first visit to $b$ across the line $y=b$. Summing over $a>0$ and rearranging gives the second claim.

We record another related result that will be useful later.
Theorem 3.1.11 (Ballot theorem). In an election with $n$ voters, candidate $A$ gets $\alpha$ votes and candidate $B$ gets $\beta<\alpha$ votes. The probability that $A$ leads $B$ throughout the counting is $\frac{\alpha-\beta}{n}$.

Recurrence Let $\left(X_{t}\right)$ be a Markov chain on a countable state space $V$. The time of the $k$-th return to $y$ is (letting $\tau_{y}^{0}:=0$ )

$$
\tau_{y}^{k}:=\inf \left\{t>\tau_{y}^{k-1}: X_{t}=y\right\} .
$$

In particular, $\tau_{y}^{1}=\tau_{y}^{+}$. Define $\rho_{x y}:=\mathbb{P}_{x}\left[\tau_{y}^{+}<+\infty\right]$. Then by the strong Markov property (and induction)

$$
\begin{equation*}
\mathbb{P}_{x}\left[\tau_{y}^{k}<+\infty\right]=\rho_{x y} \rho_{y y}^{k-1} . \tag{3.1.1}
\end{equation*}
$$

(Exercise 3.2 asks for a more formal proof.) Letting

$$
N_{y}:=\sum_{t>0} \mathbf{1}_{\left\{X_{t}=y\right\}}=\sum_{k \geq 1} \mathbf{1}_{\left\{\tau_{y}^{k}<+\infty\right\}},
$$

be the number of visits to $y$ after time 0 , by linearity

$$
\begin{equation*}
\mathbb{E}_{x}\left[N_{y}\right]=\frac{\rho_{x y}}{1-\rho_{y y}} . \tag{3.1.2}
\end{equation*}
$$

When $\rho_{y y}<1$, we have $\mathbb{E}_{y}\left[N_{y}\right]<+\infty$ by (3.1.2), and in particular $\tau_{y}^{k}=+\infty$ for some $k$. Or $\rho_{y y}=1$ and, starting at $x=y$, we have $\tau_{y}^{k}<+\infty$ almost surely for all $k$ by (3.1.1). That leads us to the following dichotomy.

Definition 3.1.12 (Recurrence). A state $x$ is recurrent if $\rho_{x x}=1$. Otherwise it is transient. We refer to the recurrence or transience of a state as its type. Let $x$ be recurrent. If in addition $\mathbb{E}_{x}\left[\tau_{x}^{+}\right]<+\infty$, we say that $x$ is positive recurrent; otherwise we say that it is null recurrent. A chain is recurrent (or transient, or positive recurrent, or null recurrent) if all its states are.

Recurrence is "contagious" in the following sense.
Lemma 3.1.13. If $x$ is recurrent and $\rho_{x y}>0$ then $y$ is recurrent and $\rho_{y x}=\rho_{x y}=$ 1.

A subset $C \subseteq V$ is closed if $x \in C$ and $\rho_{x y}>0$ implies $y \in C$. A subset $D \subseteq V$ is irreducible if $x, y \in D$ implies $\rho_{x y}>0$. This definition is consistent with (and generalizes to sets) the one we gave in Section 1.1.2. Recall that we have the following decomposition theorem.

Theorem 3.1.14 (Decomposition theorem). Let $R:=\left\{x: \rho_{x x}=1\right\}$ be the recurrent states of the chain. Then $R$ can be written as a disjoint union $\cup_{j} R_{j}$ where each $R_{j}$ is closed and irreducible.
Example 3.1.15 (Simple random walk on $\mathbb{Z}$ ). Consider simple random walk $\left(S_{t}\right)$ on $\mathbb{Z}$ started at 0 . The chain is clearly irreducible so it suffices to check the type of state 0 by Lemma 3.1.13. First note the periodicity of this chain. So we look at $S_{2 t}$. Then by Stirling's formula (see Appendix A)

$$
\mathbb{P}\left[S_{2 t}=0\right]=\binom{2 t}{t} 2^{-2 t} \sim 2^{-2 t} \frac{(2 t)^{2 t}}{\left(t^{t}\right)^{2}} \frac{\sqrt{2 t}}{\sqrt{2 \pi t}} \sim \frac{1}{\sqrt{\pi t}} .
$$

Thus

$$
\mathbb{E}\left[N_{0}\right]=\sum_{t>0} \mathbb{P}\left[S_{t}=0\right]=+\infty
$$

and the chain is recurrent.
Return times are closely related to stationary measures. We recall the following standard results without proof. We gave an alternative proof of the existence of a unique stationary distribution in the finite, irreducible case in Theorem 1.1.24.

Theorem 3.1.16. Let $x$ be a recurrent state. Then the following defines a stationary measure

$$
\mu_{x}(y):=\mathbb{E}_{x}\left[\sum_{0 \leq t<\tau_{x}^{+}} \mathbf{1}_{\left\{X_{t}=y\right\}}\right] .
$$

Theorem 3.1.17. If $\left(X_{t}\right)$ is irreducible and recurrent, then the stationary measure is unique up to a constant multiple.

Theorem 3.1.18. If there is a stationary distribution $\pi$ then all states $y$ that have $\pi(y)>0$ are recurrent.

Theorem 3.1.19. If $\left(X_{t}\right)$ is irreducible and has a stationary distribution $\pi$, then

$$
\pi(x)=\frac{1}{\mathbb{E}_{x} \tau_{x}^{+}}
$$

Theorem 3.1.20. If $\left(X_{t}\right)$ is irreducible, then the following are equivalent.
(i) There is a stationary distribution.
(ii) All states are positive recurrent.
(iii) There is a positive recurrent state.

We have seen previously that, in the irreducible, positive recurrent, aperiodic case, there is convergence to stationarity (see Theorem 1.1.33). In the transient and null recurrent cases, there is no stationary distribution to converge to by Theorem 3.1.20. Instead, we have the following.

Theorem 3.1.21 (Convergence of $P^{t}$ : transient and null recurrent cases). If $P$ is an irreducible chain which is either transient or null recurrent, we have for all $x, y$ that

$$
\lim _{t} P^{t}(x, y)=0
$$

Proof. We only prove the transient case. In that case, we showed in (3.1.2) that

$$
\sum_{t} P^{t}(x, y)=\mathbb{E}_{x}\left[\sum_{t} \mathbf{1}_{\left\{X_{t}=y\right\}}\right]=\mathbb{E}_{x}\left[N_{y}\right]<+\infty
$$

Hence $P^{t}(x, y) \rightarrow 0$.

A useful identity A slight generalization of the "cycle trick" used in the proof of Theorem 3.1.16 gives a useful identity.

Definition 3.1.22 (Green function). Let $\sigma$ be a stopping time for a Markov chain $\left(X_{t}\right)$. The Green function of the chain stopped at $\sigma$ is given by

$$
\begin{equation*}
\mathscr{G}_{\sigma}(x, y)=\mathbb{E}_{x}\left[\sum_{0 \leq t<\sigma} \mathbf{1}_{\left\{X_{t}=y\right\}}\right], \quad x, y \in V, \tag{3.1.3}
\end{equation*}
$$

that is, it is the expected number of visits to $y$ before $\sigma$ when started at $x$.
Lemma 3.1.23 (Occupation measure identity). Consider an irreducible, positive recurrent Markov chain $\left(X_{t}\right)_{t \geq 0}$ with transition matrix $P$ and stationary distribution $\pi$. Let $x$ be a state and $\sigma$ be a stopping time such that $\mathbb{E}_{x}[\sigma]<+\infty$ and $\mathbb{P}_{x}\left[X_{\sigma}=x\right]=1$. For any $y$,

$$
\mathscr{G}_{\sigma}(x, y)=\pi_{y} \mathbb{E}_{x}[\sigma] .
$$

Proof. By the uniqueness of the stationary measure up to constant multiple (Theorem 3.1.17), it suffices to show that $\mathscr{G}_{\sigma}(x, y)$ satisfies the system for a stationary measure as a function of $y$

$$
\begin{equation*}
\sum_{y} \mathscr{G}_{\sigma}(x, y) P(y, z)=\mathscr{G}_{\sigma}(x, z), \quad \forall z, \tag{3.1.4}
\end{equation*}
$$

and use the fact that

$$
\sum_{y} \mathscr{G}_{\sigma}(x, y)=\sum_{y} \mathbb{E}_{x}\left[\sum_{0 \leq t<\sigma} \mathbf{1}_{\left\{X_{t}=y\right\}}\right]=\mathbb{E}_{x}[\sigma] .
$$

To check (3.1.4), because $X_{\sigma}=X_{0}$ almost surely, observe that

$$
\begin{aligned}
\mathscr{G}_{\sigma}(x, z) & =\mathbb{E}_{x}\left[\sum_{0 \leq t<\sigma} \mathbf{1}_{X_{t}=z}\right] \\
& =\mathbb{E}_{x}\left[\sum_{0 \leq t<\sigma} \mathbf{1}_{X_{t+1}=z}\right] \\
& =\sum_{t \geq 0} \mathbb{P}_{x}\left[X_{t+1}=z, \sigma>t\right] .
\end{aligned}
$$

Since $\{\sigma>t\} \in \mathcal{F}_{t}$, applying the Markov property we get

$$
\begin{aligned}
\mathscr{G}_{\sigma}(x, z) & =\sum_{t \geq 0} \sum_{y} \mathbb{P}_{x}\left[X_{t}=y, X_{t+1}=z, \sigma>t\right] \\
& =\sum_{t \geq 0} \sum_{y} \mathbb{P}_{x}\left[X_{t+1}=z \mid X_{t}=y, \sigma>t\right] \mathbb{P}_{x}\left[X_{t}=y, \sigma>t\right] \\
& =\sum_{t \geq 0} \sum_{y} P(y, z) \mathbb{P}_{x}\left[X_{t}=y, \sigma>t\right] \\
& =\sum_{y} \mathscr{G}_{\sigma}(x, y) P(y, z),
\end{aligned}
$$

which establishes (3.1.4) and proves the claim.
Here is a typical application of this lemma.
Corollary 3.1.24. In the setting of Lemma 3.1.23, for all $x \neq y$,

$$
\mathbb{P}_{x}\left[\tau_{y}<\tau_{x}^{+}\right]=\frac{1}{\pi_{x}\left(\mathbb{E}_{x}\left[\tau_{y}\right]+\mathbb{E}_{y}\left[\tau_{x}\right]\right)}
$$

Proof. Let $\sigma$ be the time of the first visit to $x$ after the first visit to $y$. Then $\mathbb{E}_{x}[\sigma]=$ $\mathbb{E}_{x}\left[\tau_{y}\right]+\mathbb{E}_{y}\left[\tau_{x}\right]<+\infty$, where we used that the chain is irreducible and positive recurrent. By the strong Markov property, the number of visits to $x$ before the first visit to $y$ is geometric with success probability $\mathbb{P}_{x}\left[\tau_{y}<\tau_{x}^{+}\right]$(where, here, a visit to $x$ is a "failed trial"). Moreover the number of visits to $x$ after the first visit to $y$ but before $\sigma$ is 0 by definition. Hence $\mathscr{G}_{\sigma}(x, x)$ is the mean of the geometric distribution, namely $1 / \mathbb{P}_{x}\left[\tau_{y}<\tau_{x}^{+}\right]$. Applying the occupation measure identity gives the result.

### 3.1.2 $\triangleright$ Markov chains: exponential tail of hitting times and some cover time bounds

Tail of a hitting time On a finite state space, the tail of any hitting time converges to 0 exponentially fast.
Lemma 3.1.25. Let $\left(X_{t}\right)$ be a finite, irreducible Markov chain with state space $V$. For any subset of states $A \subseteq V$ and initial distribution $\mu$ :
(i) It holds that $\mathbb{E}_{\mu}\left[\tau_{A}\right]<+\infty$ (and, in particular, $\tau_{A}<+\infty$ a.s.).
(ii) Letting $\overline{\mathrm{t}}_{A}:=\max _{x} \mathbb{E}_{x}\left[\tau_{A}\right]$, we have the tail bound

$$
\mathbb{P}_{\mu}\left[\tau_{A}>t\right] \leq \exp \left(-\left\lfloor\frac{t}{\left\lceil e \overline{\mathrm{t}}_{A}\right\rceil}\right\rfloor\right) .
$$

Proof. For any positive integer $m$, for some distribution $\theta$ over the state space $V$, by the strong Markov property (Theorem 3.1.8)

$$
\mathbb{P}_{\mu}\left[\tau_{A}>m s \mid \tau_{A}>(m-1) s\right]=\mathbb{P}_{\theta}\left[\tau_{A}>s\right] \leq \max _{x} \mathbb{P}_{x}\left[\tau_{A}>s\right]=: \alpha_{s}
$$

Choose a positive integer $s$ large enough that, from any $x$, there is a path to $A$ of length at most $s$ of positive probability. Such an $s$ exists by irreducibility. In particular $\alpha_{s}<1$.

By the multiplication rule and the monotonicity of the events $\left\{\tau_{A}>r s\right\}$ over $r$, we have

$$
\mathbb{P}_{\mu}\left[\tau_{A}>m s\right]=\mathbb{P}_{\mu}\left[\tau_{A}>s\right] \prod_{r=2}^{m} \mathbb{P}_{\mu}\left[\tau_{A}>r s \mid \tau_{A}>(r-1) s\right] .
$$

Therefore, $\mathbb{P}_{\mu}\left[\tau_{A}>m s\right] \leq \alpha_{s}^{m}$, which in turn implies

$$
\begin{equation*}
\mathbb{P}_{\mu}\left[\tau_{A}>t\right] \leq \alpha_{s}^{\left\lfloor\frac{t}{s}\right\rfloor} \tag{3.1.5}
\end{equation*}
$$

The result for the expectation follows from

$$
\mathbb{E}_{\mu}\left[\tau_{A}\right]=\sum_{t \geq 0} \mathbb{P}_{\mu}\left[\tau_{A}>t\right] \leq \sum_{t} \alpha_{s}^{\left\lfloor\frac{t}{s}\right\rfloor}<+\infty,
$$

since $\alpha_{s}<1$.
Now that we have established that $\overline{\mathrm{t}}_{A}<+\infty$, by Markov's inequality (Theorem 2.1.1),

$$
\alpha_{s}=\max _{x} \mathbb{P}_{x}\left[\tau_{A}>s\right] \leq \frac{\overline{\mathrm{t}}_{A}}{s} .
$$

for all non-negative integers s. Plugging back into (3.1.5) gives $\mathbb{P}_{\mu}\left[\tau_{A}>t\right] \leq$ $\left(\frac{\overline{\bar{A}}_{A}}{s}\right)^{\left\lfloor\frac{t}{s}\right\rfloor}$. By differentiating with respect to $s$, it can be checked that a good choice for $s$ is $\left\lceil e \overline{\mathrm{t}}_{A}\right\rceil$. Simplifying gives the second claim.

Application to cover times We give an application of the previous bound to cover times. Let $\left(X_{t}\right)$ be a finite, irreducible Markov chain on $V$ with $n:=|V|>$ 1. Recall that the cover time is $\tau_{\text {cov }}:=\max _{y} \tau_{y}$. We bound the mean cover time in terms of

$$
\overline{\mathrm{t}}_{\text {hit }}:=\max _{x \neq y} \mathbb{E}_{x} \tau_{y}
$$

Claim 3.1.26.

$$
\max _{x} \mathbb{E}_{x}\left[\tau_{\text {cov }}\right] \leq(3+\log n)\left\lceil e \overline{\mathrm{t}}_{\text {hit }}\right\rceil .
$$

Proof. By a union bound over all states to be visited and Lemma 3.1.25,

$$
\max _{x} \mathbb{P}_{x}\left[\tau_{\text {cov }}>t\right] \leq \min \left\{1, n \exp \left(-\left\lfloor\frac{t}{\left\lceil e \bar{t}_{\text {hit }}\right\rceil}\right\rfloor\right)\right\} .
$$

Summing over $t \in \mathbb{Z}_{+}$and appealing to the sum of a geometric series,

$$
\max _{x} \mathbb{E}_{x}\left[\tau_{\mathrm{cov}}\right] \leq(\log n+1)\left\lceil e \overline{\mathrm{t}}_{\mathrm{hit}}\right\rceil+\frac{1}{1-e^{-1}}\left\lceil e \overline{\mathrm{t}}_{\mathrm{hit}}\right\rceil
$$

where the first term on the right-hand side comes from the fact that until $t \geq$ $(\log n+1)\left\lceil e \overline{\mathrm{t}}_{\mathrm{hit}}\right\rceil$ the upper bound above is 1 . The factor $\left\lceil e \overline{\mathrm{t}}_{\mathrm{hit}}\right\rceil$ in the second term on the right-hand side comes from the fact that we must break up the series into blocks of size $\left\lceil e \overline{\mathrm{t}}_{\text {hit }}\right\rceil$. Simplifying gives the claim.

The previous proof should be reminiscent of that of Theorem 2.4.21.
A clever argument gives a better constant factor as well as a lower bound.
Theorem 3.1.27 (Matthews' cover time bounds). Let

$$
\underline{\mathrm{t}}_{\mathrm{hit}}^{A}:=\min _{x, y \in A, x \neq y} \mathbb{E}_{x} \tau_{y},
$$

and $h_{n}:=\sum_{m=1}^{n} \frac{1}{m}$. Then

$$
\begin{equation*}
\max _{x} \mathbb{E}_{x}\left[\tau_{\mathrm{cov}}\right] \leq h_{n} \overline{\mathrm{t}}_{\mathrm{hit}}, \tag{3.1.6}
\end{equation*}
$$

and

$$
\begin{equation*}
\min _{x} \mathbb{E}_{x}\left[\tau_{\text {cov }}\right] \geq \max _{A \subseteq V} h_{|A|-1} \underline{\mathrm{t}}_{\mathrm{hit}}^{A} . \tag{3.1.7}
\end{equation*}
$$

Clearly, $\max _{x \neq y} \mathrm{t}_{\text {hit }}^{\{x, y\}}$ is a lower bound on the worst expected cover time. Lower bound (3.1.7) says that a tighter bound is obtained by finding a larger subset of states $A$ that are "far away" from each other.

We sketch the proof of the lower bound for $A=V$, which we assume is $[n]$ without loss of generality. The other cases are similar. Let $\left(J_{1}, \ldots, J_{n}\right)$ be a uniform random ordering of $V$, let $C_{m}:=\max _{i \leq m} \tau_{J_{i}}$, and let $L_{m}$ be the last state visited among $J_{1}, \ldots, J_{m}$. Then for $m \geq 2$

$$
\mathbb{E}_{x}\left[C_{m}-C_{m-1} \mid J_{1}, \ldots, J_{m},\left\{X_{t}, t \leq C_{m-1}\right\}\right] \geq \underline{\mathrm{t}}_{\mathrm{hit}}^{V} \mathbf{1}_{\left\{L_{m}=J_{m}\right\}} .
$$

By symmetry, $\mathbb{P}\left[L_{m}=J_{m}\right]=\frac{1}{m}$. To see this, first pick the set of vertices corresponding to $\left\{J_{1}, \ldots, J_{m}\right\}$, wait for all of those vertices to be visited, then pick the ordering. Moreover observe that $\mathbb{E}_{x} C_{1} \geq\left(1-\frac{1}{n}\right) \underline{\underline{h}}_{\text {hit }}^{V}$ where the factor of $\left(1-\frac{1}{n}\right)$ accounts for the probability that $J_{1} \neq x$. Taking expectations above and summing over $m$ gives the result.

Exercise 3.3 asks for a proof that the bounds above cannot in general be improved up to smaller order terms.

### 3.1.3 Martingales

Definition Martingales are an important class of stochastic processes that correspond intuitively to the "probabilistic version of a monotone sequence." They hide behind many processes and have properties that make them powerful tools in the analysis of processes where they have been identified. Formally:

Definition 3.1.28 (Martingale). An adapted process $\left(M_{t}\right)_{t \geq 0}$ with $\mathbb{E}\left|M_{t}\right|<+\infty$ for all $t$ is a martingale if

$$
\mathbb{E}\left[M_{t+1} \mid \mathcal{F}_{t}\right]=M_{t}, \quad \forall t \geq 0 .
$$

If equality is replaced with $\leq$ or $\geq$, we get a supermartingale or a submartingale respectively. We say that a martingale in bounded in $L^{p}$ if $\sup _{t} \mathbb{E}\left[\left|X_{t}\right|^{p}\right]<+\infty$.

Recall that adapted (Definition B.7.5) simply means that $M_{t} \in \mathcal{F}_{t}$, that is, roughly speaking $M_{t}$ is "known at time $t$." Note that for a martingale, by the tower property (Lemma B.6.16), we have $\mathbb{E}\left[M_{t} \mid \mathcal{F}_{s}\right]=M_{s}$ for all $t>s$, and similarly (with inequalities) for supermartingales and submartingales.

We start with a straightforward example.
Example 3.1.29 (Sums of i.i.d. random variables with mean 0). Let $X_{0}, X_{1}, \ldots$ be i.i.d. integrable, centered random variables, $\mathcal{F}_{t}=\sigma\left(X_{0}, \ldots, X_{t}\right), S_{0}=0$, and $S_{t}=\sum_{i=1}^{t} X_{i}$. Note that $\mathbb{E}\left|S_{t}\right|<\infty$ by the triangle inequality. By taking out what is known and the role of independence lemma (Lemma B.6.14) we obtain

$$
\mathbb{E}\left[S_{t} \mid \mathcal{F}_{t-1}\right]=\mathbb{E}\left[S_{t-1}+X_{t} \mid \mathcal{F}_{t-1}\right]=S_{t-1}+\mathbb{E}\left[X_{t}\right]=S_{t-1},
$$

which proves that $\left(S_{t}\right)$ is a martingale.
Martingales however are richer than random walks with centered steps. For instance mixtures of such random walks are also martingales.

Example 3.1.30 (Mixtures of random walks). Consider again the setting of Example 3.1.29. This time assume that $X_{0}$ is uniformly distributed in $\{1,2\}$ and define

$$
R_{t}=X_{0} S_{t}, \quad t \geq 0 .
$$

Then, because $\left(S_{t}\right)$ is a martingale,

$$
\mathbb{E}\left[R_{t} \mid \mathcal{F}_{t-1}\right]=X_{0} \mathbb{E}\left[S_{t} \mid \mathcal{F}_{t-1}\right]=X_{0} S_{t-1}=R_{t-1},
$$

so $\left(R_{t}\right)$ is also a martingale.

Further examples Martingales can also be a little more hidden. Here are two examples.
Example 3.1.31 (Variance of a sum of i.i.d. random variables). Consider again the setting of Example 3.1.29 with $\sigma^{2}:=\operatorname{Var}\left[X_{1}\right]<\infty$. Define

$$
M_{t}=S_{t}^{2}-t \sigma^{2}
$$

Note that by the triangle inequality and the fact that $S_{t}$ has mean zero and is a sum of independent random variables

$$
\mathbb{E}\left|M_{t}\right| \leq \sum_{i=1}^{t} \operatorname{Var}\left[X_{i}\right]+t \sigma^{2} \leq 2 t \sigma^{2}<+\infty .
$$

Moreover, arguing similarly to the previous example, and using the fact that both $X_{t}$ and $S_{t-1}$ are square integrable

$$
\begin{aligned}
\mathbb{E}\left[M_{t} \mid \mathcal{F}_{t-1}\right] & =\mathbb{E}\left[\left(X_{t}+S_{t-1}\right)^{2}-t \sigma^{2} \mid \mathcal{F}_{t-1}\right] \\
& =\mathbb{E}\left[X_{t}^{2}+2 X_{t} S_{t-1}+S_{t-1}^{2}-t \sigma^{2} \mid \mathcal{F}_{t-1}\right] \\
& =\sigma^{2}+0+S_{t-1}^{2}-t \sigma^{2} \\
& =M_{t-1},
\end{aligned}
$$

which proves that $\left(M_{t}\right)$ is a martingale.
Example 3.1.32 (Eigenvectors of a transition matrix). Let $\left(X_{t}\right)_{t \geq 0}$ be a finite Markov chain with state space $V$ and transition matrix $P$, and let $\left(\mathcal{F}_{t}\right)_{t \geq 0}$ be the corresponding filtration. Suppose $f: V \rightarrow \mathbb{R}$ is such that

$$
\sum_{j} P(i, j) f(j)=\lambda f(i), \forall i \in S
$$

In other words, $f$ is a (right) eigenvector of $P$ with eigenvalue $\lambda$. Define

$$
M_{t}=\lambda^{-t} f\left(X_{t}\right) .
$$

Note that by the finiteness of the state space

$$
\mathbb{E}\left|M_{t}\right|<+\infty,
$$

and that further by the Markov property

$$
\begin{aligned}
\mathbb{E}\left[M_{t} \mid \mathcal{F}_{t-1}\right] & =\lambda^{-t} \mathbb{E}\left[f\left(X_{t}\right) \mid \mathcal{F}_{t-1}\right] \\
& =\lambda^{-t} \sum_{j} P\left(X_{t-1}, j\right) f(j) \\
& =\lambda^{-t} \cdot \lambda f\left(X_{t-1}\right) \\
& =M_{t-1} .
\end{aligned}
$$

That is, $\left(M_{t}\right)$ is a martingale.

Or we can create martingales out of thin air. We give two important examples that will appear later.

Example 3.1.33 (Doob martingale: accumulating data). Let $X$ with $\mathbb{E}|X|<+\infty$. Define $M_{t}=\mathbb{E}\left[X \mid \mathcal{F}_{t}\right]$. Note that $\mathbb{E}\left|M_{t}\right| \leq \mathbb{E}|X|<+\infty$ by Jensens' inequality, and

$$
\mathbb{E}\left[M_{t} \mid \mathcal{F}_{t-1}\right]=\mathbb{E}\left[X \mid \mathcal{F}_{t-1}\right]=M_{t-1},
$$

by the tower property. This is known as a Doob martingale. Intuitively this process tracks our expectation of the unobserved $X$ as "more information becomes available." See the co-called "exposure martingales" in Section 3.2.3 for a concrete illustration of this idea.

Example 3.1.34 (Martingale transform). Let $\left(X_{t}\right)_{t \geq 1}$ be an integrable, adapted process and let $\left(C_{t}\right)_{t \geq 1}$ be a bounded, predictable process. Recall that predictable (Definition B.7.6) means $C_{t} \in \mathcal{F}_{t-1}$ for all $t$, that is, roughly speaking $C_{t}$ is "known at time $t-1$." Define

$$
N_{t}=\sum_{i \leq t}\left(X_{i}-\mathbb{E}\left[X_{i} \mid \mathcal{F}_{i-1}\right]\right) C_{i} .
$$

Then

$$
\mathbb{E}\left|N_{t}\right| \leq \sum_{i \leq t} 2 \mathbb{E}\left|X_{t}\right| K<+\infty,
$$

where we used that $\left|C_{t}\right|<K$ for all $t \geq 1$, and

$$
\begin{aligned}
\mathbb{E}\left[N_{t}-N_{t-1} \mid \mathcal{F}_{t-1}\right] & =\mathbb{E}\left[\left(X_{t}-\mathbb{E}\left[X_{t} \mid \mathcal{F}_{t-1}\right]\right) C_{t} \mid \mathcal{F}_{t-1}\right] \\
& =C_{t}\left(\mathbb{E}\left[X_{t} \mid \mathcal{F}_{t-1}\right]-\mathbb{E}\left[X_{t} \mid \mathcal{F}_{t-1}\right]\right) \\
& =0
\end{aligned}
$$

by taking out what is known. So $\left(N_{t}\right)$ is a martingale.
When $\left(X_{t}\right)$ is itself a martingale (in which case $\mathbb{E}\left[X_{i} \mid \mathcal{F}_{i-1}\right]=X_{i-1}$ in the definition of $N_{t}$ ), this is a sort of "stochastic (Stieltjes) integral." When, instead, $\left(X_{t}\right)$ is a supermartingale (respectively submartingale) and $\left(C_{t}\right)$ is nonnegative and bounded, then the same computation shows that

$$
N_{t}=\sum_{i \leq t}\left(X_{i}-X_{i-1}\right) C_{i},
$$

defines a supermartingale (respectively submartingale).

As implied by the next lemma, an immediate consequence of Jensen's inequality (in its conditional version of Lemma B.6.12), submartingales naturally arise as convex functions of martingales.

Lemma 3.1.35. If $\left(M_{t}\right)_{t \geq 0}$ is a martingale and $\phi$ is a convex function such that $\mathbb{E}\left|\phi\left(M_{t}\right)\right|<+\infty$ for all $t$, then $\left(\phi\left(M_{t}\right)\right)_{t \geq 0}$ is a submartingale. Moreover, if $\left(M_{t}\right)_{t \geq 0}$ is a submartingale and $\phi$ is an increasing convex function with $\mathbb{E}\left|\phi\left(M_{t}\right)\right|<$ $+\infty$ for all $t$, then $\left(\phi\left(M_{t}\right)\right)_{t \geq 0}$ is a submartingale.

Martingales and stopping times A fundamental reason explaining the utility of martingales in analyzing a variety of stochastic processes is that they play nicely with stopping times, in particular, through what is known as the optional stopping theorem (in its various forms). We will encounter many applications of this important result. First a definition:

Definition 3.1.36. Let $\left(M_{t}\right)$ be an adapted process and $\sigma$ be a stopping time. Then

$$
M_{t}^{\sigma}(\omega):=M_{\sigma(\omega) \wedge t}(\omega),
$$

is $M_{t}$ stopped at $\sigma$.
Lemma 3.1.37. Let $\left(M_{t}\right)$ be a supermartingale and $\sigma$ be a stopping time. Then the stopped process $\left(M_{t}^{\sigma}\right)$ is a supermartingale and in particular

$$
\mathbb{E}\left[M_{t}\right] \leq \mathbb{E}\left[M_{\sigma \wedge t}\right] \leq \mathbb{E}\left[M_{0}\right] .
$$

The same result holds with equalities if $\left(M_{t}\right)$ is a martingale, and with inequalities in the opposite direction if $\left(M_{t}\right)$ is a submartingale.

Proof. Note that

$$
M_{t}^{\sigma}-M_{0}=\sum_{i \leq t} C_{i}\left(X_{i}-X_{i-1}\right),
$$

with $C_{i}=\mathbf{1}\{i \leq \sigma\} \in \mathcal{F}_{i-1}$ (which is nonnegative and bounded) and $X_{i}=M_{i}$ for all $i$, and use Example 3.1.34 to conclude that $\mathbb{E}\left[M_{\sigma \wedge t}\right] \leq \mathbb{E}\left[M_{0}\right]$.

On the other hand,

$$
M_{t}-M_{t}^{\sigma}=\sum_{i \leq t}\left(1-C_{i}\right)\left(X_{i}-X_{i-1}\right) .
$$

So the other inequality follows from the same argument.

Theorem 3.1.38 (Doob's optional stopping theorem). Let $\left(M_{t}\right)$ be a supermartingale and $\sigma$ be a stopping time. Then $M_{\sigma}$ is integrable and

$$
\mathbb{E}\left[M_{\sigma}\right] \leq \mathbb{E}\left[M_{0}\right],
$$

if any of the following conditions hold:
(i) $\sigma$ is bounded;
(ii) $\left(M_{t}\right)$ is uniformly bounded and $\sigma$ is almost surely finite;
(iii) $\mathbb{E}[\sigma]<+\infty$ and $\left(M_{t}\right)$ has bounded increments (i.e., there is $c>0$ such that $\left|M_{t}-M_{t-1}\right| \leq$ c a.s. for all $t$ );
(iv) $\left(M_{t}\right)$ is nonnegative and $\sigma$ is almost surely finite.

The first three imply equality above if $\left(M_{t}\right)$ is a martingale.
Proof. Case (iv) is Fatou's lemma (Proposition B.4.14). We prove (iii). We leave the proof of the other claims as an exercise (see Exercise 3.5).

From Lemma 3.1.37, we have

$$
\begin{equation*}
\mathbb{E}\left[M_{\sigma \wedge t}-M_{0}\right] \leq 0 . \tag{3.1.8}
\end{equation*}
$$

Furthermore the assumption that $\mathbb{E}[\sigma]<+\infty$ implies that $\sigma<+\infty$ almost surely. Hence we seek to take a limit as $t \rightarrow+\infty$ inside the expectation. To justify swapping limit and expectation, note that by a telescoping sum

$$
\begin{aligned}
\left|M_{\sigma \wedge t}-M_{0}\right| & \leq\left|\sum_{s \leq \sigma \wedge t}\left(M_{s}-M_{s-1}\right)\right| \\
& \leq \sum_{s \leq \sigma}\left|M_{s}-M_{s-1}\right| \\
& \leq c \sigma .
\end{aligned}
$$

The claim now follows from dominated convergence (Proposition B.4.14). Equality holds if $\left(M_{t}\right)$ is a martingale.

Although the optional stopping theorem (Theorem 3.1.38) is useful, one often works directly with Lemma 3.1.37 and applies suitable limit theorems (see Proposition B.4.14). The following martingale-based proof of Wald's first identity provides an illustration.

Theorem 3.1.39 (Wald's first identity). Let $X_{1}, X_{2}, \ldots \in L^{1}$ be i.i.d. with $\mathbb{E}\left[X_{1}\right]=$ $\mu$ and let $\tau \in L^{1}$ be a stopping time. Let $S_{t}=\sum_{s=1}^{t} X_{s}$. Then

$$
\mathbb{E}\left[S_{\tau}\right]=\mu \mathbb{E}[\tau]
$$

Proof. We first prove the result for nonnegative $X_{i}$ s. By Example 3.1.29, $S_{t}-t \mu$ is a martingale and Lemma 3.1.37 implies that $\mathbb{E}\left[S_{\tau \wedge t}-\mu(\tau \wedge t)\right]=0$, or

$$
\mathbb{E}\left[S_{\tau \wedge t}\right]=\mu \mathbb{E}[\tau \wedge t] .
$$

Note that, in the nonnegative case, we have $S_{\tau \wedge t} \uparrow S_{\tau}$ and $\tau \wedge t \uparrow \tau$. Thus, by monotone convergence (Proposition B.4.14), the claim $\mathbb{E}\left[S_{\tau}\right]=\mu \mathbb{E}[\tau]$ follows in that case.

Consider now the general case. Again, $\mathbb{E}\left[S_{\tau \wedge t}\right]=\mu \mathbb{E}[\tau \wedge t]$ and $\mathbb{E}[\tau \wedge t] \uparrow \mathbb{E}[\tau]$. Applying the previous argument to the sum of nonnegative random variables $R_{t}=$ $\sum_{s=1}^{t}\left|X_{s}\right|$ shows that $\mathbb{E}\left[R_{\tau}\right]=\mathbb{E}\left|X_{1}\right| \mathbb{E}[\tau]<+\infty$ by assumption. Since $\left|S_{\tau \wedge t}\right| \leq$ $R_{\tau}$ for all $t$ by the triangle inequality, dominated convergence (Proposition B.4.14) implies $\mathbb{E}\left[S_{\tau \wedge t}\right] \rightarrow \mathbb{E}\left[S_{\tau}\right]$ and we are done.

We also recall Wald's second identity. The proof, which we omit, uses the martingale in Example 3.1.31.

Theorem 3.1.40 (Wald's second identity). Let $X_{1}, X_{2}, \ldots \in L^{2}$ be i.i.d. with $\mathbb{E}\left[X_{1}\right]=0$ and $\operatorname{Var}\left[X_{1}\right]=\sigma^{2}$ and let $\tau \in L^{1}$ be a stopping time. Let $S_{t}=$ $\sum_{s=1}^{t} X_{s}$. Then

$$
\mathbb{E}\left[S_{\tau}^{2}\right]=\sigma^{2} \mathbb{E}[\tau]
$$

We illustrate Wald's identities on the gambler's ruin problem that is characteristic of applications of stopping times in Markov chains. We consider the "unbiased" and "biased" cases separately.

Example 3.1.41 (Gambler's ruin: unbiased case). Let $\left(S_{t}\right)$ be simple random walk on $\mathbb{Z}$ started at 0 and let $\tau=\tau_{a} \wedge \tau_{b}$ where $-\infty<a<0<b<+\infty$, where the first visit time $\tau_{x}$ was defined in Definition 3.1.6.

Claim 3.1.42. We have:
(i) $\tau<+\infty$ almost surely;
(ii) $\mathbb{P}\left[\tau_{a}<\tau_{b}\right]=\frac{b}{b-a}$;
(iii) $\mathbb{E}[\tau]=-a b$;
(iv) $\tau_{a}<+\infty$ almost surely but $\mathbb{E}\left[\tau_{a}\right]=+\infty$.

Proof. We prove the claims in order.
(i) We argue that in fact $\mathbb{E}[\tau]<\infty$. That follows immediately from the exponential tail of hitting times in Lemma 3.1.25 for the chain ( $S_{\tau \wedge t}$ ) whose (effective) state space, $\{a, a+1, \ldots, b\}$, is finite.
(ii) By Wald's first identity (Theorem 3.1.39) and (i), we have $\mathbb{E}\left[S_{\tau}\right]=0$ or

$$
a \mathbb{P}\left[S_{\tau}=a\right]+b \mathbb{P}\left[S_{\tau}=b\right]=0,
$$

that is, using $\mathbb{P}\left[S_{\tau}=a\right]=1-\mathbb{P}\left[S_{\tau}=b\right]=\mathbb{P}\left[\tau_{a}<\tau_{b}\right]$,

$$
\mathbb{P}\left[\tau_{a}<\tau_{b}\right]=\frac{b}{b-a} \quad \text { and } \quad \mathbb{P}\left[\tau_{a}<+\infty\right] \geq \mathbb{P}\left[\tau_{a}<\tau_{b}\right] \rightarrow 1,
$$

where we took $b \rightarrow+\infty$ in the first expression to obtain the second one.
(iii) Because $\sigma^{2}=1$, Wald's second identity (Theorem 3.1.40) says that $\mathbb{E}\left[S_{\tau}^{2}\right]=$ $\mathbb{E}[\tau]$. Furthermore, we have by (ii)

$$
\mathbb{E}\left[S_{\tau}^{2}\right]=\frac{b}{b-a} a^{2}+\frac{-a}{b-a} b^{2}=-a b .
$$

Thus $\mathbb{E}[\tau]=-a b$.
(iv) The first claim was proved in (ii). When $b \rightarrow+\infty, \tau=\tau_{a} \wedge \tau_{b} \uparrow \tau_{a}$ and monotone convergence applied to (iii) gives that $\mathbb{E}\left[\tau_{a}\right]=+\infty$.

That concludes the proof.
Note that (iv) above shows that the $L^{1}$ condition on the stopping time in Wald's second identity (Theorem 3.1.40) is necessary. Indeed we have shown $a^{2}=\mathbb{E}\left[S_{\tau_{a}}^{2}\right] \neq$ $\sigma^{2} \mathbb{E}\left[\tau_{a}\right]=+\infty$.

Example 3.1.43 (Gambler's ruin: biased case). The biased random walk on $\mathbb{Z}$ with parameter $1 / 2<p<1$ is the process $\left(S_{t}\right)$ with $S_{0}=0$ and $S_{t}=\sum_{i=1}^{t} X_{i}$ where the $X_{i}$ s are i.i.d. in $\{-1,+1\}$ with $\mathbb{P}\left[X_{1}=1\right]=p$. Let again $\tau:=\tau_{a} \wedge \tau_{b}$ where $a<0<b$. Define $q:=1-p, \delta:=p-q>0$, and $\phi(x):=(q / p)^{x}$.

Claim 3.1.44. We have:
(i) $\tau<+\infty$ almost surely;
(ii) $\mathbb{P}\left[\tau_{a}<\tau_{b}\right]=\frac{\phi(b)-\phi(0)}{\phi(b)-\phi(a)}$;
(iii) $\mathbb{E}\left[\tau_{b}\right]=\frac{b}{2 p-1}$;
(iv) $\tau_{a}=+\infty$ with positive probability.

Proof. Let $\psi_{t}(x):=x-\delta t$. We use two martingales: $\left(\phi\left(S_{t}\right)\right)$ and $\left(\psi_{t}\left(S_{t}\right)\right)$. Observe that indeed both processes are clearly integrable and

$$
\mathbb{E}\left[\phi\left(S_{t}\right) \mid \mathcal{F}_{t-1}\right]=p(q / p)^{S_{t-1}+1}+q(q / p)^{S_{t-1}-1}=\phi\left(S_{t-1}\right)
$$

and

$$
\mathbb{E}\left[\psi_{t}\left(S_{t}\right) \mid \mathcal{F}_{t-1}\right]=p\left[S_{t-1}+1-\delta t\right]+q\left[S_{t-1}-1-\delta t\right]=\psi_{t-1}\left(S_{t-1}\right)
$$

(i) This claim follows by the same argument as in the unbiased case.
(ii) Note that $\left(\phi\left(S_{t}\right)\right)$ is a nonnegative, bounded martingale since $q<p$ by assumption. By Lemma 3.1.37 and dominated convergence (Proposition B.4.14),

$$
\phi(0)=\mathbb{E}\left[\phi\left(S_{\tau}\right)\right]=\mathbb{P}\left[\tau_{a}<\tau_{b}\right] \phi(a)+\mathbb{P}\left[\tau_{a}>\tau_{b}\right] \phi(b)
$$

or, rearranging, $\mathbb{P}\left[\tau_{a}<\tau_{b}\right]=\frac{\phi(b)-\phi(0)}{\phi(b)-\phi(a)}$. Taking $b \rightarrow+\infty$, by monotonicity

$$
\begin{equation*}
\mathbb{P}\left[\tau_{a}<+\infty\right]=\frac{1}{\phi(a)}<1 \tag{3.1.9}
\end{equation*}
$$

so that $\tau_{a}=+\infty$ with positive probability. On the other hand, $\mathbb{P}\left[\tau_{b}<\tau_{a}\right]=$ $1-\mathbb{P}\left[\tau_{a}<\tau_{b}\right]=\frac{\phi(0)-\phi(a)}{\phi(b)-\phi(a)}$, and taking $a \rightarrow-\infty$

$$
P\left[\tau_{b}<+\infty\right]=1
$$

(iii) By Lemma 3.1.37 applied to $\left(\psi_{t}\left(S_{t}\right)\right)$,

$$
\begin{equation*}
0=\mathbb{E}\left[S_{\tau_{b} \wedge t}-\delta\left(\tau_{b} \wedge t\right)\right] \tag{3.1.10}
\end{equation*}
$$

By monotone convergence (Proposition B.4.14), $\mathbb{E}\left[\tau_{b} \wedge t\right] \uparrow \mathbb{E}\left[\tau_{b}\right]$. Furthermore, observe that $-\inf _{t} S_{t} \geq 0$ almost surely since $S_{0}=0$. Moreover, for $x \geq 0$, by (3.1.9)

$$
\mathbb{P}\left[-\inf _{t} S_{t} \geq x\right]=\mathbb{P}\left[\tau_{-x}<+\infty\right]=\left(\frac{q}{p}\right)^{x}
$$

so that $\mathbb{E}\left[-\inf _{t} S_{t}\right]=\sum_{x \geq 1} \mathbb{P}\left[-\inf _{t} S_{t} \geq x\right]<+\infty$. Hence, in (3.1.10), we can use dominated convergence (Proposition B.4.14) with

$$
\left|S_{\tau_{b} \wedge t}\right| \leq \max \left\{b,-\inf _{t} S_{t}\right\}
$$

and the fact that $\tau_{b}<+\infty$ almost surely from (ii) to deduce that $\mathbb{E}\left[\tau_{b}\right]=$ $\frac{\mathbb{E}\left[S_{\tau_{b}}\right]}{p-q}=\frac{b}{2 p-1}$.
(iv) That claim was proved in (ii).

That concludes the proof.
Note that, in (iii) above, in order to apply Wald's first identity directly we would have had to prove that $\tau_{b} \in L^{1}$ first.

We also obtain the following maximal version of Markov's inequality (Theorem 2.1.1).
Theorem 3.1.45 (Doob's submartingale inequality). Let $\left(M_{t}\right)$ be a nonnegative submartingale. Then, for $b>0$,

$$
\mathbb{P}\left[\sup _{0 \leq s \leq t} M_{s} \geq b\right] \leq \frac{\mathbb{E}\left[M_{t}\right]}{b}
$$

Observe that a naive application of Markov's inequality implies only that

$$
\sup _{0 \leq s \leq t} \mathbb{P}\left[M_{s} \geq b\right] \leq \frac{\mathbb{E}\left[M_{t}\right]}{b}
$$

where we used that $\mathbb{E}\left[M_{s}\right] \leq \mathbb{E}\left[M_{t}\right]$ for all $0 \leq s \leq t$ for a submartingale. Introducing an appropriate stopping time immediately gives something stronger. (Exercise 3.6 asks for the supermartingale version of this.)

Proof. Let $\sigma$ be the first time that $M_{t} \geq b$. Then the event of interest can be characterized as

$$
\left\{\sup _{0 \leq s \leq t} M_{s} \geq b\right\}=\left\{M_{\sigma \wedge t} \geq b\right\}
$$

By Markov's inequality,

$$
\mathbb{P}\left[M_{\sigma \wedge t} \geq b\right] \leq \frac{\mathbb{E}\left[M_{\sigma \wedge t}\right]}{b}
$$

Lemma 3.1.37 implies that $\mathbb{E}\left[M_{\sigma \wedge t}\right] \leq \mathbb{E}\left[M_{t}\right]$, which concludes the proof.
One consequence of the previous bound is a strengthening of Chebyshev's inequality (Theorem 2.1.2) for sums of independent random variables.
Corollary 3.1.46 (Kolmogorov's maximal inequality). Let $X_{1}, X_{2}, \ldots$ be independent random variables with $\mathbb{E}\left[X_{i}\right]=0$ and $\operatorname{Var}\left[X_{i}\right]<+\infty$. Define $S_{t}=$ $\sum_{i \leq t} X_{i}$. Then, for $\beta>0$,

$$
\mathbb{P}\left[\max _{i \leq t}\left|S_{i}\right| \geq \beta\right] \leq \frac{\operatorname{Var}\left[S_{t}\right]}{\beta^{2}}
$$

Proof. By Example 3.1.29, $\left(S_{t}\right)$ is a martingale. By Lemma 3.1.35, $\left(S_{t}^{2}\right)$ is hence a (nonnegative) submartingale. The result follows from Doob's submartingale inequality (Theorem 3.1.45).

Convergence Finally another fundamental result about martingales is the following convergence theorem, which we state without proof. We give a quick application below.

Theorem 3.1.47 (Convergence theorem). Let $\left(M_{t}\right)$ be a supermartingale bounded in $L^{1}$. Then $\left(M_{t}\right)$ converges almost surely to a finite limit $M_{\infty}$. Moreover, letting $M_{\infty}:=\lim \sup _{t} M_{t}$, then $M_{\infty} \in \mathcal{F}_{\infty}$ and $\mathbb{E}\left|M_{\infty}\right|<+\infty$.
Corollary 3.1.48 (Convergence of non-negative supermartingales). If $\left(M_{t}\right)$ is a non-negative supermartingale then $M_{t}$ converges almost surely to a finite limit $M_{\infty}$ with $\mathbb{E}\left[M_{\infty}\right] \leq \mathbb{E}\left[M_{0}\right]$.
Proof. By the supermartingale property, $\left(M_{t}\right)$ is bounded in $L^{1}$ since

$$
\mathbb{E}\left|M_{t}\right|=\mathbb{E}\left[M_{t}\right] \leq \mathbb{E}\left[M_{0}\right], \forall t
$$

Then we use the martingale convergence theorem (Theorem 3.1.47) and Fatou's lemma (Proposition B.4.14).

Example 3.1.49 (Pólya's urn). An urn contains 1 red ball and 1 green ball. At each time, we pick one ball and put it back with an extra ball of the same color. This process is known as Pólya's urn. Let $R_{t}$ (respectively $G_{t}$ ) be the number of red balls (respectively green balls) after the $t$ th draw. Let

$$
\mathcal{F}_{t}:=\sigma\left(R_{0}, G_{0}, R_{1}, G_{1}, \ldots, R_{t}, G_{t}\right)
$$

Define $M_{t}$ to be the fraction of green balls after the $t$ th draw. Then

$$
\begin{aligned}
\mathbb{E}\left[M_{t} \mid \mathcal{F}_{t-1}\right]= & \frac{R_{t-1}}{G_{t-1}+R_{t-1}} \frac{G_{t-1}}{G_{t-1}+R_{t-1}+1} \\
& \quad+\frac{G_{t-1}}{G_{t-1}+R_{t-1}} \frac{G_{t-1}+1}{G_{t-1}+R_{t-1}+1} \\
= & \frac{G_{t-1}}{G_{t-1}+R_{t-1}} \\
= & M_{t-1} .
\end{aligned}
$$

Since $M_{t} \geq 0$ and is a martingale, we have $M_{t} \rightarrow M_{\infty}$ almost surely. In fact, Exercise 3.4 asks for a proof that

$$
\mathbb{P}\left[G_{t}=m+1\right]=\binom{t}{m} \frac{m!(t-m)!}{(t+1)!}=\frac{1}{t+1} .
$$

So taking a limit as $t \rightarrow+\infty$

$$
\mathbb{P}\left[M_{t} \leq x\right]=\frac{\lfloor x(t+2)-1\rfloor}{t+1} \rightarrow x
$$

That is, $\left(M_{t}\right)$ converges in distribution to a uniform random variable on $[0,1]$.

Convergence of the expectation in general requires stronger conditions. A simple case is boundedness in $L^{2}$. Before stating the result, we derive a key property of martingales in $L^{2}$ which will be useful later.

Lemma 3.1.50 (Orthogonality of increments). Let $\left(M_{t}\right)$ be a martingale with $M_{t} \in L^{2}$ for all $t$. Let $s \leq t \leq u \leq v$. Then,

$$
\left\langle M_{t}-M_{s}, M_{v}-M_{u}\right\rangle=0,
$$

where $\langle X, Y\rangle=\mathbb{E}[X Y]$.
Proof. Use $M_{u}=\mathbb{E}\left[M_{v} \mid \mathcal{F}_{u}\right]$ and $M_{t}-M_{s} \in \mathcal{F}_{u}$, and apply the $L^{2}$ characterization of the conditional expectation (Theorem B.6.2).

In words, martingale increments over disjoint time intervals are uncorrelated (provided the second moment exists). Note that this is weaker than the independence of increments of random walks. (See Section 3.2.1 for more discussion on this.)

Theorem 3.1.51 (Convergence of martingales bounded in $L^{2}$ ). Let $\left(M_{t}\right)$ be a martingale with $M_{t} \in L^{2}$ for all $t$. Then $\left(M_{t}\right)$ is bounded in $L^{2}$ if and only if

$$
\sum_{k \geq 1} \mathbb{E}\left[\left(M_{t}-M_{t-1}\right)^{2}\right]<+\infty
$$

When this is the case, $M_{t}$ converges almost surely and in $L^{2}$ to a finite limit $M_{\infty}$, and furthermore

$$
\mathbb{E}\left[M_{t}\right] \rightarrow \mathbb{E}\left[M_{\infty}\right]<+\infty,
$$

as $t \rightarrow+\infty$.
Proof. Writing $M_{t}$ as a telescoping sum of increments, the orthogonality of increments (Lemma 3.1.50) implies

$$
\begin{aligned}
\mathbb{E}\left[M_{t}^{2}\right] & =\mathbb{E}\left[\left(M_{0}+\sum_{s=1}^{t}\left(M_{s}-M_{s-1}\right)\right)^{2}\right] \\
& =\mathbb{E}\left[M_{0}^{2}\right]+\sum_{s=1}^{t} \mathbb{E}\left[\left(M_{s}-M_{s-1}\right)^{2}\right],
\end{aligned}
$$

proving the first claim.
By the monotonicity of norms (Lemma B.4.16), ( $M_{t}$ ) being bounded in $L^{2}$ implies that $\left(M_{t}\right)$ is bounded in $L^{1}$ which, in turn, implies that $M_{t}$ converges almost
surely to a finite limit $M_{\infty}$ with $\mathbb{E}\left|M_{\infty}\right|<+\infty$ by Theorem 3.1.47. Then using Fatou's lemma (Proposition B.4.14) in

$$
\mathbb{E}\left[\left(M_{t+s}-M_{t}\right)^{2}\right]=\sum_{t+1 \leq i \leq t+s} \mathbb{E}\left[\left(M_{i}-M_{i-1}\right)^{2}\right],
$$

gives

$$
\mathbb{E}\left[\left(M_{\infty}-M_{t}\right)^{2}\right] \leq \sum_{t+1 \leq i} \mathbb{E}\left[\left(M_{i}-M_{i-1}\right)^{2}\right]
$$

The right-hand side goes to 0 as $t \rightarrow+\infty$ since the full series is finite, which proves the second claim.

The last claim follows from Lemmas B.4.16 and B.4.17.

### 3.1.4 $\triangleright$ Percolation: critical regime on infinite d-regular tree

Consider bond percolation (see Definition 1.2.1) on the infinite $d$-regular tree $\mathbb{T}_{d}$ rooted at a vertex 0. In Section 2.3.3, we showed that

$$
p_{c}\left(\mathbb{T}_{d}\right)=\sup \left\{p \in[0,1]: \mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right]=0\right\}=\frac{1}{d-1}
$$

where recall that $\mathcal{C}_{0}$ is the open cluster of the root. Here we consider the critical case, that is, we set density $p=\frac{1}{d-1}$. (The same results apply to the infinite $b$-ary tree $\widehat{\mathbb{T}}_{b}$ with $d=b+1$.) Assume $d \geq 3$ (since $d=2$ is simply a path).

First:
Claim 3.1.52. $\left|\mathcal{C}_{0}\right|<+\infty$ almost surely.
Let $X_{n}:=\left|\partial_{n} \cap \mathcal{C}_{0}\right|$, where $\partial_{n}$ are the $n$-th level vertices. In Section 2.3.3, we proved the same claim in the subcritical case using the first moment method. It does not work here because

$$
\mathbb{E} X_{n}=d(d-1)^{n-1} p^{n}=\frac{d}{d-1} \nrightarrow 0 .
$$

Instead we use a martingale argument which will be generalized when we discuss branching processes in Section 6.1.

Proof of Claim 3.1.52. Let $b:=d-1$ be the branching ratio. Because the root has a different number of children, we consider the descendants of its children. Let $Z_{n}$ be the number of vertices in the open cluster of the first child of the root $n$ levels below it and let $\mathcal{F}_{n}=\sigma\left(Z_{0}, \ldots, Z_{n}\right)$. Then $Z_{0}=1$ and

$$
\mathbb{E}\left[Z_{n} \mid \mathcal{F}_{n-1}\right]=b p Z_{n-1}=Z_{n-1} .
$$

So $\left(Z_{n}\right)$ is a nonnegative, integer-valued martingale and it converges almost surely to a finite limit by Corollary 3.1.48. (In particular, $\mathbb{E}\left[Z_{n}\right]=1$, which will be useful below.) But, clearly, for any integer $k>0$ and $N \geq 0$

$$
\mathbb{P}\left[Z_{n}=k, \forall n \geq N\right]=0,
$$

so it must be that the limit is 0 almost surely. In other words, $Z_{n}$ is eventually 0 for all $n$ large enough. This is true for every child of the root. Hence the open cluster of the root is finite almost surely.

On the other hand:

## Claim 3.1.53.

$$
\mathbb{E}\left|\mathcal{C}_{0}\right|=+\infty .
$$

Proof. Consider the descendant subtree, $T_{1}$, of the first child of the root, which we denote by 1 . Let $\widetilde{\mathcal{C}_{1}}$ be the open cluster of 1 in $T_{1}$. As we showed in the previous claim, the expected number of vertices on any level of $T_{1}$ is 1 . So $\mathbb{E}\left|\widetilde{\mathcal{C}_{1}}\right|=+\infty$ by summing over the levels.

### 3.2 Concentration for martingales and applications

The Chernoff-Cramér method extends naturally to martingales. This observation leads to powerful new tail bounds that hold far beyond the case of sums of independent variables. In particular it will allow us to prove one version of the concentration phenomenon, which can be stated informally as: a function $f\left(X_{1}, \ldots, X_{n}\right)$ of many independent random variables that is not too sensitive to any of its coordinates tends to be close to its mean.

### 3.2.1 Azuma-Hoeffding inequality

The main result of this section is the following generalization of Hoeffding's inequality (Theorem 2.4.10).

Theorem 3.2.1 (Maximal Azuma-Hoeffding inequality). Let $\left(Z_{t}\right)_{t \in \mathbb{Z}_{+}}$be a martingale with respect to the filtration $\left(\mathcal{F}_{t}\right)_{t \in \mathbb{Z}_{+}}$. Assume that there are predictable processes $\left(A_{t}\right)$ and $\left(B_{t}\right)$ (i.e., $\left.A_{t}, B_{t} \in \mathcal{F}_{t-1}\right)$ and constants $0<c_{t}<+\infty$ such that: for all $t \geq 1$, almost surely,

$$
A_{t} \leq Z_{t}-Z_{t-1} \leq B_{t} \quad \text { and } \quad B_{t}-A_{t} \leq c_{t} .
$$

Then, for all $\beta>0$,

$$
\mathbb{P}\left[\sup _{0 \leq i \leq t}\left(Z_{i}-Z_{0}\right) \geq \beta\right] \leq \exp \left(-\frac{2 \beta^{2}}{\sum_{i \leq t} c_{i}^{2}}\right) .
$$

Applying this inequality to $\left(-Z_{t}\right)$ gives a tail bound in the other direction.
Proof of Theorem 3.2.1. As in the Chernoff-Cramér method, we start by applying Markov's inequality (Theorem 2.1.1). Here we use the maximal version for submartingales, Doob's submartingale inequality (Theorem 3.1.45). First notice that $e^{s x}$ is increasing and convex for $s>0$, so that by Lemma 3.1.35 the process $\left(e^{s\left(Z_{t}-Z_{0}\right)}\right)_{t}$ is a submartingale. Hence, for $s>0$, by Theorem 3.1.45

$$
\begin{align*}
\mathbb{P}\left[\sup _{0 \leq i \leq t}\left(Z_{i}-Z_{0}\right) \geq \beta\right] & =\mathbb{P}\left[\sup _{0 \leq i \leq t} e^{s\left(Z_{i}-Z_{0}\right)} \geq e^{s \beta}\right] \\
& \leq \frac{\mathbb{E}\left[e^{s\left(Z_{t}-Z_{0}\right)}\right]}{e^{s \beta}} \\
& =\frac{\mathbb{E}\left[e^{s \sum_{r=1}^{t}\left(Z_{r}-Z_{r-1}\right)}\right]}{e^{s \beta}} . \tag{3.2.1}
\end{align*}
$$

Unlike the Chernoff-Cramér case, however, the terms in the exponent are not independent. Instead, to exploit the martingale property, we condition on the filtration. By taking out what is known (Lemma B.6.13)

$$
\mathbb{E}\left[\mathbb{E}\left[e^{s \sum_{r=1}^{t}\left(Z_{r}-Z_{r-1}\right)} \mid \mathcal{F}_{t-1}\right]\right]=\mathbb{E}\left[e^{s \sum_{r=1}^{t-1}\left(Z_{r}-Z_{r-1}\right)} \mathbb{E}\left[e^{s\left(Z_{t}-Z_{t-1}\right)} \mid \mathcal{F}_{t-1}\right]\right] .
$$

The martingale property and the assumption in the statement imply that, conditioned on $\mathcal{F}_{t-1}$, the random variable $Z_{t}-Z_{t-1}$ is centered and lies in an interval of length $c_{t}$. Hence by Hoeffding's lemma (Lemma 2.4.12), it holds almost surely that

$$
\begin{equation*}
\mathbb{E}\left[e^{s\left(Z_{t}-Z_{t-1}\right)} \mid \mathcal{F}_{t-1}\right] \leq \exp \left(\frac{s^{2} c_{t}^{2} / 4}{2}\right)=\exp \left(\frac{c_{t}^{2} s^{2}}{8}\right) . \tag{3.2.2}
\end{equation*}
$$

Using the tower property (Lemma B.6.16) and arguing by induction, we obtain

$$
\mathbb{E}\left[e^{s\left(Z_{t}-Z_{0}\right)}\right] \leq \exp \left(\frac{s^{2} \sum_{r \leq t} c_{r}^{2}}{8}\right)
$$

Put differently, we have proved that $Z_{t}-Z_{0}$ is sub-Gaussian with variance factor $\frac{1}{4} \sum_{r \leq t} c_{r}^{2}$. By (2.4.16) (or, equivalently, by choosing $s=\beta / \frac{1}{4} \sum_{r \leq t} c_{r}^{2}$ in (3.2.1)) we get the result.

In Theorem 3.2.1 the martingale difference sequence $\left(X_{t}\right)$, where $X_{t}:=Z_{t}-$ $Z_{t-1}$, is not only "pairwise uncorrelated" by Lemma 3.1.50, that is,

$$
\mathbb{E}\left[X_{s} X_{r}\right]=0, \quad \forall r \neq s,
$$

but it is in fact "mutually uncorrelated," that is,

$$
\mathbb{E}\left[X_{j_{1}} \cdots X_{j_{k}}\right]=0, \quad \forall k \geq 1, \forall 1 \leq j_{1}<\cdots<j_{k}
$$

This stronger property helps explain why $\sum_{r \leq t} X_{r}$ is highly concentrated. This point is the subject of Exercise 3.7, which guides the reader through a slightly different proof of the Azuma-Hoeffding inequality. Compare with Exercises 2.5 and 2.6.

### 3.2.2 Method of bounded differences

The power of the maximal Azuma-Hoeffding inequality (Theorem 3.2.1) is that it produces tail inequalities for quantities other than sums of independent variables. The setting is the following. Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$ and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume that $f$ : $\mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function. Our goal is to characterize the concentration properties of $f(X)$ around its expectation in terms of its "discrete derivatives"

$$
\left.\begin{array}{rl}
D_{i} f(x):= & \sup _{y \in \mathcal{X}_{i}} f(
\end{array} x_{1}, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_{n}\right), ~\left(\inf _{y^{\prime} \in \mathcal{X}_{i}} f\left(x_{1}, \ldots, x_{i-1}, y^{\prime}, x_{i+1}, \ldots, x_{n}\right), ~ \$\right.
$$

where $x=\left(x_{1}, \ldots, x_{n}\right) \in \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n}$. We think of $D_{i} f(x)$ as a measure of the "sensitivity" of $f$ to its $i$-th coordinate.

## High-level idea

We begin with two easier bounds that we will improve below. The trick to analyzing the concentration of $f(X)$ is to consider the Doob martingale (see Example 3.1.33)

$$
\begin{equation*}
Z_{i}=\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right], \tag{3.2.3}
\end{equation*}
$$

where $\mathcal{F}_{i}=\sigma\left(X_{1}, \ldots, X_{i}\right)$, which is well-defined provided $\mathbb{E}|f(X)|<+\infty$. Note that

$$
Z_{n}=\mathbb{E}\left[f(X) \mid \mathcal{F}_{n}\right]=f(X),
$$

and

$$
Z_{0}=\mathbb{E}[f(X)],
$$

so that we can write

$$
f(X)-\mathbb{E}[f(X)]=\sum_{i=1}^{n}\left(Z_{i}-Z_{i-1}\right) .
$$

Intuitively, the martingale difference $Z_{i}-Z_{i-1}$ tracks the change in our expectation of $f(X)$ as $X_{i}$ is revealed.

In fact a clever probabilistic argument relates martingale differences directly to discrete derivatives. Let $X^{\prime}=\left(X_{1}^{\prime}, \ldots, X_{n}^{\prime}\right)$ be an independent copy of $X$ and let

$$
X^{(i)}=\left(X_{1}, \ldots, X_{i-1}, X_{i}^{\prime}, X_{i+1}, \ldots, X_{n}\right) .
$$

Then

$$
\begin{aligned}
Z_{i}-Z_{i-1} & =\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]-\mathbb{E}\left[f(X) \mid \mathcal{F}_{i-1}\right] \\
& =\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]-\mathbb{E}\left[f\left(X^{(i)}\right) \mid \mathcal{F}_{i-1}\right] \\
& =\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]-\mathbb{E}\left[f\left(X^{(i)}\right) \mid \mathcal{F}_{i}\right] \\
& =\mathbb{E}\left[f(X)-f\left(X^{(i)}\right) \mid \mathcal{F}_{i}\right] .
\end{aligned}
$$

Note that we crucially used the independence of the $X_{k} \mathrm{~s}$ in the second and third lines. But then, by Jensen's inequality (Lemma B.6.12),

$$
\begin{equation*}
\left|Z_{i}-Z_{i-1}\right| \leq\left\|D_{i} f\right\|_{\infty} \tag{3.2.4}
\end{equation*}
$$

Assume further that $\mathbb{E}\left[f(X)^{2}\right]<+\infty$. By the orthogonality of increments of martingales in $L^{2}$ (Lemma 3.1.50), we immediately obtain a bound on the variance of $f$

$$
\begin{equation*}
\operatorname{Var}[f(X)]=\mathbb{E}\left[\left(Z_{n}-Z_{0}\right)^{2}\right]=\sum_{i=1}^{n} \mathbb{E}\left[\left(Z_{i}-Z_{i-1}\right)^{2}\right] \leq \sum_{i=1}^{n}\left\|D_{i} f\right\|_{\infty}^{2} \tag{3.2.5}
\end{equation*}
$$

By the maximal Azuma-Hoeffding inequality and the fact that

$$
Z_{i}-Z_{i-1} \in\left[-\left\|D_{i} f\right\|_{\infty},\left\|D_{i} f\right\|_{\infty}\right]
$$

we also get a bound on the tail

$$
\begin{equation*}
\mathbb{P}[f(X)-\mathbb{E}[f(X)] \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2 \sum_{i \leq n}\left\|D_{i} f\right\|_{\infty}^{2}}\right) \tag{3.2.6}
\end{equation*}
$$

A more careful analysis, which we detail below, leads to a better bound.
We emphasize that, although it may not be immediately obvious, independence plays a crucial role in the bound (3.2.4), as the next example shows.

Example 3.2.2 (A counterexample). Let $f\left(x_{1}, \ldots, x_{n}\right)=x_{1}+\cdots+x_{n}$ where $x_{i} \in\{-1,1\}$ for all $i$. Then,

$$
\left\|D_{1} f\right\|_{\infty}=\sup _{x_{2}, \ldots, x_{n}}\left[\left(1+x_{2}+\cdots+x_{n}\right)-\left(-1+x_{2}+\cdots+x_{n}\right)\right]=2,
$$

and similarly $\left\|D_{i} f\right\|_{\infty}=2$ for $i=2, \ldots, n$. Let $X_{1}$ be a uniform random variable on $\{-1,1\}$. First consider the case where we set $X_{2}, \ldots, X_{n}$ all equal to $X_{1}$. Then

$$
\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right)\right]=0
$$

and

$$
\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right) \mid X_{1}\right]=n X_{1},
$$

so that

$$
\left|\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right) \mid X_{1}\right]-\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right)\right]\right|=n>2 .
$$

In particular, the corresponding Doob martingale does not have increments bounded by $\left\|D_{i} f\right\|_{\infty}=2$.

Fo a less extreme example which has support over all of $\{-1,1\}^{n}$, let

$$
U_{i}= \begin{cases}1, & \text { w.p. } 1-\varepsilon \\ -1, & \text { w.p. } \varepsilon\end{cases}
$$

for some $\varepsilon>0$ independently for all $i=1, \ldots, n-1$. Let again $X_{1}$ be a uniform random variable on $\{-1,1\}$ and, for $i=2, \ldots, n$, define the random variable $X_{i}=$ $U_{i-1} X_{i-1}$, that is, $X_{i}$ is the same as $X_{i-1}$ with probability $1-\varepsilon$ and otherwise is flipped. Then,

$$
\begin{aligned}
\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right)\right] & =\mathbb{E}\left[X_{1}+\cdots+X_{n}\right] \\
& =\mathbb{E}\left[X_{1}\left(1+\sum_{i=1}^{n-1} \prod_{j \leq i} U_{j}\right)\right] \\
& =\mathbb{E}\left[X_{1}\right] \mathbb{E}\left[1+\sum_{i=1}^{n-1} \prod_{j \leq i} U_{j}\right] \\
& =0,
\end{aligned}
$$

by the independence of $X_{1}$ and the $U_{i}$ s. Similarly

$$
\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right) \mid X_{1}\right]=X_{1} \mathbb{E}\left[1+\sum_{i=1}^{n-1} \prod_{j \leq i} U_{j}\right]=X_{1}\left(\sum_{i=0}^{n-1}(1-2 \varepsilon)^{i}\right),
$$

so that

$$
\left|\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right) \mid X_{1}\right]-\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right)\right]\right|=\left(\sum_{i=0}^{n-1}(1-2 \varepsilon)^{i}\right)>2,
$$

for $\varepsilon$ small enough and $n \geq 3$. In particular, the corresponding Doob martingale does not have increments bounded by $\left\|D_{i} f\right\|_{\infty}=2$.

## Variance bounds

We give improved bounds on the variance (compared to (3.2.5)). Our first bound explicitly decomposes the variance of $f(X)$ over the contributions of its individual entries.

Theorem 3.2.3 (Tensorization of the variance). Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$ and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume that $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function with $\mathbb{E}\left[f(X)^{2}\right]<$ $+\infty$. Define $\mathcal{F}_{i}=\sigma\left(X_{1}, \ldots, X_{i}\right), \mathcal{G}_{i}=\sigma\left(X_{1}, \ldots, X_{i-1}, X_{i+1}, \ldots, X_{n}\right)$ and $Z_{i}=\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]$. Then we have

$$
\operatorname{Var}[f(X)] \leq \sum_{i=1}^{n} \mathbb{E}\left[\operatorname{Var}\left[f(X) \mid \mathcal{G}_{i}\right]\right]
$$

Proof of Theorem 3.2.3. The key lemma is the following.

## Lemma 3.2.4.

$$
\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i}\right]=\mathbb{E}\left[f(X) \mid \mathcal{F}_{i-1}\right]
$$

Proof. By the tower property (Lemma B.6.16),

$$
\mathbb{E}\left[f(X) \mid \mathcal{F}_{i-1}\right]=\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i-1}\right] .
$$

Moreover, $\sigma\left(X_{i}\right)$ is independent of $\sigma\left(\mathcal{G}_{i}, \mathcal{F}_{i-1}\right)$ so by the role of independence (Lemma B.6.14), we have

$$
\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i-1}\right]=\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i-1}, X_{i}\right]=\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i}\right] .
$$

Combining the last two displays gives the result.
Again, we take advantage of the orthogonality of increments to write

$$
\operatorname{Var}[f(X)]=\sum_{i=1}^{n} \mathbb{E}\left[\left(Z_{i}-Z_{i-1}\right)^{2}\right]
$$

By the lemma above,

$$
\begin{aligned}
\left(Z_{i}-Z_{i-1}\right)^{2} & =\left(\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]-\mathbb{E}\left[f(X) \mid \mathcal{F}_{i-1}\right]\right)^{2} \\
& =\left(\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right]-\mathbb{E}\left[\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i}\right]\right)^{2} \\
& =\left(\mathbb{E}\left[f(X)-\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right] \mid \mathcal{F}_{i}\right]\right)^{2} \\
& \leq \mathbb{E}\left[\left(f(X)-\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right]\right)^{2} \mid \mathcal{F}_{i}\right],
\end{aligned}
$$

where we used Jensen's inequality on the last line. Taking expectations and using the tower property

$$
\begin{aligned}
\operatorname{Var}[f(X)] & =\sum_{i=1}^{n} \mathbb{E}\left[\left(Z_{i}-Z_{i-1}\right)^{2}\right] \\
& \leq \sum_{i=1}^{n} \mathbb{E}\left[\mathbb{E}\left[\left(f(X)-\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right]\right)^{2} \mid \mathcal{F}_{i}\right]\right] \\
& =\sum_{i=1}^{n} \mathbb{E}\left[\left(f(X)-\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right]\right)^{2}\right] \\
& =\sum_{i=1}^{n} \mathbb{E}\left[\mathbb{E}\left[\left(f(X)-\mathbb{E}\left[f(X) \mid \mathcal{G}_{i}\right]\right)^{2} \mid \mathcal{G}_{i}\right]\right] \\
& =\sum_{i=1}^{n} \mathbb{E}\left[\operatorname{Var}\left[f(X) \mid \mathcal{G}_{i}\right]\right] .
\end{aligned}
$$

That concludes the proof.
We derive two useful consequences of the tensorization property of the variance. The first one is the Efron-Stein inequality.

Theorem 3.2.5 (Efron-Stein inequality). Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$ and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume that $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function with $\mathbb{E}\left[f(X)^{2}\right]<+\infty$. Let $X^{\prime}=\left(X_{1}^{\prime}, \ldots, X_{n}^{\prime}\right)$ be an independent copy of $X$ and

$$
X^{(i)}=\left(X_{1}, \ldots, X_{i-1}, X_{i}^{\prime}, X_{i+1}, \ldots, X_{n}\right) .
$$

Then,

$$
\operatorname{Var}[f(X)] \leq \frac{1}{2} \sum_{i=1}^{n} \mathbb{E}\left[\left(f(X)-f\left(X^{(i)}\right)\right)^{2}\right]
$$

Proof. Observe that if $Y^{\prime}$ is an independent copy of $Y \in L^{2}$, then $\operatorname{Var}[Y]=$ $\frac{1}{2} \mathbb{E}\left[\left(Y-Y^{\prime}\right)^{2}\right]$, which can be seen by adding and subtracting the mean, expanding and using independence. Hence,

$$
\operatorname{Var}\left[f(X) \mid \mathcal{G}_{i}\right]=\frac{1}{2} \mathbb{E}\left[\left(f(X)-f\left(X^{(i)}\right)\right)^{2} \mid \mathcal{G}_{i}\right],
$$

where we used the independence of the $X_{i} \mathrm{~s}$ and $X_{i}^{\prime} \mathrm{s}$. Plugging back into Theorem 3.2.3 gives the claim.

Our second consequence of Theorem 3.2.3 is a Poincaré-type inequality which relates the variance of a function to its expected "square gradient." Compare to the much weaker (3.2.5), which involves in each term a supremum rather than an expectation.

Theorem 3.2.6 (Bounded differences inequality). Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$ and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume that $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function with $\mathbb{E}\left[f(X)^{2}\right]<$ $+\infty$. Then

$$
\operatorname{Var}[f(X)] \leq \frac{1}{4} \sum_{i=1}^{n} \mathbb{E}\left[D_{i} f(X)^{2}\right] .
$$

Proof. By Lemma 2.4.11,

$$
\operatorname{Var}\left[f(X) \mid \mathcal{G}_{i}\right] \leq \frac{1}{4} D_{i} f(X)^{2}
$$

Plugging back into Theorem 3.2.3 gives the claim.
Remark 3.2.7. For comparison, a version of the Poincaré inequality in one dimension asserts the following: let $f:[0, T] \rightarrow \mathbb{R}$ be continuously differentiable with $f(0)=$ $f(T)=0, \int_{0}^{T} f(x)^{2}+f^{\prime}(x)^{2} \mathrm{~d} x<+\infty$ and $\int_{0}^{T} f(x) \mathrm{d} x=0$, then

$$
\begin{equation*}
\int_{0}^{T} f(x)^{2} \mathrm{~d} x \leq C \int_{0}^{T} f^{\prime}(x)^{2} \mathrm{~d} x \tag{3.2.7}
\end{equation*}
$$

where the best possible C is $T^{2} / 4 \pi^{2}$ (see, e.g., [SSO3, Chapter 3, Exercise 11]; this case is also known as Wirtinger's inequality). We give a quick proof for $T=1$ with the suboptimal $C=1$. Note that $f(x)=\int_{0}^{x} f^{\prime}(y) \mathrm{d} y$ so, by Cauchy-Schwarz (Theorem B.4.8),

$$
f(x)^{2} \leq x \int_{0}^{x} f^{\prime}(y)^{2} \mathrm{~d} y \leq \int_{0}^{1} f^{\prime}(y)^{2} \mathrm{~d} y .
$$

The result follows by integration. Intuitively, for a function with mean 0 to have a large norm, it must have a large absolute derivative somewhere.

Example 3.2.8 (Longest common subsequence). Let $X_{1}, \ldots, X_{2 n}$ be independent uniform random variables in $\{-1,+1\}$. Let $Z$ be the length of the longest common subsequence in $\left(X_{1}, \ldots, X_{n}\right)$ and $\left(X_{n+1}, \ldots, X_{2 n}\right)$, that is,

$$
\begin{aligned}
Z=\max \{k: & \exists 1 \leq i_{1}<i_{2}<\cdots<i_{k} \leq n \\
& \text { and } n+1 \leq j_{1}<j_{2}<\cdots<j_{k} \leq 2 n \\
& \text { such that } \left.X_{i_{1}}=X_{j_{1}}, X_{i_{2}}=X_{i_{2}}, \ldots, X_{i_{k}}=X_{j_{k}}\right\} .
\end{aligned}
$$

Then, writing $Z=f\left(X_{1}, \ldots, X_{2 n}\right)$, it follows that $\left\|D_{i} f\right\|_{\infty} \leq 1$. Indeed, fix $\mathbf{x}=\left(x_{1}, \ldots, x_{2 n}\right)$ and let $\mathbf{x}^{i,+}$ (respectively $\mathbf{x}^{i,-}$ ) be $\mathbf{x}$ where the $i$-th component is replaced with +1 (respectively -1 ). Assume without loss of generality that $f\left(\mathbf{x}^{i,-}\right) \leq f\left(\mathbf{x}^{i,+}\right)$. Then $\left|f\left(\mathbf{x}^{i,+}\right)-f\left(\mathbf{x}^{i,-}\right)\right| \leq 1$ because removing the $i$-th component (and its match) from a longest common subsequence when $x_{i}=+1$ (if present) decreases the length by 1 . Since this is true for any $\mathbf{x}$, we have $\left\|D_{i} f\right\|_{\infty} \leq$ 1. Finally, by the bounded differences inequality (Theorem 3.2.6),

$$
\operatorname{Var}[Z] \leq \frac{1}{4} \sum_{i=1}^{2 n}\left\|D_{i} f\right\|_{\infty}^{2} \leq \frac{n}{2},
$$

which is much better than the obvious $\operatorname{Var}[Z] \leq \mathbb{E}\left[Z^{2}\right] \leq n^{2}$. Note that we did not require any information about the expectation of $Z$.

## McDiarmid's inequality

The following powerful consequence of the Azuma-Hoeffding inequality is commonly referred to as the method of bounded differences. Compare to (3.2.6).

Theorem 3.2.9 (McDiarmid's inequality). Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$, and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function such that $\left\|D_{i} f\right\|_{\infty}<+\infty$ for all $i$. Then for all $\beta>0$

$$
\mathbb{P}[f(X)-\mathbb{E} f(X) \geq \beta] \leq \exp \left(-\frac{2 \beta^{2}}{\sum_{i \leq n}\left\|D_{i} f\right\|_{\infty}^{2}}\right)
$$

Once again, applying the inequality to $-f$ gives a tail bound in the other direction.
Proof of Theorem 3.2.9. As before, we let

$$
Z_{i}=\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right],
$$

where $\mathcal{F}_{i}=\sigma\left(X_{1}, \ldots, X_{i}\right)$, we let $\mathcal{G}_{i}=\sigma\left(X_{1}, \ldots, X_{i-1}, X_{i+1}, \ldots, X_{n}\right)$. Then, it holds that $A_{i} \leq Z_{i}-Z_{i-1} \leq B_{i}$ where

$$
B_{i}=\mathbb{E}\left[\sup _{y \in \mathcal{X}_{i}} f\left(X_{1}, \ldots, X_{i-1}, y, X_{i+1}, \ldots, X_{n}\right)-f(X) \mid \mathcal{F}_{i-1}\right],
$$

and

$$
A_{i}=\mathbb{E}\left[\inf _{y \in \mathcal{X}_{i}} f\left(X_{1}, \ldots, X_{i-1}, y, X_{i+1}, \ldots, X_{n}\right)-f(X) \mid \mathcal{F}_{i-1}\right] .
$$

Indeed, since $\sigma\left(X_{i}\right)$ is independent of $\mathcal{F}_{i-1}$ and $\mathcal{G}_{i}$, by the role of independence (Lemma B.6.14)

$$
\begin{aligned}
Z_{i} & =\mathbb{E}\left[f(X) \mid \mathcal{F}_{i}\right] \\
& \leq \mathbb{E}\left[\sup _{y \in \mathcal{X}_{i}} f\left(X_{1}, \ldots, X_{i-1}, y, X_{i+1}, \ldots, X_{n}\right) \mid \mathcal{F}_{i}\right] \\
& =\mathbb{E}\left[\sup _{y \in \mathcal{X}_{i}} f\left(X_{1}, \ldots, X_{i-1}, y, X_{i+1}, \ldots, X_{n}\right) \mid \mathcal{F}_{i-1}, X_{i}\right] \\
& =\mathbb{E}\left[\sup _{y \in \mathcal{X}_{i}} f\left(X_{1}, \ldots, X_{i-1}, y, X_{i+1}, \ldots, X_{n}\right) \mid \mathcal{F}_{i-1}\right],
\end{aligned}
$$

and similarly for the other direction. Moreover, by definition, $B_{i}-A_{i} \leq\left\|D_{i} f\right\|_{\infty}:=$ $c_{i}$. The Azuma-Hoeffding inequality then gives the result.

## Examples

The moral of McDiarmid's inequality is that functions of independent variables that are smooth, in the sense that they do not depend too much on any one of their variables, are concentrated around their mean. Here are some straightforward applications.

Example 3.2.10 (Balls and bins: empty bins). Suppose we throw $m$ balls into $n$ bins independently, uniformly at random. The number of empty bins, $Z_{n, m}$, is centered at

$$
\mathbb{E} Z_{n, m}=n\left(1-\frac{1}{n}\right)^{m}
$$

Writing $Z_{n, m}$ as the sum of indicators $\sum_{i=1}^{n} \mathbf{1}_{B_{i}}$, where $B_{i}$ is the event that bin $i$ is empty, is a natural first attempt at proving concentration around the mean. However there is a problem - the $B_{i} \mathrm{~s}$ are not independent. Indeed, because there
is a fixed number of bins, the event $B_{i}$ intuitively makes the other such events less likely. Instead let $X_{j}$ be the index of the bin in which ball $j$ lands. The $X_{j} \mathrm{~s}$ are independent by construction and, moreover, letting $Z_{n, m}=f\left(X_{1}, \ldots, X_{m}\right)$ we have $\left\|D_{i} f\right\|_{\infty} \leq 1$. Indeed, moving a single ball changes the number of empty bins by at most 1 (if at all). Hence by the method of bounded differences

$$
\mathbb{P}\left[\left|Z_{n, m}-n\left(1-\frac{1}{n}\right)^{m}\right| \geq b \sqrt{m}\right] \leq 2 e^{-2 b^{2}}
$$

Example 3.2.11 (Pattern matching). Let $X=\left(X_{1}, X_{2}, \ldots, X_{n}\right)$ be i.i.d. random variables taking values uniformly at random in a finite set $S$ of size $s=|S|$. Let $a=\left(a_{1}, \ldots, a_{k}\right)$ be a fixed string of elements of $S$. We are interested in the number of occurrences of $a$ as a (consecutive) substring in $X$, which we denote by $N_{n}$. Denote by $E_{i}$ the event that the substring of $X$ starting at $i$ is $a$. Summing over the starting positions and using the linearity of expectation, the mean of $N_{n}$ is

$$
\mathbb{E} N_{n}=\mathbb{E}\left[\sum_{i=1}^{n-k+1} \mathbf{1}_{E_{i}}\right]=(n-k+1)\left(\frac{1}{s}\right)^{k} .
$$

However the $\mathbf{1}_{E_{i}} \mathrm{~s}$ are not independent. So we cannot use a Chernoff bound for Poisson trials (Theorem 2.4.7). Instead we use the fact that $N_{n}=f(X)$ where $\left\|D_{i} f\right\|_{\infty} \leq k$, as each $X_{i}$ appears in at most $k$ substrings of length $k$. By the method of bounded differences, for all $b>0$,

$$
\mathbb{P}\left[\left|N_{n}-(n-k+1)\left(\frac{1}{s}\right)^{k}\right| \geq b k \sqrt{n}\right] \leq 2 e^{-2 b^{2}}
$$

The last two examples are perhaps not surprising in that they involve "sums of weakly independent" indicator variables. One might reasonably expect a sub-Gaussian-type inequality in that case. The next application is more striking and hints at connections to isoperimetric considerations (which we will not explore here).

Example 3.2.12 (Concentration of measure on the hypercube). For $A \subseteq\{0,1\}^{n}$ a subset of the hypercube and $r>0$, we let

$$
A_{r}=\left\{\mathbf{x} \in\{0,1\}^{n}: \inf _{\mathbf{a} \in A}\|\mathbf{x}-\mathbf{a}\|_{1} \leq r\right\}
$$

be the points at $\ell^{1}$ distance at most $r$ from $A$. Fix $\varepsilon \in(0,1 / 2)$ and assume that $|A| \geq \varepsilon 2^{n}$. Let $\lambda_{\varepsilon}$ be such that $e^{-2 \lambda_{\varepsilon}^{2}}=\varepsilon$. The following application of the method of bounded differences indicates that much of the uniform measure on the high-dimensional hypercube lies in a close neighborhood of any such set $A$. This is an example of the concentration of measure phenomenon.

Claim 3.2.13.

$$
r>2 \lambda_{\varepsilon} \sqrt{n} \Longrightarrow\left|A_{r}\right| \geq(1-\varepsilon) 2^{n} .
$$

Proof. Let $X=\left(X_{1}, \ldots, X_{n}\right)$ be uniformly distributed in $\{0,1\}^{n}$. Note that the coordinates are in fact independent. The function

$$
f(\mathbf{x})=\inf _{\mathbf{a} \in A}\|\mathbf{x}-\mathbf{a}\|_{1},
$$

has $\left\|D_{i} f\right\|_{\infty} \leq 1$. Indeed changing one coordinate of $\mathbf{x}$ can increase the $\ell^{1}$ distance to the closest point to x by at most 1 ; in the other direction, if a one-coordinate change were to decrease $f$ by more than 1 , reversing it would produce an increase of that same amount-a contradiction. Hence McDiarmid's inequality gives

$$
\mathbb{P}[\mathbb{E} f(X)-f(X) \geq \beta] \leq \exp \left(-\frac{2 \beta^{2}}{n}\right)
$$

Choosing $\beta=\mathbb{E} f(X)$ and noting that $f(\mathbf{x}) \leq 0$ if and only if $\mathbf{x} \in A$ gives

$$
\mathbb{P}[A] \leq \exp \left(-\frac{2(\mathbb{E} f(X))^{2}}{n}\right)
$$

or, rearranging and using our assumption on $A$,

$$
\mathbb{E} f(X) \leq \sqrt{\frac{1}{2} n \log \frac{1}{\mathbb{P}[A]}} \leq \sqrt{\frac{1}{2} n \log \frac{1}{\varepsilon}}=\lambda_{\varepsilon} \sqrt{n}
$$

By a second application of the method of bounded differences with $\beta=\lambda_{\varepsilon} \sqrt{n}$,

$$
\mathbb{P}\left[f(X) \geq 2 \lambda_{\varepsilon} \sqrt{n}\right] \leq \mathbb{P}[f(X)-\mathbb{E} f(X) \geq b] \leq \exp \left(-\frac{2 \beta^{2}}{n}\right)=\varepsilon
$$

The result follows by observing that, with $r>2 \lambda_{\varepsilon} \sqrt{n}$,

$$
\frac{\left|A_{r}\right|}{2^{n}} \geq \mathbb{P}\left[f(X)<2 \lambda_{\varepsilon} \sqrt{n}\right] \geq 1-\varepsilon .
$$

Claim 3.2.13 is striking for two reasons: 1) the radius $2 \lambda_{\varepsilon} \sqrt{n}$ is much smaller than $n$, the diameter of $\{0,1\}^{n}$; and 2) it applies to any $A$ (such that $|A| \geq \varepsilon 2^{n}$ ). The smallest $r$ such that $\left|A_{r}\right| \geq(1-\varepsilon) 2^{n}$ in general depends on $A$. Here are two extremes.

For $\gamma>0$, let

$$
B(\gamma):=\left\{\mathbf{x} \in\{0,1\}^{n}:\|\mathbf{x}\|_{1} \leq \frac{n}{2}-\gamma \sqrt{\frac{n}{4}}\right\} .
$$

Note that, letting for $Y_{n} \sim B\left(n, \frac{1}{2}\right)$,

$$
\begin{equation*}
\frac{1}{2^{n}}|B(\gamma)|=\sum_{\ell=0}^{\frac{n}{2}-\gamma \sqrt{\frac{n}{4}}}\binom{n}{\ell} 2^{-n}=\mathbb{P}\left[Y_{n} \leq \frac{n}{2}-\gamma \sqrt{\frac{n}{4}}\right] \tag{3.2.8}
\end{equation*}
$$

By the Berry-Esséen theorem (e.g., [Dur10, Theorem 3.4.9]), there is a $C>0$ such that, after rearranging the final quantity in (3.2.8),

$$
\left|\mathbb{P}\left[\frac{Y_{n}-n / 2}{\sqrt{n / 4}} \leq-\gamma\right]-\mathbb{P}[Z \leq-\gamma]\right| \leq \frac{C}{\sqrt{n}},
$$

where $Z \sim N(0,1)$. Let $\varepsilon<\varepsilon^{\prime}<1 / 2$ and let $\gamma_{\varepsilon^{\prime}}$ be such that $\mathbb{P}\left[Z \leq-\gamma_{\varepsilon^{\prime}}\right]=\varepsilon^{\prime}$. Then setting $A:=B\left(\gamma_{\varepsilon^{\prime}}\right)$, for $n$ large enough, we have $|A| \geq \varepsilon 2^{n}$ by (3.2.8). On the other hand, setting $r:=\gamma_{\varepsilon^{\prime}} \sqrt{n / 4}$, we have $A_{r} \subseteq B(0)$, so that $\left|A_{r}\right| \leq \frac{1}{2} 2^{n}<$ $(1-\varepsilon) 2^{n}$. We have shown that $r=\Omega(\sqrt{n})$ is in general required for Claim 3.2.13 to hold.

For an example at the other extreme, assume for simplicity that $N:=\varepsilon 2^{n}$ is an integer. Let $A \subseteq\{0,1\}^{n}$ be constructed as follows: starting from the empty set, add points in $\{0,1\}^{n}$ to $A$ independently, uniformly at random until $|A|=N$. Set $r:=2$. Each point selected in $A$ has $\binom{n}{2}$ points within $\ell^{1}$ distance 2 . By a union bound, the probability that $A_{r}$ does not cover all of $\{0,1\}^{n}$ is at most

$$
\mathbb{P}\left[\left|\{0,1\}^{n} \backslash A_{r}\right|>0\right] \leq \sum_{x \in\{0,1\}^{n}} \mathbb{P}\left[\mathbf{x} \notin A_{r}\right] \leq 2^{n}\left(1-\frac{\binom{n}{2}}{2^{n}}\right)^{\varepsilon 2^{n}} \leq 2^{n} e^{-\varepsilon\binom{n}{2}},
$$

where, in the second inequality, we considered only the first $N$ picks in the construction of $A$ (possibly with repeats), and in the third inequality we used $1-z \leq$ $e^{-z}$ for all $z \in \mathbb{R}$ (see Exercise 1.16). In particular, as $n \rightarrow+\infty$,

$$
\mathbb{P}\left[\left|\{0,1\}^{n} \backslash A_{r}\right|>0\right]<1 .
$$

So for $n$ large enough there is a set $A$ such that $A_{r}=\{0,1\}^{n}$ where $r=2$.

Remark 3.2.14. In fact, it can be shown that sets of the form $\left\{\mathrm{x}:\|\mathrm{x}\|_{1} \leq s\right\}$ have the smallest "expansion" among subsets of $\{0,1\}^{n}$ of the same size, a result known as Harper's vertex isoperimetric theorem. See, for example, [BLM13, Theorem 7.6 and Exercises 7.11-7.13].

### 3.2.3 $\triangleright$ Random graphs: exposure martingale and application to the chromatic number in Erdốs-Rényi model

Exposure martingales In the context of the Erdős-Rényi graph model (Definition 1.2.2), a common way to apply the Azuma-Hoeffding inequality (Theorem 3.2.1) is to introduce an "exposure martingale." Let $G \sim \mathbb{G}_{n, p}$ and let $F$ be any function on graphs such that $\mathbb{E}_{n, p}|F(G)|<+\infty$ for all $n, p$. Choose an arbitrary ordering of the vertices and, for $i=1, \ldots, n$, denote by $H_{i}$ the subgraph of $G$ induced by the first $i$ vertices. Then the filtration $\mathcal{H}_{i}=\sigma\left(H_{1}, \ldots, H_{i}\right)$, $i=1, \ldots, n$, corresponds to adding the vertices of $G$ one at a time (together with their edges to the previous vertices). The Doob martingale

$$
Z_{i}=\mathbb{E}_{n, p}\left[F(G) \mid \mathcal{H}_{i}\right], \quad i=1, \ldots, n,
$$

is known as a vertex exposure martingale. An alternative way to define the filtration is to consider instead the random variables $X_{i}=\left(\mathbf{1}_{\{\{i, j\} \in G\}}: 1 \leq j \leq i\right)$ for $i=2, \ldots, n$. In words, $X_{i}$ is a vector whose entries indicate the status (present or absent) of all potential edges incident with $i$ and a vertex preceding it. Hence, $\mathcal{H}_{i}=$ $\sigma\left(X_{2}, \ldots, X_{i}\right)$ for $i=1, \ldots, n$ (and $\mathcal{H}_{1}$ is trivial as it corresponds to a graph with a single vertex and no edge). This representation has an important property: the $X_{i} \mathrm{~s}$ are independent as they pertain to disjoint subsets of edges. We are then in the setting of the method of bounded differences. Re-writing $F(G)=f\left(X_{1}, \ldots, X_{n}\right)$, the vertex exposure martingale coincides with the martingale (3.2.3) used in that context.

As an example, consider the chromatic number $\chi(G)$, that is, the smallest number of colors needed in a proper coloring of $G$. Define $f_{\chi}\left(X_{1}, \ldots, X_{n}\right):=\chi(G)$. We use the following combinatorial observation to bound $\left\|D_{i} f_{\chi}\right\|_{\infty}$.

Lemma 3.2.15. Altering the status (absent or present) of edges incident to a fixed vertex $v$ changes the chromatic number by at most 1 .

Proof. Altering the status of edges incident to $v$ increases the chromatic number by at most 1 , since in the worst case one can simply use an extra color for $v$. On the other hand, if the chromatic number were to decrease by more than 1 after altering the status of edges incident to $v$, reversing the change and using the previous observation would produce a contradiction.

A fortiori, since $X_{i}$ depends on a subset of the edges incident with vertex $i$, Lemma 3.2.15 implies that $\left\|D_{i} f_{\chi}\right\|_{\infty} \leq 1$. Hence, for all $0<p<1$ and $n$, by an immediate application of the McDiarmid's inequality (Theorem 3.2.9):

## Claim 3.2.16.

$$
\mathbb{P}_{n, p}\left[\left|\chi(G)-\mathbb{E}_{n, p}[\chi(G)]\right| \geq b \sqrt{n-1}\right] \leq 2 e^{-2 b^{2}}
$$

Edge exposure martingales can be defined in a manner similar to the vertex case: reveal the edges one at a time in an arbitrary order. By Lemma 3.2.15, the corresponding function also satisfies the same $\ell^{\infty}$ bound. Observe however that, for the chromatic number, edge exposure results in a much weaker bound as the $\Theta\left(n^{2}\right)$ random variables produce only a linear in $n$ deviation for the same tail probability. (The reader may want to ponder the apparent paradox: using a larger number of independent variables seemingly leads to weaker concentration in this case.)

Remark 3.2.17. Note that Claim 3.2.16 tells us nothing about the expectation of $\chi(G)$. It turns out that, up to logarithmic factors, $\mathbb{E}_{n, p_{n}}[\chi(G)]$ is of order $n p_{n}$ when $p_{n} \sim n^{-\alpha}$ for some $0<\alpha<1$. We will not prove this result here. See the "Bibliographic remarks" at the end of this chapter for more on the chromatic number of Erdös-Rényi graphs.
$\chi(G)$ is concentrated on few values Much stronger concentration results can be obtained: when $p_{n}=n^{-\alpha}$ with $\alpha>\frac{1}{2}$, the chromatic number $\chi(G)$ is in fact concentrated on two values! We give a partial result along those lines which illustrates a less straightforward choice of martingale in the Azuma-Hoeffding inequality (Theorem 3.2.1).

Claim 3.2.18. Let $p_{n}=n^{-\alpha}$ with $\alpha>\frac{5}{6}$ and let $G_{n} \sim \mathbb{G}_{n, p_{n}}$. Then for any $\varepsilon>0$ there is $\varphi_{n}:=\varphi_{n}(\alpha, \varepsilon)$ such that

$$
\mathbb{P}_{n, p_{n}}\left[\varphi_{n} \leq \chi\left(G_{n}\right) \leq \varphi_{n}+3\right] \geq 1-\varepsilon,
$$

for all $n$ large enough.
Proof. We consider the following martingale. Let $\varphi_{n}$ be the smallest integer such that

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\chi\left(G_{n}\right) \leq \varphi_{n}\right]>\frac{\varepsilon}{3} . \tag{3.2.9}
\end{equation*}
$$

Let $F_{n}\left(G_{n}\right)$ be the minimal size of a set of vertices, $U$, in $G_{n}$ such that $G_{n} \backslash U$ is $\varphi_{n}$-colorable. Let $\left(Z_{i}\right)$ be the vertex exposure martingale associated to the quantity $F_{n}\left(G_{n}\right)$. The proof proceeds in two steps: we show that 1$)$ all but $O(\sqrt{n})$ vertices can be $\varphi_{n}$-colored and 2 ) the remaining vertices can be colored using 3 additional colors. See Figure 3.2.3 for an illustration of the proof strategy.

We claim that $\left(Z_{i}\right)$ has increments bounded by 1 .


Figure 3.1: All but $O(\sqrt{n})$ vertices are colored using $\varphi_{n}$ colors. The remaining vertices are colored using 3 additional colors.

Lemma 3.2.19. Changing the edges incident to a single vertex can change $F_{n}$ by at most 1.

Proof. Changing the edges incident to $v$ can increase $F_{n}$ by at most 1 . Indeed, if $F_{n}$ increases after such a change, it must be that $v \notin U$ since in the other case the edges incident with $v$ would not affect the colorability of $G_{n} \backslash U$-present or not. So we can add $v$ to $U$ and restore colorability. On the other hand, if $F_{n}$ were to decrease by more than 1 , reversing the change and using the previous observation would give a contradiction.

Choose $b_{\varepsilon}$ such that $e^{-b_{\varepsilon}^{2} / 2}=\frac{\varepsilon}{3}$. Then, applying the Azuma-Hoeffding inequality to $\left(-Z_{i}\right)$,

$$
\mathbb{P}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)-\mathbb{E}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)\right] \leq-b_{\varepsilon} \sqrt{n-1}\right] \leq \frac{\varepsilon}{3},
$$

which, since $\mathbb{P}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)=0\right]=\mathbb{P}_{n, p_{n}}\left[\chi\left(G_{n}\right) \leq \varphi_{n}\right]>\frac{\varepsilon}{3}$, implies that

$$
\mathbb{E}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)\right] \leq b_{\varepsilon} \sqrt{n-1}
$$

Applying the Azuma-Hoeffding inequality to $\left(Z_{i}\right)$ gives

$$
\begin{align*}
\mathbb{P}_{n, p_{n}} & {\left[F_{n}\left(G_{n}\right) \geq 2 b_{\varepsilon} \sqrt{n-1}\right] } \\
& \leq \mathbb{P}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)-\mathbb{E}_{n, p_{n}}\left[F_{n}\left(G_{n}\right)\right] \geq b_{\varepsilon} \sqrt{n-1}\right] \\
& \leq \frac{\varepsilon}{3} . \tag{3.2.10}
\end{align*}
$$

So with probability at least $1-\frac{\varepsilon}{3}$, we can color all vertices but $2 b_{\varepsilon} \sqrt{n-1}$ using $\varphi_{n}$ colors. Let $U$ be the remaining uncolored vertices.

We claim that, with high probability, we can color the vertices in $U$ using at most 3 extra colors.
Lemma 3.2.20. Fix $c>0, \alpha>\frac{5}{6}$ and $\varepsilon>0$. Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ with $p_{n}=n^{-\alpha}$. For all n large enough,

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\text { every subset of } c \sqrt{n} \text { vertices of } G_{n} \text { can be } 3 \text {-colored }\right]>1-\frac{\varepsilon}{3} \text {. } \tag{3.2.11}
\end{equation*}
$$

Proof. We use the first moment method (Theorem 2.2.6). We refer to a subset of vertices that is not 3 -colorable but such that all of its subsets are as minimal, non 3 -colorable. Let $Y_{n}$ be the number of such subsets of size at most $c \sqrt{n}$ in $G_{n}$.

Any minimal, non 3 -colorable subset $W$ must have degree at least 3 . Indeed suppose that $w \in W$ has degree less than 3 . Then $W \backslash\{w\}$ is 3 -colorable by definition. But, since $w$ has fewer than 3 neighbors, it can also be properly colored
without adding a new color-a contradiction. In particular, the subgraph of $G_{n}$ induced by $W$ must have at least $\frac{3}{2}|W|$ edges. Hence, the probability that a subset of vertices of $G_{n}$ of size $\ell$ is minimal, non 3-colorable is at most

$$
\left(\begin{array}{c}
\ell \\
2 \\
\frac{3 \ell}{2}
\end{array}\right) p_{n}^{\frac{3 \ell}{2}},
$$

by a union bound over all subsets of edges of size $\frac{3 \ell}{2}$.
By the first moment method, by the binomial bounds $\binom{n}{\ell} \leq\left(\frac{e n}{\ell}\right)^{\ell}$ (see Appendix A) and $\binom{\ell}{2} \leq \ell^{2} / 2$, for some $c^{\prime} \in(0,+\infty)$

$$
\left.\begin{array}{rl}
\mathbb{P}_{n, p_{n}}\left[Y_{n}>0\right] & \leq \mathbb{E}_{n, p_{n}} Y_{n} \\
& \leq \sum_{\ell=4}^{c \sqrt{n}}\binom{n}{\ell}\binom{\ell}{\ell} \\
\frac{3 \ell}{2}
\end{array}\right) p_{n}^{\frac{3 \ell}{2}} .
$$

as $n \rightarrow+\infty$, where we used that $\frac{5}{4}-\frac{3 \alpha}{2}<\frac{5}{4}-\frac{5}{4}=0$ when $\alpha>\frac{5}{6}$ so that the geometric series is dominated by its first term. Therefore for $n$ large enough $\mathbb{P}_{n, p_{n}}\left[Y_{n}>0\right] \leq \varepsilon / 3$, concluding the proof.

By the choice of $\varphi_{n}$ in (3.2.9),

$$
\mathbb{P}_{n, p_{n}}\left[\chi\left(G_{n}\right)<\varphi_{n}\right] \leq \frac{\varepsilon}{3} .
$$

By (3.2.10) and (3.2.11) with $c=2 b_{\varepsilon}$,

$$
\mathbb{P}_{n, p_{n}}\left[\chi\left(G_{n}\right)>\varphi_{n}+3\right] \leq \frac{2 \varepsilon}{3} .
$$

So, overall,

$$
\mathbb{P}_{n, p_{n}}\left[\varphi_{n} \leq \chi\left(G_{n}\right) \leq \varphi_{n}+3\right] \geq 1-\varepsilon .
$$

That concludes the proof.

### 3.2.4 $\triangleright$ Random graphs: degree sequence of preferential attachment graphs

Let $\left(G_{t}\right)_{t \geq 1} \sim \mathrm{PA}_{1}$ be a preferential attachment graph (Definition 1.2.3). A key feature of such graphs is a power-law degree sequence: the fraction of vertices with degree $d$ behaves like $\propto d^{-\alpha}$ for some $\alpha>0$, that is, it has a fat tail. Recall that we restrict ourselves to the tree case. In contrast, we will show in Section 4.1.4 that a (sparse) Erdős-Rényi random graph has an asymptotically Poisson-distributed degree sequence, and therefore a much thinner tail.

Power law degree sequence Let $D_{i}(t)$ be the degree of the $i$-th vertex in $G_{t}$, denoted $v_{i}$, and let

$$
N_{d}(t):=\sum_{i=0}^{t} \mathbf{1}_{\left\{D_{i}(t)=d\right\}},
$$

be the number of vertices of degree $d$ in $G_{t}$. By construction $N_{0}(t)=0$ for all $t$. Define the sequence

$$
\begin{equation*}
f_{d}:=\frac{4}{d(d+1)(d+2)}, \quad d \geq 1 . \tag{3.2.12}
\end{equation*}
$$

Our main claim is:
Claim 3.2.21.

$$
\frac{1}{t} N_{d}(t) \rightarrow_{\mathrm{p}} f_{d}, \quad \forall d \geq 1
$$

Proof. The claim is immediately implied by the following lemmas.
Lemma 3.2.22 (Convergence of the mean).

$$
\frac{1}{t} \mathbb{E} N_{d}(t) \rightarrow f_{d}, \quad \forall d \geq 1
$$

Lemma 3.2.23 (Concentration around the mean). For any $\delta>0$,

$$
\mathbb{P}\left[\left|\frac{1}{t} N_{d}(t)-\frac{1}{t} \mathbb{E} N_{d}(t)\right| \geq \sqrt{\frac{2 \log \delta^{-1}}{t}}\right] \leq 2 \delta, \quad \forall d \geq 1, \forall t .
$$

An alternative representation of the process We start with the proof of Lemma 3.2.23, which is an application of the method of bounded differences.

Proof of Lemma 3.2.23. In our description of the preferential attachment process, the random choices made at each time depend in a seemingly complicated way on previous choices. In order to establish concentration of the process around its mean, we introduce a clever, alternative construction which has the advantage that it involves independent choices.

We start with a single vertex $v_{0}$. At time 1 , we add a single vertex $v_{1}$ and an edge $e_{1}$ connecting $v_{0}$ and $v_{1}$. For bookkeeping, we orient edges away from the vertex of higher time index (but we ignore the orientations in the output). For a directed edge $(i, j)$, we refer to $i$ as its tail and $j$ as its head. For all $s \geq 2$, let $X_{s}$ be an independent, uniformly chosen edge extremity among the edges in $G_{s-1}$, that is, pick a uniform element in

$$
\mathcal{X}_{s}:=\{(1, \text { tail }),(1, \text { head }), \ldots,(s-1, \text { tail }),(s-1, \text { head })\} .
$$

To form $G_{s}$, attach a new edge $e_{s}$ to the vertex of $G_{s-1}$ corresponding to $X_{s}$. A vertex of degree $d^{\prime}$ in $G_{s-1}$ is selected with probability $\frac{d^{\prime}}{2(s-1)}$, as it should. Note that $X_{s}$ can be picked in advance independently of the sequence $\left(G_{s^{\prime}}\right)_{s^{\prime}<s}$. For instance, if $x_{2}=(1$, head $), x_{3}=(2$, tail $)$ and $x_{4}=(3$, head $)$, the graph obtained at time 4 is depicted in Figure 3.2.

We claim that $N_{d}(t)=: h\left(X_{2}, \ldots, X_{t}\right)$ as a function of $X_{2}, \ldots, X_{t}$ satisfies $\left\|D_{i} h\right\|_{\infty} \leq 2$. Indeed let $\left(x_{2}, \ldots, x_{t}\right)$ be a realization of $\left(X_{2}, \ldots, X_{t}\right)$ and let $y \in \mathcal{X}_{s}$ with $y \neq x_{s}$. Replacing $x_{s}=(i$, end $)$ with $y=\left(j\right.$, end $\left.^{\prime}\right)$ where $i, j \in$ $\{1, \ldots, s-1\}$ and end, end ${ }^{\prime} \in\{$ tail, head $\}$ has the effect of redirecting the head of edge $e_{s}$ from the end of $e_{i}$ to the end' of $e_{j}$. This redirection also brings along with it the heads of all other edges associated with the choice ( $s$, head). But, crucially, those changes only affect the degrees of the vertices ( $i$, end) and $\left(j\right.$, end $\left.^{\prime}\right)$ in the original graph. Hence the number of vertices with degree $d$ changes by at most 2, as claimed. For instance, returning to the example of Figure 3.2. If we replace $x_{3}=(2$, tail $)$ with $y=(1$, tail $)$, one obtains the graph in Figure 3.3. Note that only the degrees of vertices $v_{1}$ and $v_{2}$ are affected by this change.

By McDiarmid's inequality (Theorem 3.2.9), for all $\beta>0$,

$$
\mathbb{P}\left[\left|N_{d}(t)-\mathbb{E} N_{d}(t)\right| \geq \beta\right] \leq 2 \exp \left(-\frac{2 \beta^{2}}{(2)^{2}(t-1)}\right)
$$

which, choosing $\beta=\sqrt{2 t \log \delta^{-1}}$, we can rewrite as

$$
\mathbb{P}\left[\left|\frac{1}{t} N_{d}(t)-\frac{1}{t} \mathbb{E} N_{d}(t)\right| \geq \sqrt{\frac{2 \log \delta^{-1}}{t}}\right] \leq 2 \delta .
$$

That concludes the proof of the lemma.


Figure 3.2: Graph obtained when $x_{2}=(1$, head $), x_{3}=(2$, tail $)$ and $x_{4}=$ (3, head).


Figure 3.3: Substituting $x_{3}=(2$, tail) with $y=(1$, tail $)$ in the example of Figure 3.2 has the effect of replacing the dashed edges with the dotted edges. Note that only the degrees of vertices $v_{1}$ and $v_{2}$ are affected by this change.

Dynamics of the mean Once again the method of bounded differences tells us nothing about the mean, which must be analyzed by other means. The proof of Lemma 3.2.22 does not rely on the Azuma-Hoeffding inequality but is given for completeness (and may be skipped).

Proof of Lemma 3.2.22. The idea of the proof is to derive a recursion for $f_{d}$ by considering the evolution of $\mathbb{E} N_{d}(t)$ and taking a limit as $t \rightarrow+\infty$. Let $d \geq 1$. Observe that $\mathbb{E} N_{d}(t)=0$ for $t \leq d-1$ since we need at least $d$ edges to have a degree- $d$ vertex. Moreover, by the description of the preferential attachment process, the following recursion holds for $t \geq d-1$

$$
\begin{equation*}
\mathbb{E} N_{d}(t+1)-\mathbb{E} N_{d}(t)=\underbrace{\frac{d-1}{2 t} \mathbb{E} N_{d-1}(t)}_{\text {(a) }}-\underbrace{\frac{d}{2 t} \mathbb{E} N_{d}(t)}_{\text {(b) }}+\underbrace{\mathbf{1}_{\{d=1\}}}_{\text {(c) }} . \tag{3.2.13}
\end{equation*}
$$

Indeed: (a) $N_{d}(t)$ increases by 1 if a vertex of degree $d-1$ is picked, an event of probability $\frac{d-1}{2 t} N_{d-1}(t)$ because the sum of degrees at time $t$ is twice the number of edges (i.e., $t$ ); (b) $N_{d}(t)$ decreases by 1 if a vertex of degree $d$ is picked, an event of probability $\frac{d}{2 t} N_{d}(t)$; and (c) the last term comes from the fact that the new vertex always has degree 1 . We rewrite (3.2.13) as

$$
\begin{align*}
\mathbb{E} N_{d}(t+1) & =\mathbb{E} N_{d}(t)+\frac{d-1}{2 t} \mathbb{E} N_{d-1}(t)-\frac{d}{2 t} \mathbb{E} N_{d}(t)+\mathbf{1}_{\{d=1\}} \\
& =\left(1-\frac{d / 2}{t}\right) \mathbb{E} N_{d}(t)+\left\{\frac{d-1}{2}\left[\frac{1}{t} \mathbb{E} N_{d-1}(t)\right]+\mathbf{1}_{\{d=1\}}\right\} \\
& =:\left(1-\frac{d / 2}{t}\right) \mathbb{E} N_{d}(t)+g_{d}(t), \tag{3.2.14}
\end{align*}
$$

where $g_{d}(t)$ is defined as the expression in curly brackets on the second line. We will not solve this recursion explicitly. Instead we seek to analyze its asymptotics, specifically we show that $\frac{1}{t} \mathbb{E} N_{d}(t) \rightarrow f_{d}$.

The key is to notice that the expression for $\mathbb{E} N_{d}(t+1)$ depends on $\frac{1}{t} \mathbb{E} N_{d-1}(t)$ so we work by induction on $d$. Because of the form of the recursion, the following technical lemma is what we need to proceed.

Lemma 3.2.24. Let $f, g$ be nonnegative functions of $t \in \mathbb{N}$ satisfying the following recursion

$$
f(t+1)=\left(1-\frac{\alpha}{t}\right) f(t)+g(t), \quad \forall t \geq t_{0}
$$

with $g(t) \rightarrow g \in[0,+\infty)$ as $t \rightarrow+\infty$, and where $\alpha>0, t_{0} \geq 2 \alpha, f\left(t_{0}\right) \geq 0$ are constants. Then

$$
\frac{1}{t} f(t) \rightarrow \frac{g}{1+\alpha}
$$

as $t \rightarrow+\infty$.

The proof of this lemma is given after the proof of Claim 3.2.21. We first conclude the proof of Lemma 3.2.22. First let $d=1$. In that case, $g_{1}(t)=g_{1}:=1$, $\alpha:=1 / 2$, and $t_{0}:=1$. By Lemma 3.2.24,

$$
\frac{1}{t} \mathbb{E} N_{1}(t) \rightarrow \frac{1}{1+1 / 2}=\frac{2}{3}=f_{1} .
$$

Assuming by induction that $\frac{1}{t} \mathbb{E} N_{d-1}(t) \rightarrow f_{d-1}$ we get

$$
g_{d}(t) \rightarrow g_{d}:=\frac{d-1}{2} f_{d-1},
$$

as $t \rightarrow+\infty$. Using Lemma 3.2.24 with $\alpha:=d / 2$ and $t_{0}:=d-1$, we obtain

$$
\frac{1}{t} \mathbb{E} N_{d}(t) \rightarrow \frac{1}{1+d / 2}\left[\frac{d-1}{2} f_{d-1}\right]=\frac{d-1}{d+2} \cdot \frac{4}{(d-1) d(d+1)}=f_{d}
$$

That concludes the proof of Lemma 3.2.22.
To prove Claim 3.2.21, we combine Lemmas 3.2.22 and 3.2.23. Fix any $d, \delta, \varepsilon>$ 0 . Choose $t^{\prime}$ large enough that for all $t \geq t^{\prime}$

$$
\max \left\{\left|\frac{1}{t} \mathbb{E} N_{d}(t)-f_{d}\right|, \sqrt{\frac{2 \log \delta^{-1}}{t}}\right\} \leq \varepsilon
$$

Then

$$
\mathbb{P}\left[\left|\frac{1}{t} N_{d}(t)-f_{d}\right| \geq 2 \varepsilon\right] \leq 2 \delta,
$$

for all $t \geq t^{\prime}$. That proves convergence in probability.

Proof of the technical lemma It remains to prove Lemma 3.2.24.
Proof of Lemma 3.2.24. By induction on $t$, we have

$$
\begin{aligned}
f(t+1) & =\left(1-\frac{\alpha}{t}\right) f(t)+g(t) \\
& =\left(1-\frac{\alpha}{t}\right)\left[\left(1-\frac{\alpha}{t-1}\right) f(t-1)+g(t-1)\right]+g(t) \\
& =\left(1-\frac{\alpha}{t}\right) g(t-1)+g(t)+\left(1-\frac{\alpha}{t}\right)\left(1-\frac{\alpha}{t-1}\right) f(t-1) \\
& =\cdots \\
& =\sum_{i=0}^{t-t_{0}} g(t-i) \prod_{j=0}^{i-1}\left(1-\frac{\alpha}{t-j}\right)+f\left(t_{0}\right) \prod_{j=0}^{t-t_{0}}\left(1-\frac{\alpha}{t-j}\right)
\end{aligned}
$$

or

$$
\begin{equation*}
f(t+1)=\sum_{s=t_{0}}^{t} g(s) \prod_{r=s+1}^{t}\left(1-\frac{\alpha}{r}\right)+f\left(t_{0}\right) \prod_{r=t_{0}}^{t}\left(1-\frac{\alpha}{r}\right), \tag{3.2.15}
\end{equation*}
$$

where empty products are equal to 1 . To guess the limit note that, for large $s, g(s)$ is roughly constant and that the product in the first term behaves like

$$
\exp \left(-\sum_{r=s+1}^{t} \frac{\alpha}{r}\right) \approx \exp (-\alpha(\log t-\log s)) \approx \frac{s^{\alpha}}{t^{\alpha}}
$$

So approximating the sum by an integral we get that $f(t+1) \approx \frac{g t}{\alpha+1}$, which is indeed consistent with the claim.

Formally, we use that there is a constant $\gamma=0.577 \ldots$ such that (see e.g. [LL10, Lemma 12.1.3])

$$
\sum_{\ell=1}^{m} \frac{1}{\ell}=\log m+\gamma+\Theta\left(m^{-1}\right)
$$

and that by a Taylor expansion, for $|z| \leq 1 / 2$,

$$
\log (1-z)=-z+\Theta\left(z^{2}\right)
$$

Fix $\eta>0$ small and take $t$ large enough that $\eta t>2 \alpha$ and $|g(s)-g|<\eta$ for all $s \geq \eta t$. Then, for $s+1 \geq t_{0}$,

$$
\begin{aligned}
\sum_{r=s+1}^{t} \log \left(1-\frac{\alpha}{r}\right) & =-\sum_{r=s+1}^{t}\left\{\frac{\alpha}{r}+\Theta\left(r^{-2}\right)\right\} \\
& =-\alpha(\log t-\log s)+\Theta\left(s^{-1}\right)
\end{aligned}
$$

so, taking exponentials,

$$
\prod_{r=s+1}^{t}\left(1-\frac{\alpha}{r}\right)=\frac{s^{\alpha}}{t^{\alpha}}\left(1+\Theta\left(s^{-1}\right)\right)
$$

Hence

$$
\frac{1}{t} f\left(t_{0}\right) \prod_{r=t_{0}}^{t}\left(1-\frac{\alpha}{r}\right)=\frac{t_{0}^{\alpha}}{t^{\alpha+1}}\left(1+\Theta\left(t_{0}^{-1}\right)\right) \rightarrow 0
$$

as $t \rightarrow+\infty$. Moreover

$$
\begin{aligned}
\frac{1}{t} \sum_{s=\eta t}^{t} g(s) \prod_{r=s+1}^{t}\left(1-\frac{\alpha}{r}\right) & \leq \frac{1}{t} \sum_{s=\eta t}^{t}(g+\eta) \frac{s^{\alpha}}{t^{\alpha}}\left(1+\Theta\left(s^{-1}\right)\right) \\
& \leq O(\eta)+\left(1+\Theta\left(t^{-1}\right)\right) \frac{g}{t^{\alpha+1}} \sum_{s=\eta t}^{t} s^{\alpha} \\
& \leq O(\eta)+\left(1+\Theta\left(t^{-1}\right)\right) \frac{g}{t^{\alpha+1}} \frac{(t+1)^{\alpha+1}}{\alpha+1} \\
& \rightarrow O(\eta)+\frac{g}{\alpha+1},
\end{aligned}
$$

where we bounded the sum on the second line by an integral. Similarly,

$$
\begin{aligned}
\frac{1}{t} \sum_{s=t_{0}}^{\eta t-1} g(s) \prod_{r=s+1}^{t}\left(1-\frac{\alpha}{r}\right) & \leq \frac{1}{t} \sum_{s=t_{0}}^{\eta t-1}(g+\eta) \frac{s^{\alpha}}{t^{\alpha}}\left(1+\Theta\left(s^{-1}\right)\right) \\
& \leq \frac{\eta t}{t}(g+\eta) \frac{(\eta t)^{\alpha}}{t^{\alpha}}\left(1+\Theta\left(t_{0}^{-1}\right)\right) \\
& \rightarrow O\left(\eta^{\alpha+1}\right) .
\end{aligned}
$$

Plugging these inequalities back into (3.2.15), we get

$$
\limsup _{t} \frac{1}{t} f(t+1) \leq \frac{g}{1+\alpha}+O(\eta)
$$

A similar inequality holds in the other direction. Letting $\eta \rightarrow 0$ concludes the proof.

Remark 3.2.25. A more quantitative result (uniform in $t$ and d) can be derived. See, for example, [vdH17, Sections 8.5, 8.6]. See the same reference for a generalization beyond trees.

### 3.2.5 $\triangleright$ Data science: stochastic bandits and the slicing method

In this section, we consider an application of the maximal Azuma-Hoeffding inequality (Theorem 3.2.1) to (multi-armed) bandit problems. These are meant as a simple model of sequential decision making with limited information where a fundamental issue is trading off between exploitation of actions that have done well in the past and exploration of actions that might perform better in the future. A typical application is online advertising, where one must decide which advertisement to display to the next visitor to a website.

In the simplest version of the (two-arm) stochastic bandit problem, there are two unknown reward distributions $\nu_{1}, \nu_{2}$ over $[0,1]$ with respective means $\mu_{1} \neq \mu_{2}$. At each time $t=1, \ldots, n$, we request an independent sample from $\nu_{I_{t}}$, where we are free to choose $I_{t} \in\{1,2\}$ based on past choices and observed rewards $\left\{\left(I_{s}, Z_{s}\right)\right\}_{s<t}$. This will be referred to as pulling $\operatorname{arm} I_{t}$. We then observe the reward $Z_{t} \sim \nu_{I_{t}}$. Letting $\mu^{*}:=\mu_{1} \vee \mu_{2}$, our goal is to minimize

$$
\begin{equation*}
\bar{R}_{n}=n \mu^{*}-\mathbb{E}\left[\sum_{t=1}^{n} \mu_{I_{t}}\right], \tag{3.2.16}
\end{equation*}
$$

which is known as the pseudo-regret. That is, we seek to make choices $\left(I_{t}\right)_{t=1}^{n}$ that minimize the difference between the best achievable cumulative mean reward and the expected cumulative mean reward from our decisions. Note that the expectation in (3.2.16) is taken over the choices $\left(I_{t}\right)_{t=1}^{n}$, which themselves depend on the random rewards $\left(Z_{s}\right)_{t=1}^{n}$. As indicated above, because $\nu_{1}$ and $\nu_{2}$ are unknown, there is a fundamental friction between exploiting the arm that has done best in the past and exploring further the other arm, which might perform better in the future.

One general approach that has proved effective in this type of problem is known as optimism in the face of uncertainty. Roughly speaking, we construct a set of plausible environments (in our case, the means of the reward distributions) that are consistent with observed data; then we make an optimal decision assuming that the true environment is the most favorable among them. A concrete implementation of this principle is the Upper Confidence Bound (UCB) algorithm, which we now describe. In words, we use a concentration inequality to build a confidence interval for each reward mean, and then we pick the arm with highest upper bound.

## UCB algorithm

To state the algorithm formally, we will need some notation. For $i=1,2$, let $T_{i}(t)$ be the number of times arm $i$ is pulled up to time $t$

$$
T_{i}(t)=\sum_{s \leq t} \mathbf{1}\left\{I_{s}=i\right\},
$$

and let $X_{i, s}, s=1, \ldots, n$, be i.i.d. samples from $\nu_{i}$. Assume that the reward at time $t$ is

$$
Z_{t}= \begin{cases}X_{1, T_{1}(t-1)+1} & \text { if } I_{t}=1 \\ X_{2, T_{2}(t-1)+1} & \text { otherwise }\end{cases}
$$

In other words, $X_{i, s}$ is the $s$-th observed reward from arm $i$. Let $\hat{\mu}_{i, s}$ be the sample mean of the observed rewards after pulling $s$ times on arm $i$

$$
\hat{\mu}_{i, s}=\frac{1}{s} \sum_{r \leq s} X_{i, r} .
$$

Since the $X_{i, s}$ s are independent and $[0,1]$-valued by assumption, by Hoeffding's inequality (Theorem 2.4.10), for any $\beta>0$

$$
\mathbb{P}\left[\hat{\mu}_{i, s}-\mu_{i} \geq \beta\right] \vee \mathbb{P}\left[\mu_{i}-\hat{\mu}_{i, s} \geq \beta\right] \leq \exp \left(-2 s \beta^{2}\right)
$$

The right-hand side can be made $\leq \delta$ provided

$$
\beta \geq \sqrt{\frac{\log \delta^{-1}}{2 s}}:=\mathrm{H}(s, \delta)
$$

We are now ready to state the $\alpha$-UCB algorithm, where $\alpha>1$ is the exploration parameter. At each time $t$, we pick

$$
I_{t} \in \underset{i=1,2}{\arg \max }\left\{\hat{\mu}_{i, T_{i}(t-1)}+\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right)\right\} .
$$

The argument above implies that the true mean $\mu_{i}$ has probability less than $1 / t^{\alpha^{2}}$ of being higher than $\hat{\mu}_{i, T_{i}(t-1)}+\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right)$. The algorithm makes an "optimistic" decision: it chooses the higher of the two values.

The following theorem shows that UCB achieves a pseudo-regret of the order of $O(\log n)$. Define $\Delta_{i}=\mu^{*}-\mu_{i}$ and $\Delta_{*}=\Delta_{1} \vee \Delta_{2}$.

Theorem 3.2.26 (Pseudo-regret of UCB). In the two-arm stochastic bandit problem where the rewards are in $[0,1]$ with distinct means, $\alpha-U C B$ with $\alpha>1$ achieves

$$
\bar{R}_{n} \leq \frac{2 \alpha^{2}}{\Delta_{*}} \log n+\Delta_{*} C_{\alpha}
$$

for some constant $C_{\alpha} \in(0,+\infty)$ depending only on $\alpha$.
This bound should not come entirely as a surprise. Indeed a simple, alternative approach to UCB is to (1) first pull each arm $m_{n}=o(n)$ times and then (2) use the arm with largest estimated mean for the remainder. Assuming there is a known lower bound on $\Delta_{*}$, then Hoeffding's inequality (Theorem 2.4.10) guarantees that $m_{n}$ can be chosen of the order of $\frac{1}{\Delta_{*}^{2}} \log n$ to identify the largest mean with probability $1-1 / n$. Because the rewards are bounded by 1 , accounting for the contribution of the first phase and the probability of failure in the second phase, one gets a pseudo-regret of the order of $\Delta^{*} \frac{1}{\Delta_{*}^{2}} \log n+\frac{1}{n} \Delta_{*} n \approx \frac{1}{\Delta_{*}} \log n$. The UCB strategy, on the other hand, elegantly adapts to the gap $\Delta_{*}$ and the horizon $n$.

## Analysis of the UCB algorithm

We break down the proof into a sequence of lemmas. We first rewrite the pseudoregret as

$$
\begin{align*}
\bar{R}_{n} & =n \mu^{*}-\mathbb{E}\left[\sum_{t=1}^{n} \mu_{I_{t}}\right] \\
& =\mathbb{E}\left[\sum_{t=1}^{n}\left(\mu^{*}-\mu_{I_{t}}\right)\right] \\
& =\mathbb{E}\left[\sum_{t=1}^{n} \sum_{i=1,2} \mathbf{1}\left\{I_{t}=i\right\} \Delta_{i}\right] \\
& =\sum_{i=1,2} \Delta_{i} \mathbb{E}\left[T_{i}(n)\right] . \tag{3.2.17}
\end{align*}
$$

Hence the problem boils down to bounding $\mathbb{E}\left[T_{i}(n)\right]$, the expected number of times that arm $i$ is pulled. Note that $T_{i}(n)$ is a complicated function of the observations. To analyze it, we will use the following sufficient condition. Let $i^{*}$ be the optimal arm, that is, the one that achieves $\mu^{*}$. Intuitively, if arm $i \neq i^{*}$ is pulled, it is because: either our upper estimate of $\mu_{i^{*}}$ happens to be low or our lower estimate of $\mu_{i}$ happens to be high (i.e., our concentration inequality failed); or there is too much uncertainty in our estimate of $\mu_{i}$ (i.e., we haven't pulled arm $i$ enough).

Lemma 3.2.27. Under the $\alpha-U C B$ strategy, if arm $i \neq i^{*}$ is pulled at time $t$ then at least one of the following events hold:

$$
\begin{gather*}
\mathcal{E}_{t, 1}=\left\{\hat{\mu}_{i^{*}, T_{i^{*}}(t-1)}+\alpha \mathrm{H}\left(T_{i^{*}}(t-1), 1 / t\right) \leq \mu^{*}\right\},  \tag{3.2.18}\\
\mathcal{E}_{t, 2}=\left\{\hat{\mu}_{i, T_{i}(t-1)}-\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right)>\mu_{i}\right\},  \tag{3.2.19}\\
\mathcal{E}_{t, 3}=\left\{\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right)>\frac{\Delta_{i}}{2}\right\} . \tag{3.2.20}
\end{gather*}
$$

Proof. We argue by contradiction. Assume all the conditions above are false. Then

$$
\begin{aligned}
\hat{\mu}_{i^{*}, T_{i^{*}}(t-1)}+\alpha \mathrm{H}\left(T_{i^{*}}(t-1), 1 / t\right) & >\mu^{*} \\
& =\mu_{i}+\Delta_{i} \\
& \geq \hat{\mu}_{i, T_{i}(t-1)}+\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right) .
\end{aligned}
$$

That implies that arm $i$ would not be chosen.

We first deal with $\mathcal{E}_{t, 3}$. Let

$$
u_{n}=\frac{2 \alpha^{2} \log n}{\Delta_{*}^{2}}
$$

Using the condition in Lemma 3.2.27, we get the following bound on $\mathbb{E}\left[T_{i}(n)\right]$.
Lemma 3.2.28. Under the $\alpha$-UCB strategy, for $i \neq i^{*}$,

$$
\mathbb{E}\left[T_{i}(n)\right] \leq u_{n}+\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 1}\right]+\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 2}\right]
$$

Proof. For $i \neq i^{*}$, by definition of $T_{i}(n)$,

$$
\left.\begin{array}{rl}
\mathbb{E}\left[T_{i}(n)\right] & =\mathbb{E}\left[\sum_{t=1}^{n} \mathbf{1}_{\left\{I_{t}=i\right\}}\right] \\
& \leq \mathbb{E}\left[\sum_{t=1}^{n}\left[\mathbf{1}_{\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 1}}+\mathbf{1}_{\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 2}}+\mathbf{1}_{\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 3}}\right]\right.
\end{array}\right],
$$

where we used that by Lemma 3.2.27

$$
\left\{I_{t}=i\right\} \subseteq \mathcal{E}_{t, 1} \cup \mathcal{E}_{t, 2} \cup \mathcal{E}_{t, 3} .
$$

The condition in $\mathcal{E}_{t, 3}$ can be written equivalently as

$$
\alpha \sqrt{\frac{\log t}{2 T_{i}(t-1)}}>\frac{\Delta_{i}}{2} \Longleftrightarrow T_{i}(t-1)<\frac{2 \alpha^{2} \log t}{\Delta_{i}^{2}} .
$$

In particular, for all $t \leq n$, the event $\mathcal{E}_{t, 3}$ implies that $T_{i}(t-1)<u_{n}$. As a result, since $T_{i}(t)=T_{i}(t-1)+1$ whenever $I_{t}=i$, the event $\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 3}$ can occur at most $u_{n}$ times and

$$
\begin{aligned}
\mathbb{E}\left[T_{i}(n)\right] & \leq u_{n}+\mathbb{E}\left[\sum_{t=1}^{n}\left[\mathbf{1}_{\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 1}}+\mathbf{1}_{\left\{I_{t}=i\right\} \cap \mathcal{E}_{t, 2}}\right]\right. \\
& \leq u_{n}+\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 1}\right]+\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 2}\right]
\end{aligned}
$$

which proves the claim.

It remains to bound $\mathbb{P}\left[\mathcal{E}_{t, 1}\right]$ and $\mathbb{P}\left[\mathcal{E}_{t, 2}\right]$ from above. This is not entirely straightforward because, while $\hat{\mu}_{i, T_{i}(t-1)}$ involves a sum of independent random variables, the number of terms $T_{i}(t-1)$ is itself a random variable. Moreover $T_{i}(t-1)$ depends on the past rewards $Z_{s}, s \leq t-1$, in a complex way. So in order to apply a concentration inequality to $\hat{\mu}_{i, T_{i}(t-1)}$, we use a rather blunt approach: we bound the worst deviation over all possible (deterministic) values in the support of $T_{i}(t-1)$. That is,

$$
\begin{align*}
\mathbb{P}\left[\mathcal{E}_{t, 2}\right] & =\mathbb{P}\left[\hat{\mu}_{i, T_{i}(t-1)}-\alpha \mathrm{H}\left(T_{i}(t-1), 1 / t\right)>\mu_{i}\right] \\
& \leq \mathbb{P}\left[\bigcup_{s \leq t-1}\left\{\hat{\mu}_{i, s}-\alpha \mathrm{H}(s, 1 / t)>\mu_{i}\right\}\right] . \tag{3.2.21}
\end{align*}
$$

We reformulate the previous bound as

$$
\begin{align*}
& \mathbb{P}\left[\bigcup_{s \leq t-1}\left\{\hat{\mu}_{i, s}-\alpha \mathrm{H}(s, 1 / t)>\mu_{i}\right\}\right] \\
&=\mathbb{P}\left[\sup _{s \leq t-1}\left(\hat{\mu}_{i, s}-\mu_{i}-\alpha \mathrm{H}(s, 1 / t)\right)>0\right] \\
&=\mathbb{P}\left[\sup _{s \leq t-1}\left(\frac{1}{s} \sum_{r \leq s} X_{i, r}-\mu_{i}-\alpha \sqrt{\frac{\log t}{2 s}}\right)>0\right] \\
&=\mathbb{P}\left[\sup _{s \leq t-1} \frac{1}{\sqrt{s}}\left(\frac{1}{\sqrt{s}} \sum_{r \leq s}\left(X_{i, r}-\mu_{i}\right)-\alpha \sqrt{\frac{\log t}{2}}\right)>0\right] \\
&=\mathbb{P}\left[\sup _{s \leq t-1} \frac{\sum_{r=1}^{s}\left(X_{i, r}-\mu_{i}\right)}{\sqrt{s}}>\alpha \sqrt{\frac{\log t}{2}}\right] . \tag{3.2.22}
\end{align*}
$$

Observe that the numerator on the left-hand side of the inequality on the last line is a martingale (see Example 3.1.29) with increments in $\left[-\mu_{i}, 1-\mu_{i}\right]$. But the denominator depends on $s$.

We try two approaches:

- We could simply use that $\sqrt{s} \geq 1$ on the denominator and apply the maximal

Azuma-Hoeffding inequality (Theorem 3.2.1) to get

$$
\begin{align*}
\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 2}\right] & \leq \sum_{t=1}^{n} \mathbb{P}\left[\sup _{s \leq t-1} \sum_{r=1}^{s}\left(X_{i, r}-\mu_{i}\right)>\alpha \sqrt{\frac{\log t}{2}}\right] \\
& \leq \sum_{t=1}^{n} \exp \left(-\frac{2(\alpha \sqrt{(\log t) / 2})^{2}}{t-1}\right) \\
& \leq \sum_{t=1}^{n} \exp \left(-\alpha^{2} \frac{\log t}{t-1}\right) . \tag{3.2.23}
\end{align*}
$$

That is of order $\Theta(n)$ for any $\alpha$.

- On the other hand, we could use a union bound over $s$ and apply the maximal Azuma-Hoeffding inequality to each term to get

$$
\begin{align*}
\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 2}\right] & \leq \sum_{t=1}^{n} \sum_{s \leq t-1} \mathbb{P}\left[\sum_{r=1}^{s}\left(X_{i, r}-\mu_{i}\right)>\alpha \sqrt{\frac{s \log t}{2}}\right] \\
& \leq \sum_{t=1}^{n} \sum_{s \leq t-1} \exp \left(-\frac{2(\alpha \sqrt{(s \log t) / 2})^{2}}{s}\right) \\
& =\sum_{t=1}^{n}(t-1) \exp \left(-\alpha^{2} \log t\right) \\
& \leq \sum_{t=1}^{n} \frac{1}{t^{\alpha^{2}-1}} \tag{3.2.24}
\end{align*}
$$

The series converges for $\alpha>\sqrt{2}$. Therefore, in that case, this bound is $\Theta(1)$, which is much better than our previous attempt. For $1<\alpha \leq \sqrt{2}$ however, we get a bound of order $\Theta\left(n^{\alpha^{2}}\right)$, which is worse than before.

It turns out that doing something "in between" the two approaches above gives a bound that significantly improves over both of them in the $1<\alpha \leq \sqrt{2}$ regime. This is known as the slicing (or peeling) method.

## Slicing method

The slicing method is useful when bounding a weighted supremum. Its application is somewhat problem-specific so we will content ourselves with illustrating it in our case. Specifically, our goal is to control probabilities of the form

$$
\mathbb{P}\left[\sup _{s \leq t-1} \frac{M_{s}}{w(s)} \geq \beta\right],
$$

where $M_{s}:=\sum_{r=1}^{s}\left(X_{i, r}-\mu_{i}\right), w(s):=\sqrt{s}$, and $\beta:=\alpha \sqrt{\frac{\log t}{2}}$. The idea is to divide up the supremum into slices $\gamma^{k-1} \leq s<\gamma^{k}, k \geq 1$, where the constant $\gamma>1$ will be optimized below. That is, fixing $K_{t}=\left\lceil\frac{\log t}{\log \gamma}\right\rceil$ (which roughly solves $\gamma^{K_{t}}=t$ ), by a union bound over the slices

$$
\mathbb{P}\left[\sup _{1 \leq s<t} \frac{M_{s}}{w(s)} \geq \beta\right] \leq \sum_{k=1}^{K_{t}} \mathbb{P}\left[\sup _{\gamma^{k-1} \leq s<\gamma^{k}} \frac{M_{s}}{w(s)} \geq \beta\right] .
$$

Because $w(s)$ is increasing, on each slice separately we can bound

$$
\begin{aligned}
\mathbb{P}\left[\sup _{\gamma^{k-1} \leq s<\gamma^{k}} \frac{M_{s}}{w(s)} \geq \beta\right] & \leq \mathbb{P}\left[\sup _{\gamma^{k-1} \leq s<\gamma^{k}} \frac{M_{s}}{w\left(\gamma^{k-1}\right)} \geq \beta\right] \\
& =\mathbb{P}\left[\sup _{\gamma^{k-1} \leq s<\gamma^{k}} M_{s} \geq \beta w\left(\gamma^{k-1}\right)\right] \\
& \leq \mathbb{P}\left[\sup _{s \leq \gamma^{k}} M_{s} \geq \beta w\left(\gamma^{k-1}\right)\right] .
\end{aligned}
$$

Now we apply the maximal Azuma-Hoeffding inequality (Theorem 3.2.1) to obtain

$$
\begin{aligned}
\mathbb{P}\left[\sup _{s \leq \gamma^{k}} M_{s} \geq \beta w\left(\gamma^{k-1}\right)\right] & \leq \exp \left(-\frac{2\left(\beta w\left(\gamma^{k-1}\right)\right)^{2}}{\gamma^{k}}\right) \\
& \leq \exp \left(-\frac{2 \beta^{2}}{\gamma}\right) \\
& =t^{-\alpha^{2} / \gamma}
\end{aligned}
$$

where we used that $M_{s}-M_{s-1}=X_{i, s}-\mu_{i} \in\left[-\mu_{i}, 1-\mu_{i}\right]$, an interval of length 1. Plugging this back above we get

$$
\begin{equation*}
\mathbb{P}\left[\sup _{1 \leq s<t} \frac{M_{s}}{w(s)} \geq \beta\right] \leq\left\lceil\frac{\log t}{\log \gamma}\right\rceil t^{-\alpha^{2} / \gamma} \tag{3.2.25}
\end{equation*}
$$

Now we see the tradeoff: increasing $\gamma$ makes the slices larger and hence the tail inequality weaker, but it also makes the number of slices smaller which helps with the union bound.

Combining (3.2.21), (3.2.22), and (3.2.25), we have proved:
Lemma 3.2.29. For any $\gamma>1$, it holds that

$$
\sum_{t=1}^{n} \mathbb{P}\left[\mathcal{E}_{t, 2}\right] \leq \sum_{t=1}^{n}\left[\frac{\log t}{\log \gamma}\right] t^{-\alpha^{2} / \gamma}
$$

and similarly for $\mathbb{P}\left[\mathcal{E}_{t, 1}\right]$.

For $\alpha>1$, we can choose $\gamma>1$ such that $\alpha^{2} / \gamma>1$. In that case, the series on the right-hand side is summable. This improves over both (3.2.23) and (3.2.24).

We are ready to prove the main result.
Proof of Theorem 3.2.26. By (3.2.17) and Lemmas 3.2.27, 3.2.28 and 3.2.29, we have

$$
\bar{R}_{n}=\sum_{i=1,2} \Delta_{i} \mathbb{E}\left[T_{i}(n)\right] \leq \Delta_{*}\left(u_{n}+2 \sum_{t=1}^{n}\left\lceil\frac{\log t}{\log \gamma}\right\rceil t^{-\alpha^{2} / \gamma}\right) .
$$

Recalling that $\alpha>1$, choose $\gamma>1$ such that $\alpha^{2} / \gamma>1$. In that case, as noted above, the series on the right hand side is summable and there is $C_{\alpha} \in(0,+\infty)$ such that

$$
\bar{R}_{n} \leq \Delta_{*}\left(u_{n}+C_{\alpha}\right) .
$$

That proves the claim.
Remark 3.2.30. A slightly better-and provably optimal-multiplicative constant in the pseudo-regret bound has been obtained by [GC11] using a variant of UCB called KL-UCB. The matching lower bound is due to [LR85]. See also [BCB12, Sections 2.3-2.4]. Further improvements can be obtained by using Bernstein's rather than Hoeffding's inequality [AMS09].

### 3.2.6 Coda: Talagrand's inequality

We end this section with a celebrated concentration inequality that applies under weaker conditions than McDiarmid's inequality (Theorem 3.2.9)-but is not proved using the martingale method. It is known as Talagrand's inequality.

Bounds on $\left\|D_{i} f\right\|_{\infty}$ are often expressed in terms of a Lipschitz condition under an appropriate metric. Let $0<c_{i}<+\infty, i=1, \ldots, n$ and $\mathbf{c}=\left(c_{1}, \ldots, c_{n}\right)$. The c-weighted Hamming distance is defined as
weighted

$$
\rho_{\mathbf{c}}(\mathbf{x}, \mathbf{y}):=\sum_{i=1}^{n} c_{i} \mathbf{1}_{\left\{x_{i} \neq y_{i}\right\}},
$$

distance
for $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right), \mathbf{y}=\left(y_{1}, \ldots, y_{n}\right) \in \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n}$. The proof of the following equivalence is left as an exercise (see Exercise 3.8).

Lemma 3.2.31 (Lipschitz condition). A function $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ satisfies the Lipschitz condition

$$
\begin{equation*}
|f(\mathbf{x})-f(\mathbf{y})| \leq \rho_{\mathbf{c}}(\mathbf{x}, \mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n}, \tag{3.2.26}
\end{equation*}
$$

if and only if

$$
\left\|D_{i} f\right\|_{\infty} \leq c_{i}, \quad \forall i
$$

Consider the following relaxed version of (3.2.26):

$$
\begin{equation*}
f(\mathbf{x})-f(\mathbf{y}) \leq \sum_{i=1}^{n} c_{i}(\mathbf{x}) \mathbf{1}_{\left\{x_{i} \neq y_{i}\right\}}, \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \tag{3.2.27}
\end{equation*}
$$

where now $c_{i}(\mathbf{x})$ is a finite, positive function over $\mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n}$. Notice the "onesided" nature of this condition, in the sense that $c_{i}$ depends on $\mathbf{x}$ but not on $\mathbf{y}$. A typical example where (3.2.27) is satisfied, but (3.2.26) is not, is given below.

We state Talagrand's inequality without proof.
Theorem 3.2.32 (Talagrand's inequality). Let $X_{1}, \ldots, X_{n}$ be independent random variables where $X_{i}$ is $\mathcal{X}_{i}$-valued for all $i$, and let $X=\left(X_{1}, \ldots, X_{n}\right)$. Assume

Talagrand's inequality $f: \mathcal{X}_{1} \times \cdots \times \mathcal{X}_{n} \rightarrow \mathbb{R}$ is a measurable function such that (3.2.27) holds. Then $f(X)$ is sub-Gaussian with variance factor $\left\|\sum_{i \leq n} c_{i}^{2}\right\|_{\infty}$. In fact, for all $\beta>0$ the following upper and lower tail bounds hold

$$
\mathbb{P}[f(X)-\mathbb{E} f(X) \geq \beta] \leq \exp \left(-\frac{\beta^{2}}{2\left\|\sum_{i \leq n} c_{i}^{2}\right\|_{\infty}}\right),
$$

and

$$
\mathbb{P}[f(X)-\mathbb{E} f(X) \leq-\beta] \leq \exp \left(-\frac{\beta^{2}}{2 \mathbb{E}\left[\sum_{i \leq n} c_{i}(X)^{2}\right]}\right)
$$

Compared to McDiarmid's inequality (Theorem 3.2.9), the upper tail in Theorem 3.2.32 has the sum over the coordinates inside the supremum, potentially a major improvement; the lower tail is even better, replacing the supremum with an expectation.

Example 3.2.33 (Spectral norm of a random matrix with bounded entries). Let $A$ be an $n \times n$ random matrix. We assume that the entries $A_{i, j}, i, j=1, \ldots, n$, are independent, centered random variables in $[-1,1]$. In Theorem 2.4.28, we proved an upper tail bound on the spectral norm

$$
\|A\|_{2}=\sup _{\substack{x \in \mathbb{R}^{n} \backslash\{0\}}} \frac{\|A \mathbf{x}\|_{2}}{\|\mathbf{x}\|_{2}}=\sup _{\substack{\mathbf{x} \in \mathbb{S}^{n-1} \\ \mathbf{y} \in \mathbb{S}^{n-1}}}\langle A \mathbf{x}, \mathbf{y}\rangle,
$$

of such a matrix (in the more general sub-Gaussian case) using an $\varepsilon$-net argument. Theorem 2.4.28 also implies that $\mathbb{E}\|A\|_{2}=O(\sqrt{n})$ by (B.5.1). (See Exercise 3.9 for a lower bound on the expectation.)

Here we use Talagrand's inequality (Theorem 3.2.32) directly to show concentration around the mean. For this, we need to check (3.2.27) where we think of the spectral norm as a function of $n^{2}$ independent random variables

$$
\|A\|_{2}=f\left(\left\{A_{i, j}\right\}_{i, j}\right) .
$$

Let $\mathbf{x}^{*}(A)$ and $\mathbf{y}^{*}(A)$ be unit vectors in $\mathbb{R}^{n}$ such that

$$
\|A\|_{2}=\left\langle A \mathbf{x}^{*}(A), \mathbf{y}^{*}(A)\right\rangle,
$$

which exist by compactness.
Given two $n \times n$ matrices $A, \widetilde{A}$ with entries in $[-1,1]$, we have

$$
\begin{aligned}
\|A\|_{2}-\|\widetilde{A}\|_{2} & =\left\langle A \mathbf{x}^{*}(A), \mathbf{y}^{*}(A)\right\rangle-\sup _{\substack{\mathbf{x} \in \mathbb{S}^{n-1} \\
\mathbf{y} \in \mathbb{S}^{n-1}}}\langle\widetilde{A} \mathbf{x}, \mathbf{y}\rangle \\
& \leq\left\langle A \mathbf{x}^{*}(A), \mathbf{y}^{*}(A)\right\rangle-\left\langle\widetilde{A} \mathbf{x}^{*}(A), \mathbf{y}^{*}(A)\right\rangle \\
& =\left\langle(A-\widetilde{A}) \mathbf{x}^{*}(A), \mathbf{y}^{*}(A)\right\rangle \\
& \leq \sum_{i, j}\left|A_{i j}-\widetilde{A}_{i j}\left\|\mathbf{x}^{*}(A)_{i}\right\| \mathbf{y}^{*}(A)_{j}\right| \\
& \leq \sum_{i, j} \mathbf{1}_{A_{i j} \neq \widetilde{A}_{i j}} c_{i j}(A),
\end{aligned}
$$

where on the last line we set

$$
c_{i j}(A):=2\left|\mathbf{x}^{*}(A)_{i}\right|\left|\mathbf{y}^{*}(A)_{j}\right|,
$$

and used the fact that $\left|A_{i j}-\widetilde{A}_{i j}\right| \leq 2$. Note that

$$
\sum_{i, j} c_{i j}(A)^{2}=4 \sum_{i} \mathbf{x}^{*}(A)_{i}^{2} \sum_{j} \mathbf{y}^{*}(A)_{j}^{2}=4 .
$$

Hence Talagrand's inequality implies that $\|A\|_{2}$ is sub-Gaussian with variance factor 4.

### 3.3 Potential theory and electrical networks

In this section we develop a classical link between random walks and electrical networks. The electrical interpretation is a useful physical analogy. The mathematical substance of the connection starts with the following observation.

Let $\left(X_{t}\right)$ be a Markov chain with transition matrix $P$ on a finite or countable state space $V$. Recall from Definition 3.1.6 that $\tau_{B}$ is the first visit time to $B \subseteq V$. For two disjoint subsets $A, Z$ of $V$, the probability of hitting $A$ before $Z$

$$
\begin{equation*}
h(x)=\mathbb{P}_{x}\left[\tau_{A}<\tau_{Z}\right], \tag{3.3.1}
\end{equation*}
$$

seen as a function of the starting point $x \in V$, is harmonic (with respect to $P$ ) on $W:=(A \cup Z)^{c}:=V \backslash(A \cup Z)$ in the sense that

$$
\begin{equation*}
h(x)=\sum_{y} P(x, y) h(y), \quad \forall x \in W . \tag{3.3.2}
\end{equation*}
$$

Indeed note that $h=1$ (respectively $=0$ ) on $A$ (respectively $Z$ ) and by the Markov property (Theorem 1.1.18), after the first step of the chain, for $x \in W$

$$
\begin{align*}
\mathbb{P}_{x}\left[\tau_{A}<\tau_{Z}\right]= & \sum_{y \notin A \cup Z} P(x, y) \mathbb{P}_{y}\left[\tau_{A}<\tau_{Z}\right] \\
& \quad+\sum_{y \in A} P(x, y) \cdot 1+\sum_{y \in Z} P(x, y) \cdot 0 \\
= & \sum_{y} P(x, y) \mathbb{P}_{y}\left[\tau_{A}<\tau_{Z}\right] . \tag{3.3.3}
\end{align*}
$$

Quantities such as (3.3.1) arise naturally, for instance in the study of recurrence, and the connection to potential theory, the study of harmonic functions, proves fruitful in that context as we outline in this section. It turns out that harmonic functions and martingales are closely related. In Section 3.3.1 we elaborate on that connection.

But first we rewrite (3.3.2) to reveal the electrical interpretation. For this we switch to reversible chains. Recall that a reversible Markov chain is equivalent to a random walk on a network $\mathcal{N}=(G, c)$ where the edges of $G$ correspond to transitions of positive probability. If the chain is reversible with respect to a stationary measure $\pi$, then the edge weights are $c(x, y)=\pi(x) P(x, y)$. In this notation (3.3.2) becomes

$$
\begin{equation*}
h(x)=\frac{1}{c(x)} \sum_{y \sim x} c(x, y) h(y), \quad \forall x \in(A \cup Z)^{c}, \tag{3.3.4}
\end{equation*}
$$

where $c(x):=\sum_{y \sim x} c(x, y)=\pi(x)$. In words, $h(x)$ is the weighted average of its neighboring values. Now comes the electrical analogy: if one interprets $c(x, y)$ as a conductance, a function satisfying (3.3.4) is known as a voltage. The voltages at $A$ and $Z$ are 1 and 0 respectively. We show in the next subsection by a martingale argument that, under appropriate conditions, such a voltage exists and is unique. We develop the electrical analogy and many of its applications in Section 3.3.2.

### 3.3.1 Martingales, the Dirichlet problem and Lyapounov functions

To see why martingales come in, let $\mathcal{F}_{t}=\sigma\left(X_{0}, \ldots, X_{t}\right)$ and let $\tau^{*}:=\tau_{W^{c}}$. By a first-step calculation again, for a function $h$ satisfying (3.3.2),

$$
\begin{equation*}
h\left(X_{t \wedge \tau^{*}}\right)=\mathbb{E}\left[h\left(X_{(t+1) \wedge \tau^{*}}\right) \mid \mathcal{F}_{t}\right], \quad \forall t \geq 0, \tag{3.3.5}
\end{equation*}
$$

that is, $\left(h\left(X_{t \wedge \tau^{*}}\right)\right)_{t}$ is a martingale with respect to $\left(\mathcal{F}_{t}\right)$. Indeed, on $\left\{\tau^{*} \leq t\right\}$,

$$
\mathbb{E}\left[h\left(X_{(t+1) \wedge \tau^{*}}\right) \mid \mathcal{F}_{t}\right]=h\left(X_{\tau^{*}}\right)=h\left(X_{t \wedge \tau^{*}}\right),
$$

and on $\left\{\tau^{*}>t\right\}$

$$
\mathbb{E}\left[h\left(X_{(t+1) \wedge \tau^{*}}\right) \mid \mathcal{F}_{t}\right]=\sum_{y} P\left(X_{t}, y\right) h(y)=h\left(X_{t}\right)=h\left(X_{t \wedge \tau^{*}}\right) .
$$

Although the rest of Section 3.3 is concerned with reversible Markov chains, the current subsection applies to the non-reversible case as well. We give an overview of potential theory for general, countable-space, discrete-time Markov chains and its connections to martingales. As a major application, we introduce the concept of a Lyapounov function which is useful in bounding certain hitting times.

## Existence and uniqueness of a harmonic extension

We begin with a special case, which will be generalized below.
Theorem 3.3.1 (Harmonic extension: existence and uniqueness). Let $P$ be an irreducible transition matrix on a finite or countably infinite state space $V$. Let $W$ be a finite, proper subset of $V$ and let $h: W^{c} \rightarrow \mathbb{R}$ be a bounded function on $W^{c}$. Then there exists a unique extension of $h$ to $W$ that is harmonic on $W$, that is, which satisfies (3.3.2). The solution is given by

$$
h(x)=\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right)\right] .
$$

Proof. We first argue about uniqueness. Suppose $h$ is defined over all of $V$ and satisfies (3.3.2). Let $\tau^{*}:=\tau_{W^{c}}$. Then the process $\left(h\left(X_{t \wedge \tau^{*}}\right)\right)_{t}$ is a martingale by (3.3.5). Because $W$ is finite and the chain is irreducible, we have $\tau^{*}<+\infty$ almost surely, as implied by Lemma 3.1.25. Moreover the process is bounded because $h$ is bounded on $W^{c}$ and $W$ is finite. Hence by Doob's optional stopping theorem (Theorem 3.1.38 (ii))

$$
h(x)=\mathbb{E}_{x}\left[h\left(X_{\tau^{*}}\right)\right], \quad \forall x \in W,
$$

which implies that $h$ is unique, since the right-hand side depends only on the chain and the fixed values of $h$ on $W^{c}$.

For the existence, simply define $h(x):=\mathbb{E}_{x}\left[h\left(X_{\tau^{*}}\right)\right], \forall x \in W$, and use a first-step argument similarly to (3.3.3).

For some insights on what happens when the assumptions of Theorem 3.3.1 are not satisfied, see Exercise 3.11. For an alternative (arguably more intuitive) proof of uniqueness based on the maximum principle, see Exercise 3.12.

In the proof above it suffices to specify $h$ on the outer boundary of $W$

$$
\partial_{V} W=\{z \in V \backslash W: \exists y \in W, P(y, z)>0\} .
$$

Introduce the Laplacian associated to $P$

$$
\begin{align*}
\Delta f(x) & =\left[\sum_{y} P(x, y) f(y)\right]-f(x) \\
& =\sum_{y} P(x, y)[f(y)-f(x)] \\
& =\mathbb{E}_{x}\left[f\left(X_{1}\right)-f\left(X_{0}\right)\right], \tag{3.3.6}
\end{align*}
$$

provided the expectation exists. We have proved that, under the assumptions of Theorem 3.3.1, there exists a unique solution to

$$
\begin{cases}\Delta f(x)=0 & \forall x \in W  \tag{3.3.7}\\ f(x)=h(x) & \forall x \in \partial_{\mathrm{V}} W\end{cases}
$$

and that solution is given by $f(x)=\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right)\right]$, for $x \in W \cup \partial_{V} W$. The system (3.3.7), in reference to its counterpart in the theory of partial differential equations, is referred to as a Dirichlet problem.

Example 3.3.2 (Simple random walk on $\mathbb{Z}^{d}$ ). The Laplacian above can be inter- problem preted as a discretized version of the standard Laplacian. For instance, for simple random walk on $\mathbb{Z}$,

$$
\begin{aligned}
\Delta f(x) & =\left[\sum_{y} P(x, y) f(y)\right]-f(x) \\
& =\sum_{y} P(x, y)[f(y)-f(x)] \\
& =\frac{1}{2}\{[f(x+1)-f(x)]-[f(x)-f(x-1)]\}
\end{aligned}
$$

which is a discretized second derivative. More generally, for simple random walk on $\mathbb{Z}^{d}$, we get

$$
\begin{aligned}
\Delta f(x) & =\left[\sum_{y} P(x, y) f(y)\right]-f(x) \\
& =\sum_{y} P(x, y)[f(y)-f(x)] \\
& =\frac{1}{2 d} \sum_{i=1}^{d}\left\{\left[f\left(x+\mathbf{e}_{i}\right)-f(x)\right]-\left[f(x)-f\left(x-\mathbf{e}_{i}\right)\right]\right\},
\end{aligned}
$$

where $\mathbf{e}_{1}, \ldots, \mathbf{e}_{d}$ is the standard basis in $\mathbb{R}^{d}$.
Theorem 3.3.1 has many applications. One of its consequences is that harmonic functions on a finite state space are constant.

Corollary 3.3.3. Let $P$ be an irreducible transition matrix on a finite state space $V$. If $h$ is harmonic on all of $V$, then it is constant.

Proof. Fix the value of $h$ at an arbitrary vertex $z$ and set $W=V \backslash\{z\}$. Applying Theorem 3.3.1, for all $x \in W, h(x)=\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right)\right]=h(z)$.

As an example of application of this corollary, we prove the following surprising result: in a finite, irreducible Markov chain, the expected time to hit a target chosen at random according to the stationary distribution does not depend on the starting point.

Theorem 3.3.4 (Random target lemma). Let $\left(X_{t}\right)$ be an irreducible Markov chain on a finite state space $V$ with transition matrix $P$ and stationary distribution $\pi$. Then

$$
h(x):=\sum_{y \in V} \pi(y) \mathbb{E}_{x}\left[\tau_{y}\right]
$$

does not in fact depend on $x$.
Proof. Because the chain is irreducible and has a finite state space, $\mathbb{E}_{x}\left[\tau_{y}\right]<+\infty$ for all $x, y$. By Corollary 3.3.3, it suffices to show that $h(x):=\sum_{y} \pi(y) \mathbb{E}_{x}\left[\tau_{y}\right]$ is harmonic on all of $V$. As before, it is natural to expand $\mathbb{E}_{x}\left[\tau_{y}\right]$ according to the first step of the chain,

$$
\mathbb{E}_{x}\left[\tau_{y}\right]=\mathbf{1}_{\{x \neq y\}}\left(1+\sum_{z} P(x, z) \mathbb{E}_{z}\left[\tau_{y}\right]\right)
$$

Substituting into the definition of $h(x)$ gives

$$
\begin{aligned}
h(x) & =(1-\pi(x))+\sum_{z} \sum_{y \neq x} \pi(y) P(x, z) \mathbb{E}_{z}\left[\tau_{y}\right] \\
& =(1-\pi(x))+\sum_{z} P(x, z)\left(h(z)-\pi(x) \mathbb{E}_{z}\left[\tau_{x}\right]\right) .
\end{aligned}
$$

Rearranging, we get

$$
\begin{aligned}
\Delta h(x) & =\left[\sum_{z} P(x, z) h(z)\right]-h(x) \\
& =\pi(x)\left(1+\sum_{z} P(x, z) \mathbb{E}_{z}\left[\tau_{x}\right]\right)-1 \\
& =0,
\end{aligned}
$$

where we used $1 / \pi(x)=\mathbb{E}_{x}\left[\tau_{x}^{+}\right]=1+\sum_{z} P(x, z) \mathbb{E}_{z}\left[\tau_{x}\right]$ by Theorem 3.1.19 and a first-step argument (recall that the first return time $\tau_{x}^{+}$was defined in Definition 3.1.6).

## Potential theory for Markov chains

More generally, many quantities of interest can be expressed in the following form. Consider again a subset $W \subset V$ and the stopping time

$$
\tau_{W^{c}}=\inf \left\{t \geq 0: X_{t} \in W^{c}\right\}
$$

Let also $h: W^{c} \rightarrow \mathbb{R}_{+}$and $k: W \rightarrow \mathbb{R}_{+}$. Define the quantity

$$
\begin{equation*}
u(x):=\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{0 \leq t<\tau_{W^{c}}} k\left(X_{t}\right)\right] . \tag{3.3.8}
\end{equation*}
$$

The first term on the right-hand side is a final cost incurred when we exit $W$ (and depends on where we do), while the second term is a unit time cost incurred along the sample path. Note that, in fact, it suffices to define $h$ on $\partial_{\mathrm{V}} W$, the outer boundary of $W$ if we restrict ourselves to $x \in W$. Observe also that the function $u(x)$ may take the value $+\infty$; the expectation is well-defined (in $\mathbb{R}_{+} \cup\{+\infty\}$ ) by the nonnegativity of the terms (see Appendix B).

Example 3.3.5 (Some special cases). Here are some important special cases:

- Revisiting (3.3.1), for two disjoint subsets $A, Z$ of $V$, the probability

$$
u(x):=\mathbb{P}_{x}\left[\tau_{A}<\tau_{Z}\right],
$$

of hitting $A$ before $Z$ as a function of the starting point $x \in V$ is obtained by taking $W:=(A \cup Z)^{c}, h=1$ (respectively $=0$ ) on $A$ (respectively $Z$ ), and $k=0$ on $V$. The further special case $Z=\emptyset$ leads to the exit probability from $A$

$$
u(x):=\mathbb{P}_{x}\left[\tau_{A}<+\infty\right] .
$$

On the other hand, if $A$ and $Z$ form a disjoint partition of $W^{c}$ (or $\partial_{\mathrm{V}} W$ will suffice if $x \in W$ ), we get the exit law from $W$

$$
u(x):=\mathbb{P}_{x}\left[X_{\tau_{W^{c}}} \in A ; \tau_{W^{c}}<+\infty\right] .
$$

- The average occupation time of $A \subseteq W$ before exiting $W$

$$
u(x):=\mathbb{E}_{x}\left[\sum_{0 \leq t<\tau_{W^{c}}} \mathbf{1}_{\left\{X_{t} \in A\right\}}\right],
$$

is obtained by taking $h=0$ on $V$, and $k=1$ (respectively $=0$ ) on $A$ (respectively on $A^{c}$ ). Revisiting (3.1.3), the Green function of the chain stopped at $\tau_{W^{c}}$, that is,

$$
u(x):=\mathscr{G}_{\tau_{W}}(x, y)=\mathbb{E}_{x}\left[\sum_{0 \leq t<\tau_{W^{c}}} \mathbf{1}_{\left\{X_{t}=y\right\}}\right],
$$

is obtained by taking $A=\{y\}$. Another special case is $A=W$ where we get the mean exit time from $A$
mean exit
time

The function $u$ in (3.3.8) turns out to satisfy a generalized version of (3.3.7). The proof is usually called first-step analysis (of which we have already seen many instances).

Theorem 3.3.6 (First-step analysis). Let $P$ be a transition matrix on a finite or countable state space $V$. Let $W$ be a proper subset of $V$, and let $h: W^{c} \rightarrow \mathbb{R}_{+}$ and $k: W \rightarrow \mathbb{R}_{+}$be bounded functions. Then the function $u \geq 0$, as defined in (3.3.8), satisfies the system of equations

$$
\begin{cases}u(x)=k(x)+\sum_{y} P(x, y) u(y) & \text { for } x \in W  \tag{3.3.9}\\ u(x)=h(x) & \text { for } x \in W^{c}\end{cases}
$$

Proof. For $x \in W^{c}$, by definition $u(x)=h(x)$ since $\tau_{W^{c}}=0$. Fix $x \in W$. By taking out what is known (Lemma B.6.13), the tower property (Lemma B.6.16) and the Markov property (Theorem 1.1.18),

$$
\begin{aligned}
u(x) & =k(x)+\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{1 \leq t<\tau_{W^{c}}} k\left(X_{t}\right)\right] \\
& =k(x)+\mathbb{E}_{x}\left[\mathbb{E}\left[h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{1 \leq t<\tau_{W^{c}}} k\left(X_{t}\right) \mid \mathcal{F}_{1}\right]\right] \\
& =k(x)+\mathbb{E}_{x}\left[u\left(X_{1}\right)\right],
\end{aligned}
$$

which gives the claim.
If $u$ is finite, the system of equations (3.3.9) can be rewritten as the Poisson equation (once again as an analogue of its counterpart in the theory of partial differential equations)

$$
\begin{cases}\Delta u=-k & \text { on } W,  \tag{3.3.10}\\ u=h & \text { on } W^{c} .\end{cases}
$$

This is well-defined for instance if $W$ is a finite subset and $P$ is irreducible. Indeed, as we argued in the proof of Theorem 3.3.1, the stopping time $\tau_{W^{c}}$ then has a finite expectation. Because $h$ is bounded, it follows that

$$
\begin{aligned}
u(x) & :=\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{0 \leq t<\tau_{W^{c}}} k\left(X_{t}\right)\right] \\
& \leq \sup _{x \in W^{c}} h(x)+\sup _{x \in W} k(x) \sup _{x \in W} \mathbb{E}_{x}\left[\tau_{W^{c}}\right] \\
& <+\infty
\end{aligned}
$$

uniformly in $x$. Using (3.3.6) and rearranging (3.3.9) gives (3.3.10).
Remark 3.3.7. A more general form of the statement which can be used to study certain moment-generating functions can be found, for example, in [Ebe, Theorem 1.3].

In a generalization of Theorem 3.3.1, our next theorem allows one to establish uniqueness of the solution of the system (3.3.10) under some conditions (which we will not detail here, but see Exercise 3.13). Perhaps even more useful, it also gives an effective approach to bound the function $u$ from above. This is based on the following supermartingale.

Lemma 3.3.8 (Locally superharmonic functions). Let $P$ be a transition matrix on a finite or countable state space $V$. Let $W$ be a proper subset of $V$, and let $h: W^{c} \rightarrow \mathbb{R}_{+}$and $k: W \rightarrow \mathbb{R}_{+}$be bounded functions. Suppose the nonnegative function $\psi: V \rightarrow \mathbb{R}_{+}$satisfies

$$
\Delta \psi \leq-k \quad \text { on } W
$$

Then the process

$$
N_{t}:=\psi\left(X_{t \wedge \tau_{W^{c}}}\right)+\sum_{0 \leq s<t \wedge \tau_{W^{c}}} k\left(X_{s}\right),
$$

is a nonnegative supermartingale for any initial point $x \in V$.
Proof. Observe that: on $\left\{\tau_{W^{c}} \leq t\right\}$, we have $N_{t+1}=N_{t}$; while on $\left\{\tau_{W^{c}}>t\right\}$, we have $N_{t+1}-N_{t}=\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right)+k\left(X_{t}\right)$ by cancellations in the sum. So, since $\left\{\tau_{W^{c}}>t\right\} \in \mathcal{F}_{t}$ by definition of a stopping time, it holds by taking out what is known that

$$
\begin{aligned}
\mathbb{E}\left[N_{t+1}-N_{t} \mid \mathcal{F}_{t}\right] & =\mathbb{E}\left[\mathbf{1}\left\{\tau_{W^{c}}>t\right\}\left(\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right)+k\left(X_{t}\right)\right) \mid \mathcal{F}_{t}\right] \\
& =\mathbf{1}\left\{\tau_{W^{c}}>t\right\}\left(\mathbb{E}\left[\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right) \mid \mathcal{F}_{t}\right]+k\left(X_{t}\right)\right) \\
& =\mathbf{1}\left\{\tau_{W^{c}}>t\right\}\left(\Delta \psi\left(X_{t}\right)+k\left(X_{t}\right)\right) \\
& \leq \mathbf{1}\left\{\tau_{W^{c}}>t\right\}\left(-k\left(X_{t}\right)+k\left(X_{t}\right)\right) \\
& =0,
\end{aligned}
$$

where we used that, by (3.3.6) and the Markov property,

$$
\begin{equation*}
\mathbb{E}\left[\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right) \mid \mathcal{F}_{t}\right]=\Delta \psi\left(X_{t}\right), \tag{3.3.11}
\end{equation*}
$$

and that $X_{t} \in W$ on $\left\{\tau_{W^{c}}>t\right\}$.
Theorem 3.3.9 (Poisson equation: bounding the solution). Let $P$ be a transition matrix on a finite or countable state space $V$. Let $W$ be a proper subset of $V$, and let $h: W^{c} \rightarrow \mathbb{R}_{+}$and $k: W \rightarrow \mathbb{R}_{+}$be bounded functions. Suppose the nonnegative function $\psi: V \rightarrow \mathbb{R}_{+}$satisfies the system of inequalities

$$
\begin{cases}\Delta \psi \leq-k & \text { on } W  \tag{3.3.12}\\ \psi \geq h & \text { on } W^{c} .\end{cases}
$$

Then

$$
\begin{equation*}
\psi \geq u, \quad \text { on } V \tag{3.3.13}
\end{equation*}
$$

where $u$ is the function defined in (3.3.8).

Proof. The system (3.3.13) holds on $W^{c}$ by Theorem 3.3.6 and (3.3.12) since in that case $u(x)=h(x) \leq \psi(x)$.

Fix $x \in W$. Consider the nonnegative supermartingale $\left(N_{t}\right)$ in Lemma 3.3.8. By the convergence of nonnegative supermartingales (Corollary 3.1.48), ( $N_{t}$ ) converges almost surely to a finite limit with expectation $\leq \mathbb{E}_{x}\left[N_{0}\right]$. In particular, the limit $N_{\tau_{W} c}$ is well-defined, nonnegative and finite, including on the event that $\left\{\tau_{W^{c}}=+\infty\right\}$. As a result,

$$
\begin{aligned}
N_{\tau_{W^{c}}} & =\lim _{t}\left[\psi\left(X_{t \wedge \tau_{W^{c}}}\right)+\sum_{0 \leq s<t \wedge \tau_{W^{c}}} k\left(X_{s}\right)\right] \\
& \geq h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{0 \leq s<\tau_{W^{c}}} k\left(X_{s}\right),
\end{aligned}
$$

where we used (3.3.12). Moreover, by Lemma 3.1.37, $\mathbb{E}_{x}\left[N_{t \wedge \tau_{W^{c}}}\right] \leq \mathbb{E}_{x}\left[N_{0}\right]$ for all $t$ and Fatou's lemma (see Proposition B.4.14) gives $\mathbb{E}_{x}\left[N_{\tau_{W} c}\right] \leq \mathbb{E}_{x}\left[N_{0}\right]$.

Hence, by definition of $u$,

$$
\begin{aligned}
u(x) & =\mathbb{E}_{x}\left[h\left(X_{\tau_{W^{c}}}\right) \mathbf{1}\left\{\tau_{W^{c}}<+\infty\right\}+\sum_{0 \leq t<\tau_{W^{c}}} k\left(X_{t}\right)\right] \\
& \leq \mathbb{E}_{x}\left[N_{\tau_{W^{c}}}\right] \\
& \leq \mathbb{E}_{x}\left[N_{0}\right] \\
& =\psi(x),
\end{aligned}
$$

where, on the last line, we used that the initial state is $x \in W$. That proves the claim.

## Lyapounov functions

Here is an important application, bounding from above the hitting time $\tau_{A}$ to a set $A$ in expectation.

Theorem 3.3.10 (Controlling hitting times via Lyapounov functions). Let $P$ be a transition matrix on a finite or countably infinite state space $V$. Let $A$ be a proper subset of $V$. Suppose the nonnegative function $\psi: V \rightarrow \mathbb{R}_{+}$satisfies the system of inequalities

$$
\begin{equation*}
\Delta \psi \leq-1, \quad \text { on } A^{c} . \tag{3.3.14}
\end{equation*}
$$

Then

$$
\mathbb{E}_{x}\left[\tau_{A}\right] \leq \psi(x)
$$

for all $x \in V$.
Proof. Indeed, by (3.3.14) and nonnegativity (in particular on $A$ ), the function $\psi$ satisfies the assumptions of Theorem 3.3.9 with $W=A^{c}, h=0$ on $A$, and $k=1$ on $A^{c}$. Hence, by definition of $u$ and the claim in Theorem 3.3.9,

$$
\begin{aligned}
\mathbb{E}_{x}\left[\tau_{A}\right] & =\mathbb{E}_{x}\left[h\left(X_{\tau_{A}}\right) \mathbf{1}\left\{\tau_{A}<+\infty\right\}+\sum_{0 \leq t<\tau_{A}} k\left(X_{t}\right)\right] \\
& =u(x) \\
& \leq \psi(x) .
\end{aligned}
$$

That establishes the claim.
Recalling (3.3.11), condition (3.3.14) is equivalent to the following conditional expected decrease in $\psi$ outside $A$ :

$$
\begin{equation*}
\mathbb{E}\left[\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right) \mid \mathcal{F}_{t}\right] \leq-1, \quad \text { on }\left\{X_{t} \in A^{c}\right\} . \tag{3.3.15}
\end{equation*}
$$

A nonnegative function satisfying an inequality of this type, also known as drift condition, is often referred to as a Lyapounov function. Intuitively, it tends to decrease along the sample path outside of $A$. Because it is non-negative, it cannot decrease forever and therefore the chain eventually enters $A$. We consider a simple example next.

Example 3.3.11 (A Markov chain on the nonnegative integers). Let $\left(Z_{t}\right)_{t \geq 1}$ be i.i.d. integrable random variables taking values in $\mathbb{Z}$ such that $\mathbb{E}\left[Z_{1}\right]<0$. Let $\left(X_{t}\right)_{t \geq 0}$ be the chain defined by $X_{0}=x$ for some $x \in \mathbb{Z}_{+}$and

$$
X_{t+1}=\left(X_{t}+Z_{t+1}\right)^{+},
$$

where recall that $z^{+}=\max \{0, z\}$. In particular $X_{t} \in \mathbb{Z}_{+}$for all $t$. Let $\left(\mathcal{F}_{t}\right)$ be the corresponding filtration. When $X_{t}$ is large, the "local drift" is close to $\mathbb{E}\left[Z_{1}\right]<0$. By analogy to the biased case of the gambler's ruin (Example 3.1.43), we might expect that, from a large starting point $x$, it will take time roughly $x /\left|\mathbb{E}\left[Z_{1}\right]\right|$ in expectation to "return to a neighborhood of 0 ." We prove something along those lines here using a Lyapounov function.

Observe that, for any $y \in \mathbb{Z}_{+}$, we have on the event $\left\{X_{t}=y\right\}$ by the Markov property

$$
\begin{align*}
\mathbb{E}_{x}\left[X_{t+1}-X_{t} \mid \mathcal{F}_{t}\right] & =\mathbb{E}\left[\left(y+Z_{t+1}\right)^{+}-y\right] \\
& =\mathbb{E}\left[-y \mathbf{1}\left\{Z_{t+1} \leq-y\right\}+Z_{t+1} \mathbf{1}\left\{Z_{t+1}>-y\right\}\right] \\
& \leq \mathbb{E}\left[Z_{t+1} \mathbf{1}\left\{Z_{t+1}>-y\right\}\right] \\
& =\mathbb{E}\left[Z_{1} \mathbf{1}\left\{Z_{1}>-y\right\}\right] . \tag{3.3.16}
\end{align*}
$$

For all $y$, the random variable $\left|Z_{1} \mathbf{1}\left\{Z_{1}>-y\right\}\right|$ is bounded by $\left|Z_{1}\right|$, itself an integrable random variable. Moreover, $Z_{1} 1\left\{Z_{1}>-y\right\} \rightarrow Z_{1}$ as $y \rightarrow+\infty$ almost surely. Hence, the dominated convergence theorem (Proposition B.4.14) implies that

$$
\lim _{y \rightarrow+\infty} \mathbb{E}\left[Z_{1} \mathbf{1}\left\{Z_{1}>-y\right\}\right]=\mathbb{E}\left[Z_{1}\right]<0
$$

So for any $0<\varepsilon<-\mathbb{E}\left[Z_{1}\right]$, there is $y_{\varepsilon} \in \mathbb{Z}_{+}$large enough that $\mathbb{E}\left[Z_{1} \mathbf{1}\left\{Z_{1}>\right.\right.$ $-y\}]<-\varepsilon$ for all $y>y_{\varepsilon}$. Fix $\varepsilon$ as above and define

$$
A:=\left\{0,1, \ldots, y_{\varepsilon}\right\} .
$$

We use Theorem 3.3.10 to bound $\tau_{A}$ in expectation. Define the Lyapounov function

$$
\psi(x)=\frac{x}{\varepsilon}, \quad \forall x \in \mathbb{Z}_{+}
$$

On the event $\left\{X_{t}=y\right\}$, we rewrite (3.3.16) as

$$
\begin{aligned}
\mathbb{E}\left[\psi\left(X_{t+1}\right)-\psi\left(X_{t}\right) \mid \mathcal{F}_{t}\right] & \leq \frac{\mathbb{E}\left[Z_{1} \mathbf{1}\left\{Z_{1}>-y\right\}\right]}{\varepsilon} \\
& \leq-1,
\end{aligned}
$$

for $y \in A^{c}$. This is the same as (3.3.15). Hence, we can apply Theorem 3.3.10 to get

$$
\mathbb{E}_{x}\left[\tau_{A}\right] \leq \psi(x)=\frac{x}{\varepsilon},
$$

for all $x \geq y_{\varepsilon}$.
A well-known, closely related result gives a criterion for positive recurrence. We state it without proof.

Theorem 3.3.12 (Foster's theorem). Let $P$ be an irreducible transition matrix on a countable state space $V$. Let $A$ be a finite, proper subset of $V$. Suppose the nonnegative function $\psi: V \rightarrow \mathbb{R}_{+}$satisfies the system of inequalities

$$
\Delta \psi \leq-1, \quad \text { on } A^{c},
$$

as well as the condition

$$
\sum_{y \in V} P(x, y) \psi(y)<+\infty, \quad \forall x \in A
$$

Then $P$ is positive recurrent.

### 3.3.2 Basic electrical network theory

We now develop the basic theory of electrical networks for the analysis of random walks. All results in this subsection (and the next one) concern reversible Markov chains, or random walks on networks (see Definition 1.2.7). We begin with a few definitions. Throughout, we will use the notation $\left.h\right|_{B}$ for the function $h$ restricted to the subset $B$. We also write $h \equiv c$ if $h$ is identically equal to the constant $c$.

## Definitions

Let $\mathcal{N}=(G, c)$ be a finite or countable network with $G=(V, E)$. Throughout this section we assume that $\mathcal{N}$ is connected and locally finite. In the context of electrical networks, edge weights are called conductances. The reciprocal of the conductances are called resistances and are denoted by $r(e):=1 / c(e)$, for all $e \in E$. For an edge $e=\{x, y\}$ we overload $c(x, y):=c(e)$ and $r(x, y):=r(e)$. Both $c$ and $r$ are symmetric as functions of $x, y$. Recall that the transition matrix of the random walk on $\mathcal{N}$ satisfies

$$
P(x, y)=\frac{c(x, y)}{c(x)}
$$

where

$$
c(x)=\sum_{z: z \sim x} c(x, z) .
$$

Let $A, Z$ be disjoint, non-empty subsets of $V$ such that $W:=(A \cup Z)^{c}$ is finite. For our purposes it will suffice to take $A$ to be a singleton, that is, $A=\{a\}$ for some $a$. Then $a$ is called the source and $Z$ is called the sink-set, or sink for short. As an immediate corollary of Theorem 3.3.1, we obtain the existence and uniqueness of a voltage function, defined formally in the next corollary. It will be useful to consider voltages taking an arbitrary value at $a$, but we always set the voltage on $Z$ to 0 .

Corollary 3.3.13 (Voltage). Fix $v_{0}>0$. Let $\mathcal{N}=(G, c)$ be a finite or countable, connected network with $G=(V, E)$. Let $A:=\{a\}, Z$ be disjoint non-empty subsets of $V$ such that $W=(A \cup Z)^{c}$ is non-empty and finite. Then there exists a unique voltage defined as follows: a function $v$ on $V$ such that $v$ is harmonic on $W$, that is,

$$
\begin{equation*}
v(x)=\frac{1}{c(x)} \sum_{y: y \sim x} c(x, y) v(y), \quad \forall x \in W \tag{3.3.17}
\end{equation*}
$$

where

$$
\begin{equation*}
v(a)=v_{0} \quad \text { and }\left.\quad v\right|_{Z} \equiv 0 \tag{3.3.18}
\end{equation*}
$$

Moreover

$$
\begin{equation*}
\frac{v(x)}{v_{0}}=\mathbb{P}_{x}\left[\tau_{a}<\tau_{Z}\right], \tag{3.3.19}
\end{equation*}
$$

for the corresponding random walk on $\mathcal{N}$.
Proof. Set $h(x)=v(x)$ on $A \cup Z$. Theorem 3.3.1 gives the result.
Note in the definition above that if $v$ is a voltage with value $v_{0}$ at $a$, then $\tilde{v}(x)=$ $v(x) / v_{0}$ is a voltage with value 1 at $a$.

Let $v$ be a voltage function on $\mathcal{N}$ with source $a$ and $\operatorname{sink} Z$. The Laplacianbased formulation of harmonicity, (3.3.7), can be interpreted in terms of flows (see Definition 1.1.13). We define the current function

$$
\begin{equation*}
i(x, y):=c(x, y)[v(x)-v(y)], \tag{3.3.20}
\end{equation*}
$$

or, equivalently, $v(x)-v(y)=r(x, y) i(x, y)$. The latter definition is usually referred to as Ohm's "law." Notice that the current is defined on ordered pairs of vertices and is anti-symmetric, that is, $i(x, y)=-i(y, x)$. In terms of the current, the harmonicity of $v$ is then expressed as

$$
\begin{equation*}
\sum_{y: y \sim x} i(x, y)=0, \quad \forall x \in W, \tag{3.3.21}
\end{equation*}
$$

that is, $i$ is a flow on $W$ (without capacity constraints). This set of equations is known as Kirchhoff's node law. We also refer to these constraints as flowconservation constraints. To be clear, the current is not just any flow. It is a flow that can be written as a potential difference according to Ohm's law. Such a current also satisfies Kirchhoff's cycle law: if $x_{1} \sim x_{2} \sim \cdots \sim x_{k} \sim x_{k+1}=x_{1}$ is a cycle, then

$$
\sum_{j=1}^{k} i\left(x_{j}, x_{j+1}\right) r\left(x_{j}, x_{j+1}\right)=0
$$

as can be seen by substituting Ohm's law.
The strength of the current is defined as

Kirchhoff's
node law

Kirchhoff's
cycle law
strength

$$
\|i\|:=\sum_{y: y \sim a} i(a, y) .
$$

Because $a \notin W$, it does not satisfy Kirchhoff's node law and the strength is not 0 in general. The definition of $i(x, y)$ ensures that the flow out of the source is nonnegative as $\mathbb{P}_{y}\left[\tau_{a}<\tau_{Z}\right] \leq 1=\mathbb{P}_{a}\left[\tau_{a}<\tau_{Z}\right]$ for all $y \sim a$ so that

$$
i(a, y)=c(a, y)[v(a)-v(y)]=c(a, y)\left[v_{0} \mathbb{P}_{a}\left[\tau_{a}<\tau_{Z}\right]-v_{0} \mathbb{P}_{y}\left[\tau_{a}<\tau_{Z}\right]\right] \geq 0
$$

Note that by multiplying the voltage by a constant we obtain a current which is similarly scaled. Up to that scaling, the current is unique from the uniqueness of the voltage. We will often consider the unit current where we scale $v$ and $i$ so as to enforce that $\|i\|=1$.

Summing up the previous paragraphs, to determine the voltage it suffices to find functions $v$ and $i$ that simultaneously satisfy Ohm's law and Kirchhoff's node law. Here is an example.

Example 3.3.14 (Network reduction: birth-death chain). Let $\mathcal{N}$ be the line on $\{0,1, \ldots, n\}$ with $j \sim k \Longleftrightarrow|j-k|=1$ and arbitrary (positive) conductances on the edges. Let $\left(X_{t}\right)$ be the corresponding walk. We use the principle above to compute $\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right]$ for $1 \leq x \leq n-1$. Consider the voltage function $v$ when $v(0)=1$ and $v(n)=0$ with current $i$, which exists and is unique by Corollary 3.3.13. The desired quantity is $v(x)$.

Note that because $i$ is a flow on $\mathcal{N}$, the flow into every vertex equals the flow out of that vertex, and we must have $i(y, y+1)=i(0,1)=\|i\|$ for all $y$. To compute $v(x)$, we note that it remains the same if we replace the path $0 \sim 1 \sim \cdots \sim x$ with a single edge of resistance $R_{0, x}=r(0,1)+\cdots+r(x-1, x)$. Indeed leave the voltage unchanged on the remaining nodes (to the right of $x$ ) and define the current on the new edge as $\|i\|$. Kirchhoff's node law is automatically satisfied by the argument above. To check Ohm's law on the new "super-edge," note that on the original network $\mathcal{N}$ (with the original voltage function)

$$
\begin{aligned}
v(0)-v(x) & =(v(0)-v(1))+\cdots+(v(x-1)-v(x)) \\
& =r(x-1, x) i(x-1, x)+\cdots+r(0,1) i(0,1) \\
& =[r(0,1)+\cdots+r(x-1, x)]\|i\| \\
& =R_{0, x}\|i\| .
\end{aligned}
$$

Ohm's law is also satisfied on every other edge (to the right of $x$ ) because nothing has changed there. That proves the claim.

We do the same reduction on the other side of $x$ by replacing $x \sim x+1 \sim$ $\cdots \sim n$ with a single edge of resistance $R_{x, n}=r(x, x+1)+\cdots+r(n-1, n)$. See Figure 3.4.

Because the voltage at $x$ was not changed by this transformation, we can compute $v(x)=\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right]$ directly on the reduced network, where it is now a straightforward computation. Indeed, starting at $x$, the reduced walk jumps to 0 with probability proportional to the conductance on the new super-edge $0 \sim x$ (or


Figure 3.4: Reduced network.
the reciprocal of the resistance), that is,

$$
\begin{aligned}
\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right] & =\frac{R_{0, x}^{-1}}{R_{0, x}^{-1}+R_{x, n}^{-1}} \\
& =\frac{R_{x, n}}{R_{x, n}+R_{0, x}} \\
& =\frac{r(x, x+1)+\cdots+r(n-1, n)}{r(0,1)+\cdots+r(n-1, n)} .
\end{aligned}
$$

Some special cases:

- Simple random walk. In the case of simple random walk, all resistances are equal and we get

$$
\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right]=\frac{n-x}{n} .
$$

- Gambler's ruin. The gambler's ruin example (see Examples 3.1.41 and 3.1.43) corresponds to taking $c(j, j+1)=(p / q)^{j}$ or $r(j, j+1)=(q / p)^{j}$, for some $0<p<1$ and $q=1-p$. In this case we obtain

$$
\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right]=\frac{\sum_{j=x}^{n-1}(q / p)^{j}}{\sum_{j=0}^{n-1}(q / p)^{j}}=\frac{(q / p)^{x}\left(1-(q / p)^{n-x}\right)}{1-(q / p)^{n}}=\frac{(p / q)^{n-x}-1}{(p / q)^{n}-1}
$$

when $p \neq q$ (otherwise we get back the simple random walk case).z

The above example illustrates the series law: resistances in series add up. There is a similar parallel law: conductances in parallel add up. To formalize these laws, one needs to introduce multigraphs. This is straightforward, although to avoid complicating the notation further we will not do this here. (But see Example 3.3.22 for a simple case.)

Another useful network reduction technique is illustrated in the next example.
Example 3.3.15 (Network reduction: binary tree). Let $\mathcal{N}$ be the rooted binary tree with $n$ levels $\widehat{\mathbb{T}}_{2}^{n}$ and equal conductances on all edges. Let 0 be the root. Pick an arbitrary leaf and denote it by $n$. The remaining vertices on the path between 0 and $n$, which we refer to as the main path, will be denoted by $1, \ldots, n-1$ moving away from the root. We claim that, for all $0<x<n$, it holds that

$$
\mathbb{P}_{x}\left[\tau_{0}<\tau_{n}\right]=(n-x) / n
$$

Indeed let $v$ be the voltage with values 1 and 0 at $a=0$ and $Z=\{n\}$ respectively. Let $i$ be the corresponding current. Notice that, for each $0 \leq y<n$, the current-as a flow-has "nowhere to go" on the subtree $T_{y}$ hanging from $y$ away from the main path. The leaves of the subtree are dead ends. Hence the current must be 0 on $T_{y}$ and by Ohm's law the voltage must be constant on it, that is, every vertex in $T_{y}$ has voltage $v(y)$.

Imagine collapsing all vertices in $T_{y}$, including $y$, into a single vertex (and removing the self-loops so created). Doing this for every vertex on the main path results in a new reduced network which is formed of a single path as in Example 3.3.14. Note that the voltage and the current can be taken to be the same as they were previously on the main path. Indeed, with this choice, Ohm's law is automatically satisfied. Moreover, because there is no current on the hanging subtrees in the original network, Kirchhoff's node law is also satisfied on the reduced network, as no current is "lost."

Hence the answer can be obtained from Example 3.3.14. That proves the claim. (You should convince yourself that this result is obvious from a probabilistic point of view.)

We gave a probabilistic interpretation of the voltage. What about the current? The following result says that, roughly speaking, $i(x, y)$ is the net traffic on the edge $\{x, y\}$ from $x$ to $y$. We start with an important formula for the voltage at $a$. For the walk started at $a$, we use the shorthand

$$
\mathbb{P}[a \rightarrow Z]:=\mathbb{P}_{a}\left[\tau_{Z}<\tau_{a}^{+}\right],
$$

for the escape probability. The next lemma can be interpreted as a sort of Ohm's law between $a$ and $Z$, where $c(a) \mathbb{P}[a \rightarrow Z]$ is the "effective conductance." (We will be more formal in Definition 3.3.19 below.)

Lemma 3.3.16 (Effective Ohm's Law). Let $v$ be a voltage on $\mathcal{N}$ with source $a$ and sink Z. Let $i$ be the associated current. Then

$$
\begin{equation*}
\frac{v(a)}{\|i\|}=\frac{1}{c(a) \mathbb{P}[a \rightarrow Z]} . \tag{3.3.22}
\end{equation*}
$$

Proof. Using the usual first-step analysis,

$$
\begin{aligned}
\mathbb{P}[a \rightarrow Z] & =\sum_{x: x \sim a} P(a, x) \mathbb{P}_{x}\left[\tau_{Z}<\tau_{a}\right] \\
& =\sum_{x: x \sim a} \frac{c(a, x)}{c(a)}\left(1-\frac{v(x)}{v(a)}\right) \\
& =\frac{1}{c(a) v(a)} \sum_{x: x \sim a} c(a, x)[v(a)-v(x)] \\
& =\frac{1}{c(a) v(a)} \sum_{x: x \sim a} i(a, x),
\end{aligned}
$$

where we used Corollary 3.3.13 on the second line and Ohm's law on the last line. Rearranging gives the result.

Recall the Green function from (3.1.3).
Theorem 3.3.17 (Probabilistic interpretation of the current). For $x \sim y$, let $N_{x \rightarrow y}^{Z}$ be the number of one-step transitions from $x$ to $y$ up to the time of the first visit to the sink $Z$ for the random walk on $\mathcal{N}$ started at $a$. Let $v$ be the voltage corresponding to the unit current $i$. Then the following formulas hold:

$$
\begin{equation*}
v(x)=\frac{\mathscr{G}_{\tau_{Z}}(a, x)}{c(x)}, \quad \forall x \tag{3.3.23}
\end{equation*}
$$

and

$$
i(x, y)=\mathbb{E}_{a}\left[N_{x \rightarrow y}^{Z}-N_{y \rightarrow x}^{Z}\right], \quad \forall x \sim y .
$$

Proof. We prove the formula for the voltage by showing that $v(x)$ as defined above is harmonic on $W=V \backslash(\{a\} \cup Z)$. Note first that, for all $z \in Z$, the expected number of visits to $z$ before reaching $Z$ (i.e., $\mathscr{G}_{\tau_{Z}}(a, z)$ ) is 0 . Or, put differently, $0=v(z)=\frac{\mathscr{G}_{\tau_{Z}}(a, z)}{c(z)}$. Moreover, to compute $\mathscr{G}_{\tau_{Z}}(a, a)$, note that the number of
visits to $a$ before the first visit to $Z$ is geometric with success probability $\mathbb{P}[a \rightarrow Z]$ by the strong Markov property (Theorem 3.1.8) and hence

$$
\mathscr{G}_{\tau_{Z}}(a, a)=\frac{1}{\mathbb{P}[a \rightarrow Z]},
$$

and, by Lemma 3.3.16 and the fact that we are using the unit current, $v(a)=$ $\frac{\mathscr{\varphi}_{T_{Z}}(a, a)}{c(a)}$, as required.

To establish the formula for $x \in W$, we compute the quantity

$$
\frac{1}{c(x)} \sum_{y: y \sim x} \mathbb{E}_{a}\left[N_{y \rightarrow x}^{Z}\right],
$$

in two ways. First, because each visit to $x \in W$ must enter through one of $x$ 's neighbors (including itself in the presence of a self-loop), we get

$$
\begin{equation*}
\frac{1}{c(x)} \sum_{y: y \sim x} \mathbb{E}_{a}\left[N_{y \rightarrow x}^{Z}\right]=\frac{\mathscr{G}_{\tau_{Z}}(a, x)}{c(x)} . \tag{3.3.24}
\end{equation*}
$$

On the other hand, by the Markov property (Theorem 1.1.18)

$$
\begin{align*}
& \mathbb{E}_{a}\left[N_{y \rightarrow x}^{Z}\right] \\
& =\mathbb{E}_{a}\left[\sum_{0 \leq t<\tau_{Z}} \mathbf{1}_{\left\{X_{t}=y, X_{t+1}=x\right\}}\right] \\
& =\sum_{t \geq 0} \mathbb{P}_{a}\left[X_{t}=y, X_{t+1}=x, \tau_{Z}>t\right] \\
& =\sum_{t \geq 0} \mathbb{P}_{a}\left[\tau_{Z}>t\right] \mathbb{P}_{a}\left[X_{t}=y \mid \tau_{Z}>t\right] \mathbb{P}_{a}\left[X_{t+1}=x \mid X_{t}=y, \tau_{Z}>t\right] \\
& =\sum_{t \geq 0} \mathbb{P}_{a}\left[\tau_{Z}>t\right] \mathbb{P}_{a}\left[X_{t}=y \mid \tau_{Z}>t\right] P(y, x) \\
& =\sum_{t \geq 0} \mathbb{P}_{a}\left[X_{t}=y, \tau_{Z}>t\right] P(y, x) \\
& =P(y, x) \mathbb{E}_{a}\left[\sum_{0 \leq t<\tau_{Z}} \mathbf{1}_{\left\{X_{t}=y\right\}}\right] \\
& =P(y, x) \mathscr{G}_{\tau_{Z}}(a, y) \tag{3.3.25}
\end{align*}
$$

so that, summing over $y$, we obtain this time

$$
\begin{align*}
\frac{1}{c(x)} \sum_{y: y \sim x} \mathbb{E}_{a}\left[N_{y \rightarrow x}^{Z}\right] & =\frac{1}{c(x)} \sum_{y: y \sim x} P(y, x) \mathscr{G}_{\tau_{Z}}(a, y) \\
& =\sum_{y: y \sim x} P(x, y) \frac{\mathscr{G}_{\tau_{Z}}(a, y)}{c(y)}, \tag{3.3.26}
\end{align*}
$$

where we used that $c(x, y)=c(x) P(x, y)=c(y) P(y, x)$ (see Definition 1.2.7). Equating (3.3.24) and (3.3.26) shows that $\frac{\mathscr{\varphi}_{Z}(a, x)}{c(x)}$ is harmonic on $W$ and hence must be equal to the voltage function by Corollary 3.3.13.

Finally, by (3.3.25),

$$
\begin{aligned}
\mathbb{E}_{a}\left[N_{x \rightarrow y}^{Z}-N_{y \rightarrow x}^{Z}\right] & =P(x, y) \mathscr{G}_{\tau_{Z}}(a, x)-P(y, x) \mathscr{G}_{T_{Z}}(a, y) \\
& =P(x, y) v(x) c(x)-P(y, x) v(y) c(y) \\
& =c(x, y)[v(x)-v(y)] \\
& =i(x, y) .
\end{aligned}
$$

That concludes the proof.
Example 3.3.18 (Network reduction: binary tree (continued)). Recall the setting of Example 3.3.15. We argued that the current on side edges, that is, edges of subtrees hanging from the main path, is 0 . This is clear from the probabilistic interpretation of the current: in a walk from $a$ to $z$, any traversal of a side edge must be undone at a later time.

The network reduction techniques illustrated above are useful. But the power of the electrical network perspective is more apparent in what comes next: the definition of the effective resistance and, especially, its variational characterization.

## Effective resistance

Before proceeding further, let us recall our original motivation. Let $\mathcal{N}=(G, c)$ be a countable, locally finite, connected network and let $\left(X_{t}\right)$ be the corresponding walk. Recall that a vertex $a$ in $G$ is transient if $\mathbb{P}_{a}\left[\tau_{a}^{+}<+\infty\right]<1$.

To relate this to our setting, consider an exhaustive sequence of induced subgraphs $G_{n}$ of $G$ which for our purposes is defined as: $G_{0}$ contains only $a, G_{n} \subseteq$ $G_{n+1}, G=\bigcup_{n} G_{n}$, and every $G_{n}$ is finite and connected. Such a sequence always exists by iteratively adding the neighbors of the previous vertices and using that $G$ is locally finite and connected. Let $Z_{n}$ be the set of vertices of $G$ not in $G_{n}$. Then,
exhaustive
sequence
by Lemma 3.1.25, $\mathbb{P}_{a}\left[\tau_{Z_{n}} \wedge \tau_{a}^{+}=+\infty\right]=0$ for all $n$ by our assumptions on $\left(G_{n}\right)$. Hence, the remaining possibilities are

$$
\begin{aligned}
1 & =\mathbb{P}_{a}\left[\exists n, \tau_{a}^{+}<\tau_{Z_{n}}\right]+\mathbb{P}_{a}\left[\forall n, \tau_{Z_{n}}<\tau_{a}^{+}\right] \\
& =\mathbb{P}_{a}\left[\tau_{a}^{+}<+\infty\right]+\lim _{n} \mathbb{P}\left[a \rightarrow Z_{n}\right] .
\end{aligned}
$$

Therefore $a$ is transient if and only if $\lim _{n} \mathbb{P}\left[a \rightarrow Z_{n}\right]>0$. Note that the limit exists because the sequence of events $\left\{\tau_{Z_{n}}<\tau_{a}^{+}\right\}$is decreasing by construction. By a sandwiching argument the limit also does not depend on the exhaustive sequence. Hence we define

$$
\mathbb{P}[a \rightarrow \infty]:=\lim _{n} \mathbb{P}\left[a \rightarrow Z_{n}\right] .
$$

We use Lemma 3.3.16 to characterize this limit using electrical network concepts.
But, first, here comes the key definition. In Lemma 3.3.16, $v(a)$ can be thought of as the potential difference between the source and the sink, and $\|i\|$ can be thought of as the total current flowing through the network from the source to the sink. Hence, viewing the network as a single "super-edge," Equation (3.3.22) is the analogue of Ohm's law if we interpret $c(a) \mathbb{P}[a \rightarrow Z]$ as an "effective conductance."

Definition 3.3.19 (Effective resistance and conductance). Let $\mathcal{N}=(G, c)$ be a finite or countable, locally finite, connected network. Let $A=\{a\}$ and $Z$ be disjoint non-empty subsets of the vertex set $V$ such that $W:=V \backslash(A \cup Z)$ is finite. Let $v$ be a voltage from source a to sink $Z$ and let $i$ be the corresponding current. The effective resistance between $a$ and $Z$ is defined as

$$
\mathscr{R}(a \leftrightarrow Z):=\frac{1}{c(a) \mathbb{P}[a \rightarrow Z]}=\frac{v(a)}{\|i\|},
$$

where the rightmost equality holds by Lemma 3.3.16. The reciprocal is called the effective conductance and is denoted by $\mathscr{C}(a \leftrightarrow Z):=1 / \mathscr{R}(a \leftrightarrow Z)$.

Going back to recurrence, for an exhaustive sequence $\left(G_{n}\right)$ with $\left(Z_{n}\right)$ as above, it is natural to define

$$
\mathscr{R}(a \leftrightarrow \infty):=\lim _{n} \mathscr{R}\left(a \leftrightarrow Z_{n}\right),
$$

where, once again, the limit does not depend on the choice of exhaustive sequence.
Theorem 3.3.20 (Recurrence and resistance). Let $\mathcal{N}=(G, c)$ be a countable, locally finite, connected network. Vertex a (and hence all vertices) in $\mathcal{N}$ is transient if and only if $\mathscr{R}(a \leftrightarrow \infty)<+\infty$.

Proof. This follows immediately from the definition of the effective resistance. Recall that, on a connected network, all states have the same type (recurrent or transient).

Note that the network reduction techniques we discussed previously leave both the voltage and the current strength unchanged on the reduced network. Hence they also leave the effective resistance unchanged.

Example 3.3.21 (Gambler's ruin chain revisited). Extend the gambler's ruin chain of Example 3.3.14 to all of $\mathbb{Z}_{+}$. We determine when this chain is transient. Because it is irreducible, all states have the same type and it suffices to look at 0 . Consider the exhaustive sequence obtained by letting $G_{n}$ be the graph restricted to $\{0,1, \ldots, n-1\}$ and letting $Z_{n}=\{n, n+1 \ldots\}$. To compute the effective resistance $\mathscr{R}\left(0 \leftrightarrow Z_{n}\right)$, we use the same reduction as in Example 3.3.14. The "super-edge" between 0 and $n$ has resistance

$$
\mathscr{R}\left(0 \leftrightarrow Z_{n}\right)=\sum_{j=0}^{n-1} r(j, j+1)=\sum_{j=0}^{n-1}(q / p)^{j}=\frac{(q / p)^{n}-1}{(q / p)-1},
$$

when $p \neq q$, and similarly it has resistance $n$ in the $p=q$ case. Hence, taking a limit as $n \rightarrow+\infty$,

$$
\mathscr{R}(0 \leftrightarrow \infty)= \begin{cases}+\infty, & p \leq 1 / 2 \\ \frac{p}{2 p-1}, & p>1 / 2\end{cases}
$$

So 0 is transient if and only if $p>1 / 2$.
Example 3.3.22 (Biased walk on the $b$-ary tree). Fix $\lambda \in(0,+\infty)$. Consider the rooted, infinite $b$-ary tree with conductance $\lambda^{j}$ on all edges between level $j-1$ and $j$, for $j \geq 1$. We determine when this chain is transient. Because it is irreducible, all states have the same type and it suffices to look at the root. Denote the root by 0 . For an exhaustive sequence, let $G_{n}$ be the root together with the first $n-1$ levels. Let $Z_{n}$ be as before. To compute $\mathscr{R}\left(0 \leftrightarrow Z_{n}\right)$ : (i) glue together all vertices of $Z_{n}$; (ii) glue together all vertices on the same level of $G_{n}$; (iii) replace parallel edges with a single edge whose conductance is the sum of the conductances; (iv) let the current on this edge be the sum of the currents; and (v) leave the voltages unchanged. It can be checked that Ohm's law and Kirchhoff's node law are still satisfied, and that hence we have not changed the effective resistance. (This is an application of the parallel law.)

The reduced network is now a line. Denote the new vertices $0,1, \ldots, n$. The conductance on the edge between $j$ and $j+1$ is $b^{j+1} \lambda^{j}=b(b \lambda)^{j}$. So this is
the chain from the previous example with $(p / q)=b \lambda$ where all conductances are scaled by a factor of $b$. Hence

$$
\mathscr{R}(0 \leftrightarrow \infty)= \begin{cases}+\infty, & b \lambda \leq 1, \\ \frac{1}{b\left(1-(b \lambda)^{-1}\right)}, & b \lambda>1 .\end{cases}
$$

So the root is transient if and only if $b \lambda>1$.
A generalization is provided in Example 3.3.27.

### 3.3.3 Bounding the effective resistance via variational principles

The examples we analyzed so far were atypical in that it was possible to reduce the network down to a single edge using simple rules and read off the effective resistance. In general, we need more robust techniques to bound the effective resistance. The following two variational principles provide a powerful approach for this purpose. We derive them for finite networks, but will later on apply them to exhaustive sequences.

## Variational principles

Recall from Definition 1.1.13 that a flow $\theta$ from source $a$ to $\operatorname{sink} Z$ on a countable, locally finite, connected network $\mathcal{N}=(G, c)$ is a function on pairs of adjacent vertices such that: $\theta$ is anti-symmetric, that is, $\theta(x, y)=-\theta(y, x)$ for all $x \sim y$; and it satisfies the flow-conservation constraint $\sum_{y: y \sim x} \theta(x, y)=0$ on all vertices $x$ except those in $\{a\} \cup Z$. The strength of the flow is $\|\theta\|=\sum_{y: y \sim a} \theta(a, y)$. The current is a special flow-one that can be written as a potential difference according to Ohm's law. As we show next, it can also be characterized as a flow minimizing a certain energy. Specifically, the energy of a flow $\theta$ is defined as

$$
\mathscr{E}(\theta)=\frac{1}{2} \sum_{x, y} r(x, y) \theta(x, y)^{2} .
$$

The proof of the variational principle we present here employs a neat trick, convex duality. In particular, it reveals that the voltage and current are dual in the sense of convex analysis.

Theorem 3.3.23 (Thomson's principle). Let $\mathcal{N}=(G, c)$ be a finite, connected network. The effective resistance between source a and sink $Z$ is characterized by

$$
\begin{equation*}
\mathscr{R}(a \leftrightarrow Z)=\inf \{\mathscr{E}(\theta): \theta \text { is a unit flow from a to } Z\} . \tag{3.3.27}
\end{equation*}
$$

The unique minimizer is the unit current.

Proof. It will be convenient to work in vector form. Let $1, \ldots, n$ be the vertices of $G$ and order the edges arbitrarily as $e_{1}, \ldots, e_{m}$. (We ignore any self-loops, which have no flow.) Choose an arbitrary orientation of $\mathcal{N}$, that is, replace each edge $e_{i}=\{x, y\}$ with either $\vec{e}_{i}=(x, y)$ or $(y, x)$. Let $\vec{G}$ be the corresponding directed graph. Think of the flow $\boldsymbol{\theta}$ as a vector with one coordinate for each oriented edge. Then the flow constraint can be written as a linear system $B \boldsymbol{\theta}=\boldsymbol{b}$. Here the matrix $B$ has a column for each directed edge and a row for each vertex except those in $Z$. The entries of $B$ are $B_{x,(x, y)}=1, B_{y,(x, y)}=-1$, and 0 otherwise. We have already encountered this matrix: it is an oriented incidence matrix of $G$ (see Definition 1.1.16) restricted to the rows in $V \backslash Z$. The vector $\boldsymbol{b}$ has 0 s everywhere except for $b_{a}=1$. Let $r$ be the vector of resistances and let $R$ be the diagonal matrix with diagonal $\boldsymbol{r}$. In vector form, $\mathscr{E}(\boldsymbol{\theta})=\boldsymbol{\theta}^{T} R \boldsymbol{\theta}$ and the optimization problem (3.3.27) reads

$$
\mathscr{E}^{*}=\inf \left\{\boldsymbol{\theta}^{T} R \boldsymbol{\theta}: B \boldsymbol{\theta}=\boldsymbol{b}\right\} .
$$

We first characterize the optimal flow. We introduce the Lagrangian

$$
\mathscr{L}(\boldsymbol{\theta} ; \boldsymbol{h}):=\boldsymbol{\theta}^{T} R \boldsymbol{\theta}-2 \boldsymbol{h}^{T}(B \boldsymbol{\theta}-\boldsymbol{b}),
$$

where $\boldsymbol{h}$ has an entry for all vertices except those in $Z$. For all $\boldsymbol{h}$,

$$
\mathscr{E}^{*} \geq \inf _{\boldsymbol{\theta}} \mathscr{L}(\boldsymbol{\theta} ; \boldsymbol{h})
$$

because those $\boldsymbol{\theta}$ s with $B \boldsymbol{\theta}=\boldsymbol{b}$ make the second term vanish in $\mathscr{L}(\boldsymbol{\theta} ; \boldsymbol{h})$. Since $\mathscr{L}(\boldsymbol{\theta} ; \boldsymbol{h})$ is strictly convex as a function of $\boldsymbol{\theta}$, the solution to its minimization is characterized by the usual optimality conditions which in this case read $2 R \boldsymbol{\theta}$ $2 B^{T} \boldsymbol{h}=0$, or

$$
\begin{equation*}
\boldsymbol{\theta}=R^{-1} B^{T} \boldsymbol{h} . \tag{3.3.28}
\end{equation*}
$$

Substituting into the Lagrangian and simplifying, we have proved that

$$
\begin{equation*}
\mathscr{E}(\boldsymbol{\theta}) \geq \mathscr{E}^{*} \geq-\boldsymbol{h}^{T} B R^{-1} B^{T} \boldsymbol{h}+2 \boldsymbol{h}^{T} \boldsymbol{b}=: \mathscr{L}^{*}(\boldsymbol{h}) \tag{3.3.29}
\end{equation*}
$$

for all $\boldsymbol{h}$ and flow $\boldsymbol{\theta}$. This inequality is a statement of weak duality. To show that a flow $\boldsymbol{\theta}$ is optimal it suffices to find $\boldsymbol{h}$ such that $\mathscr{E}(\boldsymbol{\theta})=\mathscr{L}^{*}(\boldsymbol{h})$.

Let $\boldsymbol{\theta}=\boldsymbol{i}$ be the unit current in vector form, which satisfies $B \boldsymbol{\theta}=\boldsymbol{b}$ by our choice of $\boldsymbol{b}$ and Kirchhoff's node law (i.e., (3.3.21)). The suitable dual turns out to be the corresponding voltage $\boldsymbol{h}=\boldsymbol{v}$ in vector form restricted to $V \backslash Z$. To see this, observe that $B^{T} \boldsymbol{h}$ is the vector of neighboring node differences

$$
\begin{equation*}
B^{T} \boldsymbol{h}=(h(x)-h(y))_{(x, y) \in \vec{G}}, \tag{3.3.30}
\end{equation*}
$$

where implicitly $\left.h\right|_{Z} \equiv 0$. Hence the optimality condition (3.3.28) is nothing but Ohm's law (i.e., (3.3.20)) in vector form. Therefore, if $\boldsymbol{i}$ is the unit current and $\boldsymbol{v}$ is the associated voltage in vector form, it holds that

$$
\mathscr{L}^{*}(\boldsymbol{v})=\mathscr{L}(\boldsymbol{i} ; \boldsymbol{v})=\mathscr{E}(\boldsymbol{i}),
$$

where the first equality follows from the fact that $\boldsymbol{i}$ minimizes $\mathscr{L}(\boldsymbol{i} ; \boldsymbol{v})$ by (3.3.28) and the second equality follows from the fact that $B \boldsymbol{i}=\boldsymbol{b}$. So we must have $\mathscr{E}(\boldsymbol{i})=\mathscr{E}^{*}$ by weak duality (i.e., (3.3.29)).

As for uniqueness, it can be checked that two minimizers $\boldsymbol{\theta}, \boldsymbol{\theta}^{\prime}$ satisfy

$$
\mathscr{E}^{*}=\frac{\mathscr{E}(\boldsymbol{\theta})+\mathscr{E}\left(\boldsymbol{\theta}^{\prime}\right)}{2}=\mathscr{E}\left(\frac{\boldsymbol{\theta}+\boldsymbol{\theta}^{\prime}}{2}\right)+\mathscr{E}\left(\frac{\boldsymbol{\theta}-\boldsymbol{\theta}^{\prime}}{2}\right),
$$

by definition of the energy. The first term in the rightmost expression is greater or equal to $\mathscr{E}^{*}$ since the average of two unit flows is still a unit flow. The second term is nonnegative by definition. Hence the latter must be zero and the only way for this to happen is if $\boldsymbol{\theta}=\boldsymbol{\theta}^{\prime}$.

To conclude the proof, it remains to compute the optimal value. The matrix $B R^{-1} B^{T}$ is related to the Laplacian associated to random walk on $\mathcal{N}$ (see Section 3.3.1) up to a row scaling. Multiplying by row $x \in V \backslash Z$ involves taking a conductance-weighted average of the neighboring values and subtracting the value at $x$, that is,

$$
\begin{aligned}
\left(B R^{-1} B^{T} \boldsymbol{v}\right)_{x}= & \sum_{y:(x, y) \in \vec{G}}[c(x, y)(v(x)-v(y))] \\
& -\sum_{y:(y, x) \in \vec{G}}[c(y, x)(v(y)-v(x))] \\
= & \sum_{y: y \sim x}[c(x, y)(v(x)-v(y))]
\end{aligned}
$$

where we used (3.3.30) and the facts that $r(x, y)^{-1}=c(x, y)$ and $c(x, y)=$ $c(y, x)$, and it is assumed implicitly that $\left.v\right|_{Z} \equiv 0$. By Corollary 3.3.13, this is zero except for the row $x=a$ where it is

$$
\sum_{y: y \sim a} c(a, y)[v(a)-v(y)]=\sum_{y: y \sim a} i(a, y)=1,
$$

where we used Ohm's law and the fact that the current has unit strength. We have
finally

$$
\begin{aligned}
\mathscr{E}^{*} & =\mathscr{L}^{*}(\boldsymbol{v}) \\
& =-\boldsymbol{v}^{T} B R^{-1} B^{T} \boldsymbol{v}+2 \boldsymbol{v}^{T} \boldsymbol{b} \\
& =-v(a)+2 v(a) \\
& =v(a) \\
& =\mathscr{R}(a \leftrightarrow Z),
\end{aligned}
$$

by (3.3.16). That concludes the proof.
Observe that the convex combination $\boldsymbol{\alpha}$ minimizing the sum of squares $\sum_{j} \alpha_{j}^{2}$ is constant. In a similar manner, Thomson's principle (Theorem 3.3.23) stipulates roughly speaking that the more the flow can be spread out over the network, the lower is the effective resistance (penalizing flow on edges with higher resistance). Pólya's theorem below provides a vivid illustration. Here is a simple example suggesting that, in a sense, the current is indeed a well-distributed flow.

Example 3.3.24 (Random walk on the complete graph). Let $\mathcal{N}$ be the complete graph on $\{1, \ldots, n\}$ with unit resistances, and let $a=1$ and $Z=\{n\}$. Assume $n>2$. The effective resistance is straightforward to compute in this case. Indeed, the escape probability (with a slight abuse of notation) is

$$
\mathbb{P}[1 \rightarrow n]=\frac{1}{n-1}+\frac{1}{2}\left(1-\frac{1}{n-1}\right)=\frac{n}{2(n-1)},
$$

as we either jump to $n$ immediately or jump to one of the remaining nodes, in which case we reach $n$ first with probability $1 / 2$ by symmetry. Hence, since $c(1)=n-1$, we get

$$
\mathscr{R}(1 \leftrightarrow n)=\frac{2}{n},
$$

from the definition of the effective resistance (Definition 3.3.19).
We now look for the optimal flow in Thomson's principle. Pushing a flow of 1 through the edge $\{1, n\}$ gives an upper bound of 1 , which is far from the optimal $\frac{2}{n}$. Spreading the flow a bit more by pushing $1 / 2$ through the edge $\{1, n\}$ and $1 / 2$ through the path $1 \sim 2 \sim n$ gives the slightly better bound $3 \cdot(1 / 2)^{2}=3 / 4$. Taking this further, pushing a flow of $\frac{1}{n-1}$ through $\{1, n\}$ as well as through each two-edge path to $n$ via the remaining neighbors of 1 gives the yet improved bound

$$
\left(\frac{1}{n-1}\right)^{2}+2(n-2)\left(\frac{1}{n-1}\right)^{2}=\frac{2 n-3}{(n-1)^{2}}=\frac{2}{n} \cdot \frac{2 n^{2}-3 n}{2 n^{2}-4 n+2}>\frac{2}{n}
$$

when $n>2$. Because the direct path from 1 to $n$ has a somewhat lower resistance, the optimal flow is obtained by increasing the flow on that edge slightly. Namely, for a flow $\alpha$ on $\{1, n\}$ (and the rest divided up evenly among the two-edge paths), we get an energy of $\alpha^{2}+2(n-2)\left[\frac{1-\alpha}{n-2}\right]^{2}$ which is minimized at $\alpha=\frac{2}{n}$ where it is indeed

$$
\left(\frac{2}{n}\right)^{2}+\frac{2}{n-2}\left(\frac{n-2}{n}\right)^{2}=\frac{2}{n}\left(\frac{2}{n}+\frac{n-2}{n}\right)=\frac{2}{n} .
$$

As we noted above, the matrix $B R^{-1} B^{T}$ in the proof of Thomson's principle is related to the Laplacian. Because $B^{T} \boldsymbol{h}$ is the vector of neighboring node differences, we have

$$
\boldsymbol{h}^{T} B R^{-1} B^{T} \boldsymbol{h}=\frac{1}{2} \sum_{x, y} c(x, y)[h(y)-h(x)]^{2},
$$

where we implicitly fix $\left.h\right|_{Z} \equiv 0$, which is called the Dirichlet energy. Thinking of $B^{T}$ as a "discrete gradient," the Dirichlet energy can be interpreted as the weighted norm of the gradient of $\boldsymbol{h}$. The following is a "dual" to Thomson's principle. Exercise 3.15 asks for a proof.

Theorem 3.3.25 (Dirichlet's principle). Let $\mathcal{N}=(G, c)$ be a finite, connected network. The effective conductance between source a and sink $Z$ is characterized by

$$
\left.\mathscr{C}(a \leftrightarrow Z)=\inf \left\{\frac{1}{2} \sum_{x, y} c(x, y)[h(y)-h(x)]^{2}: h(a)=1,\left.h\right|_{Z} \equiv 0\right\}\right\} .
$$

The unique minimizer is the voltage $v$ with $v(a)=1$.
The following lower bound is a typical application of Thomson’s principle. See Pólya's theorem below for an example of its use. Recall from Section 1.1.1 that, on a finite graph, a cutset separating $a$ from $Z$ is a set of edges $\Pi$ such that any path between $a$ and $Z$ must include at least one edge in $\Pi$. Similarly, as defined in Section 2.3.3, on a countable, locally finite network, a cutset separating $a$ from $\infty$ is a finite set of edges that must be crossed by any infinite (self-avoiding) path from $a$.

Corollary 3.3.26 (Nash-Williams inequality). Let $\mathcal{N}$ be a finite, connected network and let $\left\{\Pi_{j}\right\}_{j=1}^{n}$ be a collection of disjoint cutsets separating source a from sink $Z$. Then

$$
\mathscr{R}(a \leftrightarrow Z) \geq \sum_{j=1}^{n}\left(\sum_{e \in \Pi_{j}} c(e)\right)^{-1}
$$

Similarly, if $\mathcal{N}$ is a countable, locally finite, connected network, then for any collection $\left\{\Pi_{j}\right\}_{j}$ of finite, disjoint cutsets separating a from $\infty$,

$$
\mathscr{R}(a \leftrightarrow \infty) \geq \sum_{j}\left(\sum_{e \in \Pi_{j}} c(e)\right)^{-1}
$$

Proof. Consider the case where $\mathcal{N}$ is finite first. We will need the following claim, which follows immediately from Lemma 1.1.14: for any unit flow $\theta$ between $a$ and $Z$ and any cutset $\Pi_{j}$ separating $a$ from $Z$, it holds that

$$
\sum_{e \in \Pi_{j}}|\theta(e)| \geq\|\theta\|=1 .
$$

By Cauchy-Schwarz (Theorem B.4.8),

$$
\begin{aligned}
\sum_{e \in \Pi_{j}} c(e) \sum_{e^{\prime} \in \Pi_{j}} r\left(e^{\prime}\right) \theta\left(e^{\prime}\right)^{2} & \geq\left(\sum_{e \in \Pi_{j}} \sqrt{c(e) r(e)}|\theta(e)|\right)^{2} \\
& =\left(\sum_{e \in \Pi_{j}}|\theta(e)|\right)^{2} \\
& \geq 1
\end{aligned}
$$

Rearranging, summing over $j$ and using the disjointness of the cutsets,

$$
\mathscr{E}(\theta)=\frac{1}{2} \sum_{x, y} r(x, y) \theta(x, y)^{2} \geq \sum_{j=1}^{n} \sum_{e^{\prime} \in \Pi_{j}} r\left(e^{\prime}\right) \theta\left(e^{\prime}\right)^{2} \geq \sum_{j=1}^{n}\left(\sum_{e \in \Pi_{j}} c(e)\right)^{-1} .
$$

Thomson's principle gives the result.
The infinite case follows from a similar argument using an exhaustive sequence.

The following example is an application of Nash-Williams (Corollary 3.3.26) and Thomson's principle to recurrence.

Example 3.3.27 (Biased walk on general trees). Let $\mathscr{T}$ be a locally finite tree with root 0 . Consider again the biased walk from Example 3.3.22, that is, the conductance is $\lambda^{j}$ on all edges between level $j-1$ and $j$. Recall the branching number $\operatorname{br}(\mathscr{T})$ from Definition 2.3.10.

Assume $\lambda>\operatorname{br}(\mathscr{T})$. For any $\varepsilon>0$, there is a cutset $\Pi$ such that $\sum_{e \in \Pi} \lambda^{-|e|} \leq$ ع. By Nash-Williams,

$$
\mathscr{R}(0 \leftrightarrow \infty) \geq\left(\sum_{e \in \Pi} c(e)\right)^{-1} \geq \varepsilon^{-1}
$$

Since $\varepsilon$ is arbitrary, the walk is recurrent by Theorem 3.3.20.
Suppose instead that $\lambda<\operatorname{br}(\mathscr{T})$ and let $\lambda<\lambda_{*}<\operatorname{br}(\mathscr{T})$. By the proof of Claim 2.3.11, for all $n \geq 1$, there exist $\varepsilon>0$ and a unit flow $\phi_{n}$ from 0 to the $n$-level vertices $\partial_{n}$ with capacity constraints $\left|\phi_{n}(x, y)\right| \leq \varepsilon^{-1} \lambda_{*}^{-|e|}$ for all edges $e=\{x, y\}$, where $|e|$ is the graph distance from the root to the endvertex of $e$ furthest from it. Then, letting $F_{m}=\{e:|e|=m\}$, the energy of the flow is

$$
\begin{aligned}
\mathscr{E}\left(\phi_{n}\right) & =\frac{1}{2} \sum_{x, y} r(x, y) \phi_{n}(x, y)^{2} \\
& \leq \sum_{m=1}^{n} \lambda^{m} \sum_{e=\{x, y\} \in F_{m}}\left|\phi_{n}(x, y)\right| \varepsilon^{-1} \lambda_{*}^{-|e|} \\
& =\varepsilon^{-1} \sum_{m=1}^{n}\left(\frac{\lambda}{\lambda_{*}}\right)^{m} \sum_{e=\{x, y\} \in F_{m}}\left|\phi_{n}(x, y)\right| \\
& \leq \varepsilon^{-1} \sum_{m=1}^{+\infty}\left(\frac{\lambda}{\lambda_{*}}\right)^{m} \\
& <+\infty
\end{aligned}
$$

where, on the fourth line, we used Lemma 1.1.14 together with the fact that $\phi_{n}$ is a unit flow and $F_{m}$ is a cutset separating 0 and $\partial_{n}$. Thomson's principle implies that $\mathscr{R}\left(0 \leftrightarrow \partial_{n}\right)$ is uniformly bounded in $n$. The walk is transient by Theorem 3.3.20.

Another typical application of Thomson's principle is the following monotonicity property (which is not obvious from a probabilistic point of view).

Corollary 3.3.28. Adding an edge to a finite, connected network cannot increase the effective resistance between a source a and a sink $Z$. In particular, if the added edge is not incident with $a$, then $\mathbb{P}[a \rightarrow Z]$ cannot decrease.

Proof. The additional edge enlarges the space of possible flows, so by Thomson's principle it can only lower the resistance or leave it as is. The second statement follows from the definition of the effective resistance.

More generally:
Corollary 3.3.29 (Rayleigh's principle). Let $\mathcal{N}$ and $\mathcal{N}^{\prime}$ be two networks on the same finite, connected graph $G$ such that, for each edge in $G$, the resistance in $\mathcal{N}^{\prime}$ is greater than it is in $\mathcal{N}$. Then, for any source $a$ and $\operatorname{sink} Z$,

$$
\mathscr{R}_{\mathcal{N}}(a \leftrightarrow Z) \leq \mathscr{R}_{\mathcal{N}^{\prime}}(a \leftrightarrow Z) .
$$

Proof. Compare the energies of an arbitrary flow on $\mathcal{N}$ and $\mathcal{N}^{\prime}$, and apply Thomson's principle.

Note that this corollary implies the previous one by thinking of an absent edge as one with infinite resistance.

## Flows to infinity

Combining Theorem 3.3.20 and Thomson's principle, we derive a flow-based criterion for recurrence. To state the result, it is convenient to introduce the notion of a unit flow $\theta$ from source a to $\infty$ on a countable, locally finite network: $\theta$ is anti-symmetric, it satisfies the flow-conservation constraint on all vertices but $a$, and $\|\theta\|:=\sum_{y \sim a} \theta(a, y)=1$. Note that the energy $\mathscr{E}(\theta)$ of such a flow is well defined in $[0,+\infty]$.

Theorem 3.3.30 (Recurrence and finite-energy flows). Let $\mathcal{N}=(G, c)$ be a countable, locally finite, connected network. Vertex a (and hence all vertices) in $\mathcal{N}$ is transient if and only if there is a unit flow from a to $\infty$ of finite energy.

Proof. Suppose such a flow exists and has energy bounded by $B<+\infty$. Let $\left(G_{n}\right)$ be an exhaustive sequence with associated sinks $\left(Z_{n}\right)$. A unit flow from $a$ to $\infty$ on $\mathcal{N}$ yields, by projection, a unit flow from $a$ to $Z_{n}$. This projected flow also has energy bounded by $B$. Hence Thomson's principle implies $\mathscr{R}\left(a \leftrightarrow Z_{n}\right) \leq B$ for all $n$ and transience follows from Theorem 3.3.20.

Proving the other direction involves producing a flow to $\infty$. Suppose $a$ is transient and let $\left(G_{n}\right)$ be an exhaustive sequence as above. Then Theorem 3.3.20 implies that $\mathscr{R}\left(a \leftrightarrow Z_{n}\right) \leq \mathscr{R}(a \leftrightarrow \infty)<B$ for some $B<+\infty$ and Thomson's principle guarantees in turn the existence of a flow $\theta_{n}$ from $a$ to $Z_{n}$ with energy bounded by $B$. In particular there is a unit current $i_{n}$, and associated voltage $v_{n}$, of energy bounded by $B$. So it remains to use the sequence of current flows $\left(i_{n}\right)$ to construct a flow to $\infty$ on the infinite network. The technical point is to show that the limit of $\left(i_{n}\right)$ exists and is indeed a flow. For this, consider the random walk on $\mathcal{N}$ started at $a$. Let $Y_{n}(x)$ be the number of visits to $x$ before hitting $Z_{n}$ the first time. By the monotone convergence theorem (Proposition B.4.14),
$\mathbb{E}_{a} Y_{n}(x) \rightarrow \mathbb{E}_{a} Y_{\infty}(x)$ where $Y_{\infty}(x)$ is the total number of visits to $x$. Moreover $\mathbb{E}_{a} Y_{\infty}(x)<+\infty$ by transience and (3.1.2). By (3.3.23), $\mathbb{E}_{a} Y_{n}(x)=c(x) v_{n}(x)$. So we can now define

$$
v_{\infty}(x):=\lim _{n} v_{n}(x)<+\infty,
$$

and then

$$
\begin{aligned}
i_{\infty}(x, y) & :=c(x, y)\left[v_{\infty}(x)-v_{\infty}(y)\right] \\
& =\lim _{n} c(x, y)\left[v_{n}(x)-v_{n}(y)\right] \\
& =\lim _{n} i_{n}(x, y),
\end{aligned}
$$

by Ohm's law (when $n$ is large enough that both $x$ and $y$ are in $G_{n}$ ). Because $i_{n}$ is a flow for all $n$, by taking limits in the flow-conservation constraints we see that so is $i_{\infty}$. Note that by construction of $i_{\ell}$

$$
\begin{aligned}
\frac{1}{2} \sum_{x, y \in G_{n}} c(x, y) i_{\infty}(x, y)^{2} & =\lim _{\ell \geq n} \frac{1}{2} \sum_{x, y \in G_{n}} c(x, y) i_{\ell}(x, y)^{2} \\
& \leq \lim _{\ell \geq n} \sup \mathscr{E}\left(i_{\ell}\right) \\
& <B
\end{aligned}
$$

uniformly in $n$. Because the left-hand side converges to the energy of $i_{\infty}$ as $n \rightarrow$ $+\infty$, we are done.

We give an application to Pólya's theorem in Section 3.3.4.
Finally we derive a useful general result illustrating the robustness reaped from Thomson's principle. At a high level, a rough embedding from $\mathcal{N}$ to $\mathcal{N}^{\prime}$ is a mapping of the edges of $\mathcal{N}$ to paths of $\mathcal{N}^{\prime}$ of comparable overall resistance that do not overlap much. The formal definition follows. As we will see, the purpose of a rough embedding is to allow a flow on $\mathcal{N}$ to be morphed into a flow on $\mathcal{N}^{\prime}$ of comparable energy.
Definition 3.3.31 (Rough embedding). Let $\mathcal{N}=(G, c)$ and $\mathcal{N}^{\prime}=\left(G^{\prime}, c^{\prime}\right)$ be networks with resistances $r$ and $r^{\prime}$ respectively. We say that a map $\phi$ from the vertices of $G$ to the vertices of $G^{\prime}$ is a rough embedding if there are constants $\alpha, \beta<+\infty$ and a map $\Phi$ defined on the edges of $G$ such that:

1. for every edge $e=\{x, y\}$ in $G, \Phi(e)$ is a non-empty path of edges of $G^{\prime}$ between $\phi(x)$ and $\phi(y)$ such that

$$
\sum_{e^{\prime} \in \Phi(e)} r^{\prime}\left(e^{\prime}\right) \leq \alpha r(e) ;
$$

2. for every edge $e^{\prime}$ in $G^{\prime}$, there are no more than $\beta$ edges in $G$ whose image under $\Phi$ contains $e^{\prime}$.

The map $\phi$ need not in general be a bijection.
We say that two networks are roughly equivalent if there exist rough embeddings between them, one in each direction.
roughly
equivalent

Example 3.3.32 (Independent-coordinate random walk). Let $\mathcal{N}=\mathbb{L}^{d}$ with unit resistances and let $\mathcal{N}^{\prime}$ be the network corresponding to the independent-coordinate random walk

$$
\left(Y_{t}^{(1)}, \ldots, Y_{t}^{(d)}\right)
$$

where each coordinate $\left(Y_{t}^{(i)}\right)$ is an independent simple random walk on $\mathbb{Z}$ started at 0 . For example the neighborhood of the origin in $\mathcal{N}^{\prime}$ is $\left\{\left(x_{1}, \ldots, x_{d}\right): x_{i} \in\right.$ $\{-1,1\}, \forall i\}$. Note that $\mathcal{N}^{\prime}$ contains only those points of $\mathbb{Z}^{d}$ with coordinates of identical parities.

Despite encoding quite different random walks, we claim that the networks $\mathcal{N}$ and $\mathcal{N}^{\prime}$ are roughly equivalent.

- $\mathcal{N}$ to $\mathcal{N}^{\prime}$ : Consider the map $\phi$ which associates to each $x \in \mathcal{N}$ a closest point in $\mathcal{N}^{\prime}$ chosen in some arbitrary manner. For $\Phi$, associate to each edge $e=\{x, y\} \in \mathcal{N}$ a shortest path in $\mathcal{N}^{\prime}$ between $\phi(x)$ and $\phi(y)$, again chosen arbitrarily. If $\phi(x)=\phi(y)$, choose an arbitrary, non-empty, shortest cycle through $\phi(x)$.
- $\mathcal{N}^{\prime}$ to $\mathcal{N}$ : Consider the map $\phi$ which associates to each $x \in \mathcal{N}^{\prime}$ the corresponding point $x$ in $\mathcal{N}$. Construct $\Phi$ similarly to the previous case.

Exercise 3.19 asks for a rigorous proof of rough equivalence. See also Exercise 3.20 for an important generalization of this example.

Our main result about roughly equivalent networks is that they have the same type.

Theorem 3.3.33 (Recurrence and rough equivalence). Let $\mathcal{N}$ and $\mathcal{N}^{\prime}$ be roughly equivalent, locally finite, connected networks. Then $\mathcal{N}$ is transient if and only if $\mathcal{N}^{\prime}$ is transient.

Proof. Assume $\mathcal{N}$ is transient and let $\theta$ be a unit flow from some $a$ to $\infty$ of finite energy. The existence of this flow is guaranteed by Theorem 3.3.30. Let $\phi, \Phi$ be a rough embedding from $\mathcal{N}$ to $\mathcal{N}^{\prime}$ with parameters $\alpha$ and $\beta$.

The basic idea of the proof is to map the flow $\theta$ onto $\mathcal{N}^{\prime}$ using $\Phi$. Because flows are directional, it will be convenient to think of edges as being directed. For


Figure 3.5: The flow on $\left(x^{\prime}, y^{\prime}\right)$ is the sum of the flows on $\left(x_{1}, y_{1}\right),\left(x_{2}, y_{2}\right)$, and $\left(x_{3}, y_{3}\right)$.
$e=\{x, y\}$ in $\mathcal{N}$, let $\vec{\Phi}(x, y)$ be the path $\Phi(e)$ oriented from $\phi(x)$ to $\phi(y)$. So $\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)$ means that $\left\{x^{\prime}, y^{\prime}\right\} \in \Phi(e)$ and that $x^{\prime}$ is visited before $y^{\prime}$ in the path $\Phi(e)$ from $\phi(x)$ to $\phi(y)$. (If $\phi(x)=\phi(y)$, choose an arbitrary orientation of the cycle $\Phi(e)$ for $\vec{\Phi}(x, y)$ and the reversed orientation for $\vec{\Phi}(y, x)$.) Then define, for $x^{\prime}, y^{\prime}$ with $\left\{x^{\prime}, y^{\prime}\right\}$ in $\mathcal{N}^{\prime}$,

$$
\begin{equation*}
\theta^{\prime}\left(x^{\prime}, y^{\prime}\right):=\sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} \theta(x, y) . \tag{3.3.31}
\end{equation*}
$$

See Figure 3.5.
We claim that $\theta^{\prime}$ is a flow to $\infty$ of finite energy on $\mathcal{N}^{\prime}$. We first check that $\theta^{\prime}$ is a flow.

1. (Anti-symmetry) By construction, $\theta^{\prime}\left(y^{\prime}, x^{\prime}\right)=-\theta^{\prime}\left(x^{\prime}, y^{\prime}\right)$, that is, $\theta^{\prime}$ is antisymmetric, because $\theta$ itself is anti-symmetric. We used the fact that $\vec{\Phi}(y, x)$ is $\vec{\Phi}(x, y)$ oriented in the opposite direction.
2. (Flow conservation) Next we check the flow-conservation constraints. Fix $z^{\prime}$ in $\mathcal{N}^{\prime}$. By Condition 2 in Definition 3.3.31, there are finitely many edges $e$ in $\mathcal{N}$ such that $\Phi(e)$ visits $z^{\prime}$. Let $e=\{x, y\}$ be such an edge. There are two cases:

- Assume first that $\phi(x), \phi(y) \neq z^{\prime}$ and let $\left(u^{\prime}, z^{\prime}\right),\left(z^{\prime}, w^{\prime}\right)$ be the directed edges incident with $z^{\prime}$ on $\vec{\Phi}(x, y)$. Observe that, in the definition of $\theta^{\prime},(y, x)$ contributes $\theta(y, x)=-\theta(x, y)$ to $\theta^{\prime}\left(z^{\prime}, u^{\prime}\right)$ and $(x, y)$ contributes $\theta(x, y)$ to $\theta^{\prime}\left(z^{\prime}, w^{\prime}\right)$. So these contributions cancel out in the flow-conservation constraint for $z^{\prime}$, that is, in the sum $\sum_{v^{\prime}: v^{\prime} \sim z^{\prime}} \theta^{\prime}\left(z^{\prime}, v^{\prime}\right)$.
- If instead $e=\{x, y\}$ is such that $\phi(x)=z^{\prime}$, let $\left(z^{\prime}, w^{\prime}\right)$ be the first edge on the path $\vec{\Phi}(x, y)$. Edge $(x, y)$ contributes $\theta(x, y)$ to $\theta^{\prime}\left(z^{\prime}, w^{\prime}\right)$. A similar statement applies to $\phi(y)=z^{\prime}$ by changing the role of $x$ and $y$. This case also applies to $\phi(x)=\phi(y)=z^{\prime}$.

From the two cases above, summing over all paths visiting $z^{\prime}$ gives

$$
\sum_{v^{\prime}: v^{\prime} \sim z^{\prime}} \theta^{\prime}\left(z^{\prime}, v^{\prime}\right)=\sum_{z: \phi(z)=z^{\prime}}\left(\sum_{v: v \sim z} \theta(z, v)\right) .
$$

Because $\theta$ is a flow, the sum in parentheses is 0 if $z \neq a$ and 1 otherwise. So the right-hand side is 0 unless $a \in \phi^{-1}\left(\left\{z^{\prime}\right\}\right)$ in which case it is 1 .

We have shown that $\theta^{\prime}$ is a unit flow from $\phi(a)$ to $\infty$. It remains to bound the energy of $\theta^{\prime}$. By (3.3.31), Cauchy-Schwarz, and Condition 2 in Definition 3.3.31,

$$
\begin{aligned}
\theta^{\prime}\left(x^{\prime}, y^{\prime}\right)^{2} & =\left[\sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} \theta(x, y)\right]^{2} \\
& \leq\left[\sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} 1\right]\left[\sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} \theta(x, y)^{2}\right] \\
& \leq \beta \sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} \theta(x, y)^{2} .
\end{aligned}
$$

Summing over all pairs and using Condition 1 in Definition 3.3.31 gives

$$
\begin{aligned}
\frac{1}{2} \sum_{x^{\prime}, y^{\prime}} r^{\prime}\left(x^{\prime}, y^{\prime}\right) \theta^{\prime}\left(x^{\prime}, y^{\prime}\right)^{2} & \leq \beta \frac{1}{2} \sum_{x^{\prime}, y^{\prime}} r^{\prime}\left(x^{\prime}, y^{\prime}\right) \sum_{(x, y):\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} \theta(x, y)^{2} \\
& =\beta \frac{1}{2} \sum_{x, y} \theta(x, y)^{2} \sum_{\left(x^{\prime}, y^{\prime}\right) \in \vec{\Phi}(x, y)} r^{\prime}\left(x^{\prime}, y^{\prime}\right) \\
& \leq \alpha \beta \frac{1}{2} \sum_{x, y} r(x, y) \theta(x, y)^{2},
\end{aligned}
$$

which is finite by assumption. That concludes the proof.
As an application, we give a second proof of Pólya's theorem in Section 3.3.4.

## Other applications

So far we have emphasized applications to recurrence. Here we show that electrical network theory can also be used to bound commute times. In Section 3.3.5, we give further applications beyond random walks on graphs.

An application of Corollary 3.1.24 gives another probabilistic interpretation of the effective resistance-and a useful formula.

Theorem 3.3.34 (Commute time identity). Let $\mathcal{N}=(G, c)$ be a finite, connected network with vertex set $V$. For $x \neq y$, let the commute time $\tau_{x, y}$ be the time of the first return to $x$ after the first visit to $y$. Then

$$
\mathbb{E}_{x}\left[\tau_{x, y}\right]=\mathbb{E}_{x}\left[\tau_{y}\right]+\mathbb{E}_{y}\left[\tau_{x}\right]=c_{\mathcal{N}} \mathscr{R}(x \leftrightarrow y),
$$

where $c_{\mathcal{N}}=2 \sum_{e=\{x, y\} \in \mathcal{N}} c(e)$.
Proof. This follows immediately from Corollary 3.1.24 and the definition of the effective resistance (Definition 3.3.19). Specifically,

$$
\begin{aligned}
\mathbb{E}_{x}\left[\tau_{y}\right]+\mathbb{E}_{y}\left[\tau_{x}\right] & =\frac{1}{\pi_{x} \mathbb{P}_{x}\left[\tau_{y}<\tau_{x}^{+}\right]} \\
& =\frac{1}{\left(2 \sum_{e=\{x, y\} \in \mathcal{N}} c(e)\right)^{-1} c(x) \mathbb{P}_{x}\left[\tau_{y}<\tau_{x}^{+}\right]} \\
& =c_{\mathcal{N}} \mathscr{R}(x \leftrightarrow y) .
\end{aligned}
$$

Example 3.3.35 (Random walk on the torus). Consider random walk on the $d$ dimensional torus $\mathbb{L}_{n}^{d}$ with unit resistances. We use the commute time identity to lower bound the mean hitting time $\mathbb{E}_{x}\left[\tau_{y}\right]$ for arbitrary vertices $x \neq y$ at graph distance $k$ on $\mathbb{L}_{n}^{d}$. To use the commute time identity (Theorem 3.3.34), note that by symmetry $\mathbb{E}_{x}\left[\tau_{y}\right]=\mathbb{E}_{y}\left[\tau_{x}\right]$ so that

$$
\begin{equation*}
\mathbb{E}_{x}\left[\tau_{y}\right]=\frac{1}{2} c_{\mathcal{N}} \mathscr{R}(x \leftrightarrow y)=d n^{d} \mathscr{R}(x \leftrightarrow y) . \tag{3.3.32}
\end{equation*}
$$

where we used that the number of vertices is $n^{d}$ and the graph is $2 d$-regular.
To simplify, assume $n$ is odd and identify the vertices of $\mathbb{L}_{n}^{d}$ with the box

$$
B:=\{-(n-1) / 2, \ldots,(n-1) / 2\}^{d},
$$

in $\mathbb{L}^{d}$ centered at $x=0$. Let $\partial B_{j}^{\infty}=\left\{z \in \mathbb{L}^{d}:\|z\|_{\infty}=j\right\}$ and let $\Pi_{j}$ be the set of edges between $\partial B_{j}^{\infty}$ and $\partial B_{j+1}^{\infty}$. Note that on $B$ the $\ell^{1}$ norm of $y$ is at most $k$ (the graph distance between $x=0$ and $y$ ). Since the $\ell^{\infty}$ norm is at least $1 / d$ times the $\ell^{1}$ norm on $\mathbb{L}^{d}$, there exists $J=O(k)$ such that all $\Pi_{j} \mathrm{~s}, j \leq J$, are cutsets separating $x$ from $y$. By the Nash-Williams inequality

$$
\mathscr{R}(x \leftrightarrow y) \geq \sum_{0 \leq j \leq J}\left|\Pi_{j}\right|^{-1}=\sum_{0 \leq j \leq J} \Omega\left(j^{-(d-1)}\right)= \begin{cases}\Omega(\log k), & d=2 \\ \Omega(1), & d \geq 3\end{cases}
$$

From (3.3.32), we get:

## Claim 3.3.36.

$$
\mathbb{E}_{x}\left[\tau_{y}\right]= \begin{cases}\Omega\left(n^{d} \log k\right), & d=2 \\ \Omega\left(n^{d}\right), & d \geq 3\end{cases}
$$

Remark 3.3.37. The bounds in the previous example are tight up to constants. See [LPW06, Proposition 10.13]. Note that the case $d \geq 3$ does not in fact depend on the distance $k$.

See Exercise 3.22 for an application of the commute time identity to cover times.

### 3.3.4 $\triangleright$ Random walks: Pólya's theorem, two ways

The following is a classical result.
Theorem 3.3.38 (Pólya's theorem). Random walk on $\mathbb{L}^{d}$ is recurrent for $d \leq 2$ and transient for $d \geq 3$.

We prove the theorem for $d=2,3$ using the tools developed in the previous subsection. The other cases follow by Rayleigh's principle (Corollary 3.3.29). There are elementary proofs of this result. But we showed above that the electrical network approach has the advantage of being robust to the details of the lattice. For a different argument, see Exercise 2.10.

The case $d=2$ follows from the Nash-Williams inequality (Corollary 3.3.26) by letting $\Pi_{j}$ be the set of edges connecting vertices of $\ell^{\infty}$ norm $j$ and $j+1$. Using the fact that all conductances are 1 , that $\left|\Pi_{j}\right|=O(j)$, and that $\sum_{j} j^{-1}$ diverges, recurrence is established by Theorem 3.3.20.

## First proof

Now consider the case $d=3$ and let $a=0$ be the origin.
We construct a finite-energy flow to $\infty$ using the method of random paths. Note that a simple way to produce a unit flow to $\infty$ is to push a flow of 1 through an infinite path (which, recall, are self-avoiding by definition). Taking this a step further, let $\mu$ be a probability measure on infinite paths and define the anti-symmetric function

$$
\theta(x, y):=\mathbb{E}\left[\mathbf{1}_{(x, y) \in \Gamma}-\mathbf{1}_{(y, x) \in \Gamma}\right]=\mathbb{P}[(x, y) \in \Gamma]-\mathbb{P}[(y, x) \in \Gamma],
$$

where $\Gamma$ is a random path distributed according to $\mu$, oriented away from 0 . (We will give an explicit construction below where the appropriate formal probability space will be clear.) Observe that $\sum_{y \sim x}\left[\mathbf{1}_{(x, y) \in \Gamma}-\mathbf{1}_{(y, x) \in \Gamma}\right]=0$ for any $x \neq 0$ because vertices visited by $\Gamma$ are entered and exited exactly once. That same sum is 1 at $x=0$. Hence $\theta$ is a unit flow to $\infty$. For edge $e=\{x, y\}$, consider the following "edge marginal" of $\mu$ :

$$
\mu(e):=\mathbb{P}[(x, y) \in \Gamma \text { or }(y, x) \in \Gamma]=\mathbb{P}[(x, y) \in \Gamma]+\mathbb{P}[(y, x) \in \Gamma] \geq \theta(x, y),
$$

where we used that a path $\Gamma$ cannot visit both $(x, y)$ and $(y, x)$ by definition. Then we get the following bound.

Claim 3.3.39 (Method of random paths).

$$
\begin{equation*}
\mathscr{E}(\theta) \leq \sum_{e} \mu(e)^{2} \tag{3.3.33}
\end{equation*}
$$

For a measure $\mu$ concentrated on a single path, the sum above is infinite. To obtain a useful bound, what we need is a large collection of spread out paths. On the lattice $\mathbb{L}^{3}$, we construct $\mu$ as follows. Let $\boldsymbol{U}$ be a uniformly random point on the unit sphere in $\mathbb{R}^{3}$ and let $\gamma$ be the ray from 0 to $\infty$ going through $\boldsymbol{U}$. Imagine centering a unit cube around each point in $\mathbb{Z}^{3}$ whose edges are aligned with the axes. Then $\gamma$ traverses an infinite number of such cubes. Let $\Gamma$ be the corresponding path in the lattice $\mathbb{L}^{3}$. To see that this procedure indeed produces a path observe that $\gamma$, upon exiting a cube around a point $z \in \mathbb{Z}^{3}$, enters the cube of a neighboring point $z^{\prime} \in \mathbb{Z}^{3}$ through a face corresponding to the edge between $z$ and $z^{\prime}$ on the lattice $\mathbb{L}^{3}$ (unless it goes through a corner of the cube, but this has probability 0 ). To argue that $\mu$ distributes its mass among sufficiently spread out paths, we bound the probability that a vertex is visited by $\Gamma$. Let $z$ be an arbitrary vertex in $\mathbb{Z}^{3}$. Because the sphere of radius $\|z\|_{2}$ around the origin in $\mathbb{R}^{3}$ has area $O\left(\|z\|_{2}^{2}\right)$ and its intersection with the unit cube centered around $z$ has area $O(1)$, it follows that

$$
\mathbb{P}[z \in \Gamma]=O\left(1 /\|z\|_{2}^{2}\right)
$$

That immediately implies a similar bound on the probability that an edge is visited by $\Gamma$. Moreover:

Lemma 3.3.40. There are $O\left(j^{2}\right)$ edges with an endpoint at $\ell^{2}$ distance within $[j, j+1]$ from the origin.

Proof. Consider an open ball of $\ell^{2}$ radius $1 / 2$ centered around each vertex of $\ell^{2}$ norm within $[j, j+1]$. Those balls are non-intersecting and have total volume $\Theta\left(N_{j}\right)$, where $N_{j}$ is the number of such vertices. On the other hand, the volume of the shell of $\ell^{2}$ inner and outer radii $j-1 / 2$ and $j+3 / 2$ centered around the origin (where all those balls lie) is

$$
\frac{4}{3} \pi(j+3 / 2)^{3}-\frac{4}{3} \pi(j-1 / 2)^{3}=O\left(j^{2}\right) .
$$

Hence $N_{j}=O\left(j^{2}\right)$. Finally note that each vertex has 6 incident edges.
Plugging those bounds into (3.3.33), we get

$$
\mathscr{E}(\theta) \leq \sum_{j} O\left(j^{2}\right) \cdot\left[O\left(1 / j^{2}\right)\right]^{2}=O\left(\sum_{j} j^{-2}\right)<+\infty .
$$

Transience follows from Theorem 3.3.30. (This argument clearly does not work on $\mathbb{L}$ where there are only two rays. You should convince yourself that it does not work on $\mathbb{L}^{2}$ either. But see Exercise 3.17.)

## Second proof

We briefly describe a second proof based on the independent-coordinate random walk. Consider the networks $\mathcal{N}$ and $\mathcal{N}^{\prime}$ in Example 3.3.32. Because they are roughly equivalent (Definition 3.3.31), they have the same type by Theorem 3.3.33. Recall that, because the number of returns to 0 is geometric with success probability equal to the escape probability, random walk on $\mathcal{N}^{\prime}$ is transient if and only if the expected number of visits to 0 is finite (see (3.1.2)). By independence of the coordinates, this expectation can be written as

$$
\sum_{t \geq 0}\left(\mathbb{P}\left[Y_{2 t}^{(1)}=0\right]\right)^{d}=\sum_{t \geq 0}\left(\binom{2 t}{t} 2^{-2 t}\right)^{d}=\sum_{t \geq 0} \Theta\left(t^{-d / 2}\right),
$$

where we used Stirling's formula (see Appendix A). The rightmost sum is finite if and only if $d \geq 3$. That implies random walk on $\mathcal{N}^{\prime}$ is transient under that condition. By rough equivalence, the same is true of $\mathcal{N}$.

### 3.3.5 $\triangleright$ Randomized algorithms: Wilson's method for generating uniform spanning trees

In this section, we describe an application of electrical network theory to spanning trees.

With a slight abuse of notation, we use $e \in G$ to indicate that $e$ is an edge of $G$.

Uniform spanning trees Let $G=(V, E)$ be a finite connected graph. Recall that a spanning tree is a subtree of $G$ containing all its vertices. Such a tree has $|V|-1$ edges. A uniform spanning tree is a spanning tree $T$ chosen uniformly at random among all spanning trees of $G$.

We make some simple observations first. Because $G$ is connected, it has at least one spanning tree by Corollary 1.1.6. Moreover, for any edge $e \in G$, there always exists at least one spanning tree including it. To see this, let $T^{\prime}$ be any spanning tree of $G$, which exists by the previous observation. If $e \notin T^{\prime}$, then we obtain a new spanning tree by adding $e$ to $T^{\prime}$ and removing one edge $\neq e$ in the cycle created. As a consequence, the probability of inclusion $\mathbb{P}[e \in T]$ in a uniform spanning tree $T$ cannot be 0 . It is however possible for $\mathbb{P}[e \in T]$ to equal to 1 if removing $e$ disconnects the graph. Such an edge is called a bridge.

A fundamental property of uniform spanning trees is the following negative correlation between edges.

Claim 3.3.41. For a uniform spanning tree $T$ of a connected graph $G$,

$$
\mathbb{P}\left[e \in T \mid e^{\prime} \in T\right] \leq \mathbb{P}[e \in T], \quad \forall e \neq e^{\prime} \in G .
$$

This property is perhaps not surprising. For one, the number of edges in a spanning tree is fixed, so the inclusion of $e^{\prime}$ makes it seemingly less likely for other edges to be present. Yet proving Claim 3.3.41 is not trivial. The proof relies on the electrical network perspective. The key is a remarkable formula for the inclusion of an edge in a uniform spanning tree.

Theorem 3.3.42 (Kirchhoff's resistance formula). Let $G=(V, E)$ be a finite, connected graph and let $\mathcal{N}$ be the network on $G$ with unit resistances. If $T$ is a uniform spanning tree on $G$, then for all $e=\{x, y\}$

$$
\mathbb{P}[e \in T]=\mathscr{R}(x \leftrightarrow y) .
$$

Before explaining how this formula arises, we show that it implies Claim 3.3.41.

Proof of Claim 3.3.41. Recall that $\mathbb{P}\left[e^{\prime} \in T\right] \neq 0$. By the law of total probability,

$$
\mathbb{P}[e \in T]=\mathbb{P}\left[e \in T \mid e^{\prime} \in T\right] \mathbb{P}\left[e^{\prime} \in T\right]+\mathbb{P}\left[e \in T \mid e^{\prime} \notin T\right] \mathbb{P}\left[e^{\prime} \notin T\right],
$$

so, since $\mathbb{P}\left[e^{\prime} \in T\right]+\mathbb{P}\left[e^{\prime} \notin T\right]=1$, we can instead prove

$$
\begin{equation*}
\mathbb{P}\left[e \in T \mid e^{\prime} \notin T\right] \geq \mathbb{P}[e \in T] . \tag{3.3.34}
\end{equation*}
$$

Picking a uniform spanning tree on $\mathcal{N}$ conditioned on $\left\{e^{\prime} \notin T\right\}$ is the same as picking a uniform spanning tree on the modified network $\mathcal{N}^{\prime}$ where $e^{\prime}$ is removed. By Rayleigh's principle (in the form of Corollary 3.3.28),

$$
\mathscr{R}_{\mathcal{N}^{\prime}}(x \leftrightarrow y) \geq \mathscr{R}_{\mathcal{N}}(x \leftrightarrow y),
$$

and Kirchhoff's resistance formula (Theorem 3.3.42) gives (3.3.34).
Remark 3.3.43. More generally, thinking of a uniform spanning tree $T$ as a random subset of edges, the law of $T$ has the property of negative associations, defined as follows. An event $\mathcal{A} \subseteq 2^{E}$ is said to be increasing if $\omega \cup\{e\} \in \mathcal{A}$ whenever $\omega \in \mathcal{A}$. The event $\mathcal{A}$ is said to depend only on $F \subseteq E$ if for all $\omega_{1}, \omega_{2} \in 2^{E}$ that agree on $F$, either both are in $\mathcal{A}$ or neither is. The law $\mathbb{P}_{T}$ of $T$ has negative associations in the sense that for any two increasing events $\mathcal{A}$ and $\mathcal{B}$ that depend only on disjoint sets of edges, we have $\mathbb{P}_{T}[\mathcal{A} \cap \mathcal{B}] \leq \mathbb{P}_{T}[\mathcal{A}] \mathbb{P}_{T}[\mathcal{B}]$. See [LP16, Exercise 4.6].

Let $e=\{x, y\}$. To get some insight into Kirchhoff's resistance formula, we first note that, if $i$ is the unit current from $x$ to $y$ and $v$ is the associated voltage, by definition of the effective resistance

$$
\begin{equation*}
\mathscr{R}(x \leftrightarrow y)=\frac{v(x)}{\|i\|}=c(e)(v(x)-v(y))=i(x, y) \tag{3.3.35}
\end{equation*}
$$

where we used Ohm's law (i.e., (3.3.20)) as well as the fact that $c(e)=1, v(y)=0$, and $\|i\|=1$. Note that $\|i\|$ and $i(x, y)$ are not the same quantity: although $\|i\|=1$, $i(x, y)$ is only the current along the edge to $y$. Furthermore by the probabilistic interpretation of the current (Theorem 3.3.17), with $Z=\{y\}$,

$$
\begin{equation*}
i(x, y)=\mathbb{E}_{x}\left[N_{x \rightarrow y}^{Z}-N_{y \rightarrow x}^{Z}\right]=\mathbb{P}_{x}\left[(x, y) \text { is traversed before } \tau_{y}\right] . \tag{3.3.36}
\end{equation*}
$$

Indeed, started at $x, N_{y \rightarrow x}^{Z}=0$ and $N_{x \rightarrow y}^{Z} \in\{0,1\}$. Kirchhoff's resistance formula is then established by relating the random walk on $\mathcal{N}$ to the probability that $e$ is present in a uniform spanning tree $T$. To do this we introduce a random-walkbased algorithm for generating uniform spanning trees. This rather miraculous procedure, known as Wilson's method, is of independent interest. (For a classical connection between random walks and spanning trees, see also Exercise 3.23.)

Wilson's method It will be somewhat easier to work in a more general context. Let $\mathcal{N}=(G, c)$ be a finite, connected network on $G$ with arbitrary conductances and define the weight of a spanning tree $T$ on $\mathcal{N}$ as

$$
W(T)=\prod_{e \in T} c(e) .
$$

With a slight abuse, we continue to call a tree $T$ picked at random among all spanning trees of $G$ with probability proportional to $W(T)$ a "uniform" spanning tree on $\mathcal{N}$.

To state Wilson's method, we need the notion of loop erasure. Let $\mathcal{P}=x_{0} \sim$ $\ldots \sim x_{k}$ be a walk in $\mathcal{N}$. The loop erasure of $\mathcal{P}$ is obtained by removing cycles in the order they appear. That is, let $j^{*}$ be the smallest $j$ such that $x_{j}=x_{\ell}$ for some $\ell<j$. Remove the subwalk $x_{\ell+1} \sim \cdots \sim x_{j}$ from $\mathcal{P}$, and repeat. The result is self-avoiding, that is, a path, and is denoted by $\operatorname{LE}(\mathcal{P})$.

Let $v_{0}$ be an arbitrary vertex of $G$, which we refer to as the root, and let $T_{0}$ be the subtree made up of $v_{0}$ alone. Starting with the root, order arbitrarily the vertices of $G$ as $v_{0}, \ldots, v_{n-1}$. Wilson's method constructs an increasing sequence of subtrees as follows. See Figure 3.6. Let $T:=T_{0}$.

1. Let $v$ be the vertex of $G$ not in $T$ with lowest index. Perform random walk on $\mathcal{N}$ started at $v$ until the first visit to a vertex of $T$. Let $\mathcal{P}$ be the resulting walk.
2. Add the loop erasure $\operatorname{LE}(\mathcal{P})$ to $T$.
3. Repeat until all vertices of $G$ are in $T$.

Let $T_{0}, \ldots, T_{m}$ be the sequence of subtrees produced by Wilson's method.
Claim 3.3.44. Forgetting the root, $T_{m}$ is a uniform spanning tree on $\mathcal{N}$.
This claim is far from obvious. Before proving it, we finish the proof of Kirchhoff's resistance formula.

Proof of Theorem 3.3.42. From (3.3.35) and (3.3.36), it suffices to prove that, for $e=\{x, y\}$,

$$
\mathbb{P}_{x}\left[(x, y) \text { is traversed before } \tau_{y}\right]=\mathbb{P}[e \in T],
$$

where the probability on the left-hand side refers to random walk on $\mathcal{N}$ with unit resistances started at $x$ and the probability on the right-hand side refers to a uniform spanning tree $T$ on $\mathcal{N}$. Generate $T$ using Wilson's method started at root $v_{0}=y$ with the choice $v_{1}=x$. If the walk from $x$ to $y$ during the first iteration of Wilson's


Figure 3.6: An illustration of Wilson's method. The dashed lines indicate erased loops.
method includes $(x, y)$, then the loop erasure is simply $x \sim y$ and $e$ is in $T$. On the other hand, if the walk from $x$ to $y$ does not include $(x, y)$, then $e$ cannot be used at a later stage because it would create a cycle. That immediately proves the theorem.

It remains to prove the claim.
Proof of Claim 3.3.44. The idea of the proof is to cast Wilson's method in the more general framework of cycle popping algorithms. We begin by explaining how such algorithms work.

Let $P$ be the transition matrix corresponding to random walk on $\mathcal{N}=(G, c)$ with $G=(V, E)$ and root $v_{0}$. To each vertex $x \neq v_{0}$ in $V$, we assign an independent stack of "colored directed edges"

$$
\mathcal{S}_{0}^{x}:=\left(\left\langle x, Y_{1}^{x}\right\rangle_{1},\left\langle x, Y_{2}^{x}\right\rangle_{2}, \ldots\right)
$$

where each $Y_{j}^{x}$ is chosen independently at random from the distribution $P(x, \cdot)$. In particular all $Y_{j}^{x}$ s are neighbors of $x$ in $\mathcal{N}$. The index $j$ in $\left\langle x, Y_{j}^{x}\right\rangle_{j}$ is the color of the edge. It keeps track of the position of the edge in the original stack. (Picture $\mathcal{S}^{x}$ as a spring-loaded plate dispenser located on vertex $x$.)

We consider a process which involves popping edges off the stacks. We use the notation $\mathcal{S}^{x}$ to denote the current stack at $x$. The initial assignment of the stack is $\mathcal{S}^{x}:=\mathcal{S}_{0}^{x}$ as above. Given the current stacks $\left(\mathcal{S}^{x}\right)_{x}$, we call visible graph the (colored) directed graph over $V$ with edges $\operatorname{Top}\left(\mathcal{S}^{x}\right)$ for all $x \neq v_{0}$, where $\operatorname{Top}\left(\mathcal{S}^{x}\right)$ is the first edge in the current stack $\mathcal{S}^{x}$. The latter are referred to as visible edges. We denote the current visible graph by $\vec{G}_{\odot}$.

Note that $\vec{G}_{\odot}$ has out-degree 1 for all $x \neq v_{0}$ and the root has out-degree 0 . In particular all undirected cycles in $\vec{G}_{\odot}$ are in fact directed cycles, and we refer to them simply as cycles. (Indeed a set of edges forming an undirected cycle that is not directed must have a vertex of out-degree 2.) Also recall the following characterization from Corollary 1.1.8: an acyclic, undirected subgraph with $|V|$ vertices and $|V|-1$ edges is a spanning tree of $G$. Hence if there is no cycle in $\vec{G}_{\odot}$, then it must be a spanning tree (as an undirected graph) where, furthermore, all edges point towards the root. Such a tree is also known as a spanning arborescence. Once that happens, we are done.

As the name suggests, a cycle popping algorithm proceeds by popping cycles in $\vec{G}_{\odot}$ off the tops of the stacks until a spanning arborescence is produced. That is, at every iteration, if $\vec{G}_{\odot}$ contains at least one cycle, then a cycle $\vec{C}$ is picked according to some rule, the top of each stack in $\vec{C}$ is popped, and a new visible graph $\vec{G}_{\odot}$ is revealed. See Figure 3.7 for an illustration.


Figure 3.7: A realization of a cycle popping algorithm (from top to bottom). In all three figures, the underlying graph is $G$ while the arrows depict the visible edges.

With these definitions in place, the proof of the claim involves the following steps.
(i) Wilson's method is a cycle popping algorithm. Recasting Wilson's method, we can think of the initial stacks $\left(\mathcal{S}_{0}^{x}\right)$ as corresponding to picking-ahead of time-all potential transitions in the random walks. With this representation, the algorithm boils down to a recipe for choosing which cycle to pop next. Indeed, at each iteration, we start from a vertex $v$ not in the current tree $T$. A key observation: following the visible edges from $v$ traces a path whose distribution is that of random walk on $\mathcal{N}$. Loop erasure then corresponds to popping cycles as they are closed. We pop only those visible edges on the removed cycles, as they originate from vertices that will be visited again by the algorithm and for which a new transition will then be needed. Those visible edges in the resulting loop-erased path are not popped-note that they are part of the final arborescence.
(ii) The popping order does not matter. We just argued that Wilson's method is a cycle popping algorithm. In fact we claim that any cycle popping algorithm, that is, no matter what popping choices are made along the way, produces the same final arborescence. To make this precise, we identify the popped cycles uniquely. This is where the colors come in. A colored cycle is a directed cycle over $V$ made of colored edges from the stacks (not necessarily of the same color and not necessarily in the current visible graph). We say that a colored cycle $\vec{C}$ is poppable for a visible graph $\vec{G}_{\odot}$ if there exists a sequence of colored cycles $\vec{C}_{1}, \ldots, \vec{C}_{r}=\vec{C}$ that can be popped in that order starting from $\vec{G}_{\odot}$. Note that, by this definition, $\vec{C}_{1}$ is a cycle in $\vec{G}_{\odot}$. Now we claim that if $\vec{C}_{1}^{\prime}$ were popped first instead of $\vec{C}_{1}$, producing the new visible graph $\vec{G}_{\odot}^{\prime}$, then $\vec{C}$ would still be poppable for $\vec{G}_{\odot}^{\prime}$. This claim implies that, in any cycle popping algorithm, either an infinite number of cycles are popped or eventually all poppable cycles are popped-independently of the order-producing the same outcome. (Note that, while the same cycle may be popped more than once, the same colored cycle cannot.)
To prove the claim, note first that if $\vec{C}_{1}^{\prime}=\vec{C}$ or if $\vec{C}_{1}^{\prime}$ does not share a vertex with any of $\vec{C}_{1}, \ldots, \vec{C}_{r}$ there is nothing to prove. So let $\vec{C}_{j}$ be the first cycle in the sequence sharing a vertex with $\vec{C}_{1}^{\prime}$, say $x$. Let $\langle x, y\rangle_{c}$ and $\left\langle x, y^{\prime}\right\rangle_{c^{\prime}}$ be the colored edges emanating from $x$ in $\vec{C}_{j}$ and $\vec{C}_{1}^{\prime}$ respectively. By definition, $x$ is not on any of $\vec{C}_{1}, \ldots, \vec{C}_{j-1}$ so the edge originating from $x$ is not popped by that sequence and we must have $\langle x, y\rangle_{c}=\left\langle x, y^{\prime}\right\rangle_{c^{\prime}}$ as colored edges. In particular, the vertex $y$ is also a shared vertex of $\vec{C}_{j}$ and $\vec{C}_{1}^{\prime}$, and
the same argument applies to it. Proceeding by induction leads to the conclusion that $\vec{C}_{1}^{\prime}=\vec{C}_{j}$ as colored cycles. But then $\vec{C}$ is clearly poppable for the visible graph resulting from popping $\vec{C}_{1}^{\prime}$ first, because it can be popped with the rearranged sequence $\vec{C}_{1}^{\prime}=\vec{C}_{j}, \vec{C}_{1}, \ldots, \vec{C}_{j-1}, \vec{C}_{j+1}, \ldots, \vec{C}_{r}=\vec{C}$, where we used the fact that $\vec{C}_{1}^{\prime}$ does not share a vertex with $\vec{C}_{1}, \ldots, \vec{C}_{j-1}$.
(iii) Termination occurs in finite time almost surely. We have shown so far that, in any cycle popping algorithm, either an infinite number of cycles are popped or eventually all poppable cycles are popped. But Wilson's method-a cycle popping algorithm as we have shown-stops after a finite amount of time with probability 1 . Indeed, because the network is finite and connected, the random walk started at each iteration hits the current $T$ in finite time almost surely (by Lemma 3.1.25). To sum up, all cycle popping algorithms terminate and produce the same spanning arborescence. It remains to compute the distribution of the outcome.
(iv) The arborescence has the desired distribution. Let $\mathcal{A}$ be the spanning arborescence produced by any cycle popping algorithm on the stacks $\left(\mathcal{S}_{0}^{x}\right)$. To compute the distribution of $\mathcal{A}$, we first compute the distribution of a particular cycle popping realization leading to $\mathcal{A}$. Because the popping order does not matter, by "realization" we mean a collection $\mathcal{C}$ of colored cycles together with a final spanning arborescence $\mathcal{A}$. Notice that what lies in the stacks "under" $\mathcal{A}$ is not relevant to the realization, that is, the same outcome is produced no matter what is under $\mathcal{A}$.
So, from the distribution of the stacks, the probability of observing $(\mathcal{C}, \mathcal{A})$ is simply the product of the transitions corresponding to the "popped edges" in $\mathcal{C}$ and the "final edges" in $\mathcal{A}$, that is,

$$
\prod_{\vec{e} \in \mathcal{C} \cup \mathcal{A}} P(\vec{e})=\Psi(\mathcal{A}) \prod_{\vec{C} \in \mathcal{C}} \Psi(\vec{C}),
$$

where the function $\Psi$ returns the product of the transition probabilities of a set of directed edges. Thanks to the product form on the right-hand side, summing over all possible $\mathcal{C}$ s gives that the probability of producing $\mathcal{A}$ is proportional to $\Psi(\mathcal{A})$.
For this argument to work though, there are two small details to take care of. First, note that we want the probability of the "uncolored" arborescence. But observe that, in fact, there is no need to keep track of the colors on the edges of $\mathcal{A}$ because these are determined by $\mathcal{C}$. Secondly, we need for the collection
of possible $\mathcal{C}$ not to vary with $\mathcal{A}$. But it is clear that any arborescence could lie under any $\mathcal{C}$.

To see that we are done, let $T$ be the undirected spanning tree corresponding to the outcome, $\mathcal{A}$, of Wilson's method. Then, because $P(x, y)=\frac{c(x, y)}{c(x)}$, we get

$$
\Psi(\mathcal{A})=\frac{W(T)}{\prod_{x \neq v_{0}} c(x)},
$$

where note that the denominator does not depend on $T$. So if we forget the orientation of $\mathcal{A}$ which is determined by the root (i.e., sum over all choices of root), we get a spanning tree whose distribution is proportional to $W(T)$, as required.

## Exercises

Exercise 3.1 (Reflection). Give a rigorous proof of Theorem 3.1.9 through a formal application of the strong Markov property (i.e., specify $f_{t}$ and $F_{t}$ in Theorem 3.1.8).

Exercise 3.2 (Time of $k$-th return). Give a rigorous proof of (3.1.1) through a formal application of the strong Markov property (i.e., specify $f_{t}$ and $F_{t}$ in Theorem 3.1.8).

Exercise 3.3 (Tightness of Matthews' bounds). Show that the bounds (3.1.6) and (3.1.7) are tight up to smaller order terms for the coupon collector problem (Example 2.1.4). [Hint: State the problem in terms of the cover time of a random walk on the complete graph with self-loops.]

Exercise 3.4 (Pólya's urn: a suprisingly simple formula). Consider the setting of Example 3.1.49. Prove that

$$
\mathbb{P}\left[G_{t}=m+1\right]=\binom{t}{m} \frac{m!(t-m)!}{(t+1)!} .
$$

[Hint: Consider the probability of one particular sequence of outcomes producing the desired event.]

Exercise 3.5 (Optional stopping theorem). Give a rigorous proof of the remaining cases of the optional stopping theorem (Theorem 3.1.38).

Exercise 3.6 (Supermartingale inequality). Let ( $M_{t}$ ) be a nonnegative, supermartingale. Show that, for any $b>0$,

$$
\mathbb{P}\left[\sup _{s \geq 0} M_{s} \geq b\right] \leq \frac{\mathbb{E}\left[M_{0}\right]}{b} .
$$

[Hint: Mimic the proof of the submartingale case.]
Exercise 3.7 (Azuma-Hoeffding: a second proof). This exercise leads the reader through an alternative proof of the Azuma-Hoeffding inequality.
(i) Show that for all $x \in[-1,1]$ and $a>0$

$$
e^{a x} \leq \cosh a+x \sinh a .
$$

(ii) Use a Taylor expansion to show that for all $x$

$$
\cosh x \leq e^{x^{2} / 2}
$$

(iii) Let $X_{1}, \ldots, X_{n}$ be (not necessarily independent) random variables such that, for all $i,\left|X_{i}\right| \leq c_{i}$ for some constant $c_{i}<+\infty$ and

$$
\begin{equation*}
\mathbb{E}\left[X_{i_{1}} \cdots X_{i_{k}}\right]=0, \quad \forall 1 \leq k \leq n, \forall 1 \leq i_{1}<\cdots<i_{k} \leq n \tag{3.3.37}
\end{equation*}
$$

Show, using (i) and (ii), that for all $b>0$

$$
\mathbb{P}\left[\sum_{i=1}^{n} X_{i} \geq b\right] \leq \exp \left(-\frac{b^{2}}{2 \sum_{i=1}^{n} c_{i}^{2}}\right) .
$$

(iv) Prove that (iii) implies a variant of the Azuma-Hoeffding inequality (Theorem 3.2.1) for bounded increments.
(v) Show that the random variables in Exercise 2.6 (after centering) do not satisfy (3.3.37) (without using the claim in part (ii) of that exercise).

Exercise 3.8 (Lipschitz condition). Give a rigorous proof of Lemma 3.2.31.
Exercise 3.9 (Lower bound on expected spectral norm). Let $A$ be an $n \times n$ random matrix. Assume that the entries $A_{i, j}, i, j=1, \ldots, n$, are independent, centered random variables in $[-1,1]$. Suppose further that there is $0<\sigma^{2}<+\infty$ such that $\operatorname{Var}\left[A_{i j}\right] \geq \sigma^{2}$ for all $i, j$. Show that there is $0<c<+\infty$ such that

$$
\mathbb{E}\|A\| \geq c \sqrt{n}
$$

for $n$ large enough. [Hint: Use the fact that $\|A\|^{2} \geq\left\|A \mathbf{e}_{1}\right\|^{2}$ together with Chebyshev's inequality.]

Exercise 3.10 (Kirchhoff's laws). Consider a finite, connected network with a source and a sink. Show that an anti-symmetric function on the edges satisfying Kirchhoff's two laws is a current function (i.e., it corresponds to a voltage function through Ohm's law).

Exercise 3.11 (Dirichlet problem: non-uniqueness). Let $\left(X_{t}\right)$ be the birth-anddeath chain on $\mathbb{Z}_{+}$with $P(x, x+1)=p$ and $P(x, x-1)=1-p$ for all $x \geq 1$, and $P(0,1)=1$, for some $0<p<1$. Fix $h(0)=1$.
(i) When $p>1 / 2$, show that there is more than one bounded extension of $h$ to $\mathbb{Z}_{+} \backslash\{0\}$ that is harmonic on $\mathbb{Z}_{+} \backslash\{0\}$. [Hint: Consider $\mathbb{P}_{x}\left[\tau_{0}=+\infty\right]$.]
(ii) When $p \leq 1 / 2$, show that there exists a unique bounded extension of $h$ to $\mathbb{Z}_{+} \backslash\{0\}$ that is harmonic on $\mathbb{Z}_{+} \backslash\{0\}$.

Exercise 3.12 (Maximum principle). Let $\mathcal{N}=(G, c)$ be a finite or countable, connected network with $G=(V, E)$. Let $W$ be a finite, connected, proper subset of $V$.
(i) Let $h: V \rightarrow \mathbb{R}$ be a function on $V$. Prove the maximum principle: if $h$ is harmonic on $W$, that is, it satisfies

$$
h(x)=\frac{1}{c(x)} \sum_{y \sim x} c(x, y) h(y), \quad \forall x \in W
$$

and if $h$ achieves its supremum on $W$, then $h$ is constant on $W \cup \partial_{\mathrm{V}} W$, where

$$
\partial_{\mathrm{V}} W=\{z \in V \backslash W: \exists y \in W, y \sim z\}
$$

(ii) Let $h: W^{c} \rightarrow \mathbb{R}$ be a bounded function on $W^{c}:=V \backslash W$. Let $h_{1}$ and $h_{2}$ be extensions of $h$ to $W$ that are harmonic on $W$. Use part (i) to prove that $h_{1} \equiv h_{2}$.

Exercise 3.13 (Poisson equation: uniqueness). Show that $u$ is the unique solution of the system in Theorem 3.3.6 under the conditions of Theorem 3.3.1. [Hint: Use Theorem 3.3.9 and mimic the proof of Theorem 3.3.1.]

Exercise 3.14 (Effective resistance: metric). Show that effective resistances between pairs of vertices form a metric.
Exercise 3.15 (Dirichlet principle: proof). Prove Theorem 3.3.25.
Exercise 3.16 (Martingale problem). Let $V$ be countable, let $\left(X_{t}\right)$ be a stochastic process adapted to $\left(\mathcal{F}_{t}\right)$ and taking values in $V$, and let $P$ be a transition probability on $V$ with associated Laplacian operator $\Delta$. Show that the following are equivalent:
(i) The process $\left(X_{t}\right)$ is a Markov chain with transition probability $P$.
(ii) For any bounded measurable function $f: V \rightarrow \mathbb{R}$, the process

$$
M_{t}^{f}=f\left(X_{t}\right)-\sum_{s=0}^{t-1} \Delta f\left(X_{s}\right),
$$

is a martingale with respect to $\left(\mathcal{F}_{t}\right)$.
Exercise 3.17 (Random walk on $\mathbb{L}^{2}$ : effective resistance). Consider random walk on $\mathbb{L}^{2}$, which we showed is recurrent. Let $\left(G_{n}\right)$ be the exhaustive sequence corresponding to vertices at distance at most $n$ from the origin and let $Z_{n}$ be the corresponding sink-set. Show that $\mathscr{R}\left(0 \leftrightarrow Z_{n}\right)=\Theta(\log n)$. [Hint: Use the NashWilliams inequality and the method of random paths.]

Exercise 3.18 (Random walk on regular graphs: effective resistance). Let $G$ be a $d$-regular graph with $n$ vertices and $d>n / 2$. Let $\mathcal{N}$ be the network $(G, c)$ with unit conductances. Let $a$ and $z$ be arbitrary distinct vertices.
(i) Show that there are at least $2 d-n$ vertices $x \neq a, z$ such that $a \sim x \sim z$ is a path.
(ii) Prove that

$$
\mathbb{E}_{a}\left(\tau_{a, z}\right) \leq \frac{2 d n}{2 d-n} .
$$

Exercise 3.19 (Independent-coordinate random walk). Give a rigorous proof that the two networks in Example 3.3.32 are roughly equivalent.

Exercise 3.20 (Rough isometries). Graphs $G=(V, E)$ and $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ are roughly isometric (or quasi-isometric) if there is a map $\phi: V \rightarrow V^{\prime}$ and constants $0<\alpha, \beta<+\infty$ such that for all $x, y \in V$

$$
\alpha^{-1} d(x, y)-\beta \leq d^{\prime}(\phi(x), \phi(y)) \leq \alpha d(x, y)+\beta,
$$

where $d$ and $d^{\prime}$ are the graph distances on $G$ and $G^{\prime}$ respectively, and furthermore all vertices in $G^{\prime}$ are within distance $\beta$ of the image of $V$. Let $\mathcal{N}=(G, c)$ and $\mathcal{N}^{\prime}=\left(G^{\prime}, c^{\prime}\right)$ be countable, connected networks with uniformly bounded conductances, resistances and degrees. Prove that if $G$ and $G^{\prime}$ are roughly isometric then $\mathcal{N}$ and $\mathcal{N}^{\prime}$ are roughly equivalent. [Hint: Start by proving that being roughly isometric is an equivalence relation.]

Exercise 3.21 (Random walk on the cycle: hitting time). Use the commute time identity (Theorem 3.3.34) to compute $\mathbb{E}_{x}\left[\tau_{y}\right]$ in Example 3.3.35 in the case $d=1$. Give a second proof using a direct martingale argument.

Exercise 3.22 (Random walk on the binary tree: cover time). As in Example 3.3.15, let $\mathcal{N}$ be the rooted binary tree with $n$ levels $\widehat{\mathbb{T}}_{2}^{n}$ and equal conductances on all edges.
(i) Show that the maximal hitting time $\mathbb{E}_{a} \tau_{b}$ is achieved for $a$ and $b$ such that their most recent common ancestor is the root 0 . Furthermore argue that in that case $\mathbb{E}_{a}\left[\tau_{b}\right]=\mathbb{E}_{a}\left[\tau_{a, 0}\right]$, where recall that $\tau_{a, 0}$ is the commute time between $a$ and 0 .
(ii) Use the commute time identity (Theorem 3.3.34) and Matthews' cover time bounds (Theorem 3.1.27) to give an upper bound on the mean cover time of the order of $O\left(n^{2} 2^{n}\right)$.

Exercise 3.23 (Markov chain tree theorem). Let $P$ be the transition matrix of a finite, irreducible Markov chain with stationary distribution $\pi$. Let $G$ be the directed graph corresponding to the positive transitions of $P$. For an arborescence $\mathcal{A}$ of $G$, define its weight as

$$
\Psi(\mathcal{A})=\prod_{\vec{e} \in \mathcal{A}} P(\vec{e}) .
$$

Consider the following process on spanning arborescences over $G$. Let $\rho$ be the root of the current spanning arborescence $\mathcal{A}$. Pick an outgoing edge $\vec{e}=(\rho, x)$ of $\rho$ according to $P(\rho, \cdot)$. Edge $\vec{e}$ is not in $\mathcal{A}$ by definition of an arborescence. Add $\vec{e}$ to $\mathcal{A}$. This creates a cycle. Remove the edge of this cycle originating from $x$, producing a new arborescence $\mathcal{A}^{\prime}$ with root $x$. Repeat the process.
(i) Show that this chain is irreducible.
(ii) Show that $\Psi$ is a stationary measure for this chain.
(iii) Prove the Markov chain tree theorem: The stationary distribution $\pi$ of $P$ is proportional to

$$
\pi_{x}=\sum_{\mathcal{A}: \operatorname{root}(\mathcal{A})=x} \Psi(\mathcal{A}) .
$$

## Bibliographic Remarks

Section 3.1 Picking up where Appendix B leaves off, Sections 3.1.1 and 3.1.3 largely follow the textbooks [Wi191] and [Dur10], which contain excellent introductions to martingales. The latter also covers Markov chains, and includes the proofs we skipped here. Theorem 3.1.11 is proved in [Dur10, Theorem 4.3.2]. Many more results like Corollary 3.1.24 can be derived from the occupation measure identity; see for example [AF, Chapter 2]. The upper bound in Theorem 3.1.27 was first proved by Matthews [Mat88].

Section 3.2 The Azuma-Hoeffding inequality is due to Hoeffding [Hoe63] and Azuma [Azu67]. The version of the inequality in Exercise 3.7 is from [Ste97]. The method of bounded differences has its origins in the works of Yurinskii [Yur76], Maurey [Mau79], Milman and Schechtman [MS86], Rhee and Talagrand [RT87], and Shamir and Spencer [SS87]. In its current form, it appears in [McD89]. Example 3.2.11 is taken from [MU05, Section 12.5]. The presentation in Section 3.2.3 follows [AS11, Section 7.3]. Claim 3.2.16 is due to Shamir and Spencer [SS87]. The 2-point concentration result alluded to in Section 3.2.3 is due to Alon and Krivelevich [AK97]. For the full story on the chromatic number of Erdős-Rényi graphs, see [JLR11, Chapter 7]. Claim 3.2.21 is due to Bollobás, Riordan, Spencer, and Tusnády [BRST01]. It confirmed simulations of Barabási and Albert [BA99]. The expectation was analyzed by Dorogovtsev, Mendes, and Samukhin [DMS00]. For much more on preferential attachment models see [Dur06], [CL06], or [vdH17]. Example 3.2.12 borrows from [BLM13, Section 7.1] and [Pet, Section 6.3]. General references on the concentration of measure phenomenon and concentration inequalities are [Led01] and [BLM13]. See [BCB12] or [LS20] for an introduction to bandit problems; or [AJKS22] for an introduction to the sample complexity of the more general reinforcement learning problem. The slicing argument in Section 3.2.5 is based on [Bub10]. A more general discussion of the slicing method, whose best known application is the proof of the law of the iterated logarithm (e.g., [Wil91, Section 14.7]), can be found in [vH16]. Section 3.2.6 is based on [vH16, Section 4.3]. In particular, a proof of Talagrand's inequality (Theorem 3.2.32) can be found there. See also [AS11, Chapter 7] or [BLM13, Chapter 7].

Section 3.3 Section 3.3.1 is based partly on [Nor98, Sections 4.1-2], [Ebe, Sections 0.3, 1.1-2, 3.1-2], and [Bre17, Sections 7.3, 17.1]. The material in Sections 3.3.2-3.3.5 borrows from [LPW06, Chapters 9, 10], [AF, Chapters 2, 3] and, especially, [LP16, Sections 2.1-2.6, 4.1-4.2, 5.5]. Foster's theorem (Theorem 3.3.12)
is from [Fos53]. The classical reference on potential theory and its probabilistic counterpart is [Doo01]. For the discrete case and the electrical network point of view, the book of Doyle and Snell is excellent [DS84]. In particular the series and parallel laws are defined and illustrated. See also [KSK76]. For an introduction to convex optimization and duality, see for example [BV04]. The Nash-Williams inequality is due to Nash-Williams [NW59]. The result in Example 3.3.27 is due to R. Lyons [Lyo90]. Theorem 3.3.33 is due to Kanai [Kan86]. The commute time identity was proved by Chandra, Raghavan, Ruzzo, Smolensky and Tiwari [CRR ${ }^{+} 89$ ]. An elementary proof of Pólya's theorem can be found in [Dur10, Section 4.2]. The flow we used in the proof of Pólya's theorem is essentially due to T. Lyons [Lyo83]. Wilson's method is due to Wilson [Wil96]. A related method for generating uniform spanning trees was introduced by Aldous [Ald90] and Broder [Bro89]. A connection between loop-erased random walks and uniform spanning trees had previously been established by Pemantle [Pem91] using the Aldous-Broder method. For more on negative correlation in uniform spanning trees, see for example [LP16, Section 4.2]. For a proof of the matrix tree theorem using Wilson's method, see [KRS]. For a discussion of the running time of Wilson's method and other spanning tree generation approaches, see [Wil96].

## Chapter 4

## Coupling

In this chapter we move on to coupling, another probabilistic technique with a wide range of applications (far beyond discrete stochastic processes). The idea behind the coupling method is deceptively simple: to compare two probability measures $\mu$ and $\nu$, it is sometimes useful to construct a joint probability space with marginals $\mu$ and $\nu$. For instance, in the classical application of coupling to the convergence of Markov chains (Theorem 1.1.33), one simultaneously constructs two copies of a Markov chain-one of which is already at stationarity-and shows that they can be made to coincide after a random amount of time called the coupling time. We begin in Section 4.1 by defining coupling formally and deriving its connection to the total variation distance through the coupling inequality. We illustrate the basic idea on a classical Poisson approximation result, which we apply to the degree sequence of an Erdős-Rényi graph. In Section 4.2, we introduce the concept of stochastic domination and some related correlation inequalities. We develop a key application in percolation theory. Coupling of Markov chains is the subject of Section 4.3, where it serves as a powerful tool to derive mixing time bounds. Finally, we end in Section 4.4 with the Chen-Stein method for Poisson approximations, a technique that applies in particular in some natural settings with dependent variables.

### 4.1 Background

We begin with some background on coupling. After defining the concept formally and giving a few simple examples, we derive the coupling inequality, which provides a fundamental approach to bounding the distance between two distributions.

[^5]As an application, we analyze the degree distribution in the Erdős-Rényi graph model. Throughout this chapter, $(S, \mathcal{S})$ is a measurable space. Also we will denote by $\mu_{Z}$ the law of random variable $Z$.

### 4.1.1 Basic definitions

A formal definition of coupling follows. Recall (see Appendix B) that for measurable spaces $\left(S_{1}, \mathcal{S}_{1}\right)\left(S_{2}, \mathcal{S}_{2}\right)$, we can consider the product space $\left(S_{1} \times S_{2}, \mathcal{S}_{1} \times \mathcal{S}_{2}\right)$ where

$$
S_{1} \times S_{2}:=\left\{\left(s_{1}, s_{2}\right): s_{1} \in S_{1}, s_{2} \in S_{2}\right\}
$$

is the Cartesian product of $S_{1}$ and $S_{2}$, and $\mathcal{S}_{1} \times \mathcal{S}_{2}$ is the smallest $\sigma$-algebra on $S_{1} \times S_{2}$ containing the rectangles $A_{1} \times A_{2}$ for all $A_{1} \in \mathcal{S}_{1}$ and $A_{2} \in \mathcal{S}_{2}$.
Definition 4.1.1 (Coupling). Let $\mu$ and $\nu$ be probability measures on the same measurable space $(S, \mathcal{S})$. A coupling of $\mu$ and $\nu$ is a probability measure $\gamma$ on the product space $(S \times S, \mathcal{S} \times \mathcal{S})$ such that the marginals of $\gamma$ coincide with $\mu$ and $\nu$, that is,

$$
\gamma(A \times S)=\mu(A) \quad \text { and } \quad \gamma(S \times A)=\nu(A), \quad \forall A \in \mathcal{S} .
$$

For two random variables $X$ and $Y$ taking values in $(S, \mathcal{S}), a$ coupling of $X$ and $Y$ is a joint variable $\left(X^{\prime}, Y^{\prime}\right)$ taking values in $(S \times S, \mathcal{S} \times \mathcal{S})$ whose law as a probability measure is a coupling of the laws of $X$ and $Y$. Note that, under this definition, $X$ and $Y$ need not be defined on the same probability space (but $X^{\prime}$ and $Y^{\prime}$ do need to). We also say that $\left(X^{\prime}, Y^{\prime}\right)$ is a coupling of $\mu$ and $\nu$ if the law of $\left(X^{\prime}, Y^{\prime}\right)$ is a coupling of $\mu$ and $\nu$.

We give a few examples.
Example 4.1.2 (Coupling of Bernoulli variables). Let $X$ and $Y$ be Bernoulli random variables with respective parameters $0 \leq q<r \leq 1$. That is, $\mathbb{P}[X=1]=q$ and $\mathbb{P}[Y=1]=r$. Here $S=\{0,1\}$ and $\mathcal{S}=2^{S}$.

- (Independent coupling) One coupling of $X$ and $Y$ is $\left(X^{\prime}, Y^{\prime}\right)$ where $X^{\prime} \stackrel{\text { d }}{=} X$ and $Y^{\prime} \stackrel{\mathrm{d}}{=} Y$ are independent of one another. Its law is

$$
\left(\mathbb{P}\left[\left(X^{\prime}, Y^{\prime}\right)=(i, j)\right]\right)_{i, j \in\{0,1\}}=\left(\begin{array}{cc}
(1-q)(1-r) & (1-q) r \\
q(1-r) & q r
\end{array}\right) .
$$

- (Monotone coupling) Another possibility is to pick $U$ uniformly at random in $[0,1]$, and set $X^{\prime \prime}=\mathbf{1}_{\{U \leq q\}}$ and $Y^{\prime \prime}=\mathbf{1}_{\{U \leq r\}}$. Then $\left(X^{\prime \prime}, Y^{\prime \prime}\right)$ is a coupling of $X$ and $Y$ with law

$$
\left(\mathbb{P}\left[\left(X^{\prime \prime}, Y^{\prime \prime}\right)=(i, j)\right]\right)_{i, j \in\{0,1\}}=\left(\begin{array}{cc}
1-r & r-q \\
0 & q
\end{array}\right) .
$$

Example 4.1.3 (Bond percolation: monotonicity). Let $G=(V, E)$ be a countable graph. Denote by $\mathbb{P}_{p}$ the law of bond percolation (Definition 1.2.1) on $G$ with density $p$. Let $x \in V$ and assume $0 \leq q<r \leq 1$. Using the monotone coupling in the previous example on each edge independently produces a coupling of $\mathbb{P}_{q}$ and $\mathbb{P}_{r}$. More precisely:

- Let $\left\{U_{e}\right\}_{e \in E}$ be independent uniforms on $[0,1]$.
- For $p \in[0,1]$, let $W_{p}$ be the set of edges $e$ such that $U_{e} \leq p$.

Thinking of $W_{p}$ as specifying the open edges in the percolation process on $G$ under $\mathbb{P}_{p}$, we see that $\left(W_{q}, W_{r}\right)$ is a coupling of $\mathbb{P}_{q}$ and $\mathbb{P}_{r}$ with the property that $\mathbb{P}\left[W_{q} \subseteq\right.$ $\left.W_{r}\right]=1$. Let $\mathcal{C}_{x}^{(q)}$ and $\mathcal{C}_{x}^{(r)}$ be the open clusters of $x$ under $W_{q}$ and $W_{r}$ respectively. Because $\mathcal{C}_{x}^{(q)} \subseteq \mathcal{C}_{x}^{(r)}$,

$$
\begin{aligned}
\theta(q) & :=\mathbb{P}_{q}\left[\left|\mathcal{C}_{x}\right|=+\infty\right] \\
& =\mathbb{P}\left[\left|\mathcal{C}_{x}^{(q)}\right|=+\infty\right] \\
& \leq \mathbb{P}\left[\left|\mathcal{C}_{x}^{(r)}\right|=+\infty\right] \\
& =\mathbb{P}_{r}\left[\left|\mathcal{C}_{x}\right|=+\infty\right] \\
& =\theta(r) .
\end{aligned}
$$

(We made this claim in Section 2.2.4.)
Example 4.1.4 (Biased random walk on $\mathbb{Z})$. For $p \in[0,1]$, let $\left(S_{t}^{(p)}\right)$ be nearestneighbor random walk on $\mathbb{Z}$ started at 0 with probability $p$ of jumping to the right and probability $1-p$ of jumping to the left. (See the gambler's ruin problem in Example 3.1.43.) Assume $0 \leq q<r \leq 1$. Using again the monotone coupling of Bernoulli variables above we produce a coupling of $\left(S_{t}^{(q)}\right)$ and $\left(S_{t}^{(r)}\right)$.

- Let $\left(X_{i}^{\prime \prime}, Y_{i}^{\prime \prime}\right)_{i}$ be an infinite sequence of i.i.d. monotone Bernoulli couplings with parameters $q$ and $r$ respectively.
- Define $\left(Z_{i}^{(q)}, Z_{i}^{(r)}\right):=\left(2 X_{i}^{\prime \prime}-1,2 Y_{i}^{\prime \prime}-1\right)$. Note that $\mathbb{P}\left[2 X_{1}^{\prime \prime}-1=1\right]=$ $\mathbb{P}\left[X_{1}^{\prime \prime}=1\right]=q$ and $\mathbb{P}\left[2 X_{1}^{\prime \prime}-1=-1\right]=\mathbb{P}\left[X_{1}^{\prime \prime}=0\right]=1-q$, and similarly for $Y_{i}^{\prime \prime}$.

$$
\text { - Let } \hat{S}_{t}^{(q)}=\sum_{i \leq t} Z_{i}^{(q)} \text { and } \hat{S}_{t}^{(r)}=\sum_{i \leq t} Z_{i}^{(r)} .
$$

Then $\left(\hat{S}_{t}^{(q)}, \hat{S}_{t}^{(r)}\right)$ is a coupling of $\left(S_{t}^{(q)}\right)$ and $\left(S_{t}^{(r)}\right)$ such that $\hat{S}_{t}^{(q)} \leq \hat{S}_{t}^{(r)}$ for all $t$ almost surely. In particular, we deduce that for all $y$ and all $t$

$$
\mathbb{P}\left[S_{t}^{(q)} \leq y\right]=\mathbb{P}\left[\hat{S}_{t}^{(q)} \leq y\right] \geq \mathbb{P}\left[\hat{S}_{t}^{(r)} \leq y\right]=\mathbb{P}\left[S_{t}^{(r)} \leq y\right]
$$

### 4.1.2 $\triangleright$ Random walks: harmonic functions on lattices and infinite $d$ regular trees

Let $\left(X_{t}\right)$ be a Markov chain on a finite or countably infinite state space $V$ with transition matrix $P$ and let $\mathbb{P}_{x}$ be the law of $\left(X_{t}\right)$ started at $x$. We say that a function $h: V \rightarrow \mathbb{R}$ is bounded if $\sup _{x \in V}|h(x)|<+\infty$. Recall from Section 3.3 that $h$ is harmonic (with respect to $P$ ) on $V$ if

$$
h(x)=\sum_{y \in V} P(x, y) h(y), \quad \forall x \in V .
$$

We first give a coupling-based criterion for bounded harmonic functions to be constant. Recall that we treated the finite state-space case (where boundedness is automatic) in Corollary 3.3.3.

Lemma 4.1.5 (Coupling and bounded harmonic functions). If, for all $y, z \in V$, there is a coupling $\left(\left(Y_{t}\right)_{t},\left(Z_{t}\right)_{t}\right)$ of $\mathbb{P}_{y}$ and $\mathbb{P}_{z}$ such that

$$
\lim _{t} \mathbb{P}\left[Y_{t} \neq Z_{t}\right]=0,
$$

then all bounded harmonic functions on $V$ are constant.
Proof. Let $h$ be bounded and harmonic on $V$ with $\sup _{x}|h(x)|=M<+\infty$. Let $y, z$ be any points in $V$. Then, arguing as in Section 3.3.1, $\left(h\left(Y_{t}\right)\right)$ and $\left(h\left(Z_{t}\right)\right)$ are martingales and, in particular,

$$
\mathbb{E}\left[h\left(Y_{t}\right)\right]=\mathbb{E}\left[h\left(Y_{0}\right)\right]=h(y) \quad \text { and } \quad \mathbb{E}\left[h\left(Z_{t}\right)\right]=\mathbb{E}\left[h\left(Z_{0}\right)\right]=h(z) .
$$

So by Jensen's inequality (Theorem B.4.15) and the boundedness assumption

$$
\begin{aligned}
|h(y)-h(z)| & =\left|\mathbb{E}\left[h\left(Y_{t}\right)\right]-\mathbb{E}\left[h\left(Z_{t}\right)\right]\right| \\
& \leq \mathbb{E}\left|h\left(Y_{t}\right)-h\left(Z_{t}\right)\right| \\
& \leq 2 M \mathbb{P}\left[Y_{t} \neq Z_{t}\right] \\
& \rightarrow 0 .
\end{aligned}
$$

So $h(y)=h(z)$.

Harmonic functions on $\mathbb{Z}^{d} \quad$ Consider random walk on $\mathbb{L}^{d}$ for $d \geq 1$. In that case, we show that all bounded harmonic functions are constant.

Theorem 4.1.6 (Bounded harmonic functions on $\mathbb{Z}^{d}$ ). All bounded harmonic functions on $\mathbb{L}^{d}$ are constant.

Proof. From (3.3.2), $h$ is harmonic with respect to random walk on $\mathbb{L}^{d}$ if and only if it is harmonic with respect to lazy random walk (Definition 1.1.31), that is, the walk that stays put with probability $1 / 2$ at every step. Let $\mathbb{P}_{y}$ and $\mathbb{P}_{z}$ be the laws of lazy random walk on $\mathbb{L}^{d}$ started at $y$ and $z$ respectively. We construct a coupling $\left(\left(Y_{t}\right),\left(Z_{t}\right)\right)=\left(\left(Y_{t}^{(i)}\right)_{i \in[d]},\left(Z_{t}^{(i)}\right)_{i \in[d]}\right)$ of $\mathbb{P}_{y}$ and $\mathbb{P}_{z}$ as follows: at time $t$, pick a coordinate $I \in[d]$ uniformly at random, then

- if $Y_{t}^{(I)}=Z_{t}^{(I)}$ then do nothing with probability $1 / 2$ and otherwise pick $W \in\{-1,+1\}$ uniformly at random, set $Y_{t+1}^{(I)}=Z_{t+1}^{(I)}:=Z_{t}^{(I)}+W$ and leave the other coordinates unchanged;
- if instead $Y_{t}^{(I)} \neq Z_{t}^{(I)}$, pick $W \in\{-1,+1\}$ uniformly at random, and with probability $1 / 2$ set $Y_{t+1}^{(I)}:=Y_{t}^{(I)}+W$ and leave $Z_{t}$ and the other coordinates of $Y_{t}$ unchanged, or otherwise set $Z_{t+1}^{(I)}:=Z_{t}^{(I)}+W$ and leave $Y_{t}$ and the other coordinates of $Z_{t}$ unchanged.

It is straightforward to check that $\left(\left(Y_{t}\right),\left(Z_{t}\right)\right)$ is indeed a coupling of $\mathbb{P}_{y}$ and $\mathbb{P}_{z}$. To apply the previous lemma, it remains to bound $\mathbb{P}\left[Y_{t} \neq Z_{t}\right]$.

The key is to note that, for each coordinate $i$, the difference $\left(Y_{t}^{(i)}-Z_{t}^{(i)}\right)$ is itself a nearest-neighbor random walk on $\mathbb{Z}$ started at $y^{(i)}-z^{(i)}$ with holding probability (i.e., probability of staying put) $1-\frac{1}{d}$-until it hits 0 . Simple random walk on $\mathbb{Z}$ is irreducible and recurrent (Theorem 3.3.38). The holding probability does not affect the type of the walk. So $\left(Y_{t}^{(i)}-Z_{t}^{(i)}\right)$ hits 0 in finite time with probability 1. Hence, letting $\tau^{(i)}$ be the first time $Y_{t}^{(i)}-Z_{t}^{(i)}=0$, we have $\mathbb{P}\left[Y_{t}^{(i)} \neq Z_{t}^{(i)}\right] \leq$ $\mathbb{P}\left[\tau^{(i)}>t\right] \rightarrow \mathbb{P}\left[\tau^{(i)}=+\infty\right]=0$.

By a union bound,

$$
\mathbb{P}\left[Y_{t} \neq Z_{t}\right] \leq \sum_{i \in[d]} \mathbb{P}\left[Y_{t}^{(i)} \neq Z_{t}^{(i)}\right] \rightarrow 0
$$

as desired.
Exercise 4.1 asks for an example of a non-constant (necessarily unbounded) harmonic function on $\mathbb{Z}^{d}$.

Harmonic functions on $\mathbb{T}_{d}$ On trees, the situation is different. Let $\mathbb{T}_{d}$ be the infinite $d$-regular tree with root $\rho$. For $x \in \mathbb{T}_{d}$, we let $T_{x}$ be the subtree, rooted at $x$, of descendants of $x$.

Theorem 4.1.7 (Bounded harmonic functions on $\mathbb{T}_{d}$ ). For $d \geq 3$, let $\left(X_{t}\right)$ be simple random walk on $\mathbb{T}_{d}$ and let $P$ be the corresponding transition matrix. Let a be a neighbor of the root and consider the function

$$
h(x):=\mathbb{P}_{x}\left[X_{t} \in T_{a} \text { for all but finitely many } t\right] .
$$

Then $h$ is a non-constant, bounded harmonic function on $\mathbb{T}_{d}$.
Proof. The function $h$ is bounded since it is defined as a probability, and by the usual first-step analysis

$$
h(x)=\sum_{y: y \sim x} \frac{1}{d} \mathbb{P}_{y}\left[X_{t} \in T_{a} \text { for all but finitely many } t\right]=\sum_{y} P(x, y) h(y),
$$

so $h$ is harmonic on all of $\mathbb{T}_{d}$.
Let $b \neq a$ be a neighbor of the root. The key of the proof is the following lemma.

Lemma 4.1.8.

$$
q:=\mathbb{P}_{a}\left[\tau_{\rho}=+\infty\right]=\mathbb{P}_{b}\left[\tau_{\rho}=+\infty\right]>0 .
$$

Proof. The equality of the two probabilities follows by symmetry. To see that $q>0$, let $\left(Z_{t}\right)$ be simple random walk on $\mathbb{T}_{d}$ started at $a$ until the walk hits $\rho$ and let $L_{t}$ be the graph distance between $Z_{t}$ and the root. Then $\left(L_{t}\right)$ is a biased random walk on $\mathbb{Z}$ started at 1 jumping to the right with probability $1-\frac{1}{d}$ and jumping to the left with probability $\frac{1}{d}$. The probability that $\left(L_{t}\right)$ hits 0 in finite time is $<1$ because $1-\frac{1}{d}>\frac{1}{2}$ when $d \geq 3$ by the gambler's ruin (Example 3.1.43).

Note that

$$
h(\rho) \leq\left(1-\frac{1}{d}\right)(1-q)<1 .
$$

Indeed, if on the first step the random walk started at $\rho$ moves away from $a$, an event of probability $1-\frac{1}{d}$, then it must come back to $\rho$ in finite time to reach $T_{a}$. Similarly, by the strong Markov property (Theorem 3.1.8),

$$
h(a)=q+(1-q) h(\rho) .
$$

Since $h(\rho) \neq 1$ and $q>0$, this shows that $h(a)>h(\rho)$. So $h$ is not constant.

### 4.1.3 Total variation distance and coupling inequality

In the examples of Section 4.1.1, we used coupling to prove monotonicity statements. Coupling is also useful to bound the distance between probability measures. For this, we need the coupling inequality.

Total variation distance Let $\mu$ and $\nu$ be probability measures on $(S, \mathcal{S})$. Recall the definition of the total variation distance

$$
\|\mu-\nu\|_{\mathrm{TV}}:=\sup _{A \in \mathcal{S}}|\mu(A)-\nu(A)| .
$$

As the next lemma shows in the countable case, the total variation distance can be thought of as an $\ell^{1}$ distance on probability measures as vectors (up to a constant factor).

Lemma 4.1.9 (Alternative definition of total variation distance). If $S$ is countable, then it holds that

$$
\|\mu-\nu\|_{\mathrm{TV}}=\frac{1}{2} \sum_{x \in S}|\mu(x)-\nu(x)|
$$

Proof. Let $E_{*}:=\{x: \mu(x) \geq \nu(x)\}$. Then, for any $A \subseteq S$, by definition of $E_{*}$

$$
\mu(A)-\nu(A) \leq \mu\left(A \cap E_{*}\right)-\nu\left(A \cap E_{*}\right) \leq \mu\left(E_{*}\right)-\nu\left(E_{*}\right) .
$$

Similarly, we have

$$
\begin{aligned}
\nu(A)-\mu(A) & \leq \nu\left(E_{*}^{c}\right)-\mu\left(E_{*}^{c}\right) \\
& =\left(1-\nu\left(E_{*}\right)\right)-\left(1-\mu\left(E_{*}\right)\right) \\
& =\mu\left(E_{*}\right)-\nu\left(E_{*}\right) .
\end{aligned}
$$

The two bounds above are equal so $|\mu(A)-\nu(A)| \leq \mu\left(E_{*}\right)-\nu\left(E_{*}\right)$. Equality is achieved when $A=E_{*}$. Also

$$
\begin{aligned}
\mu\left(E_{*}\right)-\nu\left(E_{*}\right) & =\frac{1}{2}\left[\mu\left(E_{*}\right)-\nu\left(E_{*}\right)+\nu\left(E_{*}^{c}\right)-\mu\left(E_{*}^{c}\right)\right] \\
& =\frac{1}{2} \sum_{x \in S}|\mu(x)-\nu(x)| .
\end{aligned}
$$

That concludes the proof.
Like the $\ell^{1}$ distance, the total variation distance is a metric. In particular, it satisfies the triangle inequality.

Lemma 4.1.10 (Total variation distance: triangle inequality). Let $\mu, \nu, \eta$ be probability measures on $(S, \mathcal{S})$. Then

$$
\|\mu-\nu\|_{\mathrm{TV}} \leq\|\mu-\eta\|_{\mathrm{TV}}+\|\eta-\nu\|_{\mathrm{TV}} .
$$

Proof. From the definition,

$$
\begin{aligned}
\sup _{A \in \mathcal{S}}|\mu(A)-\nu(A)| & \leq \sup _{A \in \mathcal{S}}\{|\mu(A)-\eta(A)|+|\eta(A)-\nu(A)|\} \\
& \leq \sup _{A \in \mathcal{S}}|\mu(A)-\eta(A)|+\sup _{A \in \mathcal{S}}|\eta(A)-\nu(A)| .
\end{aligned}
$$

Coupling inequality We come to an elementary, yet fundamental inequality.
Lemma 4.1.11 (Coupling inequality). Let $\mu$ and $\nu$ be probability measures on $(S, \mathcal{S})$. For any coupling $(X, Y)$ of $\mu$ and $\nu$,

$$
\|\mu-\nu\|_{\mathrm{TV}} \leq \mathbb{P}[X \neq Y] .
$$

Proof. For any $A \in \mathcal{S}$,

$$
\begin{aligned}
\mu(A)-\nu(A)= & \mathbb{P}[X \in A]-\mathbb{P}[Y \in A] \\
= & \mathbb{P}[X \in A, X=Y]+\mathbb{P}[X \in A, X \neq Y] \\
& -\mathbb{P}[Y \in A, X=Y]-\mathbb{P}[Y \in A, X \neq Y] \\
= & \mathbb{P}[X \in A, X \neq Y]-\mathbb{P}[Y \in A, X \neq Y] \\
\leq & \mathbb{P}[X \neq Y],
\end{aligned}
$$

and, similarly, $\nu(A)-\mu(A) \leq \mathbb{P}[X \neq Y]$. Hence

$$
|\mu(A)-\nu(A)| \leq \mathbb{P}[X \neq Y] .
$$

Taking a supremum over $A$ gives the claim.
Here is a quick example.
Example 4.1.12 (A coupling of Poisson random variables). Let $X \sim \operatorname{Poi}(\lambda)$ and $Y \sim \operatorname{Poi}(\nu)$ with $\lambda>\nu$. Recall that a sum of independent Poisson is Poisson (see Exercise 6.7). This fact leads to a natural coupling: let $\hat{Y} \sim \operatorname{Poi}(\nu), \hat{Z} \sim$ $\operatorname{Poi}(\lambda-\nu)$ independently of $\hat{Y}$, and $\hat{X}=\hat{Y}+\hat{Z}$. Then $(\hat{X}, \hat{Y})$ is a coupling of $X$ and $Y$, and by the coupling inequality (Lemma 4.1.11)

$$
\left\|\mu_{X}-\mu_{Y}\right\|_{T V} \leq \mathbb{P}[\hat{X} \neq \hat{Y}]=\mathbb{P}[\hat{Z}>0]=1-e^{-(\lambda-\nu)} \leq \lambda-\nu,
$$

where we used $1-e^{-x} \leq x$ for all $x$ (see Exercise 1.16).


Figure 4.1: Proof by picture that: $1-p=\alpha=\beta=\|\mu-\nu\|_{\mathrm{TV}}$.

Remarkably, the inequality in Lemma 4.1.11 is tight. For simplicity, we prove this in the finite case only.

Lemma 4.1.13 (Maximal coupling). Assume $S$ is finite and let $\mathcal{S}=2^{S}$. Let $\mu$ and $\nu$ be probability measures on $(S, \mathcal{S})$. Then,

$$
\|\mu-\nu\|_{\mathrm{TV}}=\inf \{\mathbb{P}[X \neq Y]: \text { coupling }(X, Y) \text { of } \mu \text { and } \nu\} .
$$

Proof. We construct a coupling which achieves equality in the coupling inequality. Such a coupling is called a maximal coupling.

Let $A=\{x \in S: \mu(x)>\nu(x)\}, B=\{x \in S: \mu(x) \leq \nu(x)\}$ and

$$
p:=\sum_{x \in S} \mu(x) \wedge \nu(x), \quad \alpha:=\sum_{x \in A}[\mu(x)-\nu(x)], \quad \beta:=\sum_{x \in B}[\nu(x)-\mu(x)] .
$$

Assume $p>0$ (otheriwse there is nothing to prove). First, two lemmas. See Figure 4.1 for a proof by picture.

## Lemma 4.1.14.

$$
\sum_{x \in S} \mu(x) \wedge \nu(x)=1-\|\mu-\nu\|_{\mathrm{TV}}
$$

Proof. We have

$$
\begin{aligned}
2\|\mu-\nu\|_{\mathrm{TV}} & =\sum_{x \in S}|\mu(x)-\nu(x)| \\
& =\sum_{x \in A}[\mu(x)-\nu(x)]+\sum_{x \in B}[\nu(x)-\mu(x)] \\
& =\sum_{x \in A} \mu(x)+\sum_{x \in B} \nu(x)-\sum_{x \in S} \mu(x) \wedge \nu(x) \\
& =2-\sum_{x \in B} \mu(x)-\sum_{x \in A} \nu(x)-\sum_{x \in S} \mu(x) \wedge \nu(x) \\
& =2-2 \sum_{x \in S} \mu(x) \wedge \nu(x),
\end{aligned}
$$

where we used that both $\mu$ and $\nu$ sum to 1 . Rearranging gives the claim.

## Lemma 4.1.15.

$$
\sum_{x \in A}[\mu(x)-\nu(x)]=\sum_{x \in B}[\nu(x)-\mu(x)]=\|\mu-\nu\|_{\mathrm{TV}}=1-p .
$$

Proof. The first equality is immediate by the fact that $\mu$ and $\nu$ are probability measures. The second equality follows from the first one together with the second line in the proof of the previous lemma. The last equality is a restatement of the last lemma.

The maximal coupling is defined as follows:

- With probability $p$, pick $X=Y$ from $\gamma_{\text {min }}$ where

$$
\gamma_{\min }(x):=\frac{1}{p} \mu(x) \wedge \nu(x), \quad x \in S
$$

- Otherwise, pick $X$ from $\gamma_{A}$ where

$$
\gamma_{A}(x):=\frac{\mu(x)-\nu(x)}{1-p}, \quad x \in A,
$$

and, independently, pick $Y$ from

$$
\gamma_{B}(x):=\frac{\nu(x)-\mu(x)}{1-p}, \quad x \in B .
$$

Note that $X \neq Y$ in that case because $A$ and $B$ are disjoint.

The marginal law of $X$ is: for $x \in A$,

$$
p \gamma_{\min }(x)+(1-p) \gamma_{A}(x)=\nu(x)+\mu(x)-\nu(x)=\mu(x),
$$

and for $x \in B$,

$$
p \gamma_{\min }(x)+(1-p) \gamma_{A}(x)=\mu(x)+0=\mu(x) .
$$

A similar calculation holds for $Y$. Finally $\mathbb{P}[X \neq Y]=1-p=\|\mu-\nu\|_{\mathrm{TV}}$.
Remark 4.1.16. A proof of this result for general Polish spaces can be found in $[d H$, Section 2.5].

We return to our coupling of Bernoulli variables.
Example 4.1.17 (Coupling of Bernoulli variables (continued)). Recall the setting of Example 4.1.2. To construct the maximal coupling as above, we note that

$$
\begin{gathered}
A:=\{0\}, \quad B:=\{1\}, \\
p:=\sum_{x} \mu(x) \wedge \nu(x)=(1-r)+q, \quad 1-p=\alpha=\beta:=r-q, \\
\left(\gamma_{\min }(x)\right)_{x=0,1}=\left(\frac{1-r}{(1-r)+q}, \frac{q}{(1-r)+q}\right), \\
\gamma_{A}(0):=1, \quad \gamma_{B}(1):=1 .
\end{gathered}
$$

The law of the maximal coupling $\left(X^{\prime \prime \prime}, Y^{\prime \prime \prime}\right)$ is given by

$$
\begin{aligned}
\left(\mathbb{P}\left[\left(X^{\prime \prime \prime}, Y^{\prime \prime \prime}\right)=(i, j)\right]\right)_{i, j \in\{0,1\}} & =\left(\begin{array}{cc}
p \gamma_{\min }(0) & (1-p) \gamma_{A}(0) \gamma_{B}(1) \\
0 & p \gamma_{\min }(1)
\end{array}\right) \\
& =\left(\begin{array}{cc}
1-r & r-q \\
0 & q
\end{array}\right) .
\end{aligned}
$$

Notice that it happens to coincide with the monotone coupling.

Poisson approximation Here is a classical application of coupling: the approximation of a sum of independent Bernoulli variables with a Poisson. It gives a quantitative bound in total variation distance. Let $X_{1}, \ldots, X_{n}$ be independent Bernoulli random variables with parameters $p_{1}, \ldots, p_{n}$ respectively. We are interested in the case where the $p_{i}$ s are "small." Let $S_{n}:=\sum_{i \leq n} X_{i}$. We approximate $S_{n}$ with a Poisson random variable $Z_{n}$ as follows: let $W_{1}, \ldots, W_{n}$ be independent Poisson random variables with means $\lambda_{1}, \ldots, \lambda_{n}$ respectively and define $Z_{n}:=\sum_{i \leq n} W_{i}$. We choose $\lambda_{i}=-\log \left(1-p_{i}\right)$ for reasons that will become clear below. Note that $Z_{n} \sim \operatorname{Poi}(\lambda)$ where $\lambda=\sum_{i \leq n} \lambda_{i}$.

Theorem 4.1.18 (Poisson approximation).

$$
\left\|\mu_{S_{n}}-\operatorname{Poi}(\lambda)\right\|_{\mathrm{TV}} \leq \frac{1}{2} \sum_{i \leq n} \lambda_{i}^{2}
$$

Proof. We couple the pairs $X_{i}, W_{i}$ independently for $i \leq n$. Let

$$
W_{i}^{\prime} \sim \operatorname{Poi}\left(\lambda_{i}\right) \quad \text { and } \quad X_{i}^{\prime}=W_{i}^{\prime} \wedge 1 .
$$

Because of our choice $\lambda_{i}=-\log \left(1-p_{i}\right)$ which implies

$$
1-p_{i}=\mathbb{P}\left[X_{i}=0\right]=\mathbb{P}\left[W_{i}=0\right]=e^{-\lambda_{i}},
$$

$\left(X_{i}^{\prime}, W_{i}^{\prime}\right)$ is indeed a coupling of $X_{i}, W_{i}$. Let $S_{n}^{\prime}:=\sum_{i \leq n} X_{i}^{\prime}$ and $Z_{n}^{\prime}:=\sum_{i \leq n} W_{i}^{\prime}$. Then $\left(S_{n}^{\prime}, Z_{n}^{\prime}\right)$ is a coupling of $S_{n}, Z_{n}$. By the coupling inequality

$$
\begin{aligned}
\left\|\mu_{S_{n}}-\mu_{Z_{n}}\right\|_{\mathrm{TV}} & \leq \mathbb{P}\left[S_{n}^{\prime} \neq Z_{n}^{\prime}\right] \\
& \leq \sum_{i \leq n} \mathbb{P}\left[X_{i}^{\prime} \neq W_{i}^{\prime}\right] \\
& =\sum_{i \leq n} \mathbb{P}\left[W_{i}^{\prime} \geq 2\right] \\
& =\sum_{i \leq n} \sum_{j \geq 2} e^{-\lambda_{i}} \frac{\lambda_{i}^{j}}{j!} \\
& \leq \sum_{i \leq n} \frac{\lambda_{i}^{2}}{2} \sum_{\ell \geq 0} e^{-\lambda_{i}} \frac{\lambda_{i}^{\ell}}{\ell!} \\
& =\sum_{i \leq n} \frac{\lambda_{i}^{2}}{2} .
\end{aligned}
$$

Mappings reduce the total variation distance The following lemma will be useful.

Lemma 4.1.19 (Mappings). Let $X$ and $Y$ be random variables taking values in $(S, \mathcal{S})$, let $h$ be a measurable map from $(S, \mathcal{S})$ to $\left(S^{\prime}, \mathcal{S}^{\prime}\right)$, and let $X^{\prime}:=h(X)$ and $Y^{\prime}:=h(Y)$. The following inequality holds

$$
\left\|\mu_{X^{\prime}}-\mu_{Y^{\prime}}\right\|_{\mathrm{TV}} \leq\left\|\mu_{X}-\mu_{Y}\right\|_{\mathrm{TV}}
$$

Proof. From the definition of the total variation distance, we seek to bound

$$
\begin{aligned}
& \sup _{A^{\prime} \in \mathcal{S}^{\prime}}\left|\mathbb{P}\left[X^{\prime} \in A^{\prime}\right]-\mathbb{P}\left[Y^{\prime} \in A^{\prime}\right]\right| \\
& =\sup _{A^{\prime} \in \mathcal{S}^{\prime}}\left|\mathbb{P}\left[h(X) \in A^{\prime}\right]-\mathbb{P}\left[h(Y) \in A^{\prime}\right]\right| \\
& =\sup _{A^{\prime} \in \mathcal{S}^{\prime}}\left|\mathbb{P}\left[X \in h^{-1}\left(A^{\prime}\right)\right]-\mathbb{P}\left[Y \in h^{-1}\left(A^{\prime}\right)\right]\right| .
\end{aligned}
$$

Since $h^{-1}\left(A^{\prime}\right) \in \mathcal{S}$ by the measurability of $h$, this last expression is less or equal than

$$
\sup _{A \in \mathcal{S}}|\mathbb{P}[X \in A]-\mathbb{P}[Y \in A]|
$$

which proves the claim.

Coupling of Markov chains In the context of Markov chains, a natural way to couple is to do so step by step. We will refer to such couplings as Markovian. An important special case is a Markovian coupling of a chain with itself.

Definition 4.1.20 (Markovian coupling). Let $P$ and $Q$ be transition matrices on the same state space $V$. A Markovian coupling of $P$ and $Q$ is a Markov chain $\left(X_{t}, Y_{t}\right)_{t}$ on $V \times V$ with transition matrix $R$ satisfying: for all $x, y, x^{\prime}, y^{\prime} \in V$,

$$
\begin{aligned}
& \sum_{z^{\prime}} R\left((x, y),\left(x^{\prime}, z^{\prime}\right)\right)=P\left(x, x^{\prime}\right), \\
& \sum_{z^{\prime}} R\left((x, y),\left(z^{\prime}, y^{\prime}\right)\right)=Q\left(y, y^{\prime}\right) .
\end{aligned}
$$

We will give many examples throughout this chapter. See also Example 4.2.14 for an example of a coupling of Markov chains that is not Markovian.

### 4.1.4 $\triangleright$ Random graphs: degree sequence in Erdös-Rényi model

Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ be an Erdős-Rényi graph with $p_{n}:=\frac{\lambda}{n}$ and $\lambda>0$ (see Definition 1.2.2). For $i \in[n]$, let $D_{i}(n)$ be the degree of vertex $i$ and define

$$
N_{d}(n):=\sum_{i=1}^{n} \mathbf{1}_{\left\{D_{i}(n)=d\right\}},
$$

the number of vertices of degree $d$.

Theorem 4.1.21 (Erdős-Rényi graph: degree sequence).

$$
\frac{1}{n} N_{d}(n) \rightarrow_{\mathrm{p}} f_{d}:=e^{-\lambda} \frac{\lambda^{d}}{d!}, \quad \forall d \geq 0
$$

Proof. We proceed in two steps:

1. we use the coupling inequality (Lemma 4.1.11) to show that the expectation of $\frac{1}{n} N_{d}(n)$ is close to $f_{d}$; and
2. we appeal to Chebyshev's inequality (Theorem 2.1.2) to show that $\frac{1}{n} N_{d}(n)$ is close to its expectation.
We justify each step as a lemma.
Lemma 4.1.22 (Convergence of the mean).

$$
\lim _{n \rightarrow+\infty} \frac{1}{n} \mathbb{E}_{n, p_{n}}\left[N_{d}(n)\right]=f_{d}, \quad \forall d \geq 1 .
$$

Proof. Note that the degrees $D_{i}(n), i \in[n]$, are identically distributed (but not independent) so

$$
\frac{1}{n} \mathbb{E}_{n, p_{n}}\left[N_{d}(n)\right]=\mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right] .
$$

Moreover, by definition, $D_{1}(n) \sim \operatorname{Bin}\left(n-1, p_{n}\right)$. Let $S_{n} \sim \operatorname{Bin}\left(n, p_{n}\right)$ and $Z_{n} \sim \operatorname{Poi}(\lambda)$. Using the Poisson approximation (Theorem 4.1.18) and a Taylor expansion,

$$
\begin{aligned}
\left\|\mu_{S_{n}}-\mu_{Z_{n}}\right\|_{\mathrm{TV}} & \leq \frac{1}{2} \sum_{i \leq n}\left(-\log \left(1-p_{n}\right)\right)^{2} \\
& =\frac{1}{2} \sum_{i \leq n}\left(\frac{\lambda}{n}+O\left(n^{-2}\right)\right)^{2} \\
& =\frac{\lambda^{2}}{2 n}+O\left(n^{-2}\right) .
\end{aligned}
$$

We can further couple $D_{1}(n)$ and $S_{n}$ as

$$
\left(\sum_{i \leq n-1} X_{i}, \sum_{i \leq n} X_{i}\right)
$$

where the $X_{i}$ s are i.i.d. $\operatorname{Ber}\left(p_{n}\right)$, that is, Bernoulli with parameter $p_{n}$. By the coupling inequality (Theorem 4.1.11),

$$
\left\|\mu_{D_{1}(n)}-\mu_{S_{n}}\right\|_{\mathrm{TV}} \leq \mathbb{P}\left[\sum_{i \leq n-1} X_{i} \neq \sum_{i \leq n} X_{i}\right]=\mathbb{P}\left[X_{n}=1\right]=p_{n}=\frac{\lambda}{n}
$$

By the triangle inequality for the total variation distance (Lemma 4.1.10) and the bounds above,

$$
\begin{aligned}
\frac{1}{2} \sum_{d \geq 0}\left|\mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right]-f_{d}\right| & =\left\|\mu_{D_{1}(n)}-\mu_{Z_{n}}\right\|_{\mathrm{TV}} \\
& \leq\left\|\mu_{D_{1}(n)}-\mu_{S_{n}}\right\|_{\mathrm{TV}}+\left\|\mu_{S_{n}}-\mu_{Z_{n}}\right\|_{\mathrm{TV}} \\
& \leq \frac{\lambda+\lambda^{2} / 2}{n}+O\left(n^{-2}\right) .
\end{aligned}
$$

Therefore, for all $d$,

$$
\left|\frac{1}{n} \mathbb{E}_{n, p_{n}}\left[N_{d}(n)\right]-f_{d}\right| \leq \frac{2 \lambda+\lambda^{2}}{n}+O\left(n^{-2}\right) \rightarrow 0,
$$

as $n \rightarrow+\infty$.
Lemma 4.1.23 (Concentration around the mean).

$$
\mathbb{P}_{n, p_{n}}\left[\left|\frac{1}{n} N_{d}(n)-\frac{1}{n} \mathbb{E}_{n, p_{n}}\left[N_{d}(n)\right]\right| \geq \varepsilon\right] \leq \frac{2 \lambda+1}{\varepsilon^{2} n}, \quad \forall d \geq 1, \forall n .
$$

Proof. By Chebyshev's inequality, for all $\varepsilon>0$

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\left|\frac{1}{n} N_{d}(n)-\frac{1}{n} \mathbb{E}_{n, p_{n}}\left[N_{d}(n)\right]\right| \geq \varepsilon\right] \leq \frac{\operatorname{Var}_{n, p_{n}}\left[\frac{1}{n} N_{d}(n)\right]}{\varepsilon^{2}} \tag{4.1.1}
\end{equation*}
$$

To compute the variance, we note that

$$
\begin{align*}
& \operatorname{Var}_{n, p_{n}}\left[\frac{1}{n} N_{d}(n)\right] \\
& =\frac{1}{n^{2}}\left\{\mathbb{E}_{n, p_{n}}\left[\left(\sum_{i \leq n} \mathbf{1}_{\left\{D_{i}(n)=d\right\}}\right)^{2}\right]-\left(n \mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right]\right)^{2}\right\} \\
& =\frac{1}{n^{2}}\left\{n(n-1) \mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d, D_{2}(n)=d\right]\right. \\
& \left.\quad \quad+n \mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right]-n^{2} \mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right]^{2}\right\} \\
& \leq  \tag{4.1.2}\\
& \leq \\
& \frac{1}{n}+\left\{\mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d, D_{2}(n)=d\right]-\mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right]^{2}\right\}
\end{align*}
$$

where we used the crude bound $\mathbb{P}_{n, p_{n}}\left[D_{1}(n)=d\right] \leq 1$. We bound the last line using a neat coupling argument. Let $Y_{1}$ and $Y_{2}$ be independent $\operatorname{Bin}\left(n-2, p_{n}\right)$, and let $X_{1}$ and $X_{2}$ be independent $\operatorname{Ber}\left(p_{n}\right)$. By separating the contribution of the edge
between 1 and 2 from those of edges to other vertices, we see that the joint degrees ( $\left.D_{1}(n), D_{2}(n)\right)$ have the same distribution as $\left(X_{1}+Y_{1}, X_{1}+Y_{2}\right)$. So the term in curly bracket above is equal to

$$
\begin{aligned}
\mathbb{P}[ & \left.\left(X_{1}+Y_{1}, X_{1}+Y_{2}\right)=(d, d)\right]-\mathbb{P}\left[X_{1}+Y_{1}=d\right]^{2} \\
\quad= & \mathbb{P}\left[\left(X_{1}+Y_{1}, X_{1}+Y_{2}\right)=(d, d)\right]-\mathbb{P}\left[\left(X_{1}+Y_{1}, X_{2}+Y_{2}\right)=(d, d)\right] \\
\quad \leq & \mathbb{P}\left[\left(X_{1}+Y_{1}, X_{1}+Y_{2}\right)=(d, d),\left(X_{1}+Y_{1}, X_{2}+Y_{2}\right) \neq(d, d)\right] \\
= & \mathbb{P}\left[\left(X_{1}+Y_{1}, X_{1}+Y_{2}\right)=(d, d), X_{2}+Y_{2} \neq d\right] \\
= & \mathbb{P}\left[X_{1}=0, Y_{1}=Y_{2}=d, X_{2}=1\right] \\
& +\mathbb{P}\left[X_{1}=1, Y_{1}=Y_{2}=d-1, X_{2}=0\right] \\
\leq & \mathbb{P}\left[X_{2}=1\right]+\mathbb{P}\left[X_{1}=1\right] \\
= & \frac{2 \lambda}{n} .
\end{aligned}
$$

Plugging back into (4.1.2) we get $\operatorname{Var}_{n, p_{n}}\left[\frac{1}{n} N_{d}(n)\right] \leq \frac{2 \lambda+1}{n}$, and (4.1.1) gives the claim.

Combining the lemmas concludes the proof of Theorem 4.1.21.

### 4.2 Stochastic domination

In comparing two probability measures, a natural relationship is that of "domination." For instance, let $\left(X_{i}\right)_{i=1}^{n}$ be independent $\mathbb{Z}_{+}$-valued random variables with

$$
\mathbb{P}\left[X_{i} \geq 1\right] \geq p,
$$

and let $S=\sum_{i=1}^{n} X_{i}$ be their sum. Now consider a separate random variable

$$
S_{*} \sim \operatorname{Bin}(n, p) .
$$

It is intuitively clear that one should be able to bound $S$ from below by analyzing $S_{*}$ instead-which may be considerably easier. Indeed, in some sense, $S$ "dominates" $S_{*}$, that is, $S$ should have a tendency to be bigger than $S_{*}$. One expects more specifically that

$$
\mathbb{P}[S>x] \geq \mathbb{P}\left[S_{*}>x\right]
$$

Coupling provides a formal characterization of this notion, as we detail in this section.

In particular we study an important special case known as positive associations. Here a measure "dominates itself" in the following sense: conditioning on certain events makes other events more likely. That concept is formalized in Section 4.2.3.


Figure 4.2: The law of $X$, represented here by its cumulative distribution function $F_{X}$ in solid, stochastically dominates the law of $Y$, in dashed. The construction of a monotone coupling, $(\hat{X}, \hat{Y}):=\left(F_{X}^{-1}(U), F_{Y}^{-1}(U)\right)$ where $U$ is uniform in $[0,1]$, is also depicted.

### 4.2.1 Definitions

We start with the simpler case of real random variables then consider partially ordered sets, a natural setting for this concept.

Ordering of real random variables Recall that, intuitively, stochastic domination captures the idea that one variable "tends to take larger values" than the other. For real random variables, it is defined in terms of tail probabilities, or equivalently in terms of cumulative distribution functions. See Figure 4.2 for an illustration.

Definition 4.2.1 (Stochastic domination). Let $\mu$ and $\nu$ be probability measures on $\mathbb{R}$. The measure $\mu$ is said to stochastically dominate $\nu$, denoted by $\mu \succeq \nu$, if for all $x \in \mathbb{R}$

$$
\mu[(x,+\infty)] \geq \nu[(x,+\infty)] .
$$

A real random variable $X$ stochastically dominates $Y$, denoted by $X \succeq Y$, if the law of $X$ dominates the law of $Y$.

Example 4.2.2 (Bernoulli vs. Poisson). Let $X \sim \operatorname{Poi}(\lambda)$ be Poisson with mean $\lambda>0$ and let $Y$ be a Bernoulli trial with success probability $p \in(0,1)$. In order
for $X$ to stochastically dominate $Y$, we need to have

$$
\mathbb{P}[X>\ell] \geq \mathbb{P}[Y>\ell], \quad \forall \ell \geq 0
$$

This is always true for $\ell \geq 1$ since $\mathbb{P}[X>\ell]>0$ but $\mathbb{P}[Y>\ell]=0$. So it remains to consider the case $\ell=0$. We have

$$
1-e^{-\lambda}=\mathbb{P}[X>0] \geq \mathbb{P}[Y>0]=p,
$$

if and only if

$$
\lambda \geq-\log (1-p) .
$$

Note that stochastic domination does not require $X$ and $Y$ to be defined on the same probability space. However the connection to coupling arises from the following characterization.

Theorem 4.2.3 (Coupling and stochastic domination). The real random variable $X$ stochastically dominates $Y$ if and only if there is a coupling $(\hat{X}, \hat{Y})$ of $X$ and $Y$ such that

$$
\begin{equation*}
\mathbb{P}[\hat{X} \geq \hat{Y}]=1 . \tag{4.2.1}
\end{equation*}
$$

We refer to $(\hat{X}, \hat{Y})$ as a monotone coupling of $X$ and $Y$.
Proof. Suppose there is such a coupling. Then for all $x \in \mathbb{R}$

$$
\mathbb{P}[Y>x]=\mathbb{P}[\hat{Y}>x]=\mathbb{P}[\hat{X} \geq \hat{Y}>x] \leq \mathbb{P}[\hat{X}>x]=\mathbb{P}[X>x] .
$$

For the other direction, define the cumulative distribution functions $F_{X}(x)=$ $\mathbb{P}[X \leq x]$ and $F_{Y}(x)=\mathbb{P}[Y \leq x]$. Assume $X \succeq Y$. The idea of the proof is to use the following standard way of generating a real random variable (see Theorem B.2.7)

$$
\begin{equation*}
X \stackrel{\mathrm{~d}}{=} F_{X}^{-1}(U), \tag{4.2.2}
\end{equation*}
$$

where $U$ is a $[0,1]$-valued uniform random variable and

$$
F_{X}^{-1}(u):=\inf \left\{x \in \mathbb{R}: F_{X}(x) \geq u\right\}
$$

is a generalized inverse. It is natural to construct a coupling of $X$ and $Y$ by simply using the same uniform random variable $U$ in this representation, that is, we define $\hat{X}=F_{X}^{-1}(U)$ and $\hat{Y}=F_{Y}^{-1}(U)$. See Figure 4.2. By (4.2.2), this is a coupling of $X$ and $Y$. It remains to check (4.2.1). Because $F_{X}(x) \leq F_{Y}(x)$ for all $x$ by definition of stochastic domination, by the definition of the generalized inverse,

$$
\mathbb{P}[\hat{X} \geq \hat{Y}]=\mathbb{P}\left[F_{X}^{-1}(U) \geq F_{Y}^{-1}(U)\right]=1
$$

as required.

Example 4.2.4. Returning to the example in the first paragraph of Section 4.2, let $\left(X_{i}\right)_{i=1}^{n}$ be independent $\mathbb{Z}_{+}$-valued random variables with $\mathbb{P}\left[X_{i} \geq 1\right] \geq p$ and consider their sum $S:=\sum_{i=1}^{n} X_{i}$. Further let $S_{*} \sim \operatorname{Bin}(n, p)$. Write $S_{*}$ as the sum $\sum_{i=1}^{n} Y_{i}$ where $\left(Y_{i}\right)$ are independent Bernoullli variables with $\mathbb{P}\left[Y_{i}=1\right]=p$. To couple $S$ and $S_{*}$, first set $\left(\hat{Y}_{i}\right):=\left(Y_{i}\right)$ and $\hat{S}_{*}:=\sum_{i=1}^{n} \hat{Y}_{i}$. Let $\hat{X}_{i}$ be 0 whenever $\hat{Y}_{i}=0$. Otherwise (i.e., if $\hat{Y}_{i}=1$ ), generate $\hat{X}_{i}$ according to the distribution of $X_{i}$ conditioned on $\left\{X_{i} \geq 1\right\}$, independently of everything else. By construction $\hat{X}_{i} \geq \hat{Y}_{i}$ almost surely for all $i$ and as a result $\sum_{i=1}^{n} \hat{X}_{i}=: \hat{S} \geq \hat{S}_{*}$ almost surely, or $S \succeq S_{*}$ by Theorem 4.2.3. That implies for instance that $\mathbb{P}[S>x] \geq \mathbb{P}\left[S_{*}>x\right]$ as we claimed earlier. A slight modification of this argument gives the following useful fact about binomials

$$
n \geq m, q \geq p \Longrightarrow \operatorname{Bin}(n, q) \succeq \operatorname{Bin}(m, p) .
$$

Exercise 4.2 asks for a formal proof.
Example 4.2.5 (Poisson distribution). Let $X \sim \operatorname{Poi}(\mu)$ and $Y \sim \operatorname{Poi}(\nu)$ with $\mu>\nu$. Recall that a sum of independent Poisson is Poisson (see Exercise 6.7). This fact leads to a natural coupling: let $\hat{Y} \sim \operatorname{Poi}(\nu), \hat{Z} \sim \operatorname{Poi}(\mu-\nu)$ independently of $Y$, and $\hat{X}=\hat{Y}+\hat{Z}$. Then $(\hat{X}, \hat{Y})$ is a coupling and $\hat{X} \geq \hat{Y}$ a.s. because $\hat{Z} \geq 0$. Hence $X \succeq Y$.

We record two useful consequences of Theorem 4.2.3.
Corollary 4.2.6. Let $X$ and $Y$ be real random variables with $X \succeq Y$ and let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a non-decreasing function. Then $f(X) \succeq f(Y)$ and furthermore, provided $\mathbb{E}|f(X)|, \mathbb{E}|f(Y)|<+\infty$, we have that

$$
\mathbb{E}[f(X)] \geq \mathbb{E}[f(Y)]
$$

Proof. Let $(\hat{X}, \hat{Y})$ be the monotone coupling of $X$ and $Y$ whose existence is guaranteed by Theorem 4.2.3. Then $f(\hat{X}) \geq f(\hat{Y})$ almost surely so that, provided the expectations exist,

$$
\mathbb{E}[f(X)]=\mathbb{E}[f(\hat{X})] \geq \mathbb{E}[f(\hat{Y})]=\mathbb{E}[f(Y)]
$$

and furthermore $(f(\hat{X}), f(\hat{Y}))$ is a monotone coupling of $f(X)$ and $f(Y)$. Hence $f(X) \succeq f(Y)$.

Corollary 4.2.7. Let $X_{1}, X_{2}$ be independent random variables. Let $Y_{1}, Y_{2}$ be independent random variables such that $X_{i} \succeq Y_{i}, i=1,2$. Then

$$
X_{1}+X_{2} \succeq Y_{1}+Y_{2} .
$$

Proof. Let $\left(\hat{X}_{1}, \hat{Y}_{1}\right)$ and $\left(\hat{X}_{2}, \hat{Y}_{2}\right)$ be independent, monotone couplings of $X_{1}, Y_{1}$ and $X_{2}, Y_{2}$ on the same probability space. Then

$$
X_{1}+X_{2} \stackrel{\mathrm{~d}}{=} \hat{X}_{1}+\hat{X}_{2} \geq \hat{Y}_{1}+\hat{Y}_{2} \stackrel{\mathrm{~d}}{=} Y_{1}+Y_{2} .
$$

Example 4.2.8 (Binomial vs. Poisson). A sum of $n$ independent Poisson variables with mean $\lambda$ is $\operatorname{Poi}(n \lambda)$. A sum of $n$ independent Bernoulli trials with success probability $p$ is $\operatorname{Bin}(n, p)$. Using Example 4.2.2 and Corollary 4.2.7, we get

$$
\begin{equation*}
\lambda \geq-\log (1-p) \quad \Longrightarrow \quad \operatorname{Poi}(n \lambda) \succeq \operatorname{Bin}(n, p) \tag{4.2.3}
\end{equation*}
$$

The following special case will be useful later. Let $0<\Lambda<1$ and let $m$ be a positive integer. Then

$$
\frac{\Lambda}{m-1} \geq \frac{\Lambda}{m-\Lambda}=\frac{m}{m-\Lambda}-1 \geq \log \left(\frac{m}{m-\Lambda}\right)=-\log \left(1-\frac{\Lambda}{m}\right)
$$

where we used that $\log x \leq x-1$ for all $x \in \mathbb{R}_{+}$(see Exercise 1.16). So, setting $\lambda:=\frac{\Lambda}{m-1}, p:=\frac{\Lambda}{m}$ and $n:=m-1$ in (4.2.3), we get

$$
\begin{equation*}
\Lambda \in(0,1) \quad \Longrightarrow \quad \operatorname{Poi}(\Lambda) \succeq \operatorname{Bin}\left(m-1, \frac{\Lambda}{m}\right) \tag{4.2.4}
\end{equation*}
$$

Ordering on partially ordered sets The definition of stochastic domination hinges on the totally ordered nature of $\mathbb{R}$. It also extends naturally to posets. Let $(\mathcal{X}, \leq)$ be a poset, that is, for all $x, y, z \in \mathcal{X}$ :

- (Reflexivity) $x \leq x$;
- (Antisymmetry) if $x \leq y$ and $y \leq x$ then $x=y$; and
- (Transitivity) if $x \leq y$ and $y \leq z$ then $x \leq z$.

Throughout, we assume that $\mathcal{X}$ is a measurable space.
For instance the set $\{0,1\}^{F}$ is a poset when equipped with the relation $\mathbf{x} \leq \mathbf{y}$ if and only if $x_{i} \leq y_{i}$ for all $i \in F$, where $\mathbf{x}=\left(x_{i}\right)_{i \in F}$ and $\mathbf{y}=\left(y_{i}\right)_{i \in F}$. Equivalently the subsets of $F$, denoted by $2^{F}$, form a poset with the inclusion relation.

A totally ordered set satisfies in addition that, for any $x, y$, we have either $x \leq y$ or $y \leq x$. That is not satisfied in the previous example.

Let $\mathcal{F}$ be a $\sigma$-algebra over the poset $\mathcal{X}$. An event $A \in \mathcal{F}$ is increasing if $x \in A$ implies that any $y \geq x$ is also in $A$. A function $f: \mathcal{X} \rightarrow \mathbb{R}$ is increasing if $x \leq y$ implies $f(x) \leq f(y)$. Some properties of increasing events are derived in Exercise 4.4.

Definition 4.2.9 (Stochastic domination for posets). Let $(\mathcal{X}, \leq)$ be a poset and let $\mathcal{F}$ be a $\sigma$-algebra on $\mathcal{X}$. Let $\mu$ and $\nu$ be probability measures on $(\mathcal{X}, \mathcal{F})$. The measure $\mu$ is said to stochastically dominate $\nu$, denoted by $\mu \succeq \nu$, if for all increasing $A \in \mathcal{F}$

$$
\mu(A) \geq \nu(A)
$$

An $\mathcal{X}$-valued random variable $X$ stochastically dominates $Y$, denoted by $X \succeq Y$, if the law of $X$ dominates the law of $Y$.
As before, a monotone coupling $(\hat{X}, \hat{Y})$ of $X$ and $Y$ is one which satisfies $\hat{X} \geq \hat{Y}$ almost surely.

Example 4.2.10 (Monotonicity of the percolation function). We briefly revisit Example 4.1.3 to illustrate our definitions. Consider bond percolation on the $d$ dimensional lattice $\mathbb{L}^{d}$ (Definition 1.2.1). Here the poset is the collection of all subsets of edges, specifying the open edges, with the inclusion relation. Recall that the percolation function is given by

$$
\theta(p):=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right],
$$

where $\mathcal{C}_{0}$ is the open cluster of the origin. We argued in Example 4.1.3 (see also Section 2.2.4) that $\theta(p)$ is nondecreasing by considering the following alternative representation of the percolation process under $\mathbb{P}_{p}$ : to each edge $e$, assign a uniform [ 0,1$]$-valued random variable $U_{e}$ and declare the edge open if $U_{e} \leq p$. Using the same $U_{e} \mathrm{~s}$ for two different values of $p$, say $p_{1}<p_{2}$, gives a monotone coupling of the processes for $p_{1}$ and $p_{2}$.

The existence of a monotone coupling is perhaps more surprising for posets. We prove the result in the finite case only, which will be enough for our purposes.

Theorem 4.2.11 (Strassen's theorem). Let $X$ and $Y$ be random variables taking values in a finite poset $(\mathcal{X}, \leq)$ with the $\sigma$-algebra $\mathcal{F}=2^{\mathcal{X}}$. Then $X \succeq Y$ if and only if there exists a monotone coupling $(\hat{X}, \hat{Y})$ of $X$ and $Y$.

Proof. Suppose there is such a coupling. Then for all increasing $A$

$$
\mathbb{P}[Y \in A]=\mathbb{P}[\hat{Y} \in A]=\mathbb{P}[\hat{X} \geq \hat{Y} \in A] \leq \mathbb{P}[\hat{X} \in A]=\mathbb{P}[X \in A]
$$

The proof in the other direction relies on the max-flow min-cut theorem (Theorem 1.1.15). To see the connection with flows, let $\mu_{X}$ and $\mu_{Y}$ be the laws of $X$ and
$Y$ respectively, and denote by $\nu$ their joint distribution under the desired coupling. Noting that we want $\nu(x, y)>0$ only if $x \geq y$, the marginal conditions on the coupling read

$$
\begin{equation*}
\sum_{y: x \geq y} \nu(x, y)=\mu_{X}(x), \quad \forall x \in \mathcal{X}, \tag{4.2.5}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{x: x \geq y} \nu(x, y)=\mu_{Y}(y), \quad \forall y \in \mathcal{X} \tag{4.2.6}
\end{equation*}
$$

These equations can be interpreted as flow-conservation constraints. Consider the following directed graph. There are two vertices, $(w, 1)$ and $(w, 2)$, for each element $w$ in $\mathcal{X}$ with edges connecting each $(x, 1)$ to those $(y, 2)$ s with $x \geq y$. These edges have capacity $+\infty$. In addition there is a source $a$ and a sink $z$. The source has a directed edge of capacity $\mu_{X}(x)$ to $(x, 1)$ for each $x \in \mathcal{X}$ and, similarly, each $(y, 2)$ has a directed edge of capacity $\mu_{Y}(y)$ to the sink. The existence of a monotone coupling will follow once we show that there is a flow of strength 1 between $a$ and $z$. Indeed, in that case, all edges from the source and all edges to the sink must be at capacity. If we let $\nu(x, y)$ be the flow on edge $\langle(x, 1),(y, 2)\rangle$, the systems in (4.2.5) and (4.2.6) encode conservation of flow on the vertices $(\mathcal{X} \times\{1\}) \cup(\mathcal{X} \times\{2\})$. Hence the flow between $\mathcal{X} \times\{1\}$ and $\mathcal{X} \times\{2\}$ yields the desired coupling. See Figure 4.3.

By the max-flow min-cut theorem (Theorem 1.1.15), it suffices to show that a minimum cut has capacity 1 . Such a cut is of course obtained by choosing all edges out of the source. So it remains to show that no cut has capacity less than 1 . This is where we use the fact that $\mu_{X}(A) \geq \mu_{Y}(A)$ for all increasing $A$. Because the edges between $\mathcal{X} \times\{1\}$ and $\mathcal{X} \times\{2\}$ have infinite capacity, they cannot be used in a minimum cut. So we can restrict our attention to those cuts containing edges from $a$ to $A_{*} \times\{1\}$ and from $Z_{*} \times\{2\}$ to $z$ for subsets $A_{*}, Z_{*} \subseteq \mathcal{X}$.

We must have

$$
A_{*} \supseteq\left\{x \in \mathcal{X}: \exists y \in Z_{*}^{c}, x \geq y\right\},
$$

to block all paths of the form $a \sim(x, 1) \sim(y, 2) \sim z$ with $x$ and $y$ as in the previous display; here $Z_{*}^{c}=\mathcal{X} \backslash Z_{*}$. In fact, for a minimum cut, we further have

$$
A_{*}=\left\{x \in \mathcal{X}: \exists y \in Z_{*}^{c}, x \geq y\right\}
$$

as adding an $x$ not satisfying this property is redundant. In particular $A_{*}$ is increasing: if $x_{1} \in A_{*}$ and $x_{2} \geq x_{1}$, then $\exists y \in Z_{*}^{c}$ such that $x_{1} \geq y$ and, since $x_{2} \geq x_{1} \geq y$, we also have $x_{2} \in A_{*}$.

Observe further that, because $y \geq y$, the set $A_{*}$ also includes $Z_{*}^{c}$. If it were the case that $A_{*} \neq Z_{*}^{c}$, then we could construct a cut with lower or equal capacity


Figure 4.3: Construction of a monotone coupling through the max-flow representation for independent Bernoulli pairs with parameters $r$ (on the left) and $q<r$ (on the right). Edge labels indicate capacity. Edges without labels have infinite capacity. The dotted edges depict a suboptimal cut. The dark vertices correspond to the sets $A_{*}$ and $Z_{*}$ for this cut. The capacity of the cut is $r^{2}+r(1-r)+(1-q)^{2}+(1-q) q=r+(1-q)>r+(1-r)=1$.
by fixing $A_{*}$ and setting $Z_{*}:=A_{*}^{c}$ : suppose $A_{*} \cap Z_{*}$ is nonempty; because $A_{*}$ is increasing, any $y \in A_{*} \cap Z_{*}$ is such that paths of the form $a \sim(x, 1) \sim(y, 2) \sim z$ with $x \geq y$ are cut by $x \in A_{*}$; so we do not need those $y$ s in $Z_{*}$. Hence, for a minimum cut, we can assume that in fact $A_{*}=Z_{*}^{c}$. The capacity of the cut is

$$
\mu_{X}\left(A_{*}\right)+\mu_{Y}\left(Z_{*}\right)=\mu_{X}\left(A_{*}\right)+1-\mu_{Y}\left(A_{*}\right)=1+\left(\mu_{X}\left(A_{*}\right)-\mu_{Y}\left(A_{*}\right)\right) \geq 1,
$$

where the term in parenthesis is nonnegative by assumption and the fact that $A_{*}$ is increasing. That concludes the proof.

Remark 4.2.12. Strassen's theorem (Theorem 4.2.11) holds more generally on Polish spaces with a closed partial order. See, e.g., [Lin02, Section IV.1.2] for the details.

The proof of Corollary 4.2.6 immediately extends to:
Corollary 4.2.13. Let $X$ and $Y$ be $\mathcal{X}$-valued random variables with $X \succeq Y$ and let $f: \mathcal{X} \rightarrow \mathbb{R}$ be an increasing function. Then $f(X) \succeq f(Y)$ and furthermore, provided $\mathbb{E}|f(X)|, \mathbb{E}|f(Y)|<+\infty$, we have that

$$
\mathbb{E}[f(X)] \geq \mathbb{E}[f(Y)]
$$

Ordering of Markov chains Stochastic domination also arises in the context of Markov chains. We begin with an example. Recall the notion of a Markovian coupling from Definition 4.1.20. The following coupling of Markov chains is not Markovian.

Example 4.2.14 (Lazier chain). Consider a random walk $\left(X_{t}\right)$ on the network $\mathcal{N}=((V, E), c)$ where $V=\{0,1, \ldots, n\}$ and $i \sim j$ if and only if $|i-j| \leq 1$ (including self-loops). Let $\mathcal{N}^{\prime}=\left((V, E), c^{\prime}\right)$ be a modified version of $\mathcal{N}$ on the same graph where, for all $i, c(i, i) \leq c^{\prime}(i, i)$. That is, if $\left(X_{t}^{\prime}\right)$ is random walk on $\mathcal{N}^{\prime}$, then $\left(X_{t}^{\prime}\right)$ is "lazier" than $\left(X_{t}\right)$ in that it is more likely to stay put. To simplify the calculations, assume $c(i, i)=0$ for all $i$.

Assume that both $\left(X_{t}\right)$ and $\left(X_{t}^{\prime}\right)$ start at $i_{0}$ and define $M_{s}:=\max _{t \leq s} X_{t}$ and $M_{s}^{\prime}:=\max _{t \leq s} X_{t}^{\prime}$. Since ( $X_{t}^{\prime}$ ) "travels less" than $\left(X_{t}\right)$ the following claim is intuitively obvious:

## Claim 4.2.15.

$$
M_{s} \succeq M_{s}^{\prime} .
$$

We prove this by producing a monotone coupling. First set $\left(\hat{X}_{t}\right)_{t \in \mathbb{Z}_{+}}:=\left(X_{t}\right)_{t \in \mathbb{Z}_{+}}$. We then generate $\left(\hat{X}_{t}^{\prime}\right)_{t \in \mathbb{Z}_{+}}$as a "sticky" version of $\left(\hat{X}_{t}\right)_{t \in \mathbb{Z}_{+}}$. That is, $\left(\hat{X}_{t}^{\prime}\right)$ follows
exactly the same transitions as $\left(\hat{X}_{t}\right)$ (including the self-loops), but at each time it opts to stay where it currently is, say state $j$, for an extra time step with probability

$$
\alpha_{j}:=\frac{c^{\prime}(j, j)}{\sum_{i: i \sim j} c^{\prime}(i, j)},
$$

which is in $[0,1]$ by assumption. Marginally, $\left(\hat{X}_{t}^{\prime}\right)$ is a random walk on $\mathcal{N}^{\prime}$. Indeed, we have by construction of the coupling that the probability of staying put when in state $j$ is

$$
\alpha_{j}=\frac{c^{\prime}(j, j)}{\sum_{i: i \sim j} c^{\prime}(i, j)},
$$

and, for $k \neq j$ with $k \sim j$, the probability of moving to state $k$ when in state $j$ is

$$
\begin{aligned}
\left(1-\alpha_{j}\right) \frac{c(j, k)}{\sum_{i: i \sim j} c(i, j)} & =\left(\frac{\left[\sum_{i: i \sim j} c^{\prime}(i, j)\right]-c^{\prime}(j, j)}{\sum_{i: i \sim j} c^{\prime}(i, j)}\right) \frac{c(j, k)}{\sum_{i: i \sim j} c(i, j)} \\
& =\left(\frac{\sum_{i: i \sim j} c(i, j)}{\sum_{i: i \sim j} c^{\prime}(i, j)}\right) \frac{c^{\prime}(j, k)}{\sum_{i: i \sim j} c(i, j)} \\
& =\frac{c^{\prime}(j, k)}{\sum_{i: i \sim j} c^{\prime}(i, j)},
\end{aligned}
$$

where, on the second line, we used that $c^{\prime}(i, j)=c(i, j)$ for $i \neq j$ and $i \sim j$. This coupling satisfies almost surely

$$
\widehat{M}_{s}:=\max _{t \leq s} \hat{X}_{t} \geq \max _{t \leq s} \hat{X}_{t}^{\prime}=: \widehat{M}_{s}^{\prime}
$$

because $\left(\hat{X}_{t}^{\prime}\right)_{t \leq s}$ visits a subset of the states visited by $\left(\hat{X}_{t}\right)_{t \leq s}$. In other words $\left(\widehat{M}_{s}, \widehat{M}_{s}^{\prime}\right)$ is a monotone coupling of $\left(M_{s}, M_{s}^{\prime}\right)$ and this proves the claim.

As we indicated, the previous example involved an "asynchronous" coupling of the chains. Often a simpler step-by-step approach-that is, through the construction of a Markovian coupling-is possible. We specialize the notion of stochastic domination to that important case.

Definition 4.2.16 (Stochastic domination of Markov chains). Let $P$ and $Q$ be transition matrices on a finite or countably infinite poset $(\mathcal{X}, \leq)$. The transition matrix $Q$ is said to stochastically dominate the transition matrix $P$ if

$$
\begin{equation*}
x \leq y \Longrightarrow P(x, \cdot) \preceq Q(y, \cdot) . \tag{4.2.7}
\end{equation*}
$$

If the above condition is satisfied for $P=Q$, we say that $P$ is stochastically monotone

The analogue of Strassen's theorem in this case is the following theorem, which we prove in the finite case only again.

Theorem 4.2.17. Let $\left(X_{t}\right)_{t \in \mathbb{Z}_{+}}$and $\left(Y_{t}\right)_{t \in \mathbb{Z}_{+}}$be Markov chains on a finite poset $(\mathcal{X}, \leq)$ with transition matrices $P$ and $Q$ respectively. Assume that $Q$ stochastically dominates $P$. Then for all $x_{0} \leq y_{0}$ there is a coupling $\left(\hat{X}_{t}, \hat{Y}_{t}\right)$ of $\left(X_{t}\right)$ started at $x_{0}$ and $\left(Y_{t}\right)$ started at $y_{0}$ such that almost surely

$$
\hat{X}_{t} \leq \hat{Y}_{t}, \quad \forall t
$$

Furthermore, if the chains are irreducible and have stationary distributions $\pi$ and $\mu$ respectively, then $\pi \preceq \mu$.

Observe that, for a Markovian, monotone coupling to exist, it is not generally enough for the weaker condition $P(x, \cdot) \preceq Q(x, \cdot)$ to hold for all $x$, as should be clear from the proof. See also Exercise 4.3.

Proof of Theorem 4.2.17. Let

$$
\mathcal{W}:=\{(x, y) \in \mathcal{X} \times \mathcal{X}: x \leq y\}
$$

For all $(x, y) \in \mathcal{W}$, let $R((x, y), \cdot)$ be the joint law of a monotone coupling of $P(x, \cdot)$ and $Q(y, \cdot)$. Such a coupling exists by Strassen's theorem and Condition (4.2.7). Let $\left(\hat{X}_{t}, \hat{Y}_{t}\right)$ be a Markov chain on $\mathcal{W}$ with transition matrix $R$ started at $\left(x_{0}, y_{0}\right)$. By construction, $\hat{X}_{t} \leq \hat{Y}_{t}$ for all $t$ almost surely. That proves the first half of the theorem.

For the second half, let $A$ be increasing on $\mathcal{X}$. Note that the first half implies that for all $s \geq 1$

$$
P^{s}\left(x_{0}, A\right)=\mathbb{P}\left[\hat{X}_{s} \in A\right] \leq \mathbb{P}\left[\hat{Y}_{s} \in A\right]=Q^{s}\left(y_{0}, A\right),
$$

because $\hat{X}_{s} \leq \hat{Y}_{s}$ and $A$ is increasing. Then, by a standard convergence result for irreducible Markov chains (i.e., (1.1.5)),

$$
\pi(A)=\lim _{t \rightarrow+\infty} \frac{1}{t} \sum_{s \leq t} P^{s}\left(x_{0}, A\right) \leq \lim _{t \rightarrow+\infty} \frac{1}{t} \sum_{s \leq t} Q^{s}\left(y_{0}, A\right)=\mu(A) .
$$

This proves the claim by definition of stochastic domination.
An example of application of this theorem is given in the next subsection.

### 4.2.2 $\triangleright$ Ising model: boundary conditions

Consider the $d$-dimensional lattice $\mathbb{L}^{d}$. Let $\Lambda$ be a finite subset of vertices in $\mathbb{L}^{d}$ and define $\mathcal{X}:=\{-1,+1\}^{\Lambda}$, which is a poset when equipped with the relation $\sigma \leq \sigma^{\prime}$ if and only if $\sigma_{i} \leq \sigma_{i}^{\prime}$ for all $i \in \Lambda$. Generalizing Example 1.2.5, for $\xi \in\{-1,+1\}^{\mathbb{L}^{d}}$, the (ferromagnetic) Ising model on $\Lambda$ with boundary conditions $\xi$ and inverse temperature $\beta$ is the probability distribution over spin configurations $\sigma \in \mathcal{X}$ given by

$$
\mu_{\beta, \Lambda}^{\xi}(\sigma):=\frac{1}{\mathcal{Z}_{\Lambda, \xi}(\beta)} e^{-\beta \mathcal{H}_{\Lambda, \xi}(\sigma)}
$$

where

$$
\mathcal{H}_{\Lambda, \xi}(\sigma):=-\sum_{\substack{i \sim j \\ i, j \in \Lambda}} \sigma_{i} \sigma_{j}-\sum_{\substack{i \sim j \\ i \in \Lambda, j \notin \Lambda}} \sigma_{i} \xi_{j},
$$

is the Hamiltonian and

$$
\mathcal{Z}_{\Lambda, \xi}(\beta):=\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}_{\Lambda, \xi}(\sigma)}
$$

is the partition function. For shorthand, we occasionally write + and - instead of +1 and -1 .

For the all- $(+1)$ and all- $(-1)$ boundary conditions we denote the measure above by $\mu_{\beta, \Lambda}^{+}(\sigma)$ and $\mu_{\beta, \Lambda}^{-}(\sigma)$ respectively. In this section, we show that these two measures are "extreme" in the following sense.
Claim 4.2.18. For all boundary conditions $\xi \in\{-1,+1\}^{\mathbb{L}^{d}}$,

$$
\mu_{\beta, \Lambda}^{+} \succeq \mu_{\beta, \Lambda}^{\xi} \succeq \mu_{\beta, \Lambda}^{-}
$$

Intuitively, because the ferromagnetic Ising model favors spin agreement, the all$(+1)$ boundary condition tends to produce more +1 s which in turn makes increasing events more likely. And vice versa.

The idea of the proof is to use Theorem 4.2.17 with a suitable choice of Markov chain.

Stochastic domination Recall that, in this context, vertices are often referred to as sites. Adapting Definition 1.2.8, we consider the single-site Glauber dynamics, which is the Markov chain on $\mathcal{X}$ which, at each time, selects a site $i \in \Lambda$ uniformly at random and updates the spin $\sigma_{i}$ according to $\mu_{\beta, \Lambda}^{\xi}(\sigma)$ conditioned on agreeing with $\sigma$ at all sites in $\Lambda \backslash\{i\}$. Specifically, for $\gamma \in\{-1,+1\}, i \in \Lambda$, and $\sigma \in \mathcal{X}$, let
$\sigma^{i, \gamma}$ be the configuration $\sigma$ with the state at $i$ being set to $\gamma$. Then, letting $n=|\Lambda|$, the transition matrix of the Glauber dynamics is

$$
Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i, \gamma}\right):=\frac{1}{n} \cdot \frac{e^{\gamma \beta S_{i}^{\xi}(\sigma)}}{e^{-\beta S_{i}^{\xi}(\sigma)}+e^{\beta S_{i}^{\xi}(\sigma)}},
$$

where

$$
S_{i}^{\xi}(\sigma):=\sum_{\substack{j: j \sim i \\ j \in \Lambda}} \sigma_{j}+\sum_{\substack{j: j \sim i \\ j \notin \Lambda}} \xi_{j} .
$$

All other transitions have probability 0 . It is straightforward to check that $Q_{\beta, \Lambda}^{\xi}$ is a stochastic matrix.

This chain is clearly irreducible. It is also reversible with respect to $\mu_{\beta, \Lambda}^{\xi}$. Indeed, for all $\sigma \in \mathcal{X}$ and $i \in \Lambda$, let

$$
S_{\neq i}^{\xi}(\sigma):=\mathcal{H}_{\Lambda, \xi}\left(\sigma^{i,+}\right)+S_{i}^{\xi}(\sigma)=\mathcal{H}_{\Lambda, \xi}\left(\sigma^{i,-}\right)-S_{i}^{\xi}(\sigma),
$$

Arguing as in Theorem 1.2.9, we have

$$
\begin{aligned}
\mu_{\beta, \Lambda}^{\xi}\left(\sigma^{i,-}\right) Q_{\beta, \Lambda}^{\xi}\left(\sigma^{i,-}, \sigma^{i,+}\right) & =\frac{e^{-\beta S_{\neq i}^{\xi}(\sigma)} e^{-\beta S_{i}^{\xi}(\sigma)}}{\mathcal{Z}_{\Lambda, \xi}(\beta)} \cdot \frac{e^{\beta S_{i}^{\xi}(\sigma)}}{n\left[e^{-\beta S_{i}^{\xi}(\sigma)}+e^{\beta S_{i}^{\xi}(\sigma)}\right]} \\
& =\frac{e^{-\beta S_{\neq i}^{\xi}(\sigma)}}{n \mathcal{Z}_{\Lambda, \xi}(\beta)\left[e^{-\beta S_{i}^{\xi}(\sigma)}+e^{\beta S_{i}^{\xi}(\sigma)}\right]} \\
& =\frac{e^{-\beta S_{\neq i}^{\xi}(\sigma)} e^{\beta S_{i}^{\xi}(\sigma)}}{\mathcal{Z}_{\Lambda, \xi}(\beta)} \cdot \frac{e^{-\beta S_{i}^{\xi}(\sigma)}}{n\left[e^{-\beta S_{i}^{\xi}(\sigma)}+e^{\beta S_{i}^{\xi}(\sigma)}\right]} \\
& =\mu_{\beta, \Lambda}^{\xi}\left(\sigma^{i,+}\right) Q_{\beta, \Lambda}^{\xi}\left(\sigma^{i,+}, \sigma^{i,-}\right)
\end{aligned}
$$

In particular $\mu_{\beta, \Lambda}^{\xi}$ is the stationary distribution of $Q_{\beta, \Lambda}^{\xi}$.

## Claim 4.2.19.

$$
\begin{equation*}
\xi^{\prime} \geq \xi \Longrightarrow Q_{\beta, \Lambda}^{\xi^{\prime}} \text { stochastically dominates } Q_{\beta, \Lambda}^{\xi} . \tag{4.2.8}
\end{equation*}
$$

Proof. Because the Glauber dynamics updates a single site at a time, establishing stochastic domination reduces to checking simple one-site inequalities.
Lemma 4.2.20. To establish (4.2.8), it suffices to show that for all i and all $\sigma \leq \tau$

$$
\begin{equation*}
Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,+}\right) \leq Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right) . \tag{4.2.9}
\end{equation*}
$$

Proof. Assume (4.2.9) holds. Let $A$ be increasing in $\mathcal{X}$ and let $\sigma \leq \tau$. Then, for the single-site Glauber dynamics, we have

$$
\begin{equation*}
Q_{\beta, \Lambda}^{\xi}(\sigma, A)=Q_{\beta, \Lambda}^{\xi}\left(\sigma, A \cap B_{\sigma}\right) \tag{4.2.10}
\end{equation*}
$$

where

$$
B_{\sigma}:=\left\{\sigma^{i, \gamma}: i \in \Lambda, \gamma \in\{-1,+1\}\right\},
$$

and similarly for $\tau, \xi^{\prime}$. Moreover, because $A$ is increasing and $\tau \geq \sigma$,

$$
\begin{equation*}
\sigma^{i, \gamma} \in A \Longrightarrow \tau^{i, \gamma} \in A, \tag{4.2.11}
\end{equation*}
$$

and

$$
\begin{equation*}
\sigma^{i,-} \in A \Longrightarrow \sigma^{i,+} \in A \tag{4.2.12}
\end{equation*}
$$

Letting

$$
I_{\sigma, A}^{ \pm}:=\left\{i \in \Lambda: \sigma^{i,-} \in A\right\}, \quad I_{\sigma, A}^{+}:=\left\{i \in \Lambda: \sigma^{i,-} \notin A, \sigma^{i,+} \in A\right\},
$$

and similarly for $\tau$, we have by (4.2.9), (4.2.10), (4.2.11), and (4.2.12),

$$
\begin{aligned}
Q_{\beta, \Lambda}^{\xi}(\sigma, A) & =Q_{\beta, \Lambda}^{\xi}\left(\sigma, A \cap B_{\sigma}\right) \\
& =\sum_{i \in I_{\sigma, A}^{+}} Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,+}\right)+\sum_{i \in I_{\sigma, A}^{ \pm}}\left[Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,-}\right)+Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,+}\right)\right] \\
& \leq \sum_{i \in I_{\sigma, A}^{+}} Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right)+\sum_{i \in I_{\sigma, A}^{ \pm}}\left[Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,-}\right)+Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,+}\right)\right] \\
& =\sum_{i \in I_{\sigma, A}^{+}} Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right)+\sum_{i \in I_{\sigma, A}^{ \pm}} \frac{1}{n} \\
& \leq \sum_{i \in I_{\tau, A}^{+}} Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right)+\sum_{i \in I_{\tau, A}^{ \pm}}\left[Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,-}\right)+Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right)\right] \\
& =Q_{\beta, \Lambda}^{\xi^{\prime}}(\tau, A),
\end{aligned}
$$

as claimed, where on the fifth line we used that $I_{\sigma, A}^{+} \subseteq I_{\tau, A}^{+} \cup I_{\tau, A}^{ \pm}$by (4.2.11) and that $Q_{\beta, \Lambda}^{\xi^{\prime}}\left(\tau, \tau^{i,+}\right) \leq 1 / n$ for all $i$ (in particular for $\left.i \in I_{\tau, A}^{+} \backslash I_{\sigma, A}^{+}\right)$.
Returning to the proof of Claim 4.2.19, observe that

$$
Q_{\beta, \Lambda}^{\xi}\left(\sigma, \sigma^{i,+}\right)=\frac{1}{n} \cdot \frac{e^{\beta S_{i}^{\xi}(\sigma)}}{e^{-\beta S_{i}^{\xi}(\sigma)}+e^{\beta S_{i}^{\xi}(\sigma)}}=\frac{1}{n} \cdot \frac{1}{e^{-2 \beta S_{i}^{\xi}(\sigma)}+1},
$$

which is increasing in $S_{i}^{\xi}(\sigma)$. Now $\sigma \leq \tau$ and $\xi \leq \xi^{\prime}$ imply that $S_{i}^{\xi}(\sigma) \leq S_{i}^{\xi^{\prime}}(\tau)$. That proves the claim by Lemma 4.2.20.

Finally:
Proof of Claim 4.2.18. Combining Theorem 4.2.17 and Claim 4.2.19 gives the result.

Remark 4.2.21. One can make sense of the limit of $\mu_{\beta, \Lambda}^{+}$and $\mu_{\beta, \Lambda}^{-}$when $|\Lambda| \rightarrow+\infty$, which is known as an infinite-volume Gibbs measure. For more, see for example [RAS15, Chapters 7-10].

Observe that we have not used any special property of the $d$-dimensional lattice. Indeed Claim 4.2.18 in fact holds for any countable, locally finite graph with positive coupling constants. We give another proof in Example 4.2.33.

### 4.2.3 Correlation inequalities: FKG and Holley's inequalities

A special case of stochastic domination is positive associations. In this section, we restrict ourselves to posets of the form $\{0,1\}^{F}$ for $F$ finite. We begin with an example.

Example 4.2.22 (Erdős-Rényi graph: positive associations). Consider an ErdősRényi graph $G \sim \mathbb{G}_{n, p}$. Let $\mathcal{E}=\{\{x, y\}: x, y \in[n], x \neq y\}$. Think of $G$ as taking values in the poset $\left(\{0,1\}^{\mathcal{E}}, \leq\right)$ where a 1 indicates that the corresponding edge is present. In fact observe that the law of $G$, which we denote as usual by $\mathbb{P}_{n, p}$, is a product measure on $\{0,1\}^{\mathcal{E}}$. The event $\mathcal{A}$ that $G$ is connected is increasing because adding edges cannot disconnect an already connected graph. So is the event $\mathcal{B}$ of having a chromatic number larger than 4 . Intuitively then, conditioning on $\mathcal{A}$ makes $\mathcal{B}$ more likely: the occurrence of $\mathcal{A}$ tends to be accompanied with a larger number of edges which in turn makes $\mathcal{B}$ more probable.

This is an example of a more general phenomenon. That is, for any non-empty increasing events $\mathcal{A}$ and $\mathcal{B}$, we have:

Claim 4.2.23.

$$
\begin{equation*}
\mathbb{P}_{n, p}[\mathcal{B} \mid \mathcal{A}] \geq \mathbb{P}_{n, p}[\mathcal{B}] . \tag{4.2.13}
\end{equation*}
$$

Or, put differently, the conditional measure $\mathbb{P}_{n, p}[\cdot \mid \mathcal{A}]$ stochastically dominates the unconditional measure $\mathbb{P}_{n, p}[\cdot]$. This is a special case of what is known as Harris' inequality, proved below. Note that (4.2.13) is equivalent to $\mathbb{P}_{n, p}[\mathcal{A} \cap \mathcal{B}] \geq$ $\mathbb{P}_{n, p}[\mathcal{A}] \mathbb{P}_{n, p}[\mathcal{B}]$, that is, to the fact that $\mathcal{A}$ and $\mathcal{B}$ are positively correlated.

## More generally:

Definition 4.2.24 (Positive associations). Let $\mu$ be a probability measure on $\{0,1\}^{F}$ where $F$ is finite. Then $\mu$ is said to have positive associations, or is positively as-
sociated, if for all increasing functions $f, g:\{0,1\}^{F} \rightarrow \mathbb{R}$

$$
\mu(f g) \geq \mu(f) \mu(g),
$$

where

$$
\mu(h):=\sum_{\omega \in\{0,1\}^{F}} \mu(\omega) h(\omega) .
$$

In particular, for any increasing events $A$ and $B$ it holds that

$$
\mu(A \cap B) \geq \mu(A) \mu(B),
$$

that is, $A$ and $B$ are positively correlated. Denoting by $\mu(A \mid B)$ the conditional
probability of $A$ given $B$, this is equivalent to

$$
\mu(A \mid B) \geq \mu(A)
$$

Remark 4.2.25. Note that positive associations is concerned only with "monotone" events. See Remark 4.2.45.

Remark 4.2.26. A notion of negative associations, which is a somewhat more delicate concept, was defined in Remark 3.3.43. See also [Pem00].

Let $\mu$ be positively associated. Note that if $A$ and $B$ are decreasing, that is, their complements are increasing (see Exercise 4.4), then

$$
\begin{aligned}
\mu(A \cap B) & =1-\mu\left(A^{c} \cup B^{c}\right) \\
& =1-\mu\left(A^{c}\right)-\mu\left(B^{c}\right)+\mu\left(A^{c} \cap B^{c}\right) \\
& \geq 1-\mu\left(A^{c}\right)-\mu\left(B^{c}\right)+\mu\left(A^{c}\right) \mu\left(B^{c}\right) \\
& =\mu(A) \mu(B),
\end{aligned}
$$

or $\mu(A \mid B) \geq \mu(A)$. Similarly, if $A$ is increasing and $B$ is decreasing, we have $\mu(A \cap B) \leq \mu(A) \mu(B)$, or

$$
\begin{equation*}
\mu(A \mid B) \leq \mu(A) \tag{4.2.14}
\end{equation*}
$$

Harris' inequality states that product measures on $\{0,1\}^{F}$ have positive associations. We prove a more general result known as the $F K G$ inequality. For two configurations $\omega, \omega^{\prime}$ in $\{0,1\}^{F}$, we let $\omega \wedge \omega^{\prime}$ and $\omega \vee \omega^{\prime}$ be the coordinatewise minimum and maximum of $\omega$ and $\omega^{\prime}$.

Definition 4.2.27 (FKG condition). Let $\mathcal{X}=\{0,1\}^{F}$ where $F$ is finite. A positive probability measure $\mu$ on $\mathcal{X}$ satisfies the FKG condition if

$$
\begin{equation*}
\mu\left(\omega \vee \omega^{\prime}\right) \mu\left(\omega \wedge \omega^{\prime}\right) \geq \mu(\omega) \mu\left(\omega^{\prime}\right), \quad \forall \omega, \omega^{\prime} \in \mathcal{X} \tag{4.2.15}
\end{equation*}
$$

This property is also known as log-supermodularity. We call such a measure an FKG measure.

Theorem 4.2.28 (FKG inequality). Let $\mathcal{X}=\{0,1\}^{F}$ where $F$ is finite. Suppose $\mu$ is a positive probability measure on $\mathcal{X}$ satisfying the $F K G$ condition. Then $\mu$ has positive associations.

Remark 4.2.29. Strict positivity is not in fact needed [FKG71]. The FKG condition is equivalent to a strong form of positive associations. See Exercise 4.8.

Note that product measures satisfy the FKG condition with equality. Indeed if $\mu(\omega)$ is of the form $\prod_{f \in F} \mu_{f}\left(\omega_{f}\right)$ then

$$
\begin{aligned}
\mu\left(\omega \vee \omega^{\prime}\right) \mu\left(\omega \wedge \omega^{\prime}\right) & =\prod_{f} \mu_{f}\left(\omega_{f} \vee \omega_{f}^{\prime}\right) \mu_{f}\left(\omega_{f} \wedge \omega_{f}^{\prime}\right) \\
& =\prod_{f: \omega_{f}=\omega_{f}^{\prime}} \mu_{f}\left(\omega_{f}\right)^{2} \prod_{f: \omega_{f} \neq \omega_{f}^{\prime}} \mu_{f}\left(\omega_{f}\right) \mu_{f}\left(\omega_{f}^{\prime}\right) \\
& =\prod_{f: \omega_{f}=\omega_{f}^{\prime}} \mu_{f}\left(\omega_{f}\right) \mu_{f}\left(\omega_{f}^{\prime}\right) \prod_{f: \omega_{f} \neq \omega_{f}^{\prime}} \mu_{f}\left(\omega_{f}\right) \mu_{f}\left(\omega_{f}^{\prime}\right) \\
& =\mu(\omega) \mu\left(\omega^{\prime}\right) .
\end{aligned}
$$

So the FKG inequality (Theorem 4.2.28) applies, for instance, to bond percolation and the Erdős-Rényi random graph model. The pointwise nature of the FKG condition also makes it relatively easy to check for measures which are defined explicitly up to a normalizing constant, such as the Ising model.

Example 4.2.30 (Ising model with boundary conditions: checking FKG). Consider again the setting of Section 4.2.2. We work on the space $\mathcal{X}:=\{-1,+1\}^{\Lambda}$ rather than $\{0,1\}^{F}$. Fix a finite $\Lambda \subseteq \mathbb{L}^{d}, \xi \in\{-1,+1\}^{\mathbb{L}^{d}}$ and $\beta>0$.

Claim 4.2.31. The measure $\mu_{\beta, \Lambda}^{\xi}$ satisfies the $F K G$ condition and therefore has positive associations.

Intuitively, taking the minimum (or maximum) of two spin configurations tends to increase agreement and therefore leads to a higher likelihood. For $\sigma, \sigma^{\prime} \in \mathcal{X}$, let $\bar{\tau}=\sigma \vee \sigma^{\prime}$ and $\underline{\tau}=\sigma \wedge \sigma^{\prime}$. By taking logarithms in the FKG condition and rearranging, we arrive at

$$
\begin{equation*}
\mathcal{H}_{\Lambda, \xi}(\bar{\tau})+\mathcal{H}_{\Lambda, \xi}(\underline{\tau}) \leq \mathcal{H}_{\Lambda, \xi}(\sigma)+\mathcal{H}_{\Lambda, \xi}\left(\sigma^{\prime}\right), \tag{4.2.16}
\end{equation*}
$$

and we see that proving the claim boils down to checking an inequality for each term in the Hamiltonian (which, confusingly, has a negative sign in it).

When $i \in \Lambda$ and $j \notin \Lambda$ such that $i \sim j$, we have

$$
\begin{equation*}
\bar{\tau}_{i} \xi_{j}+\underline{\tau}_{i} \xi_{j}=\left(\bar{\tau}_{i}+\underline{\tau}_{i}\right) \xi_{j}=\left(\sigma_{i}+\sigma_{i}^{\prime}\right) \xi_{j}=\sigma_{i} \xi_{j}+\sigma_{i}^{\prime} \xi_{j} . \tag{4.2.17}
\end{equation*}
$$

For $i, j \in \Lambda$ with $i \sim j$, note first that the case $\sigma_{j}=\sigma_{j}^{\prime}$ reduces to the previous calculation (with $\sigma_{j}=\sigma_{j}^{\prime}$ playing the role of $\xi_{j}$ ), so we assume $\sigma_{i} \neq \sigma_{i}^{\prime}$ and $\sigma_{j} \neq \sigma_{j}^{\prime}$. Then

$$
\bar{\tau}_{i} \bar{\tau}_{j}+\underline{\tau}_{i} \underline{\tau}_{j}=(+1)(+1)+(-1)(-1)=2 \geq \sigma_{i} \sigma_{j}+\sigma_{i}^{\prime} \sigma_{j}^{\prime},
$$

since 2 is the largest value the rightmost expression ever takes. We have established (4.2.16), which implies the claim.

Again, we have not used any special property of the lattice and the same result holds for countable, locally finite graphs with positive coupling constants. Note however that in the anti-ferromagnetic case, that is, if we multiply the Hamiltonian by -1 , the above argument does not work. Indeed there is no reason to expect positive associations in that case.

The FKG inequality in turn follows from a more general result known as Holley's inequality.

Theorem 4.2.32 (Holley's inequality). Let $\mathcal{X}=\{0,1\}^{F}$ where $F$ is finite. Suppose $\mu_{1}$ and $\mu_{2}$ are positive probability measures on $\mathcal{X}$ satisfying

$$
\begin{equation*}
\mu_{2}\left(\omega \vee \omega^{\prime}\right) \mu_{1}\left(\omega \wedge \omega^{\prime}\right) \geq \mu_{2}(\omega) \mu_{1}\left(\omega^{\prime}\right), \quad \forall \omega, \omega^{\prime} \in \mathcal{X} \tag{4.2.18}
\end{equation*}
$$

Then $\mu_{1} \preceq \mu_{2}$.
Before proving Holley's inequality (Theorem 4.2.32), we check that it indeed implies the FKG inequality. See Exercise 4.5 for an elementary proof in the independent case, that is, of Harris' inequality.

Proof of Theorem 4.2.28. Assume that $\mu$ satisfies the FKG condition and let $f, g$ be increasing functions. Because of our restriction to positive measures in Holley's inequality, we will work with positive functions. This is done without loss of generality. Indeed, letting $\mathbf{0}$ be the all- 0 vector, note that $f$ and $g$ are increasing if and only if $f^{\prime}:=f-f(\mathbf{0})+1>0$ and $g^{\prime}:=g-g(\mathbf{0})+1>0$ are increasing and that, moreover,

$$
\begin{aligned}
\mu\left(f^{\prime} g^{\prime}\right)-\mu\left(f^{\prime}\right) \mu\left(g^{\prime}\right) & =\mu\left(\left[f^{\prime}-\mu\left(f^{\prime}\right)\right]\left[g^{\prime}-\mu\left(g^{\prime}\right)\right]\right) \\
& =\mu([f-\mu(f)][g-\mu(g)]) \\
& =\mu(f g)-\mu(f) \mu(g) .
\end{aligned}
$$

In Holley's inequality, we let $\mu_{1}:=\mu$ and define the positive probability measure

$$
\mu_{2}(\omega):=\frac{g(\omega) \mu(\omega)}{\mu(g)} .
$$

We check that $\mu_{1}$ and $\mu_{2}$ satisfy the conditions of Theorem 4.2.32. Note that $\omega^{\prime} \leq$ $\omega \vee \omega^{\prime}$ for any $\omega$ so that, because $g$ is increasing, we have $g\left(\omega^{\prime}\right) \leq g\left(\omega \vee \omega^{\prime}\right)$. Hence, for any $\omega, \omega^{\prime}$,

$$
\begin{aligned}
\mu_{1}(\omega) \mu_{2}\left(\omega^{\prime}\right) & =\mu(\omega) \frac{g\left(\omega^{\prime}\right) \mu\left(\omega^{\prime}\right)}{\mu(g)} \\
& =\mu(\omega) \mu\left(\omega^{\prime}\right) \frac{g\left(\omega^{\prime}\right)}{\mu(g)} \\
& \leq \mu\left(\omega \wedge \omega^{\prime}\right) \mu\left(\omega \vee \omega^{\prime}\right) \frac{g\left(\omega \vee \omega^{\prime}\right)}{\mu(g)} \\
& =\mu_{1}\left(\omega \wedge \omega^{\prime}\right) \mu_{2}\left(\omega \vee \omega^{\prime}\right),
\end{aligned}
$$

where on the third line we used the FKG condition satisfied by $\mu$.
So Holley's inequality implies that $\mu_{2} \succeq \mu_{1}$. Hence, since $f$ is increasing, by Corollary 4.2.13

$$
\mu(f)=\mu_{1}(f) \leq \mu_{2}(f)=\frac{\mu(f g)}{\mu(g)},
$$

and the theorem is proved.
Proof of Theorem 4.2.32. The idea of the proof is to use Theorem 4.2.17. This is similar to what was done in Section 4.2.2. Again we use a single-site dynamic. For $x \in \mathcal{X}$ and $\gamma \in\{0,1\}$, we let $x^{i, \gamma}$ be $x$ with coordinate $i$ set to $\gamma$. We write $x \sim y$ if $\|x-y\|_{1}=1$. Let $n=|F|$. We use a scheme analogous to the Metropolis algorithm (see Example 1.1.30). A natural symmetric chain on $\mathcal{X}$ is to pick a coordinate uniformly at random, and flip its value. We modify it to guarantee reversibility with respect to the desired stationary distributions, namely $\mu_{1}$ and $\mu_{2}$.

For $\alpha, \beta>0$ small enough, the following transition matrix over $\mathcal{X}$ is irreducible and reversible with respect to its stationary distribution $\mu_{2}$ : for all $i \in F$, $y \in \mathcal{X}$,

$$
\begin{gathered}
Q\left(y^{i, 0}, y^{i, 1}\right)=\frac{1}{n} \alpha\{\beta\}, \\
Q\left(y^{i, 1}, y^{i, 0}\right)=\frac{1}{n} \alpha\left\{\beta \frac{\mu_{2}\left(y^{i, 0}\right)}{\mu_{2}\left(y^{i, 1}\right)}\right\}, \\
Q(y, y)=1-\sum_{z: z \sim y} Q(y, z) .
\end{gathered}
$$

Let $P$ be similarly defined with respect to $\mu_{1}$ with the same values of $\alpha$ and $\beta$. For reasons that will be clear below, the value of $0<\beta<1$ is chosen small enough that the sum of the two expressions in brackets above is smaller than 1 for all $y, i$ in both $P$ and $Q$. The value of $\alpha>0$ is then chosen small enough that $P(y, y), Q(y, y) \geq 0$ for all $y$. Reversibility follows immediately from the first two equations. We call the first transition above an upward transition and the second one a downward transition.

By Theorem 4.2.17, it remains to show that $Q$ stochastically dominates $P$. That is, for any $x \leq y$, we want to show that $P(x, \cdot) \preceq Q(y, \cdot)$. We produce a monotone coupling $(\hat{X}, \hat{Y})$ of these two distributions. Because $x \leq y$, our goal is never to perform an upward transition in $x$ simultaneously with a downward transition in $y$. Observe that

$$
\begin{equation*}
\frac{\mu_{1}\left(x^{i, 0}\right)}{\mu_{1}\left(x^{i, 1}\right)} \geq \frac{\mu_{2}\left(y^{i, 0}\right)}{\mu_{2}\left(y^{i, 1}\right)} \tag{4.2.19}
\end{equation*}
$$

by taking $\omega=y^{i, 0}$ and $\omega^{\prime}=x^{i, 1}$ in Condition (4.2.18).
The coupling works as follows. Fix $x \leq y$. With probability $1-\alpha$, set $(\hat{X}, \hat{Y}):=(x, y)$. Otherwise, pick a coordinate $i \in F$ uniformly at random. There are several cases to consider depending on the coordinates $x_{i}, y_{i}$ (with $x_{i} \leq y_{i}$ by assumption):

- $\left(x_{i}, y_{i}\right)=(0,0)$ : With probability $\beta$, perform an upward transition in both coordinates, that is, set $\hat{X}:=x^{i, 1}$ and $\hat{Y}:=y^{i, 1}$. With probability $1-\beta$, set $(\hat{X}, \hat{Y}):=(x, y)$ instead.
- $\left(x_{i}, y_{i}\right)=(1,1)$ : With probability $\beta \frac{\mu_{2}\left(y^{i, 0}\right)}{\mu_{2}\left(y^{i, 1}\right)}$, perform a downward transition in both coordinates, that is, set $\hat{X}:=x^{i, 0}$ and $\hat{Y}:=y^{i, 0}$. With probability

$$
\beta\left(\frac{\mu_{1}\left(x^{i, 0}\right)}{\mu_{1}\left(x^{i, 1}\right)}-\frac{\mu_{2}\left(y^{i, 0}\right)}{\mu_{2}\left(y^{i, 1}\right)}\right)
$$

perform a downward transition in $x$ only, that is, set $\hat{X}:=x^{i, 0}$ and $\hat{Y}:=$ $y$. With the remaining probability, set $(\hat{X}, \hat{Y}):=(x, y)$ instead. Note that (4.2.19) and our choice of $\beta$ guarantees that this step is well-defined.

- $\left(x_{i}, y_{i}\right)=(0,1)$ : With probability $\beta$, perform an upward transition in $x$ only, that is, set $\hat{X}:=x^{i, 1}$ and $\hat{Y}:=y$. With probability $\beta \frac{\mu_{2}\left(y^{i, 0}\right)}{\mu_{2}\left(y^{i, 1}\right)}$, perform a downward transition in $y$ only, that is, set $\hat{X}:=x$ and $\hat{Y}:=y^{i, 0}$. With the remaining probability, set $(\hat{X}, \hat{Y}):=(x, y)$ instead. Again our choice of $\beta$ guarantees that this step is well-defined.

A little accounting shows that this is indeed a coupling of $P(x, \cdot)$ and $Q(y, \cdot)$. By construction, this coupling satisfies $\hat{X} \leq \hat{Y}$ almost surely. An application of Theorem 4.2.17 concludes the proof.

Example 4.2.33 (Ising model revisited). Holley's inequality implies Claim 4.2.18. To see this, just repeat the calculations of Example 4.2.30, where now (4.2.17) is replaced with an inequality. See Exercise 4.6.

### 4.2.4 $\triangleright$ Random graphs: Janson's inequality and application to the clique number in the Erdö́s-Rényi model

Let $G=(V, E) \sim \mathbb{G}_{n, p_{n}}$ be an Erdős-Rényi graph (see Definition 1.2.2). By Claim 2.3.5, the property of being triangle-free has threshold $n^{-1}$. That is, the probability that $G$ contains a triangle goes to 0 or 1 as $n \rightarrow+\infty$ according to whether $p_{n} \ll n^{-1}$ or $p_{n} \gg n^{-1}$ respectively. In this section, we investigate what happens "at the threshold," by which we mean that we take $p_{n}=\lambda / n$ for some $\lambda>0$ not depending on $n$.

For any subset $S$ of three distinct vertices of $G$, let $B_{S}$ be the event that $S$ forms a triangle in $G$. So

$$
\begin{equation*}
\varepsilon:=\mathbb{P}_{n, p_{n}}\left[B_{S}\right]=p_{n}^{3} \rightarrow 0 . \tag{4.2.20}
\end{equation*}
$$

Denoting the unordered triples of distinct vertices by $\binom{V}{3}$, let $X_{n}=\sum_{S \in\binom{V}{3}} \mathbf{1}_{B_{S}}$ be the number of triangles in $G$. By the linearity of expectation, the mean number of triangles is

$$
\mathbb{E}_{n, p_{n}} X_{n}=\binom{n}{3} p_{n}^{3}=\frac{n(n-1)(n-2)}{6}\left(\frac{\lambda}{n}\right)^{3} \rightarrow \frac{\lambda^{3}}{6},
$$

as $n \rightarrow+\infty$. If the events $\left\{B_{S}\right\}_{S}$ were mutually independent, $X_{n}$ would be binomially distributed and the event that $G$ is triangle-free would have probability

$$
\begin{equation*}
\prod_{S \in\binom{V}{3}} \mathbb{P}_{n, p_{n}}\left[B_{S}^{c}\right]=\left(1-p_{n}^{3}\right)^{\binom{n}{3}} \rightarrow e^{-\lambda^{3} / 6} . \tag{4.2.21}
\end{equation*}
$$

In fact, by the Poisson approximation to the binomial (e.g., Theorem 4.1.18), we would have that the number of triangles converges weakly to $\operatorname{Poi}\left(\lambda^{3} / 6\right)$.

In reality, of course, the events $\left\{B_{S}\right\}$ are not mutually independent. Observe however that, for most pairs $S, S^{\prime}$, the events $B_{S}$ and $B_{S^{\prime}}$ are in fact pairwise independent. That is the case whenever $\left|S \cap S^{\prime}\right| \leq 1$, that is, whenever the edges connecting $S$ are disjoint from those connecting $S^{\prime}$. Write $S \sim S^{\prime}$ if $S \neq S^{\prime}$ are
not independent, that is, if $\left|S \cap S^{\prime}\right|=2$. The expected number of unordered pairs $S \sim S^{\prime}$ both forming a triangle is

$$
\Delta:=\frac{1}{2} \sum_{\substack{S, S^{\prime} \in\left(\begin{array}{l}
V \\
3 \\
S \sim S^{\prime} \\
\hline \tag{4.2.22}
\end{array}\right.}} \mathbb{P}_{n, p_{n}}\left[B_{S} \cap B_{S^{\prime}}\right]=\frac{1}{2}\binom{n}{3}\binom{3}{2}(n-3) p_{n}^{5}=\Theta\left(n^{4} p_{n}^{5}\right) \rightarrow 0,
$$

where the $\binom{n}{3}$ comes from the number of ways of choosing $S$, the $\binom{3}{2}$ comes from the number of ways of choosing the vertices in common between $S$ and $S^{\prime}$, and the $n-3$ comes from the number of ways of choosing the third vertex of $S^{\prime}$. Given that the events $\left\{B_{S}\right\}_{S \in\binom{V}{3}}$ are "mostly" independent, it is natural to expect that $X_{n}$ behaves asymptotically as it would in the independent case. Indeed we prove:

## Claim 4.2.34.

$$
\mathbb{P}_{n, p_{n}}\left[X_{n}=0\right] \rightarrow e^{-\lambda^{3} / 6}
$$

Remark 4.2.35. In fact, $X_{n} \xrightarrow{\mathrm{~d}} \operatorname{Poi}\left(\lambda^{3} / 6\right)$. See Exercises 2.18 and 4.9.
The FKG inequality (Theorem 4.2.28) immediately gives one direction. Recall that $\mathbb{P}_{n, p_{n}}$, as a product measure over edge sets, satisfies the FKG condition and therefore has positive associations by the FKG inequality. Moreover the events $B_{S}^{c}$ are decreasing for all $S$. Hence, applying positive associations inductively,

$$
\mathbb{P}_{n, p_{n}}\left[\bigcap_{S \in\binom{V}{3}} B_{S}^{c}\right] \geq \prod_{S \in\binom{V}{3}} \mathbb{P}_{n, p_{n}}\left[B_{S}^{c}\right] \rightarrow e^{-\lambda^{3} / 6}
$$

where the limit follows from (4.2.21). As it turns out, the FKG inequality also gives a bound in the other direction. This is known as Janson's inequality, which we state in a more general context.

Janson's inequality Let $\mathcal{X}:=\{0,1\}^{F}$ where $F$ is finite. Let $B_{i}, i \in I$, be a finite collection of events of the form

$$
B_{i}:=\left\{\omega \in \mathcal{X}: \omega \geq \beta^{(i)}\right\},
$$

for some $\beta^{(i)} \in \mathcal{X}$. Think of these as "bad events" corresponding to a certain subset of coordinates being set to 1 . By definition, the $B_{i} \mathrm{~s}$ are increasing. Assume $\mathbb{P}$ is a positive product measure on $\mathcal{X}$. Write $i \sim j$ if $\beta_{r}^{(i)}=\beta_{r}^{(j)}=1$ for at least one $r$ and note that $B_{i}$ is independent of $B_{j}$ if $i \nsim j$. Set

$$
\Delta:=\sum_{\substack{\{i, j\} \\ i \sim j}} \mathbb{P}\left[B_{i} \cap B_{j}\right] .
$$

Theorem 4.2.36 (Janson's inequality). Let $\mathcal{X}:=\{0,1\}^{F}$ where $F$ is finite and $\mathbb{P}$ be a positive product measure on $\mathcal{X}$. Let $\left\{B_{i}\right\}_{i \in I}$ and $\Delta$ be as above. Assume further that there is $\varepsilon>0$ such that $\mathbb{P}\left[B_{i}\right] \leq \varepsilon$ for all $i \in I$. Then

$$
\prod_{i \in I} \mathbb{P}\left[B_{i}^{c}\right] \leq \mathbb{P}\left[\cap_{i \in I} B_{i}^{c}\right] \leq e^{\frac{\Delta}{1-\varepsilon}} \prod_{i \in I} \mathbb{P}\left[B_{i}^{c}\right] .
$$

Before proving the theorem, we show that it implies Claim 4.2.34. We have already shown in (4.2.20) and (4.2.22) that $\varepsilon \rightarrow 0$ and $\Delta \rightarrow 0$. Janson's inequality (Theorem 4.2.36) immediately implies the claim by (4.2.21).

Proof of Theorem 4.2.36. The lower bound is the FKG inequality.
In the other direction, assume without loss of generality that $I=[m]$. The first step is to apply the chain rule to obtain

$$
\mathbb{P}\left[\cap_{i \in I} B_{i}^{c}\right]=\prod_{i=1}^{m} \mathbb{P}\left[B_{i}^{c} \mid \cap_{j \in[i-1]} B_{j}^{c}\right] .
$$

The rest is clever manipulation. For $i \in[m]$, let $N(i):=\{\ell \in[m]: \ell \sim i\}$ and $N_{<}(i):=N(i) \cap[i-1]$. Note that $B_{i}$ is independent of $\left\{B_{\ell}: \ell \in[i-1] \backslash N_{<}(i)\right\}$. Hence,

$$
\begin{aligned}
\mathbb{P}\left[B_{i} \mid \cap_{j \in[i-1]} B_{j}^{c}\right] & =\frac{\mathbb{P}\left[B_{i} \cap\left(\cap_{j \in[i-1]} B_{j}^{c}\right)\right]}{\mathbb{P}\left[\cap_{j \in[i-1]} B_{j}^{c}\right]} \\
& \geq \frac{\mathbb{P}\left[B_{i} \cap\left(\cap_{j \in N_{<}(i)} B_{j}^{c}\right) \cap\left(\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right)\right]}{\mathbb{P}\left[\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right]} \\
& =\mathbb{P}\left[B_{i} \cap\left(\cap_{j \in N_{<}(i)} B_{j}^{c}\right) \mid \cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right] \\
& =\mathbb{P}\left[B_{i} \mid \cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right] \\
& \times \mathbb{P}\left[\cap_{j \in N_{<}(i)} B_{j}^{c} \mid B_{i} \cap\left(\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right)\right] \\
& =\mathbb{P}\left[B_{i}\right] \mathbb{P}\left[\cap_{j \in N_{<}(i)} B_{j}^{c} \mid B_{i} \cap\left(\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right)\right]
\end{aligned}
$$

where we used independence for the first term on the last line. By a union bound, the second term on the last line is

$$
\begin{aligned}
\mathbb{P}[ & \left.\cap_{j \in N_{<}(i)} B_{j}^{c} \mid B_{i} \cap\left(\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right)\right] \\
& \geq 1-\sum_{j \in N_{<}(i)} \mathbb{P}\left[B_{j} \mid B_{i} \cap\left(\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}\right)\right] \\
& \geq 1-\sum_{j \in N_{<}(i)} \mathbb{P}\left[B_{j} \mid B_{i}\right],
\end{aligned}
$$

where the last line follows from the FKG inequality. This requires some explanations:

- On the event $B_{i}$, all coordinates $\ell$ with $\beta_{\ell}^{(i)}=1$ are fixed to 1 , and the other ones are free. So we can think of $\mathbb{P}\left[\cdot \mid B_{i}\right]$ as a positive product measure on $\{0,1\}^{F^{\prime}}$ with $F^{\prime}:=\left\{\ell \in[m]: \beta_{\ell}^{(i)}=0\right\}$.
- The event $B_{j}$ is increasing, while the event $\cap_{j \in[i-1] \backslash N_{<}(i)} B_{j}^{c}$ is decreasing as the intersection of decreasing events (see Exercise 4.4).
- So we can apply the FKG inequality in the form (4.2.14) to $\mathbb{P}\left[\cdot \mid B_{i}\right]$.

Combining the last three displays and using $1+x \leq e^{x}$ for all $x$ (see Exercise 1.16), we get

$$
\begin{aligned}
\mathbb{P}\left[\cap_{i \in I} B_{i}^{c}\right] & \leq \prod_{i=1}^{m}\left[\mathbb{P}\left[B_{i}^{c}\right]+\sum_{j \in N_{<}(i)} \mathbb{P}\left[B_{i} \cap B_{j}\right]\right] \\
& \leq \prod_{i=1}^{m} \mathbb{P}\left[B_{i}^{c}\right]\left[1+\frac{1}{1-\varepsilon} \sum_{j \in N_{<}(i)} \mathbb{P}\left[B_{i} \cap B_{j}\right]\right] \\
& \leq \prod_{i=1}^{m} \mathbb{P}\left[B_{i}^{c}\right] \exp \left(\frac{1}{1-\varepsilon} \sum_{j \in N_{<}(i)} \mathbb{P}\left[B_{i} \cap B_{j}\right]\right) .
\end{aligned}
$$

where we used the assumption $\mathbb{P}\left[B_{i}\right] \leq \varepsilon$ on the second line. By the definition of $\Delta$, we are done.

### 4.2.5 $\triangleright$ Percolation: RSW theory and a proof of Harris' theorem

Consider bond percolation (Definition 1.2.1) on the two-dimensional lattice $\mathbb{L}^{2}$. Recall that the percolation function is given by

$$
\theta(p):=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right],
$$

where $\mathcal{C}_{0}$ is the open cluster of the origin. We know from Example 4.2.10 that $\theta(p)$ is non-decreasing. Let

$$
p_{c}\left(\mathbb{L}^{2}\right):=\sup \{p \geq 0: \theta(p)=0\},
$$

be the critical value. We proved in Section 2.2.4 that there is a non-trivial transition, that is, $p_{c}\left(\mathbb{L}^{2}\right) \in(0,1)$. See Exercise 2.3 for a proof that $p_{c}\left(\mathbb{L}^{2}\right) \in[1 / 3,2 / 3]$.

Our goal in this section is to use the FKG inequality to improve this further to:

Theorem 4.2.37 (Harris' theorem).

$$
\theta(1 / 2)=0 .
$$

Or, put differently, $p_{c}\left(\mathbb{L}^{2}\right) \geq 1 / 2$.
Remark 4.2.38. This bound is tight, that is, in fact $p_{c}\left(\mathbb{L}^{2}\right)=1 / 2$. The other direction is known as Kesten's theorem. See, e.g., [BRO6a].

Here we present a proof of Harris' theorem that uses an important tool in percolation theory, the Russo-Seymour-Welsh (RSW) lemma, an application of the FKG inequality.

## Harris' theorem

To motivate the RSW lemma, we start with the proof of Harris' theorem.
Proof of Theorem 4.2.37. Fix $p=1 / 2$. We use the dual lattice $\widetilde{\mathbb{L}}^{2}$ as we did in Section 2.2.4. Consider the annulus

$$
\operatorname{Ann}(\ell):=[-3 \ell, 3 \ell]^{2} \backslash[-\ell, \ell]^{2} .
$$

The existence of a closed dual cycle inside $\operatorname{Ann}(\ell)$, an event we denote by $\mathrm{O}_{\mathrm{d}}(\ell)$, prevents the possibility of an infinite open path from the origin in the primal lattice $\mathbb{L}^{2}$. See Figure 4.4. That is,

$$
\begin{equation*}
\mathbb{P}_{1 / 2}\left[\left|\mathcal{C}_{0}\right|=+\infty\right] \leq \prod_{k=0}^{K}\left\{1-\mathbb{P}_{1 / 2}\left[\mathrm{O}_{\mathrm{d}}\left(3^{k}\right)\right]\right\} \tag{4.2.23}
\end{equation*}
$$

for all $K$, where we took powers of 3 to make the annuli disjoint and therefore independent. To prove the theorem, it suffices to show that there is a constant $c^{*}>0$ such that, for all $\ell, \mathbb{P}_{1 / 2}\left[\mathrm{O}_{\mathrm{d}}(\ell)\right] \geq c^{*}$. Then the right-hand side of (4.2.23) tends to 0 as $K \rightarrow+\infty$.

To simplify further, thinking of $\operatorname{Ann}(\ell)$ as a union of four rectangles $[-3 \ell,-\ell) \times$ $[-3 \ell, 3 \ell],[-3 \ell, 3 \ell] \times(\ell, 3 \ell]$, etc., it suffices to consider the event $\mathrm{O}_{\mathrm{d}}^{\#}(\ell)$ that each one of these rectangles contains a closed dual path connecting its two shorter sides. To be more precise, for the first rectangle above for instance, the path connects $[-3 \ell+1 / 2,-\ell-1 / 2] \times\{3 \ell-1 / 2\}$ to $[-3 \ell+1 / 2,-\ell-1 / 2] \times\{-3 \ell+1 / 2\}$ and stays inside the rectangle. See Figure 4.4. By symmetry the probability that such a path exists is the same for all four rectangles. Denote that probability by $\rho_{\ell}$. Moreover the event that such a path exists is increasing so, although the four events


Figure 4.4: Top: the event $\mathrm{O}_{\mathrm{d}}(\ell)$. Bottom: the event $\mathrm{O}_{\mathrm{d}}^{\#}(\ell)$.
are not independent, we can apply the FKG inequality (Theorem 4.2.28). Hence, since $\mathrm{O}_{\mathrm{d}}^{\#}(\ell) \subseteq \mathrm{O}_{\mathrm{d}}(\ell)$, we finally get the bound

$$
\mathbb{P}_{1 / 2}\left[\mathrm{O}_{\mathrm{d}}(\ell)\right] \geq \rho_{\ell}^{4} .
$$

The RSW lemma and some symmetry arguments, both of which are detailed below, imply:

Claim 4.2.39. There is some $c>0$ such that, for all $\ell$,

$$
\rho_{\ell} \geq c .
$$

That concludes the proof.
It remains to prove Claim 4.2.39. We first state the RSW lemma.

## RSW theory

We have reduced the proof of Harris' theorem to bounding the probability that certain closed paths exist in the dual lattice. To be consistent with the standard RSW notation, we switch to the primal lattice and consider open paths. We also let $p$ take any value in $(0,1)$.

Let $R_{n, \alpha}(p)$ be the probability that the rectangle

$$
B(\alpha n, n):=[-n,(2 \alpha-1) n] \times[-n, n],
$$

has an open path connecting its left and right sides with the path remaining inside the rectangle. Such a path is called an (open) left-right crossing. The event that a left-right crossing exists in a rectangle $B$ is denoted by $\operatorname{LR}(B)$. We similarly define the event, $\mathrm{TB}(B)$, that a top-bottom crossing exists in $B$. In essence, the RSW lemma says this: if there is a significant probability that a left-right crossing exists in the square $B(n, n)$, then there is a significant probability that a left-right crossing exists in the rectangle $B(3 n, n)$. More precisely, here is a version of the theorem that will be enough for our purposes. (See Exercise 4.10 for a generalization.)

Lemma 4.2.40 (RSW lemma). For all $n \geq 2$ (divisible by 4) and $p \in(0,1)$,

$$
\begin{equation*}
R_{n, 3}(p) \geq \frac{1}{256} R_{n, 1}(p)^{11} R_{n / 2,1}(p)^{12} \tag{4.2.24}
\end{equation*}
$$

The right-hand side of (4.2.24) depends only on the probability of crossing a square from left to right. By a duality argument, at $p=1 / 2$, it turns out that this probability is at least $1 / 2$ independently of $n$. Before presenting a proof of the RSW lemma, we detail this argument and finish the proof of Harris' theorem.

Proof of Claim 4.2.39. The point of (4.2.24) is that, if $R_{n, 1}(1 / 2)$ is bounded away from 0 uniformly in $n$, then so is the left-hand side. By the argument in the proof of Harris' theorem, this then implies that a closed cycle exists in $\operatorname{Ann}(n)$ with a probability bounded away from 0 as well. Hence to prove Claim 4.2.39 it suffices to give a lower bound on $R_{n, 1}(1 / 2)$. It is crucial that this bound not depend on the "scale" $n$.

As it turns out, a simple duality-based symmetry argument does the trick. The following fact about $\mathbb{L}^{2}$ is a variant of the contour lemma (Lemma 2.2.14). Its proof is similar and Exercise 4.11 asks for the details (the "if" direction being the non-trivial implication).

Lemma 4.2.41. There is an open left-right crossing in the primal rectangle $[0, n+$ $1] \times[0, n]$ if and only if there is no closed top-bottom crossing in the dual rectangle $[1 / 2, n+1 / 2] \times[-1 / 2, n+1 / 2]$.

By symmetry, when $p=1 / 2$, the two events in Lemma 4.2.41 have equal probability. So they must have probability $1 / 2$ because they form a partition of the space of outcomes. By monotonicity, that implies $R_{n, 1}(1 / 2) \geq 1 / 2$ for all $n$. The RSW lemma then implies the required bound.

The proof of the RSW lemma involves a clever choice of event that relates the existence of crossings in squares and rectangles. (Combining crossings of squares into crossings of rectangles is not as trivial as it might look. Try it before reading the proof.)

Proof of Lemma 4.2.40. There are several steps in the proof.

Step 1: it suffices to bound $R_{n, 3 / 2}(p)$ We first reduce the proof to finding a bound on $R_{n, 3 / 2}(p)$. Let $B_{1}^{\prime}:=B(2 n, n)$ and $B_{2}^{\prime}:=[n, 5 n] \times[-n, n]$. Note that $B_{1}^{\prime} \cup B_{2}^{\prime}=B(3 n, n)$ and $B_{1}^{\prime} \cap B_{2}^{\prime}=[n, 3 n] \times[-n, n]$. Then we have the implication

$$
\operatorname{LR}\left(B_{1}^{\prime}\right) \cap \mathrm{TB}\left(B_{1}^{\prime} \cap B_{2}^{\prime}\right) \cap \mathrm{LR}\left(B_{2}^{\prime}\right) \subseteq \operatorname{LR}(B(3 n, n))
$$

See Figure 4.5. Each event on the left-hand side is increasing so the FKG inequality gives

$$
R_{n, 3}(p) \geq R_{n, 2}(p)^{2} R_{n, 1}(p) .
$$

A similar argument over $B(2 n, n)$ gives

$$
R_{n, 2}(p) \geq R_{n, 3 / 2}(p)^{2} R_{n, 1}(p)
$$

Combining the two, we have proved:


Figure 4.5: Illustration of the implication $\operatorname{LR}\left(B_{1}^{\prime}\right) \cap \mathrm{TB}\left(B_{1}^{\prime} \cap B_{2}^{\prime}\right) \cap \operatorname{LR}\left(B_{2}^{\prime}\right) \subseteq$ $\operatorname{LR}(B(3 n, n))$.

Lemma 4.2.42 (Proof of RSW: step 1).

$$
\begin{equation*}
R_{n, 3}(p) \geq R_{n, 3 / 2}(p)^{4} R_{n, 1}(p)^{3} . \tag{4.2.25}
\end{equation*}
$$

Step 2: bounding $R_{n, 3 / 2}(p)$ The heart of the proof is to bound $R_{n, 3 / 2}(p)$ using an event involving crossings of squares. Let

$$
\begin{aligned}
B_{1} & :=B(n, n)=[-n, n] \times[-n, n], \\
B_{2} & :=[0,2 n] \times[-n, n], \\
B_{12} & :=B_{1} \cap B_{2}=[0, n] \times[-n, n], \\
S & :=[0, n] \times[0, n] .
\end{aligned}
$$

Let $\Gamma_{1}$ be the event that there are paths $P_{1}, P_{2}$, where $P_{1}$ is a top-bottom crossing of $S$ and $P_{2}$ is an open path connecting the left side of $B_{1}$ to $P_{1}$ and stays inside $B_{1}$. Similarly let $\Gamma_{2}^{\prime}$ be the event that there are paths $P_{1}^{\prime}, P_{2}^{\prime}$, where $P_{1}^{\prime}$ is a topbottom crossing of $S$ and $P_{2}^{\prime}$ is an open path connecting the right side of $B_{2}$ to $P_{1}^{\prime}$ and stays inside $B_{2}$. Then we have the implication

$$
\Gamma_{1} \cap \operatorname{LR}(S) \cap \Gamma_{2}^{\prime} \subseteq \operatorname{LR}(B(3 n / 2, n))
$$

See Figure 4.6. By symmetry $\mathbb{P}_{p}\left[\Gamma_{1}\right]=\mathbb{P}_{p}\left[\Gamma_{2}^{\prime}\right]$. Moreover, the events on the lefthand side are increasing so by the FKG inequality:

Lemma 4.2.43 (Proof of RSW: step 2).

$$
\begin{equation*}
R_{n, 3 / 2}(p) \geq \mathbb{P}_{p}\left[\Gamma_{1}\right]^{2} R_{n / 2,1}(p) . \tag{4.2.26}
\end{equation*}
$$

Step 3: bounding $\mathbb{P}_{p}\left[\Gamma_{1}\right]$ It remains to bound $\mathbb{P}_{p}\left[\Gamma_{1}\right]$. That requires several additional definitions. Let $P_{1}$ and $P_{2}$ be top-bottom crossings of $S$. There is a natural partial order over such crossings. The path $P_{1}$ divides $S$ into two subgraphs: $\left[P_{1}\right\}$ which includes the left side of $S$ (including edges on the left incident with $P_{1}$ but not those edges on $P_{1}$ itself) and $\left\{P_{1}\right\}$ which includes the right side of $S$ (and $P_{1}$ itself). Then we write $P_{1} \preceq P_{2}$ if $\left\{P_{1}\right] \subseteq\left\{P_{2}\right]$. Assuming $\operatorname{TB}(S)$ holds, one also gets the existence of a unique rightmost crossing. Roughly speaking, take the union of all top-bottom crossings of $S$ as sets of edges; then the "right boundary" of this set is a top-bottom crossing $P_{S}^{*}$ such that $P_{S}^{*} \preceq P$ for all top-bottom crossings $P$ of $S$. (We accept as a fact the existence of a unique rightmost crossing. See Exercise 4.11 for a related construction.)

Let $I_{S}$ be the set of (not necessarily open) paths connecting the top and bottom of $S$ and stay inside $S$. For $P \in I_{S}$, we let $P^{\prime}$ be the reflection of $P$ in $B_{12} \backslash S$ through the $x$-axis and we let $\frac{P}{P^{\prime}}$ be the union of $P$ and $P^{\prime}$. Define $\left\{\frac{P}{P^{\prime}}\right\}$ to be the


Figure 4.6: Top: illustration of the implication $\Gamma_{1} \cap \mathrm{LR}(S) \cap \Gamma_{2}^{\prime} \subseteq$ $\operatorname{LR}(B(3 n / 2, n))$. Bottom: the event $\mathrm{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right) \cap\left\{P=P_{S}^{*}\right\}$; the dashed path is the mirror image of the rightmost top-bottom crossing in $S$; the shaded region on the right is the complement in $B_{1}$ of the set $\left[\frac{P}{P}\right\}$. Note that, because in the bottom figure the left-right path must stay within $\left\{\frac{P}{P^{\prime}}\right\}$ by definition of $P_{S}^{*}$, the configuration shown in the top figure where a left-right path (dotted) "travels behind" the top-bottom crossing of $S$ cannot occur.
subgraph of $B_{1}$ to the left of $\frac{P}{P^{\prime}}$ (including edges on the left incident with $\frac{P}{P^{\prime}}$ but not those edges on $\frac{P}{P^{\prime}}$ itself). Let $\mathrm{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right)$ be the event that there is a left-right crossing of $\left\{\frac{P}{P^{\prime}}\right\}$ ending on $P$, that is, that there is an open path connecting the left side of $B_{1}$ and $P$ that stays within $\left[\frac{P}{P^{\prime}}\right\}$. See Figure 4.6. Note that the existence of a left-right crossing of $B_{1}$ implies the existence of an open path connecting the left side of $B_{1}$ to $\frac{P}{P^{\prime}}$. By symmetry we then get

$$
\begin{equation*}
\mathbb{P}_{p}\left[\mathrm{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right)\right] \geq \frac{1}{2} \mathbb{P}_{p}\left[\mathrm{LR}\left(B_{1}\right)\right]=\frac{1}{2} R_{n, 1}(p) \tag{4.2.27}
\end{equation*}
$$

Now comes a subtle point. We turn to the rightmost crossing of $S$-for two reasons:

- First, by uniqueness of the rightmost crossing, $\left\{P_{S}^{*}=P\right\}_{P \in I_{S}}$ forms a partition of $\mathrm{TB}(S)$. Recall that we are looking to bound a probability from below, and therefore we have to be careful not to "double count."
- Second, the rightmost crossing has a Markov-like property. Observe that, for $P \in I_{S}$, the event that $\left\{P_{S}^{*}=P\right\}$ depends only the bonds in $\{P]$. In particular it is independent of the bonds in $\left\{\frac{P}{P^{\prime}}\right\}$, for example, of the event $\mathrm{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right)$. Hence

$$
\begin{equation*}
\mathbb{P}_{p}\left[\left.\operatorname{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right) \right\rvert\, P_{S}^{*}=P\right]=\mathbb{P}_{p}\left[\operatorname{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right)\right] . \tag{4.2.28}
\end{equation*}
$$

Note that the event $\left\{P_{S}^{*}=P\right\}$ is not increasing, as adding more open bonds can shift the rightmost crossing rightward. Therefore, we cannot use the FKG inequality here.

Combining (4.2.27) and (4.2.28), we get

$$
\begin{aligned}
\mathbb{P}_{p}\left[\Gamma_{1}\right] & \geq \sum_{P \in I_{S}} \mathbb{P}_{p}\left[P_{S}^{*}=P\right] \mathbb{P}_{p}\left[\left.\operatorname{LR}^{+}\left(\left[\frac{P}{P^{\prime}}\right\}\right) \right\rvert\, P_{S}^{*}=P\right] \\
& \geq \frac{1}{2} R_{n, 1}(p) \sum_{P \in I_{S}} \mathbb{P}_{p}\left[P_{S}^{*}=P\right] \\
& =\frac{1}{2} R_{n, 1}(p) \mathbb{P}_{p}[\mathrm{~TB}(S)] \\
& =\frac{1}{2} R_{n, 1}(p) R_{n / 2,1}(p) .
\end{aligned}
$$

We have proved:
Lemma 4.2.44 (Proof of RSW: step 3).

$$
\begin{equation*}
\mathbb{P}_{p}\left[\Gamma_{1}\right] \geq \frac{1}{2} R_{n, 1}(p) R_{n / 2,1}(p) . \tag{4.2.29}
\end{equation*}
$$

Step 4: putting everything together Combining (4.2.25), (4.2.26) and (4.2.29) gives that

$$
\begin{aligned}
R_{n, 3}(p) & \geq R_{n, 3 / 2}(p)^{4} R_{n, 1}(p)^{3} \\
& \geq\left[\mathbb{P}_{p}\left[\Gamma_{1}\right]^{2} R_{n / 2,1}(p)\right]^{4} R_{n, 1}(p)^{3} \\
& \geq\left[\left(\frac{1}{2} R_{n, 1}(p) R_{n / 2,1}(p)\right)^{2} R_{n / 2,1}(p)\right]^{4} R_{n, 1}(p)^{3} .
\end{aligned}
$$

Collecting the terms concludes the proof of the RSW lemma.
Remark 4.2.45. This argument is quite subtle. It is instructive to read the remark after [Gri97, Theorem 9.3].

### 4.3 Coupling of Markov chains and application to mixing

As we have seen, coupling is useful to bound total variation distance. In this section we apply the technique to bound the mixing time of Markov chains.

### 4.3.1 Bounding the mixing time via coupling

Let $P$ be an irreducible, aperiodic Markov transition matrix on the finite state space $V$ with stationary distribution $\pi$. Recall from Definition 1.1.35 that, for a fixed $0<\varepsilon<1 / 2$, the mixing time of $P$ is

$$
\mathrm{t}_{\text {mix }}(\varepsilon):=\min \{t: d(t) \leq \varepsilon\},
$$

where

$$
d(t):=\max _{x \in V}\left\|P^{t}(x, \cdot)-\pi\right\|_{\mathrm{TV}} .
$$

It will be easier to work with

$$
\bar{d}(t):=\max _{x, y \in V}\left\|P^{t}(x, \cdot)-P^{t}(y, \cdot)\right\|_{\mathrm{TV}} .
$$

The quantities $d(t)$ and $\bar{d}(t)$ are related in the following way.

## Lemma 4.3.1.

$$
d(t) \leq \bar{d}(t) \leq 2 d(t), \quad \forall t
$$

Proof. The second inequality follows from an application of the triangle inequality.

For the first inequality, note that by definition of the total variation distance and the stationarity of $\pi$

$$
\begin{aligned}
\left\|P^{t}(x, \cdot)-\pi\right\|_{\mathrm{TV}} & =\sup _{A \subseteq V}\left|P^{t}(x, A)-\pi(A)\right| \\
& =\sup _{A \subseteq V}\left|\sum_{y \in V} \pi(y)\left[P^{t}(x, A)-P^{t}(y, A)\right]\right| \\
& \leq \sup _{A \subseteq V} \sum_{y \in V} \pi(y)\left|P^{t}(x, A)-P^{t}(y, A)\right| \\
& \leq \sum_{y \in V} \pi(y)\left\{\sup _{A \subseteq V}\left|P^{t}(x, A)-P^{t}(y, A)\right|\right\} \\
& \leq \sum_{y \in V} \pi(y)\left\|P^{t}(x, \cdot)-P^{t}(y, \cdot)\right\|_{\mathrm{TV}} \\
& \leq \max _{x, y \in V}\left\|P^{t}(x, \cdot)-P^{t}(y, \cdot)\right\|_{\mathrm{TV}} .
\end{aligned}
$$

Coalescence Recall that a Markovian coupling of $P$ with itself is a Markov chain $\left(X_{t}, Y_{t}\right)_{t}$ on $V \times V$ with transition matrix $Q$ satisfying: for all $x, y, x^{\prime}, y^{\prime} \in V$,

$$
\begin{aligned}
& \sum_{z^{\prime}} Q\left((x, y),\left(x^{\prime}, z^{\prime}\right)\right)=P\left(x, x^{\prime}\right) \\
& \sum_{z^{\prime}} Q\left((x, y),\left(z^{\prime}, y^{\prime}\right)\right)=P\left(y, y^{\prime}\right) .
\end{aligned}
$$

We say that a Markovian coupling is coalescing if further: for all $z \in V$,

$$
x^{\prime} \neq y^{\prime} \Longrightarrow Q\left((z, z),\left(x^{\prime}, y^{\prime}\right)\right)=0 .
$$

Let $\left(X_{t}, Y_{t}\right)$ be a coalescing Markovian coupling of $P$. By the coalescing condition, if $X_{s}=Y_{s}$ then $X_{t}=Y_{t}$ for all $t \geq s$. That is, once $\left(X_{t}\right)$ and $\left(Y_{t}\right)$ meet, they remain equal. Let $\tau_{\text {coal }}$ be the coalescence time (also called coupling time), that is,

$$
\tau_{\text {coal }}:=\inf \left\{t \geq 0: X_{t}=Y_{t}\right\} .
$$

The key to the coupling approach to mixing times is the following immediate consequence of the coupling inequality (Lemma 4.1.11). For any starting point $(x, y)$,

$$
\begin{equation*}
\left\|P^{t}(x, \cdot)-P^{t}(y, \cdot)\right\|_{\mathrm{TV}} \leq \mathbb{P}_{(x, y)}\left[X_{t} \neq Y_{t}\right]=\mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t\right] . \tag{4.3.1}
\end{equation*}
$$

Combining (4.3.1) and Lemma 4.3.1, we get the main tool of this section.

Theorem 4.3.2 (Bounding the mixing time: coupling method). Let $\left(X_{t}, Y_{t}\right)$ be a coalescing Markovian coupling of an irreducible transition matrix $P$ on a finite state space $V$ with stationary distribution $\pi$. Then

$$
d(t) \leq \max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t\right]
$$

In particular

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \inf \left\{t \geq 0: \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t\right] \leq \varepsilon, \forall x, y\right\} .
$$

Note that a Markovian coupling can be made coalescing by modifying it as follows: when $X_{t}=Y_{t}$, perform one step of the chain to determine $X_{t+1}$ and set $Y_{t+1}:=$ $X_{t+1}$. That modification does not affect the coalescence time.

We give a few simple examples of the coupling method in the next subsection. First, we discuss a classical result.

Example 4.3.3 (Doeblin's condition). Let $P$ be a transition matrix on a countable space $V$. One form of Doeblin's condition (also called a minorization condition) is: there is $s \in \mathbb{Z}_{+}$and $\delta>0$ such that

$$
\sup _{z \in V} \inf _{w \in V} P^{s}(w, z)>\delta
$$

In words there is a state $z_{0} \in V$ such that, starting from any state $w \in V$, the probability of reaching $z_{0}$ in exactly $s$ steps is at least $\delta$ (which does not depend on $w)$. Assume such a $z_{0}$ exists.

We construct a coalescing Markovian coupling $\left(X_{t}, Y_{t}\right)$ of $P$. Assume first that $s=1$ and let

$$
\tilde{P}(w, z)=\frac{1}{1-\delta}\left[P(w, z)-\delta \mathbf{1}\left\{z=z_{0}\right\}\right] .
$$

It can be checked that $\tilde{P}$ is a stochastic matrix on $V$ provided $z_{0}$ satisfies the condition above (see Exercise 4.13). We use a technique known as splitting. While $X_{t} \neq Y_{t}$, at the next time step: (i) with probability $\delta$ we set $X_{t+1}=Y_{t+1}=z_{0}$, (ii) otherwise we pick $X_{t+1} \sim \tilde{P}\left(X_{t}, \cdot\right)$ and $Y_{t+1} \sim \tilde{P}\left(Y_{t}, \cdot\right)$ independently. On the other hand, if $X_{t}=Y_{t}$, we maintain the equality and pick the next state according to $P$. Put differently, the coupling $Q$ is defined as: if $x \neq y$,

$$
Q\left((x, y),\left(x^{\prime}, y^{\prime}\right)\right)=\delta \mathbf{1}\left\{x^{\prime}=y^{\prime}=z_{0}\right\}+(1-\delta) \tilde{P}\left(x, x^{\prime}\right) \tilde{P}\left(y, y^{\prime}\right),
$$

while if $x=y$,

$$
Q\left((x, x),\left(x^{\prime}, x^{\prime}\right)\right)=P\left(x, x^{\prime}\right)
$$

Observe that, in case (i) above, coalescence occurs at time $t+1$. In case (ii), coalescence may or may not occur at time $t+1$. In other words, while $X_{t} \neq$
$Y_{t}$, coalescence occurs at the next step with probability at least $\delta$. So $\tau_{\text {coal }}$ is stochastically dominated by a geometric random variable with success probability $\delta$, or

$$
\max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t\right] \leq(1-\delta)^{t}
$$

By Theorem 4.3.2,

$$
\max _{x \in V}\left\|P^{t}(x, \cdot)-\pi\right\|_{\mathrm{TV}} \leq(1-\delta)^{t} .
$$

Exponential decay of the worst-case total variation distance to the stationary distribution is referred to as uniform geometric ergodicity.

Suppose now that $s>1$. We apply the argument above to the chain $P^{s}$ this time. We get

$$
\max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t s\right] \leq(1-\delta)^{t}
$$

uniform
so that, after a change of variable,

$$
\max _{x \in V}\left\|P^{t}(x, \cdot)-\pi\right\|_{\mathrm{TV}} \leq(1-\delta)^{\lfloor t / s\rfloor}
$$

So, we have shown that uniform geometric ergodicity is implied by Doeblin's condition.

We note however that the rate of decay derived from this technique can be very slow. For instance the condition always holds when $P$ is finite, irreducible and aperiodic (as follows from Lemma 1.1.32), but a straight application of the technique may lead to a bound depending badly on the size of the state space $V$ (see Exercise 4.14).

### 4.3.2 $\triangleright$ Random walks: mixing on cycles, hypercubes, and trees

In this section, we consider lazy simple random walk on various graphs. By this we mean that the walk stays put with probability $1 / 2$ and otherwise picks an adjacent vertex uniformly at random. In each case, we construct a coupling to bound the mixing time. As a reference, we compare our upper bounds to the diameter-based lower bound we will derive in Section 5.2.3. Specifically, by Claim 5.2.25, for a finite, reversible Markov chain with stationary distribution $\pi$ and diameter $\Delta$ we have the lower bound

$$
\mathrm{t}_{\mathrm{mix}}(\varepsilon)=\Omega\left(\frac{\Delta^{2}}{\log \left(n \vee \pi_{\min }^{-1}\right)}\right)
$$

where $\pi_{\text {min }}$ is the smallest value taken by $\pi$.

## Cycle

Let $\left(Z_{t}\right)$ be lazy simple random walk on the cycle of size $n, \mathbb{Z}_{n}:=\{0,1 \ldots, n-1\}$, where $i \sim j$ if $|j-i|=1(\bmod n)$. For any starting points $x, y$, we construct a Markovian coupling $\left(X_{t}, Y_{t}\right)$ of this chain. Set $\left(X_{0}, Y_{0}\right):=(x, y)$. At each time, flip a fair coin. On heads, $Y_{t}$ stays put and $X_{t}$ moves one step, the direction of which is uniform at random. On tails, proceed similarly with the roles of $X_{t}$ and $Y_{t}$ reversed. Let $D_{t}$ be the clockwise distance between $X_{t}$ and $Y_{t}$. Observe that, by construction, $\left(D_{t}\right)$ is simple random walk on $\{0, \ldots, n\}$ and $\tau_{\text {coal }}=\tau_{\{0, n\}}^{D}$, the first time $\left(D_{t}\right)$ hits $\{0, n\}$.

We use Markov's inequality (Theorem 2.1.1) to bound $\mathbb{P}_{(x, y)}\left[\tau_{\{0, n\}}^{D}>t\right]$. Denote by $D_{0}=d_{x, y}$ the starting distance. By Wald's second identity (Theorem 3.1.40),

$$
\mathbb{E}_{(x, y)}\left[\tau_{\{0, n\}}^{D}\right]=d_{x, y}\left(n-d_{x, y}\right) .
$$

Applying Theorem 4.3.2 and Markov's inequality we get

$$
\begin{aligned}
d(t) & \leq \max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>t\right] \\
& \leq \max _{x, y \in V} \frac{\mathbb{E}_{(x, y)}\left[\tau_{\{0, n\}}^{D}\right]}{t} \\
& =\max _{x, y \in V} \frac{d_{x, y}\left(n-d_{x, y}\right)}{t} \\
& \leq \frac{n^{2}}{4 t},
\end{aligned}
$$

or:

## Claim 4.3.4.

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \frac{n^{2}}{4 \varepsilon}
$$

By the diameter-based lower bound on mixing in Section 5.2.3, this bound gives the correct order of magnitude in $n$ up to logarithmic factors. Indeed, the diameter is $\Delta=n / 2$ and $\pi_{\min }=1 / n$ so that Claim 5.2.25 gives

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{n^{2}}{64 \log n},
$$

for $n$ large enough. Exercise 4.15 sketches a tighter lower bound.

## Hypercube

Let $\left(Z_{t}\right)_{t \in \mathbb{Z}_{+}}$be lazy simple random walk on the $n$-dimensional hypercube $\mathbb{Z}_{2}^{n}:=$ $\{0,1\}^{n}$ where $i \sim j$ if $\|i-j\|_{1}=1$. We denote the coordinates of $Z_{t}$ by $\left(Z_{t}^{(1)}, \ldots, Z_{t}^{(n)}\right)$. This is equivalent to performing the Glauber dynamics chain on an empty graph (see Definition 1.2.8): at each step, we first pick a coordinate uniformly at random, then refresh its value. Because of the way the updating is done, the chain stays put with probability $1 / 2$ at each time as required.

Inspired by this observation, the coupling $\left(X_{t}, Y_{t}\right)$ started at $(x, y)$ is the following. At each time $t$, pick a coordinate $i$ uniformly at random in $[n]$, pick a bit value $b$ in $\{0,1\}$ uniformly at random independently of the coordinate choice. Set both $i$ coordinates to $b$, that is, $X_{t}^{(i)}=Y_{t}^{(i)}=b$. By design we reach coalescence when all coordinates have been updated at least once.

The following standard bound from the coupon collector's problem (see Example 2.1.4) is what is needed to conclude.

Lemma 4.3.5. Let $\tau_{\text {coll }}$ be the time it takes to update each coordinate at least once.
Then, for any $c>0$,

$$
\mathbb{P}\left[\tau_{\text {coll }}>\lceil n \log n+c n\rceil\right] \leq e^{-c}
$$

Proof. Let $B_{i}$ be the event that the $i$-th coordinate has not been updated by time $\lceil n \log n+c n\rceil$. Then, using that $1-x \leq e^{-x}$ for all $x$ (see Exercise 1.16),

$$
\begin{aligned}
\mathbb{P}\left[\tau_{\text {coll }}>\lceil n \log n+c n\rceil\right] & \leq \sum_{i} \mathbb{P}\left[B_{i}\right] \\
& =\sum_{i}\left(1-\frac{1}{n}\right)^{\lceil n \log n+c n\rceil} \\
& \leq n \exp \left(-\frac{n \log n+c n}{n}\right) \\
& =e^{-c}
\end{aligned}
$$

Applying Theorem 4.3.2, we get

$$
\begin{aligned}
d(\lceil n \log n+c n\rceil) & \leq \max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau_{\text {coal }}>\lceil n \log n+c n\rceil\right] \\
& \leq \mathbb{P}\left[\tau_{\text {coll }}>\lceil n \log n+c n\rceil\right] \\
& \leq e^{-c .}
\end{aligned}
$$

Hence for $c_{\varepsilon}>0$ large enough:

## Claim 4.3.6.

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq\left\lceil n \log n+c_{\varepsilon} n\right\rceil .
$$

Again we get a quick lower bound using the diameter-based result from Section 5.2.3. Here $\Delta=n$ and $\pi_{\min }=1 / 2^{n}$ so that Claim 5.2.25 gives

$$
\mathrm{t}_{\mathrm{mix}}(\varepsilon) \geq \frac{n^{2}}{12 \log n+(4 \log 2) n}=\Omega(n)
$$

for $n$ large enough. So the upper bound we derived above is off at most by a logarithmic factor in $n$. In fact:

## Claim 4.3.7.

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{1}{2} n \log n-O(n) .
$$

Proof. For simplicity, we assume that $n$ is odd. Let $W_{t}$ be the number of 1 s , or Hamming weight, at time $t$. Let $A$ be the event that the Hamming weight is $\leq n / 2$. To bound the mixing time, we use the fact that for any $z_{0}$

$$
\begin{equation*}
d(t) \geq\left\|P^{t}\left(z_{0}, \cdot\right)-\pi\right\|_{\mathrm{TV}} \geq\left|P^{t}\left(z_{0}, A\right)-\pi(A)\right| . \tag{4.3.2}
\end{equation*}
$$

Under the stationary distribution, the Hamming weight is equal in distribution to a $\operatorname{Bin}(n, 1 / 2)$. In particular the probability that a majority of coordinates are 0 is $1 / 2$. That is, $\pi(A)=1 / 2$.

On the other hand, let $\left(Z_{t}\right)$ start at $z_{0}$, the all-1 vector. By the definition of $A$,

$$
\begin{equation*}
\left|P^{t}\left(z_{0}, A\right)-\pi(A)\right|=\left|\mathbb{P}\left[W_{t} \leq n / 2\right]-1 / 2\right| . \tag{4.3.3}
\end{equation*}
$$

We use Chebyshev's inequality (Theorem 2.1.2) to bound the probability on the right-hand side. So we need to compute the expectation and variance of $W_{t}$.

Let $U_{t}$ be the number of (distinct) updated coordinates up to time $t$ in the Glauber dynamics representation of the chain discussed above. Observe that, conditioned on $U_{t}$, the Hamming weight $W_{t}$ is equal in distribution to $\operatorname{Bin}\left(U_{t}, 1 / 2\right)+$ $\left(n-U_{t}\right)$ as the updated coordinates are uniform and the other ones are 1 . Thus we
have

$$
\begin{align*}
\mathbb{E}\left[W_{t}\right] & =\mathbb{E}\left[\mathbb{E}\left[W_{t} \mid U_{t}\right]\right] \\
& =\mathbb{E}\left[\frac{1}{2} U_{t}+\left(n-U_{t}\right)\right] \\
& =\mathbb{E}\left[n-\frac{1}{2} U_{t}\right] \\
& =n-\frac{1}{2} n\left[1-\left(1-\frac{1}{n}\right)^{t}\right] \\
& =\frac{n}{2}\left[1+\left(1-\frac{1}{n}\right)^{t}\right], \tag{4.3.4}
\end{align*}
$$

where on the fourth line we used the fact that $\mathbb{E}\left[U_{t}\right]=n\left[1-\left(1-\frac{1}{n}\right)^{t}\right]$ by summing over the coordinates and using linearity of expectation.

As to the variance, using again the observation above about the distribution of $W_{t}$ given $U_{t}$,

$$
\begin{align*}
\operatorname{Var}\left[W_{t}\right] & =\mathbb{E}\left[\operatorname{Var}\left[W_{t} \mid U_{t}\right]\right]+\operatorname{Var}\left[\mathbb{E}\left[W_{t} \mid U_{t}\right]\right] \\
& =\frac{1}{4} \mathbb{E}\left[U_{t}\right]+\frac{1}{4} \operatorname{Var}\left[U_{t}\right] . \tag{4.3.5}
\end{align*}
$$

It remains to compute $\operatorname{Var}\left[U_{t}\right]$. Let $I_{t}^{(i)}$ be 1 if coordinate $i$ has not been updated up to time $t$ and 0 otherwise. Note that for $i \neq j$

$$
\begin{aligned}
\operatorname{Cov}\left[I_{t}^{(i)}, I_{t}^{(j)}\right] & =\mathbb{E}\left[I_{t}^{(i)} I_{t}^{(j)}\right]-\mathbb{E}\left[I_{t}^{(i)}\right] \mathbb{E}\left[I_{t}^{(j)}\right] \\
& =\left(1-\frac{2}{n}\right)^{t}-\left(1-\frac{1}{n}\right)^{2 t} \\
& =\left(1-\frac{2}{n}\right)^{t}-\left(1-\frac{2}{n}+\frac{1}{n^{2}}\right)^{t} \\
& \leq 0,
\end{aligned}
$$

that is, $I_{t}^{(i)}$ and $I_{t}^{(j)}$ are negatively correlated, while

$$
\operatorname{Var}\left[I_{t}^{(i)}\right]=\mathbb{E}\left[\left(I_{t}^{(i)}\right)^{2}\right]-\left(\mathbb{E}\left[I_{t}^{(i)}\right]\right)^{2} \leq \mathbb{E}\left[I_{t}^{(i)}\right]=\left(1-\frac{1}{n}\right)^{t}
$$

Then, writing $n-U_{t}$ as the sum of these indicators, we have

$$
\begin{aligned}
\operatorname{Var}\left[U_{t}\right] & =\operatorname{Var}\left[n-U_{t}\right] \\
& =\sum_{i=1}^{n} \operatorname{Var}\left[I_{t}^{(i)}\right]+2 \sum_{i<j} \operatorname{Cov}\left[I_{t}^{(i)}, I_{t}^{(j)}\right] \\
& \leq n\left(1-\frac{1}{n}\right)^{t} .
\end{aligned}
$$

Plugging this back into (4.3.5), we get

$$
\operatorname{Var}\left[W_{t}\right] \leq \frac{n}{4}\left[1-\left(1-\frac{1}{n}\right)^{t}\right]+\frac{n}{4}\left(1-\frac{1}{n}\right)^{t}=\frac{n}{4}
$$

For $t_{\alpha}=\frac{1}{2} n \log n-n \log \alpha$ with $\alpha>0$, by (4.3.4),

$$
\mathbb{E}\left[W_{t_{\alpha}}\right]=\frac{n}{2}+\frac{n}{2} e^{t_{\alpha}\left(-1 / n+\Theta\left(1 / n^{2}\right)\right)}=\frac{n}{2}+\frac{\alpha}{2} \sqrt{n}+o(1),
$$

where we used that by a Taylor expansion, for $|z| \leq 1 / 2, \log (1-z)=-z+$ $\Theta\left(z^{2}\right)$. Fix $0<\varepsilon<1 / 2$. By Chebyshev's inequality, for $t_{\alpha}=\frac{1}{2} n \log n-n \log \alpha$ and $n$ large enough,

$$
\mathbb{P}\left[W_{t_{\alpha}} \leq n / 2\right] \leq \mathbb{P}\left[\left|W_{t_{\alpha}}-\mathbb{E}\left[W_{t_{\alpha}}\right]\right| \geq(\alpha / 2) \sqrt{n}\right] \leq \frac{n / 4}{(\alpha / 2)^{2} n} \leq \frac{1}{2}-\varepsilon,
$$

for $\alpha$ large enough. By (4.3.2) and (4.3.3), that implies $d\left(t_{\alpha}\right) \geq \varepsilon$ and we are done.

The previous proof relies on a "distinguishing statistic." Recall from Lemma 4.1.19 that for any random variables $X, Y$ and mapping $h$ it holds that

$$
\left\|\mu_{h(X)}-\mu_{h(Y)}\right\|_{\mathrm{TV}} \leq\left\|\mu_{X}-\mu_{Y}\right\|_{\mathrm{TV}}
$$

where $\mu_{Z}$ is the law of $Z$. The mapping used in the proof of the claim is the Hamming weight. In essence, we gave a lower bound on the total variation distance between the laws of the Hamming weight at stationarity and under $P^{t}\left(z_{0}, \cdot\right)$. See Exercise 4.16 for a more general treatment of the distinguishing statistic approach.

Remark 4.3.8. The upper bound in Claim 4.3.6 is indeed off by a factor of 2. See [LPW06, Theorem 18.3] for an improved upper bound and a discussion of the so-called cutoff phenomenon. The latter refers to the fact that for all $0<\varepsilon<1 / 2$ it can be shown in this case that

$$
\lim _{n \rightarrow+\infty} \frac{\mathrm{t}_{\text {mix }}^{(n)}(\varepsilon)}{\mathrm{t}_{\text {mix }}^{(n)}(1-\varepsilon)}=1,
$$

where $\mathrm{t}_{\text {mix }}^{(n)}(\varepsilon)$ is the mixing time on the $n$-dimensional hypercube. In words, for large $n$, the total variation distance drops from 1 to 0 in a short time window. See Exercise 5.10 for a necessary condition for cutoff.

## $b$-ary tree

Let $\left(Z_{t}\right)_{t \in \mathbb{Z}_{+}}$be lazy simple random walk on the $\ell$-level rooted $b$-ary tree, $\widehat{\mathbb{T}}$, with $\ell \geq 2$. The root, 0 , is on level 0 and the leaves, $L$, are on level $\ell$. All vertices have degree $b+1$, except for the root which has degree $b$ and the leaves which have degree 1. By Example 1.1.29 (noting that laziness makes no difference), the stationary distribution is

$$
\pi(x):=\frac{\delta(x)}{2(n-1)}
$$

where $n$ is the number of vertices and $\delta(x)$ is the degree of $x$. We used that a tree on $n$ vertices has $n-1$ edges (Corollary 1.1.7). We construct a coupling ( $X_{t}, Y_{t}$ ) of this chain started at $(x, y)$. Assume without loss of generality that $x$ is no further from the root than $y$, which we denote by $x \preccurlyeq y$ (which, here, does not mean that $y$ is a descendant of $x$ ). The coupling has two stages:

- In the first stage, at each time, flip a fair coin. On heads, $Y_{t}$ stays put and $X_{t}$ moves one step chosen uniformly at random among its neighbors. Similarly, on tails, reverse the roles of $X_{t}$ and $Y_{t}$. Do this until $X_{t}$ and $Y_{t}$ are on the same level.
- In the second stage, that is, once the two chains are on the same level, at each time first let $X_{t}$ move as a lazy simple random walk on $\widehat{\mathbb{T}}_{b}^{\ell}$. Then let $Y_{t}$ move in the same direction as $X_{t}$, that is, if $X_{t}$ moves closer to the root, so does $Y_{t}$ and so on.

By construction, $X_{t} \preccurlyeq Y_{t}$ for all $t$. The key observation is the following. Let $\tau^{*}$ be the first time $\left(X_{t}\right)$ visits the root after visiting the leaves. By time $\tau^{*}$, the two chains have necessarily met: because $X_{t} \preccurlyeq Y_{t}$, when $X_{t}$ reaches the leaves, so does $Y_{t}$; after that time, the coupling is in the second stage so $X_{t}$ and $Y_{t}$ remain on the same level; in particular, when $X_{t}$ reaches the root (after visiting the leaves), so does $Y_{t}$. Hence $\tau_{\text {coal }} \leq \tau^{*}$. Intuitively, the mixing time is indeed dominated by the time it takes to reach the root from the worst starting point, a leaf. See Figure 4.7 and the corresponding lower bound argument.

To estimate $\mathbb{P}_{(x, y)}\left[\tau^{*}>t\right]$, we use Markov's inequality (Theorem 2.1.1), for which we need a bound on $\mathbb{E}_{(x, y)}\left[\tau^{*}\right]$. We note that $\mathbb{E}_{(x, y)}\left[\tau^{*}\right]$ is less than the mean time for the walk to go from the root to the leaves and back. Let $L_{t}$ be the level of $X_{t}$ and let $\mathcal{N}$ be the corresponding network (where the conductances are equal to the number of edges on each level of the tree). In terms of $L_{t}$, the quantity we seek to bound is the mean of $\tau_{0, \ell}$, the commute time of the chain $\left(L_{t}\right)$ between the states 0 and $\ell$. By the commute time identity (Theorem 3.3.34),

$$
\begin{equation*}
\mathbb{E}\left[\tau_{0, \ell}\right]=c_{\mathcal{N}} \mathscr{R}(0 \leftrightarrow \ell), \tag{4.3.6}
\end{equation*}
$$

where

$$
c_{\mathcal{N}}=2 \sum_{e=\{x, y\} \in \mathcal{N}} c(e)=4(n-1),
$$

where we simply counted the number of edges in $\widehat{\mathbb{T}}_{b}^{\ell}$ and the extra factor of 2 accounts for self-loops. Using network reduction techniques, we computed the effective resistance $\mathscr{R}(0 \leftrightarrow \ell)$ in Examples 3.3.21 and 3.3.22-without self-loops. Of course adding self-loops does not affect the effective resistance as we can use the same voltage and current. So, ignoring them, we get

$$
\begin{equation*}
\mathscr{R}(0 \leftrightarrow \ell)=\sum_{j=0}^{\ell-1} r(j, j+1)=\sum_{j=0}^{\ell-1} b^{-(j+1)}=\frac{1}{b} \cdot \frac{1-b^{-\ell}}{1-b^{-1}}, \tag{4.3.7}
\end{equation*}
$$

which implies

$$
\frac{1}{b} \leq \mathscr{R}(0 \leftrightarrow \ell) \leq \frac{1}{b-1} \leq 1 .
$$

Finally, applying Theorem 4.3.2 and Markov's inequality and using (4.3.6), we get

$$
\begin{aligned}
d(t) & \leq \max _{x, y \in V} \mathbb{P}_{(x, y)}\left[\tau^{*}>t\right] \\
& \leq \max _{x, y \in V} \frac{\mathbb{E}_{(x, y)}\left[\tau^{*}\right]}{t} \\
& \leq \frac{\mathbb{E}\left[\tau_{0, \ell}\right]}{t} \\
& \leq \frac{4 n}{t}
\end{aligned}
$$

or:

## Claim 4.3.9.

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \frac{4 n}{\varepsilon}
$$

This time the diameter-based bound is far off. We have $\Delta=2 \ell=\Theta(\log n)$ and $\pi_{\text {min }}=1 / 2(n-1)$ so that Claim 5.2.25 gives

$$
\mathrm{t}_{\mathrm{mix}}(\varepsilon) \geq \frac{(2 \ell)^{2}}{12 \log n+4 \log (2(n-1))}=\Omega(\log n)
$$

for $n$ large enough.
Here is a better lower bound. We take $b=2$ to simplify. Intuitively the mixing time is significantly greater than the squared diameter because the chain tends to be pushed away from the root. Consider the time it takes to go from the leaves


Figure 4.7: Setup for the lower bound on the mixing time on a $b$-ary tree. (Here $b=2$.)
on one side of the root to the leaves on the other, both of which have substantial weight under the stationary distribution. That typically takes time exponential in the diameter-that is, linear in $n$. Indeed one first has to reach the root, which by the gambler's ruin problem (Example 3.1.43), takes an exponential in $\ell$ number of "excursions" (see Claim 3.1.44 (ii)).

Formally let $x_{0}$ be a leaf of $\widehat{\mathbb{T}}_{b}^{\ell}$ and let $A$ be the set of vertices "on the other side of root (inclusively)," that is, vertices whose graph distance from $x_{0}$ is at least $\ell$. See Figure 4.7. Then $\pi(A) \geq 1 / 2$ by symmetry. We use the fact that

$$
\left\|P^{t}\left(x_{0}, \cdot\right)-\pi\right\|_{\mathrm{TV}} \geq\left|P^{t}\left(x_{0}, A\right)-\pi(A)\right|,
$$

to bound the mixing time from below. We claim that, started at $x_{0}$, the walk takes time linear in $n$ to reach $A$ with nontrivial probability.

Consider again the level $L_{t}$ of $X_{t}$. Using definition of the effective resistance (Definition 3.3.19) as well as the expression for it in (4.3.7), we have

$$
\mathbb{P}_{\ell}\left[\tau_{0}<\tau_{\ell}^{+}\right]=\frac{1}{c(\ell) \mathscr{R}(0 \leftrightarrow \ell)}=\frac{1}{b^{\ell}} \cdot \frac{b-1}{1-b^{-\ell}}=\frac{b-1}{b^{\ell}-1}=O\left(\frac{1}{n}\right) .
$$

Hence, started from the leaves, the number of excursions back to the leaves needed to reach the root for the first time is geometric with success probability $O\left(n^{-1}\right)$. Each such excursion takes time at least 2 (which corresponds to going right back
to the leaves after the first step). So $P^{t}\left(x_{0}, A\right)$ is bounded above by the probability that at least one such excursion was successful among the first $t / 2$ attempts. That is,

$$
P^{t}\left(x_{0}, A\right) \leq 1-\left(1-O\left(n^{-1}\right)\right)^{t / 2}<\frac{1}{2}-\varepsilon,
$$

for all $t \leq \alpha_{\varepsilon} n$ with $\alpha_{\varepsilon}>0$ small enough and

$$
\left\|P^{\alpha_{\varepsilon} n}\left(x_{0}, \cdot\right)-\pi\right\|_{\mathrm{TV}} \geq\left|P^{\alpha_{\varepsilon} n}\left(x_{0}, A\right)-\pi(A)\right|>\varepsilon .
$$

We have proved that $\mathrm{t}_{\text {mix }}(\varepsilon) \geq \alpha_{\varepsilon} n$.

### 4.3.3 Path coupling

Path coupling is a method for constructing Markovian couplings from "simpler" couplings. The building blocks are one-step couplings starting from pairs of initial states that are close in some "dissimilarity graph."

Let $\left(X_{t}\right)$ be an irreducible Markov chain on a finite state space $V$ with transition matrix $P$ and stationary distribution $\pi$. Assume that we have a dissimilarity graph $H_{0}=\left(V_{0}, E_{0}\right)$ on $V_{0}:=V$ with edge weights $w_{0}: E_{0} \rightarrow \mathbb{R}_{+}$. This graph need not have the same edges as the transition graph of $\left(X_{t}\right)$. We extend $w_{0}$ to the path metric

$$
w_{0}(x, y):=\inf \left\{\sum_{i=0}^{m-1} w_{0}\left(x_{i}, x_{i+1}\right): x=x_{0}, x_{1}, \ldots, x_{m}=y \text { is a path in } H_{0}\right\}
$$

where the infimum is over all paths connecting $x$ and $y$ in $H_{0}$. We call a path achieving the infimum a minimum-weight path. It is straightforward to check that $w_{0}$ satisfies the triangle inequality. Let

$$
\Delta_{0}:=\max _{x, y} w_{0}(x, y)
$$

be the weighted diameter of $H_{0}$.
Theorem 4.3.10 (Path coupling method). Assume that

$$
w_{0}(u, v) \geq 1,
$$

for all $\{u, v\} \in E_{0}$. Assume further that there exists $\kappa \in(0,1)$ such that:

- (Local couplings) For all $x, y$ with $\{x, y\} \in E_{0}$, there is a coupling $\left(X^{*}, Y^{*}\right)$ of $P(x, \cdot)$ and $P(y, \cdot)$ satisfying the following contraction property

$$
\begin{equation*}
\mathbb{E}\left[w_{0}\left(X^{*}, Y^{*}\right)\right] \leq \kappa w_{0}(x, y) \tag{4.3.8}
\end{equation*}
$$

Then

$$
d(t) \leq \Delta_{0} \kappa^{t}
$$

or

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq\left\lceil\frac{\log \Delta_{0}+\log \varepsilon^{-1}}{\log \kappa^{-1}}\right\rceil .
$$

Proof. The crux of the proof is to extend (4.3.8) to arbitrary pairs of vertices.
Claim 4.3.11 (Global coupling). For all $x, y \in V$, there is a coupling $\left(X^{*}, Y^{*}\right)$ of $P(x, \cdot)$ and $P(y, \cdot)$ such that (4.3.8) holds.

Iterating the coupling in this last claim immediately implies the existence of a coalescing Markovian coupling $\left(X_{t}, Y_{t}\right)$ of $P$ such that

$$
\begin{aligned}
\mathbb{E}_{(x, y)}\left[w_{0}\left(X_{t}, Y_{t}\right)\right] & =\mathbb{E}_{(x, y)}\left[\mathbb{E}\left[w_{0}\left(X_{t}, Y_{t}\right) \mid X_{t-1}, Y_{t-1}\right]\right] \\
& \leq \mathbb{E}_{(x, y)}\left[\kappa w_{0}\left(X_{t-1}, Y_{t-1}\right)\right] \\
& \leq \cdots \\
& \leq \kappa^{t} \mathbb{E}_{(x, y)}\left[w_{0}\left(X_{0}, Y_{0}\right)\right] \\
& =\kappa^{t} w_{0}(x, y) \\
& \leq \kappa^{t} \Delta_{0} .
\end{aligned}
$$

By assumption, $\mathbf{1}_{\{x \neq y\}} \leq w_{0}(x, y)$ so that by the coupling inequality and Lemma 4.3.1, we have

$$
d(t) \leq \bar{d}(t) \leq \max _{x, y} \mathbb{P}_{(x, y)}\left[X_{t} \neq Y_{t}\right] \leq \max _{x, y} \mathbb{E}_{(x, y)}\left[w_{0}\left(X_{t}, Y_{t}\right)\right] \leq \kappa^{t} \Delta_{0},
$$

which implies the theorem.
Remark 4.3.12. In essence, $w_{0}$ satisfies a form of Lyapounov condition (i.e., (3.3.15)) with a "geometric drift." See, e.g., [MT09, Chapter 15].

It remains to prove Claim 4.3.11.
Proof of Claim 4.3.11. Fix $x^{\prime}, y^{\prime} \in V$ such that $\left\{x^{\prime}, y^{\prime}\right\}$ is not an edge in the dissimilarity graph $H_{0}$. The idea is to combine the local couplings on a minimumweight path between $x^{\prime}$ and $y^{\prime}$ in $H_{0}$. Let $x^{\prime}=x_{0} \sim \cdots \sim x_{m}=y^{\prime}$ be such a path. For all $i=0, \ldots, m-1$, let $\left(Z_{i, 0}^{*}, Z_{i, 1}^{*}\right)$ be a coupling of $P\left(x_{i}, \cdot\right)$ and $P\left(x_{i+1}, \cdot\right)$ satisfying the contraction property (4.3.8).

Then we proceed as follows. Set $Z^{(0)}:=Z_{0,0}^{*}$ and $Z^{(1)}:=Z_{0,1}^{*}$. Then iteratively pick $Z^{(i+1)}$ according to the law $\mathbb{P}\left[Z_{i, 1}^{*} \in \cdot \mid Z_{i, 0}^{*}=Z^{(i)}\right]$. By induction on $i,\left(X^{*}, Y^{*}\right):=\left(Z^{(0)}, Z^{(m)}\right)$ is then a coupling of $P\left(x^{\prime}, \cdot\right)$ and $P\left(y^{\prime}, \cdot\right)$. See Figure 4.8.


Figure 4.8: Coupling of $P\left(x^{\prime}, \cdot\right)$ and $P\left(y^{\prime}, \cdot\right)$ constructed from a sequence of local couplings $\left(Z_{0,0}^{*}, Z_{0,1}^{*}\right), \ldots,\left(Z_{0, m-1}^{*}, Z_{0, m-1}^{*}\right)$.

To be more formal, define the transition matrix

$$
R_{i}\left(z^{(i)}, z^{(i+1)}\right):=\mathbb{P}\left[Z_{i, 1}^{*}=z^{(i+1)} \mid Z_{i, 0}^{*}=z^{(i)}\right] .
$$

Observe that

$$
\begin{equation*}
\sum_{z^{(i+1)}} R_{i}\left(z^{(i)}, z^{(i+1)}\right)=1, \tag{4.3.9}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{z^{(i)}} P\left(x_{i}, z^{(i)}\right) R_{i}\left(z^{(i)}, z^{(i+1)}\right)=P\left(x_{i+1}, z^{(i+1)}\right), \tag{4.3.10}
\end{equation*}
$$

by construction of the coupling $\left(Z_{i, 0}^{*}, Z_{i, 1}^{*}\right)$ and the definition of $R_{i}$. The law of the full coupling

$$
\left(Z^{(0)}, \ldots, Z^{(m)}\right)
$$

is

$$
\begin{aligned}
& \mathbb{P}\left[\left(Z^{(0)}, \ldots, Z^{(m)}\right)=\left(z^{(0)}, \ldots, z^{(m)}\right)\right] \\
& \quad=P\left(x_{0}, z^{(0)}\right) R_{0}\left(z^{(0)}, z^{(1)}\right) \cdots R_{m-1}\left(z^{(m-1)}, z^{(m)}\right) .
\end{aligned}
$$

Using (4.3.9) and (4.3.10) inductively gives respectively

$$
\begin{gathered}
\mathbb{P}\left[X^{*}=z^{(0)}\right]=\mathbb{P}\left[Z^{(0)}=z^{(0)}\right]=P\left(x_{0}, z^{(0)}\right), \\
\mathbb{P}\left[Y^{*}=z^{(m)}\right]=\mathbb{P}\left[Z^{(m)}=z^{(m)}\right]=P\left(x_{m}, z^{(m)}\right),
\end{gathered}
$$

as required.
By the triangle inequality for $w_{0}$, the coupling $\left(X^{*}, Y^{*}\right)$ satisfies

$$
\begin{aligned}
\mathbb{E}\left[w_{0}\left(X^{*}, Y^{*}\right)\right] & =\mathbb{E}\left[w_{0}\left(Z^{(0)}, Z^{(m)}\right)\right] \\
& \leq \sum_{i=0}^{m-1} \mathbb{E}\left[w_{0}\left(Z^{(i)}, Z^{(i+1)}\right)\right] \\
& \leq \sum_{i=0}^{m-1} \kappa w_{0}\left(x_{i}, x_{i+1}\right) \\
& =\kappa w_{0}\left(x^{\prime}, y^{\prime}\right),
\end{aligned}
$$

where, on the third line, we used (4.3.8) for adjacent pairs and the last line follows from the fact that we chose a minimum-weight path.

That concludes the proof of the theorem.
We illustrate the path coupling method in the next subsection. See Exercise 4.17 for an optimal transport perspective on the path coupling method.

### 4.3.4 $\triangleright$ Ising model: Glauber dynamics at high temperature

Let $G=(V, E)$ be a finite, connected graph with maximal degree $\bar{\delta}$. Define $\mathcal{X}:=$ $\{-1,+1\}^{V}$. Recall from Example 1.2.5 that the (ferromagnetic) Ising model on $V$ with inverse temperature $\beta$ is the probability distribution over spin configurations $\sigma \in \mathcal{X}$ given by

$$
\mu_{\beta}(\sigma):=\frac{1}{\mathcal{Z}(\beta)} e^{-\beta \mathcal{H}(\sigma)},
$$

where

$$
\mathcal{H}(\sigma):=-\sum_{i \sim j} \sigma_{i} \sigma_{j}
$$

is the Hamiltonian and

$$
\mathcal{Z}(\beta):=\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}(\sigma)},
$$

is the partition function. In this context, recall that vertices are often referred to as sites. The single-site Glauber dynamics (Definition 1.2.8) of the Ising model is the Markov chain on $\mathcal{X}$ which, at each time, selects a site $i \in V$ uniformly at random and updates the spin $\sigma_{i}$ according to $\mu_{\beta}(\sigma)$ conditioned on agreeing with $\sigma$ at all sites in $V \backslash\{i\}$. Specifically, for $\gamma \in\{-1,+1\}, i \in V$, and $\sigma \in \mathcal{X}$, let $\sigma^{i, \gamma}$ be the configuration $\sigma$ with the state at $i$ being set to $\gamma$. Then, letting $n=|V|$, the transition matrix of the Glauber dynamics is

$$
\begin{align*}
Q_{\beta}\left(\sigma, \sigma^{i, \gamma}\right) & :=\frac{1}{n} \cdot \frac{e^{\gamma \beta S_{i}(\sigma)}}{e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}} \\
& =\frac{1}{n}\left\{\frac{1}{2}+\frac{1}{2} \tanh \left(\gamma \beta S_{i}(\sigma)\right)\right\}, \tag{4.3.11}
\end{align*}
$$

where

$$
S_{i}(\sigma):=\sum_{j \sim i} \sigma_{j} .
$$

All other transitions have probability 0 . Recall that this chain is irreducible and reversible with respect to $\mu_{\beta}$. In particular $\mu_{\beta}$ is the stationary distribution of $Q_{\beta}$.

In this section we give an upper bound on the mixing time, $\mathrm{t}_{\text {mix }}(\varepsilon)$, of $Q_{\beta}$ using path coupling. We say that the Glauber dynamics is fast mixing if $\mathrm{t}_{\text {mix }}(\varepsilon)=$ $O(n \log n)$. We first make a simple observation:

Claim 4.3.13 (Glauber dynamics: lower bound on mixing).

$$
\mathrm{t}_{\text {mix }}(\varepsilon)=\Omega(n), \quad \forall \beta>0
$$

Proof. Similarly to what we did in Section 4.3.2 in the context of random walk on the hypercube (but for a lower bound this time), we use a coupon collecting argument (see Example 2.1.4). Let $\bar{\sigma}$ be the all-( -1 ) configuration and let $A$ be the set of configurations where at least half of the sites are +1 . Then, by symmetry, $\mu_{\beta}(A)=\mu_{\beta}\left(A^{c}\right)=1 / 2$ where we assumed for simplicity that $n$ is odd. By definition of the total variation distance,

$$
\begin{align*}
d(t) & \geq\left\|Q_{\beta}^{t}(\bar{\sigma}, \cdot)-\mu_{\beta}(\cdot)\right\|_{\mathrm{TV}} \\
& \geq\left|Q_{\beta}^{t}(\bar{\sigma}, A)-\mu_{\beta}(A)\right| \\
& =\left|Q_{\beta}^{t}(\bar{\sigma}, A)-1 / 2\right| . \tag{4.3.12}
\end{align*}
$$

So it remains to show that by time $c n$, for $c>0$ small, the chain is unlikely to have reached $A$. That happens if, say, fewer than a third of the sites have been updated. Using the notation of Example 2.1.4, we are seeking a bound on $T_{n, n / 3}$, that is, the time to collect $n / 3$ coupons out of $n$.

We can write this random variable as a sum of $n / 3$ independent geometric variables $T_{n, n / 3}=\sum_{i=1}^{n / 3} \tau_{n, i}$, where $\mathbb{E}\left[\tau_{n, i}\right]=\left(1-\frac{i-1}{n}\right)^{-1}$ and $\operatorname{Var}\left[\tau_{n, i}\right] \leq$ $\left(1-\frac{i-1}{n}\right)^{-2}$. Hence, approximating the Riemann sums below by integrals,

$$
\begin{equation*}
\mathbb{E}\left[T_{n, n / 3}\right]=\sum_{i=1}^{n / 3}\left(1-\frac{i-1}{n}\right)^{-1}=n \sum_{j=2 n / 3+1}^{n} j^{-1}=\Theta(n), \tag{4.3.13}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Var}\left[T_{n, n / 3}\right] \leq \sum_{i=1}^{n / 3}\left(1-\frac{i-1}{n}\right)^{-2}=n^{2} \sum_{j=2 n / 3+1}^{n} j^{-2}=\Theta(n) . \tag{4.3.14}
\end{equation*}
$$

So by Chebyshev's inequality (Theorem 2.1.2)

$$
\mathbb{P}\left[\left|T_{n, n / 3}-\mathbb{E}\left[T_{n, n / 3}\right]\right| \geq \varepsilon n\right] \leq \frac{\operatorname{Var}\left[T_{n, n / 3}\right]}{(\varepsilon n)^{2}} \rightarrow 0
$$

by (4.3.14). In view of (4.3.13), taking $\varepsilon>0$ small enough and $n$ large enough, we have shown that for $t \leq c_{\varepsilon} n$ for some $c_{\varepsilon}>0$

$$
Q_{\beta}^{t}(\bar{\sigma}, A) \leq 1 / 3,
$$

which proves the claim by (4.3.12) and the definition of the mixing time (Definition 1.1.35).

Remark 4.3.14. In fact, Ding and Peres proved that $\mathrm{t}_{\text {mix }}(\varepsilon)=\Omega(n \log n)$ for any graph on $n$ vertices [DP11]. In Claim 4.3.7, we treated the special case of the empty graph, which is equivalent to lazy random walk on the hypercube. See also Section 5.3.4 for a much stronger lower bound at low temperature for certain graphs with good "expansion properties."

In our main result of this section, we show that the Glauber dynamics of the Ising model is fast mixing when the inverse temperature $\beta$ is small enough as a function of the maximum degree.

Claim 4.3.15 (Glauber dynamics: fast mixing at high temperature).

$$
\beta<\bar{\delta}^{-1} \Longrightarrow \mathrm{t}_{\text {mix }}(\varepsilon)=O(n \log n) .
$$

Proof. We use path coupling. Let $H_{0}=\left(V_{0}, E_{0}\right)$ where $V_{0}:=\mathcal{X}$ and $\{\sigma, \omega\} \in E_{0}$ if $\frac{1}{2}\|\sigma-\omega\|_{1}=1$ (i.e., they differ in exactly one coordinate) with unit weight on all edges. To avoid confusion, we reserve the notation $\sim$ for adjacency in $G$.

Let $\{\sigma, \omega\} \in E_{0}$ differ at coordinate $i$. We construct a coupling $\left(X^{*}, Y^{*}\right)$ of $Q_{\beta}(\sigma, \cdot)$ and $Q_{\beta}(\omega, \cdot)$. We first pick the same coordinate $i_{*}$ to update. If $i_{*}$ is such that all its neighbors in $G$ have the same state in $\sigma$ and $\omega$, that is, if $\sigma_{j}=\omega_{j}$ for all $j \sim i_{*}$, we update $X^{*}$ from $\sigma$ according to the Glauber rule and set $Y^{*}:=X^{*}$. Note that this includes the case $i_{*}=i$. Otherwise, that is, if $i_{*} \sim i$, we proceed as follows. From the state $\sigma$, the probability of updating site $i_{*}$ to state $\gamma \in\{-1,+1\}$ is given by the expression in brackets in (4.3.11), and similarly for $\omega$. Unlike the previous case, we cannot guarantee that the update is identical in both chains. In order to minimize the chance of increasing the distance between the two chains, we use a monotone coupling, which recall from Example 4.1.17 is maximal in the two-state case. Specifically, we pick a uniform random variable $U$ in $[-1,1]$ and set

$$
X_{i_{*}}^{*}:= \begin{cases}+1 & \text { if } U \leq \tanh \left(\beta S_{i_{*}}(\sigma)\right) \\ -1 & \text { otherwise }\end{cases}
$$

and

$$
Y_{i_{*}}^{*}:= \begin{cases}+1 & \text { if } U \leq \tanh \left(\beta S_{i_{*}}(\omega)\right) \\ -1 & \text { otherwise }\end{cases}
$$

We set $X_{j}^{*}:=\sigma_{j}$ and $Y_{j}^{*}:=\omega_{j}$ for all $j \neq i^{*}$. The expected distance between $X^{*}$ and $Y^{*}$ is then

$$
\begin{align*}
& \mathbb{E}\left[w_{0}\left(X^{*}, Y^{*}\right)\right] \\
& =1-\underbrace{\frac{1}{n}}_{\text {(a) }}+\underbrace{\frac{1}{n} \sum_{j: j \sim i} \frac{1}{2}\left|\tanh \left(\beta S_{j}(\sigma)\right)-\tanh \left(\beta S_{j}(\omega)\right)\right|}_{\text {(b) }}, \tag{4.3.15}
\end{align*}
$$

where: (a) corresponds to $i_{*}=i$ in which case $w_{0}\left(X^{*}, Y^{*}\right)=0$; (b) corresponds to $i_{*} \sim i$ in which case $w_{0}\left(X^{*}, Y^{*}\right)=2$ with probability

$$
\frac{1}{2}\left|\tanh \left(\beta S_{i_{*}}(\sigma)\right)-\tanh \left(\beta S_{i_{*}}(\omega)\right)\right|
$$

by our coupling; and otherwise $w_{0}\left(X^{*}, Y^{*}\right)=w_{0}(\sigma, \omega)=1$. To bound (b), we note that for any $j \sim i$

$$
\begin{equation*}
\left|\tanh \left(\beta S_{j}(\sigma)\right)-\tanh \left(\beta S_{j}(\omega)\right)\right|=\tanh (\beta(s+2))-\tanh (\beta s), \tag{4.3.16}
\end{equation*}
$$

where

$$
s:=S_{j}(\sigma) \wedge S_{j}(\omega)
$$

The derivative of $\tanh$ is maximized at 0 where it is equal to 1 . So the right-hand side of (4.3.16) is $\leq \beta(s+2)-\beta s=2 \beta$. Plugging this back into (4.3.15) and using $1-x \leq e^{-x}$ for all $x$ (see Exercise 1.16), we get

$$
\mathbb{E}\left[w_{0}\left(X^{*}, Y^{*}\right)\right] \leq 1-\frac{1-\bar{\delta} \beta}{n} \leq \exp \left(-\frac{1-\bar{\delta} \beta}{n}\right)=\kappa w_{0}(\sigma, \omega),
$$

where

$$
\kappa:=\exp \left(-\frac{1-\bar{\delta} \beta}{n}\right)<1,
$$

by our assumption on $\beta$. The diameter of $H_{0}$ is $\Delta_{0}=n$. By Theorem 4.3.10,

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq\left\lceil\frac{\log \Delta_{0}+\log \varepsilon^{-1}}{\log \kappa^{-1}}\right\rceil=\left\lceil\frac{n\left(\log n+\log \varepsilon^{-1}\right)}{1-\bar{\delta} \beta}\right\rceil
$$

which implies the claim.
Remark 4.3.16. A slighlty more careful analysis shows that the condition $\bar{\delta} \tanh (\beta)<1$ is enough for the claim to hold. See [LPW06, Theorem 15.1].

### 4.4 Chen-Stein method

The Chen-Stein method serves to establish Poisson approximation results with quantitative bounds in certain settings with dependent variables that are common, for instance, in random graphs and string statistics.

Setting The basic setup is a sum of Bernoulli (i.e., $\{0,1\}$-valued) random variables $\left\{X_{i}\right\}_{i=1}^{n}$

$$
\begin{equation*}
W=\sum_{i=1}^{n} X_{i}, \tag{4.4.1}
\end{equation*}
$$

where the $X_{i} \mathrm{~s}$ are not assumed independent or identically distributed. Define

$$
\begin{equation*}
p_{i}=\mathbb{P}\left[X_{i}=1\right], \tag{4.4.2}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{E}[W]=\lambda:=\sum_{i=1}^{n} p_{i} . \tag{4.4.3}
\end{equation*}
$$

Letting $\mu$ denote the law of $W$ and $\pi$ be the Poisson distribution with mean $\lambda$, our goal is to bound $\|\mu-\pi\|_{\mathrm{TV}}$.

We first state the main bounds and give some examples of its use. We then motivate and prove the result, and return to further applications. Throughout the next two subsections, we use the notation in (4.4.1), (4.4.2) and (4.4.3).

### 4.4.1 Main bounds and examples

We begin with an elementary observation.
Theorem 4.4.1 (Stein equation for the Poisson distribution). Let $\lambda>0$. A nonnegative integer-valued random variable $Z$ is $\operatorname{Poi}(\lambda)$ if and only iffor all $g$ bounded

$$
\begin{equation*}
\mathbb{E}[\lambda g(Z+1)-Z g(Z)]=0 . \tag{4.4.4}
\end{equation*}
$$

The "only if" follows a direct calculation. The "if" follows from taking $g(z):=$ $\mathbf{1}_{\{z=k\}}$ for all $k \geq 1$ and deriving a recursion. Exercise 4.18 asks for the details. One might expect that if the left-hand side of (4.4.4) is "small for many $g \mathrm{~s}$," then $Z$ is approximately Poisson.

The following key result in some sense helps to formalize this intuition. We prove it by constructing a Markov chain that "interpolates" between $\mu$ and $\pi$, where (4.4.4) will arise naturally (see Section 4.4.2).

Theorem 4.4.2 (Chen-Stein method). Let $W \sim \mu$ and $\pi \sim \operatorname{Poi}(\lambda)$. Then there exists a function $h:\{0,1, \ldots, n+1\} \rightarrow \mathbb{R}$ such

$$
\begin{equation*}
\|\mu-\pi\|_{\mathrm{TV}}=\mathbb{E}[-\lambda h(W+1)+W h(W)] . \tag{4.4.5}
\end{equation*}
$$

Moreover $h$ satisfies the following Lipschitz condition: for all $y, y^{\prime} \in\{0,1, \ldots, n+$ $1\}$,

$$
\begin{equation*}
\left|h\left(y^{\prime}\right)-h(y)\right| \leq\left(1 \wedge \lambda^{-1}\right)\left|y^{\prime}-y\right| . \tag{4.4.6}
\end{equation*}
$$

By bounding the right-hand side of (4.4.5) for any function satisfying (4.4.6), we get a Poisson approximation result for $\mu$.

One way to do this is to construct a certain type of coupling. We begin with a definition, which will be justified in the corollary below. We write $X \sim Y \mid A$ to mean that $X$ is distributed as $Y$ conditioned on the event $A$.
Definition 4.4.3 (Stein coupling). A Stein coupling is a pair $\left(U_{i}, V_{i}\right)$, for each $i=1, \ldots, n$, such that

$$
U_{i} \sim W, \quad V_{i} \sim W-1 \mid X_{i}=1
$$

Each pair $\left(U_{i}, V_{i}\right)$ is defined on a joint probability space, but different pairs do not need to.

How such a coupling is constructed will become clearer in the examples below.
Corollary 4.4.4. Let $\left(U_{i}, V_{i}\right), i=1, \ldots, n$, be a Stein coupling. Then

$$
\begin{equation*}
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|U_{i}-V_{i}\right| \tag{4.4.7}
\end{equation*}
$$

Proof. By (4.4.5), using the facts that $\lambda=\sum_{i=1}^{n} p_{i}$ and $W=\sum_{i=1}^{n} X_{i}$, we get

$$
\begin{aligned}
& \|\mu-\pi\|_{\mathrm{TV}} \\
& =\mathbb{E}[-\lambda h(W+1)+W h(W)] \\
& =\mathbb{E}\left[-\left(\sum_{i=1}^{n} p_{i}\right) h(W+1)+\left(\sum_{i=1}^{n} X_{i}\right) h(W)\right] \\
& =\sum_{i=1}^{n}\left(-p_{i} \mathbb{E}[h(W+1)]+\mathbb{E}\left[X_{i} h(W)\right]\right) \\
& =\sum_{i=1}^{n}\left(-p_{i} \mathbb{E}[h(W+1)]+\mathbb{E}\left[h(W) \mid X_{i}=1\right] \mathbb{P}\left[X_{i}=1\right]\right) \\
& =\sum_{i=1}^{n} p_{i}\left(-\mathbb{E}[h(W+1)]+\mathbb{E}\left[h(W) \mid X_{i}=1\right]\right) .
\end{aligned}
$$

Let $\left(U_{i}, V_{i}\right), i=1, \ldots, n$, be a Stein coupling (Definition 4.4.3). Then, we can rewrite this last expression is

$$
\begin{aligned}
& =\sum_{i=1}^{n} p_{i}\left(-\mathbb{E}\left[h\left(U_{i}+1\right)\right]+\mathbb{E}\left[h\left(V_{i}+1\right)\right]\right) \\
& \leq \sum_{i=1}^{n} p_{i} \mathbb{E}\left[\left|h\left(U_{i}+1\right)-h\left(V_{i}+1\right)\right|\right]
\end{aligned}
$$

By (4.4.6), we finally get

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|U_{i}-V_{i}\right|
$$

which concludes the proof.
As a first example, we derive a Poisson approximation result in the independent case. Compare to Theorem 4.1.18.

Example 4.4.5 (Independent $X_{i} \mathrm{~s}$ ). Assume the $X_{i} \mathrm{~s}$ are independent. We prove the following:

## Claim 4.4.6.

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i}^{2}
$$

We use the following Stein coupling. For each $i=1, \ldots, n$, we let

$$
U_{i}=W
$$

and

$$
V_{i}=\sum_{j: j \neq i} X_{j} .
$$

By independence,

$$
V_{i}=W-X_{i} \sim W-1 \mid X_{i}=1,
$$

as desired. Plugging into (4.4.7), we obtain the bound

$$
\begin{aligned}
\|\mu-\pi\|_{\mathrm{TV}} & \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|U_{i}-V_{i}\right| \\
& \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|W-\sum_{j \neq i} X_{j}\right| \\
& \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|X_{i}\right| \\
& \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i}^{2} .
\end{aligned}
$$

Here is a less straightforward example.
Example 4.4.7 (Balls in boxes). Suppose we throw $k$ balls uniformly at random in $n$ boxes independently. Let

$$
X_{i}=\mathbf{1}\{\text { box } i \text { is empty }\},
$$

and let $W=\sum_{i=1}^{n} X_{i}$ be the number of empty boxes. Note that the $X_{i}$ s are not independent. In particular, we cannot use Theorem 4.1.18. Note that

$$
p_{i}=\left(1-\frac{1}{n}\right)^{k}
$$

for all $i$ and, hence,

$$
\lambda=n\left(1-\frac{1}{n}\right)^{k} .
$$

For each $i=1, \ldots, n$, we generate the coupling $\left(U_{i}, V_{i}\right)$ in the following way. We let $U_{i}=W$. If box $i$ is empty then $V_{i}=W-1$. Otherwise, we re-distribute all balls in box $i$ among the remaining boxes and let $V_{i}$ count the number of empty boxes $\neq i$. By construction, both conditions of the Stein coupling are satisfied. Moreover we have almost surely $V_{i} \leq U_{i}$ so that

$$
\sum_{i=1}^{n} p_{i} \mathbb{E}\left|U_{i}-V_{i}\right|=\sum_{i=1}^{n} p_{i} \mathbb{E}\left[U_{i}-V_{i}\right]=\lambda^{2}-\sum_{i=1}^{n} p_{i} \mathbb{E}\left[V_{i}\right] .
$$

By the fact that $V_{i} \sim W-1 \mid X_{i}=1$ and Bayes' rule,

$$
\begin{aligned}
\sum_{i=1}^{n} p_{i} \mathbb{E}\left[V_{i}\right] & =\sum_{i=1}^{n} \mathbb{P}\left[X_{i}=1\right] \sum_{k=1}^{n}(k-1) \mathbb{P}\left[V_{i}=k-1\right] \\
& =\sum_{i=1}^{n} \sum_{k=1}^{n}(k-1) \mathbb{P}\left[W=k \mid X_{i}=1\right] \mathbb{P}\left[X_{i}=1\right] \\
& =\sum_{i=1}^{n} \sum_{k=1}^{n}(k-1) \mathbb{P}\left[X_{i}=1 \mid W=k\right] \mathbb{P}[W=k] .
\end{aligned}
$$

Now we use the fact that $\mathbb{P}\left[X_{i}=1 \mid W=k\right]=\mathbb{E}\left[X_{i} \mid W=k\right]$ because $X_{i}$ is an
indicator variable. So the last line above is

$$
\begin{aligned}
& =\sum_{i=1}^{n} \sum_{k=1}^{n}(k-1) \mathbb{E}\left[X_{i} \mid W=k\right] \mathbb{P}[W=k] \\
& =\sum_{k=1}^{n}(k-1) \mathbb{E}\left[\sum_{i=1}^{n} X_{i} \mid W=k\right] \mathbb{P}[W=k] \\
& =\sum_{k=1}^{n}(k-1) k \mathbb{P}[W=k] \\
& =\mathbb{E}\left[W^{2}\right]-\mathbb{E}[W] .
\end{aligned}
$$

It remains to compute $\mathbb{E}\left[W^{2}\right]$. We have by symmetry

$$
\begin{aligned}
\mathbb{E}\left[W^{2}\right] & =n \mathbb{E}\left[X_{1}^{2}\right]+n(n-1) \mathbb{E}\left[X_{1} X_{2}\right] \\
& =\lambda+n(n-1)\left(1-\frac{2}{n}\right)^{k},
\end{aligned}
$$

so by Corollary 4.4.4

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right)\left\{n^{2}\left(1-\frac{1}{n}\right)^{2 k}-n(n-1)\left(1-\frac{2}{n}\right)^{k}\right\}
$$

When $k=n \log n+C n$ for instance, it can be checked that $\|\mu-\pi\|_{\mathrm{TV}}=$ $O(\log n / n)$.

This last example is generalized in Exercise 4.22.
In special settings, one can give useful general bounds by constructing an appropriate Stein coupling. We give an important example next. Recall that $[n]=$ $\{1, \ldots, n\}$.

Theorem 4.4.8 (Chen-Stein method: dissociated case). Suppose that for each $i$ there is a neighborhood $\mathcal{N}_{i} \subseteq[n] \backslash\{i\}$ such that

$$
X_{i} \text { is independent of }\left\{X_{j}: j \notin \mathcal{N}_{i} \cup\{i\}\right\} .
$$

Then

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n}\left\{p_{i}^{2}+\sum_{j \in \mathcal{N}_{i}}\left(p_{i} p_{j}+\mathbb{E}\left[X_{i} X_{j}\right]\right)\right\}
$$

Proof. We use the following Stein coupling. Let

$$
U_{i}=W .
$$

Then generate

$$
\left(Y_{j}^{(i)}\right)_{j \in \mathcal{N}_{i}} \sim\left(X_{j}\right)_{j \in \mathcal{N}_{i}} \mid\left\{X_{k}: k \notin \mathcal{N}_{i} \cup\{i\}\right\}, X_{i}=1,
$$

and set

$$
V_{i}=\sum_{k \notin \mathcal{N}_{i} \cup\{i\}} X_{k}+\sum_{j \in \mathcal{N}_{i}} Y_{j}^{(i)} .
$$

Because the law of $\left\{X_{k}: k \notin \mathcal{N}_{i} \cup\{i\}\right\}$ (and therefore of the first term in $V_{i}$ ) is independent of the event $\left\{X_{i}=1\right\}$, the above scheme satisfies the conditions of the Stein coupling.

Hence we can apply Corollary 4.4.4. The construction of $\left(U_{i}, V_{i}\right)$ guarantees that $U_{i}-V_{i}$ depends only on " $i$ and its neighborhood." Specifically, we get

$$
\begin{aligned}
\|\mu-\pi\|_{\mathrm{TV}} & \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|U_{i}-V_{i}\right| \\
& =\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|\sum_{j=1}^{n} X_{j}-\sum_{k \notin \mathcal{N}_{i} \cup\{i\}} X_{k}-\sum_{j \in \mathcal{N}_{i}} Y_{j}^{(i)}\right| \\
& =\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i} \mathbb{E}\left|X_{i}+\sum_{j \in \mathcal{N}_{i}}\left(X_{j}-Y_{j}^{(i)}\right)\right| \\
& \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i}\left(\mathbb{E}\left|X_{i}\right|+\sum_{j \in \mathcal{N}_{i}}\left(\mathbb{E}\left|X_{j}\right|+\mathbb{E}\left|Y_{j}^{(i)}\right|\right)\right),
\end{aligned}
$$

where we used the triangle inequality. Recalling that $p_{i}=\mathbb{P}\left[X_{i}=1\right]=\mathbb{E}\left[X_{i}\right]=$ $\mathbb{E}\left|X_{i}\right|$ and the definition of $Y_{j}^{(i)}$, the last expression above is

$$
\begin{aligned}
& =\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n} p_{i}\left(p_{i}+\sum_{j \in \mathcal{N}_{i}}\left[p_{j}+\mathbb{E}\left[\left|X_{j}\right| \mid X_{i}=1\right]\right)\right. \\
& =\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n}\left\{p_{i}^{2}+\sum_{j \in \mathcal{N}_{i}}\left(p_{i} p_{j}+p_{i} \mathbb{E}\left[X_{j} \mid X_{i}=1\right]\right)\right\} \\
& =\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n}\left\{p_{i}^{2}+\sum_{j \in \mathcal{N}_{i}}\left(p_{i} p_{j}+\mathbb{E}\left[X_{i} X_{j}\right]\right)\right\} .
\end{aligned}
$$

That concludes the proof.
Next we give an example of the previous theorem.
Example 4.4.9 (Longest head run). Let $0<q<1$ and let $Z_{1}, Z_{2}, \ldots$ be i.i.d. Bernoulli random variables with success probability $q=\mathbb{P}\left[Z_{i}=1\right]$. We are interested in the distribution of $R$, the length of the longest run of 1 s starting in the first $n$ tosses. For any positive integer $t$, let $X_{1}^{(t)}:=Z_{1} \cdots Z_{t}$ and

$$
X_{i}^{(t)}:=\left(1-Z_{i-1}\right) Z_{i} \cdots Z_{i+t-1}, \quad i \geq 2
$$

The event $\left\{X_{i}^{(t)}=1\right\}$ indicates that a head run of length at least $t$ starts at the $i$-th toss. Now define

$$
W^{(t)}:=\sum_{i=1}^{n} X_{i}^{(t)} .
$$

The key observation is that

$$
\begin{equation*}
\{R<t\}=\left\{W^{(t)}=0\right\} . \tag{4.4.8}
\end{equation*}
$$

Notice that, for fixed $t$, the $X_{i}^{(t)} \mathrm{s}$ are neither independent nor identically distributed. However, they exhibit a natural neighborhood structure as in Theorem 4.4.8. Indeed let

$$
\mathcal{N}_{i}^{(t)}:=\{\alpha \in[n]:|\alpha-i| \leq t\} \backslash\{i\} .
$$

Then, $X_{i}^{(t)}$ is independent of $\left\{X_{j}^{(t)}: j \notin \mathcal{N}_{i} \cup\{i\}\right\}$. For example,

$$
X_{i}^{(t)}=\left(1-Z_{i-1}\right) Z_{i} \cdots Z_{i+t-1},
$$

and

$$
X_{i+t+1}^{(t)}=\left(1-Z_{i+t}\right) Z_{i+t+1} \cdots Z_{i+2 t},
$$

do not depend on any common $Z_{j}$, while $X_{i}^{(t)}$ and

$$
X_{i+t}^{(t)}=\left(1-Z_{i+t-1}\right) Z_{i+t} \cdots Z_{i+2 t-1},
$$

both depend on $Z_{i+t-1}$.
We compute the quantities needed to apply Theorem 4.4.8. We have

$$
p_{1}^{(t)}:=\mathbb{E}\left[Z_{1} \cdots Z_{t}\right]=\prod_{j=1}^{t} \mathbb{E}\left[Z_{j}\right]=q^{t}
$$

and, for $i \geq 2$,

$$
\begin{aligned}
p_{i}^{(t)} & :=\mathbb{E}\left[\left(1-Z_{i-1}\right) Z_{i} \cdots Z_{i+t-1}\right] \\
& =\mathbb{E}\left[1-Z_{i-1}\right] \prod_{j=i}^{i+t-1} \mathbb{E}\left[Z_{j}\right] \\
& =(1-q) q^{t} \\
& \leq q^{t} .
\end{aligned}
$$

For $i \geq 1$ and $j \in \mathcal{N}_{i}^{(t)}$, observe that a head run of length at least $t$ cannot start simultaneously at $i$ and $j$. So $\mathbb{E}\left[X_{i}^{(t)} X_{j}^{(t)}\right]=0$ in that case. We also have

$$
\lambda^{(t)}:=\mathbb{E}\left[W^{(t)}\right]=q^{t}+(n-1)(1-q) q^{t} \in\left[n(1-q) q^{t}, n q^{t}\right],
$$

and

$$
\left|\mathcal{N}_{i}^{(t)}\right| \leq 2 t
$$

We are ready to apply Theorem 4.4.8. We get, letting $\mu^{(t)}$ denote the law of $W^{(t)}$ and $\pi^{(t)}$ be the Poisson distribution with mean $\lambda^{(t)}$,

$$
\begin{aligned}
\left\|\mu^{(t)}-\pi^{(t)}\right\|_{\mathrm{TV}} & \leq\left(1 \wedge\left(\lambda^{(t)}\right)^{-1}\right) \sum_{i=1}^{n}\left\{\left(p_{i}^{(t)}\right)^{2}+\sum_{j \in \mathcal{N}_{i}^{(t)}}\left(p_{i}^{(t)} p_{j}^{(t)}+\mathbb{E}\left[X_{i}^{(t)} X_{j}^{(t)}\right]\right)\right\} \\
& \leq\left(1 \wedge\left(n(1-q) q^{t}\right)^{-1}\right)\left[n q^{2 t}+2 t n q^{2 t}\right] \\
& \leq \frac{1}{(1-q) n}\left(1 \wedge\left(n q^{t}\right)^{-1}\right)[2 t+1]\left(n q^{t}\right)^{2}
\end{aligned}
$$

This bound is non-asymptotic-it holds for any $q, n, t$. One special regime of note is $t=\log _{1 / q} n+C$ with large $n$. In that case, we have $n q^{t} \rightarrow C^{\prime}$ as $n \rightarrow+\infty$ for some $0<C^{\prime}<+\infty$ and the total variation above is of the order to $O(\log n / n)$.

Going back to (4.4.8), we finally obtain when $t=\log _{1 / q} n+C$ that

$$
\left|\mathbb{P}[R<t]-e^{-\lambda^{(t)}}\right|=O\left(\frac{\log n}{n}\right)
$$

where recall that $R$ and $\lambda^{(t)}$ implicitly depend on $n$.

### 4.4.2 Some motivation and proof

The idea behind the Chen-Stein method is to interpolate between $\mu$ and $\pi$ in Theorem 4.4.2 by constructing a Markov chain with initial distribution $\mu$ and stationary distribution $\pi$. Here we use a discrete-time, finite Markov chain.

Proof of Theorem 4.4.2. We seek a bound on

$$
\begin{align*}
\|\mu-\pi\|_{\mathrm{TV}} & =\sup _{A \subseteq \mathbb{Z}_{+}}|\mu(A)-\pi(A)| \\
& =\mu\left(A^{*}\right)-\pi\left(A^{*}\right) \\
& =\sum_{z \in A^{*}}(\mu(z)-\pi(z)), \tag{4.4.9}
\end{align*}
$$

where $A^{*}=\left\{z \in \mathbb{Z}_{+}: \mu(z)>\pi(z)\right\}$, by Lemma 4.1.15. Since $W \leq n$ almost surely, $\mu(z)=0$ for all $z>n$ which implies that $A^{*} \subseteq\{0,1, \ldots, n\}$. In particular, it will suffice to bound $\mu(z)-\pi(z)$ for $0 \leq z \leq n$. We also assume $\lambda<n$ (the case $\lambda=n$ being uninteresting).

Constructing the Markov chain It will be convenient to truncate $\pi$ at $n$, that is, we define

$$
\bar{\pi}(z)= \begin{cases}\pi(z) & 0 \leq z \leq n \\ 1-\Pi(n) & z=n+1 \\ 0 & \text { otherwise }\end{cases}
$$

where $\Pi(z)=\sum_{w \leq z} \pi(w)$ is the cumulative distribution function of the Poisson distribution with mean $\lambda$. We construct a Markov chain with stationary distribution $\bar{\pi}$. We will also need the chain to be aperiodic and irreducible over $\{0,1, \ldots, n+$ $1\}$.

We choose the transition matrix $(P(x, y))_{0 \leq x, y \leq n+1}$ to be that of a birth-death chain reversible with respect to $\bar{\pi}$, that is, we require $P(x, y)=0$ unless $|x-y| \leq 1$ and

$$
\begin{equation*}
\frac{P(x, x+1)}{P(x+1, x)}=\frac{\bar{\pi}(x+1)}{\bar{\pi}(x)}, \quad \forall x \in[n] . \tag{4.4.10}
\end{equation*}
$$

For $x<n$,

$$
\frac{\bar{\pi}(x+1)}{\bar{\pi}(x)}=\frac{\pi(x+1)}{\pi(x)}=\frac{e^{-\lambda} \lambda^{x+1} /(x+1)!}{e^{-\lambda} \lambda^{x} / x!}=\frac{\lambda}{x+1} .
$$

In view of this, we want $P(x, x+1) \propto \lambda$ and $P(x, x-1) \propto x$. We choose the proportionality constant to ensure that all transition probabilities are in $[0,1]$. Specifically, for $x \neq y$, the nonzero transition probabilities take values

$$
P(x, y)= \begin{cases}\frac{1}{2 n} \lambda, & \text { if } 0 \leq x \leq n, y=x+1,  \tag{4.4.11}\\ \frac{1}{2 n} x, & \text { if } 1 \leq x \leq n, y=x-1, \\ \frac{1}{2 n} \lambda \frac{\pi(n)}{1-\Pi(n)}, & \text { if } x=n+1, y=n .\end{cases}
$$

The probability of staying put is: $1-\frac{1}{2 n} \lambda$ if $x=0,1-\frac{1}{2 n} x-\frac{1}{2 n} \lambda$ if $1 \leq x \leq n$, and $1-\frac{1}{2 n} \lambda \frac{\pi(n)}{1-\Pi(n)}$ if $x=n+1$. Those are all strictly positive when $\lambda<n$. Hence by construction $P$ is aperiodic and irreducible, and it satisfies the detailed balance conditions (4.4.10).

Recalling (3.3.6), the Laplacian is

$$
\begin{aligned}
\Delta f(x) & =\sum_{y} P(x, y)[f(y)-f(x)] \\
& =P(x, x+1)[f(x+1)-f(x)]-P(x, x-1)[f(x)-f(x-1)] \\
& =\lambda g(x+1)-x g(x),
\end{aligned}
$$

for $0 \leq x \leq n$, where we defined

$$
\begin{equation*}
g(x):=\frac{f(x)-f(x-1)}{2 n}, \quad x \in\{1, \ldots, n+1\} \tag{4.4.12}
\end{equation*}
$$

and $g(0)$ is arbitrary. At $x=n+1$,

$$
\Delta f(n+1)=-\lambda n \frac{\pi(n)}{1-\Pi(n)} g(n+1) .
$$

It is a standard fact (see Exercise 4.19) that the expectation of the Laplacian under the stationary distribution is 0 . Inverting the relationship (4.4.12), for any $g:\{0, \ldots, n+1\} \rightarrow \mathbb{R}$, there is a corresponding $f$ (unique up to an additive constant). So we have shown that if $Z \sim \bar{\pi}$ then

$$
\mathbb{E}\left[(\lambda g(Z+1)-Z g(Z)) \mathbf{1}_{\{Z \leq n\}}-\lambda n \frac{\pi(n)}{1-\Pi(n)} g(Z) \mathbf{1}_{\{Z=n+1\}}\right]=0,
$$

that is,

$$
\mathbb{E}\left[(\lambda g(Z+1)-Z g(Z)) \mathbf{1}_{\{Z \leq n\}}\right]=\lambda n \pi(n) g(n+1) .
$$

Notice that, if $g$ is extended to a bounded function on $\mathbb{Z}_{+}, \lambda$ is fixed and $Z \sim$ $\operatorname{Poi}(\lambda)$, then taking $n \rightarrow+\infty$ recovers Theorem 4.4.1 by dominated convergence (Proposition B.4.14).*

Markov chains calculations By the convergence theorem for Markov chains (Theorem 1.1.33),

$$
P^{t}(y, z) \rightarrow \bar{\pi}(z)
$$

[^6]for all $0 \leq y \leq n+1$ and $0 \leq z \leq n+1$ as $t \rightarrow+\infty$. Hence, letting $\delta_{z}(x)=$ $\mathbf{1}_{\{x=z\}}$, by telescoping
\[

$$
\begin{align*}
\delta_{z}(y)-\bar{\pi}(z) & =\lim _{t \rightarrow+\infty} \mathbb{E}_{y}\left[\delta_{z}\left(X_{0}\right)-\delta_{z}\left(X_{t}\right)\right] \\
& =\lim _{t \rightarrow+\infty} \sum_{s=0}^{t-1} \mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)-\delta_{z}\left(X_{s+1}\right)\right], \tag{4.4.13}
\end{align*}
$$
\]

where the subscript of $\mathbb{E}$ indicates the initial state. We will later take expectations over $\mu$ and use the fact that $\pi(z)=\bar{\pi}(z)$ on $\{0,1 \ldots, n\}$ to interpolate between $\mu$ and $\pi$.

First, we use standard Markov chains facts to compute (4.4.13). Define for $y \in\{1, \ldots, n+1\}$

$$
\begin{equation*}
g_{z}^{t}(y):=\frac{1}{2 n} \sum_{s=0}^{t-1}\left(\mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)\right]-\mathbb{E}_{y-1}\left[\delta_{z}\left(X_{s}\right)\right]\right), \tag{4.4.14}
\end{equation*}
$$

and $g_{z}^{t}(0):=0$. The function $g_{z}^{t}(y)$ is, up to a factor (whose purpose will be clearer below), the difference between the expected number of visits to $z$ up to time $t-1$ when started at $y$ and $y-1$ respectively. It depends on $\mu$ only through $\lambda$ and $n$. By Chapman-Kolmogorov (Theorem 1.1.20) applied to the first step of the chain,

$$
\begin{aligned}
\mathbb{E}_{y}\left[\delta_{z}\left(X_{s+1}\right)\right]= & P(y, y+1) \mathbb{E}_{y+1}\left[\delta_{z}\left(X_{s}\right)\right] \\
& +P(y, y) \mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)\right]+P(y, y-1) \mathbb{E}_{y-1}\left[\delta_{z}\left(X_{s}\right)\right]
\end{aligned}
$$

Using that $P(y, y+1)+P(y, y)+P(y, y-1)=1$ and rearranging we get for $0 \leq y \leq n$ and $0 \leq z \leq n+1$

$$
\begin{align*}
& \sum_{s=0}^{t-1} \mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)-\delta_{z}\left(X_{s+1}\right)\right] \\
& =\sum_{s=0}^{t-1}\left\{-P(y, y+1)\left(\mathbb{E}_{y+1}\left[\delta_{z}\left(X_{s}\right)\right]-\mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)\right]\right)\right. \\
& \left.\quad+P(y, y-1)\left(\mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)\right]-\mathbb{E}_{y-1}\left[\delta_{z}\left(X_{s}\right)\right]\right)\right\} \\
& =-2 n P(y, y+1) g_{z}^{t}(y+1)+2 n P(y, y-1) g_{z}^{t}(y) \\
& =-\lambda g_{z}^{t}(y+1)+y g_{z}^{t}(y) \tag{4.4.15}
\end{align*}
$$

where we used (4.4.11) on the last line.
We establish after the proof of the theorem that $g_{z}^{t}(y)$ has a well-defined limit. That fact is not immediately obvious as the limit is the "difference of two infinities." But a simple coupling argument does the trick.

Lemma 4.4.10. Let $g_{z}^{t}:\{0,1, \ldots, n+1\} \rightarrow \mathbb{R}$ be defined in (4.4.14). Then there exists a bounded function $g_{z}^{\infty}:\{0,1, \ldots, n+1\} \rightarrow \mathbb{R}$ such that for all $0 \leq z \leq n+1$ and $0 \leq y \leq n+1$,

$$
g_{z}^{\infty}(y)=\lim _{t \rightarrow+\infty} g_{z}^{t}(y)
$$

In fact, an explicit expression for $g_{z}^{\infty}$ can be derived via the following recursion. That expression will be helpful to establish the Lipschitz condition in Theorem 4.4.2.

Lemma 4.4.11. For all $0 \leq y \leq n$ and $0 \leq z \leq n+1$,

$$
\delta_{z}(y)-\bar{\pi}(z)=-\lambda g_{z}^{\infty}(y+1)+y g_{z}^{\infty}(y)
$$

Proof. Combine (4.4.13), (4.4.15), and Lemma 4.4.10.
Lemma 4.4.11 leads to the following formula for $g_{z}^{\infty}$, which we establish after the proof of the theorem.

Lemma 4.4.12. For $1 \leq y \leq n+1$ and $0 \leq z \leq n+1$,

$$
g_{z}^{\infty}(y)= \begin{cases}\frac{\Pi(y-1)}{y \pi(y)} \bar{\pi}(z), & \text { if } z \geq y  \tag{4.4.16}\\ -\frac{1-\Pi(y-1)}{y \pi(y)} \bar{\pi}(z), & \text { if } z<y\end{cases}
$$

and $g_{z}^{\infty}(0)=0$.

Interpolating between $\mu$ and $\pi \quad$ For $A \subseteq\{0,1, \ldots, n\}$, define

$$
g_{A}^{\infty}(y):=\sum_{z \in A} g_{z}^{\infty}(y)
$$

We obtain the following key bound.
Lemma 4.4.13 (Chen's equation). Let $W \sim \mu$ and $\pi \stackrel{\mathrm{d}}{=} \operatorname{Poi}(\lambda)$. Then,

$$
\begin{equation*}
\|\mu-\pi\|_{\mathrm{TV}}=\mathbb{E}\left[-\lambda g_{A^{*}}^{\infty}(W+1)+W g_{A^{*}}^{\infty}(W)\right] \tag{4.4.17}
\end{equation*}
$$

where $A^{*}=\left\{z \in \mathbb{Z}_{+}: \mu(z)>\pi(z)\right\}$.
Proof. Fix $z \in\{0,1, \ldots, n\}$. Multiplying both sides in Lemma 4.4.11 by $\mu(y)$ and summing over $y$ in $\{0,1, \ldots, n\}$ gives

$$
\mu(z)-\pi(z)=\mathbb{E}\left[-\lambda g_{z}^{\infty}(W+1)+W g_{z}^{\infty}(W)\right]
$$

Now summing over $z$ in $A^{*} \subseteq\{0,1, \ldots, n\}$ and using (4.4.9) gives the claim.

Lemma 4.4.12 can be used to derive a Lipschitz constant for $g_{A}^{\infty}$. That lemma is also established after the proof of the theorem.

Lemma 4.4.14. For $A \subseteq\{0,1, \ldots, n\}$ and $y, y^{\prime} \in\{0,1, \ldots, n+1\}$,

$$
\left|g_{A}^{\infty}\left(y^{\prime}\right)-g_{A}^{\infty}(y)\right| \leq\left(1 \wedge \lambda^{-1}\right)\left|y^{\prime}-y\right| .
$$

Lemmas 4.4.13 and 4.4.14 imply the theorem with $h:=g_{A^{*}}^{\infty}$.

Proofs of technical lemmas It remains to prove Lemmas 4.4.10, 4.4.12 and 4.4.14.

Proof of Lemma 4.4.10. We use a coupling argument. Let $\left(Y_{s}, \tilde{Y}_{s}\right)_{s=0}^{+\infty}$ be an independent Markovian coupling of $\left(Y_{s}\right)$, the chain started at $y-1$, and $\left(\tilde{Y}_{s}\right)$, the chain started at $y$. Let $\tau$ be the first time $s$ that $Y_{s}=\tilde{Y}_{s}$. Because $Y_{s}$ and $\tilde{Y}_{s}$ are independent and $P$ is a birth-death chain with strictly positive nearest-neighbor and staying-put transition probabilities, the coupled chain $\left(Y_{s}, \tilde{Y}_{s}\right)_{s=0}^{+\infty}$ is aperiodic and irreducible over $\{0,1, \ldots, n+1\}^{2}$. By the exponential tail of hitting times, Lemma 3.1.25, it holds that $\mathbb{E}[\tau]<+\infty$.

Modify the coupling $\left(Y_{s}, \tilde{Y}_{s}\right)$ to enforce $\tilde{Y}_{s}=Y_{s}$ for all $s \geq \tau$ (while not changing $\left(Y_{s}\right)$ ), that is, to make it coalescing. By the Strong Markov property (Theorem 3.1.8), the resulting chain $\left(Y_{s}^{*}, \tilde{Y}_{s}^{*}\right)$ is also a Markovian coupling of the chain started at $y-1$ and $y$ respectively. Using this coupling, we rewrite

$$
\begin{aligned}
g_{z}^{t}(y) & =\frac{1}{2 n} \sum_{s=0}^{t-1}\left(\mathbb{E}_{y}\left[\delta_{z}\left(X_{s}\right)\right]-\mathbb{E}_{y-1}\left[\delta_{z}\left(X_{s}\right)\right]\right) \\
& =\frac{1}{2 n} \sum_{s=0}^{t-1} \mathbb{E}\left[\delta_{z}\left(\tilde{Y}_{s}^{*}\right)-\delta_{z}\left(Y_{s}^{*}\right)\right] \\
& =\frac{1}{2 n} \mathbb{E}\left[\sum_{s=0}^{t-1}\left(\delta_{z}\left(\tilde{Y}_{s}^{*}\right)-\delta_{z}\left(Y_{s}^{*}\right)\right)\right] .
\end{aligned}
$$

The random variable inside the expectation is bounded in absolute value by

$$
\left|\sum_{s=0}^{t-1}\left(\delta_{z}\left(\tilde{Y}_{s}^{*}\right)-\delta_{z}\left(Y_{s}^{*}\right)\right)\right| \leq \tau
$$

uniformly in $t$. Indeed, after $s=\tau$, the terms in the sum are 0 , while before $s=\tau$ the terms are bounded by 1 in absolute value. By the integrability of $\tau$,
the dominated convergence theorem (Proposition B.4.14) allows to take the limit, leading to

$$
\begin{aligned}
g_{z}^{\infty}(y) & =\lim _{t \rightarrow+\infty} \frac{1}{2 n} \mathbb{E}\left[\sum_{s=0}^{t-1}\left(\delta_{z}\left(\tilde{Y}_{s}^{*}\right)-\delta_{z}\left(Y_{s}^{*}\right)\right)\right] \\
& =\frac{1}{2 n} \mathbb{E}\left[\sum_{s=0}^{+\infty}\left(\delta_{z}\left(\tilde{Y}_{s}^{*}\right)-\delta_{z}\left(Y_{s}^{*}\right)\right)\right] \\
& <+\infty .
\end{aligned}
$$

That concludes the proof.
Proof of Lemma 4.4.12. Our starting point is Lemma 4.4.11, from which we deduce the recursive formula

$$
\begin{equation*}
g_{z}^{\infty}(y+1)=\frac{1}{\lambda}\left\{y g_{z}^{\infty}(y)+\pi(z)-\delta_{z}(y)\right\}, \tag{4.4.18}
\end{equation*}
$$

for $0 \leq y \leq n$ and $0 \leq z \leq n$.
We guess a general formula and then check it. By (4.4.18),

$$
\begin{gather*}
g_{z}^{\infty}(1)=\frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(0)\right\},  \tag{4.4.19}\\
g_{z}^{\infty}(2)=\frac{1}{\lambda}\left\{g_{z}^{\infty}(1)+\pi(z)-\delta_{z}(1)\right\} \\
=\frac{1}{\lambda}\left\{\frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(0)\right\}+\pi(z)-\delta_{z}(1)\right\} \\
=\frac{1}{\lambda^{2}}\left\{\pi(z)-\delta_{z}(0)\right\}+\frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(1)\right\}, \\
g_{z}^{\infty}(3)=\frac{1}{\lambda}\left\{2 g_{z}^{\infty}(2)+\pi(z)-\delta_{z}(2)\right\} \\
=\frac{1}{\lambda}\left\{2 \frac{1}{\lambda^{2}}\left\{\pi(z)-\delta_{z}(0)\right\}+2 \frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(1)\right\}+\pi(z)-\delta_{z}(2)\right\} \\
=\frac{2}{\lambda^{3}}\left\{\pi(z)-\delta_{z}(0)\right\}+\frac{2}{\lambda^{2}}\left\{\pi(z)-\delta_{z}(1)\right\}+\frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(2)\right\},
\end{gather*}
$$

and so forth. We posit the general formula

$$
\begin{equation*}
g_{z}^{\infty}(y)=\frac{(y-1)!}{\lambda^{y}} \sum_{k=0}^{y-1} \frac{\lambda^{k}}{k!}\left\{\pi(z)-\delta_{z}(k)\right\}, \tag{4.4.20}
\end{equation*}
$$

for $1 \leq y \leq n+1$ and $0 \leq z \leq n$.
The formula is straightforward to confirm by induction. Indeed it holds for $y=1$ as can be seen in (4.4.19) (and recalling that $0!=1$ by convention) and, assuming it holds for $y$, we have by (4.4.18)

$$
\begin{aligned}
g_{z}^{\infty}(y+1) & =\frac{1}{\lambda}\left\{y g_{z}^{\infty}(y)+\pi(z)-\delta_{z}(y)\right\} \\
& =\frac{1}{\lambda}\left\{y \frac{(y-1)!}{\lambda^{y}} \sum_{k=0}^{y-1} \frac{\lambda^{k}}{k!}\left\{\pi(z)-\delta_{z}(k)\right\}+\pi(z)-\delta_{z}(y)\right\} \\
& =\frac{y!}{\lambda^{y+1}} \sum_{k=0}^{y-1} \frac{\lambda^{k}}{k!}\left\{\pi(z)-\delta_{z}(k)\right\}+\frac{1}{\lambda}\left\{\pi(z)-\delta_{z}(y)\right\} \\
& =\frac{y!}{\lambda^{y+1}} \sum_{k=0}^{y} \frac{\lambda^{k}}{k!}\left\{\pi(z)-\delta_{z}(k)\right\}
\end{aligned}
$$

as desired.
We rewrite (4.4.20) according to whether the term $\delta_{z}(y)=\mathbf{1}\{z=y\}$ plays a role in the equation. For $z \geq y>0$, the equation simplifies to

$$
\begin{aligned}
g_{z}^{\infty}(y) & =\frac{(y-1)!}{\lambda^{y}} \sum_{k=0}^{y-1} \frac{\lambda^{k}}{k!} \pi(z) \\
& =\frac{1}{y} \frac{y!}{e^{-\lambda} \lambda^{y}} \sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!} \pi(z) \\
& =\frac{\Pi(y-1)}{y \pi(y)} \pi(z) .
\end{aligned}
$$

For $0 \leq z<y$, we get instead

$$
\begin{aligned}
g_{z}^{\infty}(y) & =\frac{(y-1)!}{\lambda^{y}}\left\{\left(\sum_{k=0}^{y-1} \frac{\lambda^{k}}{k!} \pi(z)\right)-\frac{\lambda^{z}}{z!}\right\} \\
& =\frac{1}{y} \frac{y!}{e^{-\lambda} \lambda^{y}}\left\{\left(\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!} \pi(z)\right)-\pi(z)\right\} \\
& =\frac{\Pi(y-1)-1}{y \pi(y)} \pi(z) .
\end{aligned}
$$

The cases $z=n+1$ are analogous. That concludes the proof.

Proof of Lemma 4.4.14. It suffices to prove that, for $A \subseteq\{0,1, \ldots, n\}$ and $y \in$ $\{0,1, \ldots, n\}$,

$$
\begin{equation*}
\left|g_{A}^{\infty}(y+1)-g_{A}^{\infty}(y)\right| \leq\left(1 \wedge \lambda^{-1}\right), \tag{4.4.21}
\end{equation*}
$$

and then use the triangle inequality.
We start with the case $y \geq 1$. We use the expression derived in Lemma 4.4.12. For $1 \leq y<z \leq n+1$,

$$
\begin{aligned}
g_{z}^{\infty}(y+1)-g_{z}^{\infty}(y) & =\frac{\Pi(y)}{(y+1) \pi(y+1)} \bar{\pi}(z)-\frac{\Pi(y-1)}{y \pi(y)} \bar{\pi}(z) \\
& =\bar{\pi}(z) \frac{1}{y \pi(y)}\left\{\frac{y}{\lambda} \Pi(y)-\Pi(y-1)\right\},
\end{aligned}
$$

where we used that $\pi(y+1) / \pi(y)=\lambda /(y+1)$. We show that the expression in curly brackets is non-negative. Indeed, taking out the term $k^{\prime}=0$ in the first sum below and changing variables, we get

$$
\begin{aligned}
& \frac{y}{\lambda} \sum_{k^{\prime}=0}^{y} \frac{e^{-\lambda} \lambda^{k^{\prime}}}{\left(k^{\prime}\right)!}-\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!} \\
& =\frac{y}{\lambda} e^{-\lambda}+\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{(k+1)-1}}{(k+1)!/ y}-\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!} \\
& \geq \frac{y}{\lambda} e^{-\lambda}+\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!}-\sum_{k=0}^{y-1} \frac{e^{-\lambda} \lambda^{k}}{k!} \\
& \geq 0 .
\end{aligned}
$$

So $g_{z}^{\infty}(y+1)-g_{z}^{\infty}(y) \geq 0$ for $1 \leq y<z$. A similar calculation, which we omit, shows that the same inequality holds for $z<y \leq n$. The cases $y=0$, which are analogous, are detailed below.

For notational convenience, it will be helpful to define $g_{z}^{\infty}(n+2)=0$ for all $z$. Then, for $y=n+1$ and $z \leq n$, we get

$$
g_{z}^{\infty}(n+2)-g_{z}^{\infty}(n+1)=0+\frac{1-\Pi(n)}{(n+1) \pi(n+1)} \pi(z) \geq 0 .
$$

Moreover, by telescoping,

$$
0=g_{z}^{\infty}(n+2)-g_{z}^{\infty}(0)=\sum_{y=0}^{n+1}\left\{g_{z}^{\infty}(y+1)-g_{z}^{\infty}(y)\right\}
$$

We have argued that all the terms in this last sum are non-negative-with the sole exception of the term $y=z$. Hence, for a fixed $0 \leq z \leq n$, it must be that
the maximum of $\left|g_{z}^{\infty}(y+1)-g_{z}^{\infty}(y)\right|$ is achieved at $y=z$. The case $z=$ $n+1$ is analogous. By definition of $g_{z}^{\infty}$, for a fixed $0 \leq y \leq n$, it holds that $\sum_{z}\left\{g_{z}^{\infty}(y+1)-g_{z}^{\infty}(y)\right\}=0$ and the maximum of $\left|g_{A}^{\infty}(y+1)-g_{A}^{\infty}(y)\right|$ over $A \subseteq\{0,1, \ldots, n\}$ is achieved at $A=\{y\}$. It remains to bound that last case.

We have, using $\pi(y+1) / \pi(y)=\lambda /(y+1)$ again, that for $1 \leq y \leq n$

$$
\begin{aligned}
& \left|g_{y}^{\infty}(y+1)-g_{y}^{\infty}(y)\right| \\
& =\left|-\frac{1-\Pi(y)}{(y+1) \pi(y+1)} \pi(y)-\frac{\Pi(y-1)}{y \pi(y)} \pi(y)\right| \\
& =\frac{1}{\lambda} \sum_{k \geq y+1} e^{-\lambda} \frac{\lambda^{k}}{k!}+\frac{1}{y} \sum_{k=0}^{y-1} e^{-\lambda} \frac{\lambda^{k}}{k!} \\
& =\frac{e^{-\lambda}}{\lambda}\left\{\sum_{k^{\prime}=1}^{y} \frac{\lambda^{k^{\prime}}}{\left(k^{\prime}\right)!} \frac{k^{\prime}}{y}+\sum_{k \geq y+1} \frac{\lambda^{k}}{k!}\right\} \\
& \leq \frac{e^{-\lambda}}{\lambda}\left\{\sum_{k \geq 1} \frac{\lambda^{k}}{k!}\right\} \\
& =\frac{e^{-\lambda}}{\lambda}\left\{e^{\lambda}-1\right\} \\
& =\frac{1-e^{-\lambda}}{\lambda} .
\end{aligned}
$$

For $\lambda \geq 1$, we have $\frac{1-e^{-\lambda}}{\lambda} \leq \frac{1}{\lambda}=\left(1 \wedge \lambda^{-1}\right)$, while for $0<\lambda<1$ we have $\frac{1-e^{-\lambda}}{\lambda} \leq \frac{\lambda}{\lambda}=1=\left(1 \wedge \lambda^{-1}\right)$ by Exercise 1.16.

It remains to consider the cases $y=0$. Recall that $g_{z}^{\infty}(0)=0$. By Lemma 4.4.12, for $z \geq 1$,

$$
\begin{aligned}
g_{z}^{\infty}(1)-g_{z}^{\infty}(0) & =g_{z}^{\infty}(1) \\
& =\frac{\Pi(0)}{\pi(1)} \bar{\pi}(z) \\
& \geq 0 .
\end{aligned}
$$

And

$$
g_{0}^{\infty}(1)-g_{0}^{\infty}(0)=g_{0}^{\infty}(1)=-\frac{1-\Pi(0)}{\pi(1)} \pi(0)=-\frac{1-e^{-\lambda}}{\lambda} .
$$

So we have established (4.4.21) and that concludes the proof.

### 4.4.3 $\triangleright$ Random graphs: clique number at the threshold in the ErdôsRényi model

We revisit the subgraph containment problem of Section 2.3.2 (and Section 4.2.4). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ be an Erdős-Rényi graph with $n$ vertices and density $p_{n}$. Let $\omega(G)$ be the clique number of a graph $G$, that is, the size of its largest clique. We showed previously that the property $\omega(G) \geq 4$ has threshold function $n^{-2 / 3}$. Here we consider what happens when

$$
p_{n}=C n^{-2 / 3},
$$

for some constant $C>0$. We use the Chen-Stein method in the form of Theorem 4.4.8.

For an enumeration $S_{1}, \ldots, S_{m}$ of the 4 -tuples of vertices in $G_{n}$, let $A_{1}, \ldots, A_{m}$ be the events that the corresponding 4 -cliques are present and define $Z_{i}=\mathbf{1}_{A_{i}}$. Then $W=\sum_{i=1}^{m} Z_{i}$ is the number of 4-cliques in $G_{n}$. We argued previously (see Claim 2.3.4) that

$$
q_{i}:=\mathbb{E}\left[Z_{i}\right]=p_{n}^{6},
$$

and

$$
\lambda:=\mathbb{E}[W]=\binom{n}{4} p_{n}^{6} .
$$

In our regime of interest, $\lambda$ is of constant order.
Observe that the $Z_{i}$ s are not independent because the 4 -tuples may share potential edges. However they admit a neighborhood structure as in Theorem 4.4.8. Specifically, for $i=1, \ldots, m$, define

$$
\mathcal{N}_{i}=\left\{j: S_{i} \text { and } S_{j} \text { share at least two vertices }\right\} \backslash\{i\} .
$$

Then the conditions of Theorem 4.4.8 are satisfied, that is, $X_{i}$ is independent of $\left\{Z_{j}: j \notin \mathcal{N}_{i} \cup\{i\}\right\}$. We argued previously (again see Claim 2.3.4) that

$$
\left|\mathcal{N}_{i}\right|=\binom{4}{3}(n-4)+\binom{4}{2}\binom{n-4}{2}=\Theta\left(n^{2}\right),
$$

where the first term counts the number of $S_{j} \mathrm{~s}$ sharing exactly three vertices with $S_{i}$, in which case $\mathbb{E}\left[Z_{i} Z_{j}\right]=p_{n}^{9}$, and the second term counts those sharing two, in which case $\mathbb{E}\left[Z_{i} Z_{j}\right]=p_{n}^{11}$.

We are ready to apply the bound in Theorem 4.4.8. Let $\pi$ be the Poisson distribution with mean $\lambda$. Using the formulas above, we get when $p_{n}=C n^{-2 / 3}$

$$
\begin{aligned}
&\|\mu-\pi\|_{\mathrm{TV}} \\
& \leq\left(1 \wedge \lambda^{-1}\right) \sum_{i=1}^{n}\left\{q_{i}^{2}+\sum_{j \in \mathcal{N}_{i}}\left(q_{i} q_{j}+\mathbb{E}\left[Z_{i} Z_{j}\right]\right)\right\} \\
& \leq\left(1 \wedge \lambda^{-1}\right)\binom{n}{4} \\
& \times\left[p_{n}^{12}+\left\{\binom{4}{3}(n-4)\left(p_{n}^{12}+p_{n}^{9}\right)+\binom{4}{2}\binom{n-4}{2}\left(p_{n}^{12}+p_{n}^{11}\right)\right\}\right] \\
&=\left(1 \wedge \lambda^{-1}\right) \Theta\left(n^{4} p_{n}^{12}+n^{5} p_{n}^{9}+n^{6} p_{n}^{11}\right) \\
&=\left(1 \wedge \lambda^{-1}\right) \Theta\left(n^{4} n^{-8}+n^{5} n^{-6}+n^{6} n^{-22 / 3}\right) \\
&=\left(1 \wedge \lambda^{-1}\right) \Theta\left(n^{-1}\right),
\end{aligned}
$$

which goes to 0 as $n \rightarrow+\infty$.
See Exercise 4.21 for an improved bound.

## Exercises

Exercise 4.1 (Harmonic function on $\mathbb{Z}^{d}$ : unbounded). Give an example of an unbounded harmonic function on $\mathbb{Z}$. Give one on $\mathbb{Z}^{d}$ for general $d$. [Hint: What is the simplest function after the constant one?]

Exercise 4.2 (Binomial vs. Binomial). Use coupling to show that

$$
n \geq m, q \geq p \Longrightarrow \operatorname{Bin}(n, q) \succeq \operatorname{Bin}(m, p) .
$$

Exercise 4.3 (A chain that is not stochastically monotone). Consider random walk on a network $\mathcal{N}=((V, E), c)$ where $V=\{0,1, \ldots, n\}$ and $i \sim j$ if and only if $|i-j|=1$ (in particular, not including self-loops). Show that the transition matrix is, in general, not stochastically monotone (see Definition 4.2.16).

Exercise 4.4 (Increasing events: properties). Let $\mathcal{F}$ be a $\sigma$-algebra over the poset $\mathcal{X}$. Recall that an event $A \in \mathcal{F}$ is increasing if $x \in A$ implies that any $y \geq x$ is also in $A$ and that a function $f: \mathcal{X} \rightarrow \mathbb{R}$ is increasing if $x \leq y$ implies $f(x) \leq f(y)$.
(i) Show that an event $A \in \mathcal{F}$ is increasing if and only if the indicator function $\mathbf{1}_{A}$ is increasing.
(ii) Let $A, B \in \mathcal{F}$ be increasing. Show that $A \cap B$ and $A \cup B$ are increasing.
(iii) An event $A$ is decreasing if $x \in A$ implies that any $y \leq x$ is also in $A$. Show that $A$ is decreasing if and only if $A^{c}$ is increasing.
(iv) Let $A, B \in \mathcal{F}$ be decreasing. Show that $A \cap B$ and $A \cup B$ are decreasing.

Exercise 4.5 (Harris' inequality: alternative proof). We say that $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is coordinatewise nondecreasing if it is nondecreasing in each variable while keeping the other variables fixed.
(i) (Chebyshev's association inequality) Let $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ be coordinatewise nondecreasing and let $X$ be a real random variable. Show that

$$
\mathbb{E}[f(X) g(X)] \geq \mathbb{E}[f(X)] \mathbb{E}[g(X)] .
$$

[Hint: Consider the quantity $\left(f(X)-f\left(X^{\prime}\right)\right)\left(g(X)-g\left(X^{\prime}\right)\right)$ where $X^{\prime}$ is an independent copy of $X$.]
(ii) (Harris' inequality) Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ and $g: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be coordinatewise nondecreasing and let $X=\left(X_{1}, \ldots, X_{n}\right)$ be independent real random variables. Show by induction on $n$ that

$$
\mathbb{E}[f(X) g(X)] \geq \mathbb{E}[f(X)] \mathbb{E}[g(X)]
$$

[Hint: You may need Lemma B.6.15.]

Exercise 4.6. Provide the details for Example 4.2.33.
Exercise 4.7 (FKG: sufficient conditions). Let $\mathcal{X}:=\{0,1\}^{F}$ where $F$ is finite and let $\mu$ be a positive probability measure on $\mathcal{X}$. We use the notation introduced in the proof of Holley's inequality (Theorem 4.2.32).
(i) To check the FKG condition, show that it suffices to check that, for all $x \leq$ $y \in \mathcal{X}$ and $i \in F$,

$$
\frac{\mu\left(y^{i, 1}\right)}{\mu\left(y^{i, 0}\right)} \geq \frac{\mu\left(x^{i, 1}\right)}{\mu\left(x^{i, 0}\right)} .
$$

[Hint: Write $\mu\left(\omega \vee \omega^{\prime}\right) / \mu(\omega)$ as a telescoping product.]
(ii) To check the FKG condition, show that it suffices to check (4.2.15) only for those $\omega, \omega^{\prime} \in \mathcal{X}$ such that $\left\|\omega-\omega^{\prime}\right\|_{1}=2$ and neither $\omega \leq \omega^{\prime}$ nor $\omega^{\prime} \leq \omega$. [Hint: Use (i).]

Exercise 4.8 (FKG and strong positive associations). Let $\mathcal{X}:=\{0,1\}^{F}$ where $F$ is finite and let $\mu$ be a positive probability measure on $\mathcal{X}$. For $\Lambda \subseteq F$ and $\xi \in \mathcal{X}$, let

$$
\mathcal{X}_{\Lambda}^{\xi}:=\left\{\omega_{\Lambda} \times \xi_{\Lambda^{c}}: \omega_{\Lambda} \in\{0,1\}^{\Lambda}\right\},
$$

where $\omega_{\Lambda} \times \xi_{\Lambda^{c}}$ agrees with $\omega$ on coordinates in $\Lambda$ and with $\xi$ on coordinates in $F \backslash \Lambda$. Define the measure $\mu_{\Lambda}^{\xi}$ over $\{0,1\}^{\Lambda}$ as

$$
\mu_{\Lambda}^{\xi}\left(\omega_{\Lambda}\right):=\frac{\mu\left(\omega_{\Lambda} \times \xi_{\Lambda^{c}}\right)}{\mu\left(\mathcal{X}_{\Lambda}^{\xi}\right)}
$$

That is, $\mu_{\Lambda}^{\xi}$ is $\mu$ conditioned on agreeing with $\xi$ on $F \backslash \Lambda$. The measure $\mu$ is said to be strongly positively associated if $\mu_{\Lambda}^{\xi}\left(\omega_{\Lambda}\right)$ is positively associated for all $\Lambda$ and $\xi$. Prove that the FKG condition is equivalent to strong positive associations. [Hint: Use Exercise 4.7 as well as the FKG inequality.]

Exercise 4.9 (Triangle-freeness: a second proof). Consider again the setting of Section 4.2.4.
(i) Let $e_{t}$ be the minimum number of edges in a $t$-vertex union of $k$ not mutually vertex-disjoint triangles. Show that, for any $k \geq 2$ and $k \leq t<3 k$, it holds that $e_{t}>t$.
(ii) Use Exercise 2.18 to give a second proof of the fact that $\mathbb{P}\left[X_{n}=0\right] \rightarrow$ $e^{-\lambda^{3} / 6}$.

Exercise 4.10 (RSW lemma: general $\alpha$ ). Let $R_{n, \alpha}(p)$ be as defined in Section 4.2.5. Show that for all $n \geq 2$ (divisible by 4 ) and $p \in(0,1)$

$$
R_{n, \alpha}(p) \geq\left(\frac{1}{2}\right)^{2 \alpha-2} R_{n, 1}(p)^{6 \alpha-7} R_{n / 2,1}(p)^{6 \alpha-6}
$$

Exercise 4.11 (Primal and dual crossings). Modify the proof of Lemma 2.2.14 to prove Lemma 4.2.41.
Exercise 4.12 (Square-root trick). Let $\mu$ be an FKG measure on $\{0,1\}^{F}$ where $F$ is finite. Let $A_{1}$ and $A_{2}$ be increasing events with $\mu\left(A_{1}\right)=\mu\left(A_{2}\right)$. Show that

$$
\mu\left(A_{1}\right) \geq 1-\sqrt{1-\mu\left(A_{1} \cup A_{2}\right)} .
$$

Exercise 4.13 (Splitting: details). Show that $\tilde{P}$, as defined in Example 4.3.3, is a transition matrix on $V$ provided $z_{0}$ satisfies the condition there.

Exercise 4.14 (Doeblin's condition in finite case). Let $P$ be a transition matrix on a finite state space.
(i) Show that Doeblin's condition (see Example 4.3.3) holds when $P$ is finite, irreducible and aperiodic.
(ii) Show that Doeblin's condition holds for lazy random walk on the hypercube with $s=n$. Use it to derive a bound on the mixing time.
Exercise 4.15 (Mixing on cycles: lower bound). Let $\left(Z_{t}\right)$ be lazy, simple random walk on the cycle of size $n, \mathbb{Z}_{n}:=\{0,1 \ldots, n-1\}$, where $i \sim j$ if $|j-i|=1$ $(\bmod n)$. Assume $n$ is divisible by 4 and fix $0<\varepsilon<1 / 2$.
(i) Let $A=\{n / 2, \ldots, n-1\}$. By coupling $\left(Z_{t}\right)$ with lazy, simple random walk on $\mathbb{Z}$, show that

$$
P^{\alpha n^{2}}(n / 4, A)<\frac{1}{2}-\varepsilon
$$

for $\alpha \leq \alpha_{\varepsilon}$ for some $\alpha_{\varepsilon}>0$. [Hint: You may want to use Chebyshev's inequality (Theorem 2.1.2) or Kolmogorov's maximal inequality (Corollary 3.1.46).]
(ii) Deduce that

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \alpha_{\varepsilon} n^{2} .
$$

Exercise 4.16 (Lower bound on mixing: distinguishing statistic). Let $X$ and $Y$ be random variables on a finite state space $S$. Let $h: S \rightarrow \mathbb{R}$ be a measurable real-valued map. Assume that

$$
\mathbb{E}[h(Y)]-\mathbb{E}[h(X)] \geq r \sigma,
$$

where $r>0$ and $\sigma^{2}:=\max \{\operatorname{Var}[h(X)], \operatorname{Var}[h(Y)]\}$. Show that

$$
\left\|\mu_{X}-\mu_{Y}\right\|_{\mathrm{TV}} \geq 1-\frac{8}{r^{2}}
$$

[Hint: Consider the interval on one side of the midpoint between $\mathbb{E}[h(X)]$ and $\mathbb{E}[h(Y)]$.

Exercise 4.17 (Path coupling and optimal transport). Let $V$ be a finite state space and let $P$ be an irreducible transition matrix on $V$ with stationary distribution $\pi$. Let $w_{0}$ be a metric on $V$. For probability measures $\mu, \nu$ on $V$, let

$$
W_{0}(\mu, \nu):=\inf \left\{\mathbb{E}\left[w_{0}(X, Y)\right]:(X, Y) \text { is a coupling of } \mu \text { and } \nu\right\},
$$

be the so-called Wasserstein distance (or transportation metric) between $\mu$ and $\nu$.
(i) Show that $W_{0}$ is a metric. [Hint: See the proof of Claim 4.3.11.]
(ii) Assume that the conditions of Theorem 4.3.10 hold. Show that for any probability measures $\mu, \nu$

$$
W_{0}(\mu P, \nu P) \leq \kappa W_{0}(\mu, \nu) .
$$

(iii) Use (i) and (ii) to prove Theorem 4.3.10.

Exercise 4.18 (Stein equation for the Poisson distribution). Let $\lambda>0$. Show that a non-negative integer-valued random variable $Z$ is $\operatorname{Poi}(\lambda)$ if and only if for all $g$ bounded

$$
\mathbb{E}[\lambda g(Z+1)-Z g(Z)]=0 .
$$

Exercise 4.19 (Laplacian and stationarity). Let $P$ be an irreducible transition matrix on a finite or countably infinite spate space $V$. Recall the Laplacian operator is

$$
\Delta f(x)=\left[\sum_{y} P(x, y) f(y)\right]-f(x),
$$

provided the sum if finite. Show that a probability distribution $\mu$ over $V$ is stationary for $P$ if and only if for all bounded measurable functions

$$
\sum_{x \in V} \mu(x) \Delta f(x)=0
$$

Exercise 4.20 (Chen-Stein method for positively related variables). Using the notation in (4.4.1), (4.4.2) and (4.4.3), suppose that for each $i$ we can construct a coupling $\left\{\left(X_{j}^{(i)}: j=1, \ldots, n\right),\left(Y_{j}^{(i)}: j \neq i\right)\right\}$ with $\left(X_{j}^{(i)}\right)_{j} \sim\left(X_{j}\right)_{j}$ such that

$$
\left(Y_{j}^{(i)}, j \neq i\right) \sim\left(X_{j}^{(i)}, j \neq i\right) \mid X_{i}^{(i)}=1 \quad \text { and } \quad Y_{j}^{(i)} \geq X_{j}^{(i)}, \forall j \neq i
$$

Show that

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right)\left\{\operatorname{Var}(W)-\lambda+2 \sum_{i=1}^{n} p_{i}^{2}\right\}
$$

Exercise 4.21 (Chen-Stein and 4-cliques). Use Exercise 4.20 to give an improved asymptotic bound in the setting of Section 4.4.3.

Exercise 4.22 (Chen-Stein for negatively related variables). Using the notation in (4.4.1), (4.4.2) and (4.4.3), suppose that for each $i$ we can construct a coupling $\left\{\left(X_{j}^{(i)}: j=1, \ldots, n\right),\left(Y_{j}^{(i)}: j \neq i\right)\right\}$ with $\left(X_{j}^{(i)}\right)_{j} \sim\left(X_{j}\right)_{j}$ such that

$$
\left(Y_{j}^{(i)}, j \neq i\right) \sim\left(X_{j}^{(i)}, j \neq i\right) \mid X_{i}^{(i)}=1 \quad \text { and } \quad Y_{j}^{(i)} \leq X_{j}^{(i)}, \forall j \neq i .
$$

Show that

$$
\|\mu-\pi\|_{\mathrm{TV}} \leq\left(1 \wedge \lambda^{-1}\right)\{\lambda-\operatorname{Var}(W)\} .
$$

## Bibliographic remarks

Section 4.1 The coupling method is generally attributed to Doeblin [Doe38]. The standard reference on coupling is [Lin02]. See that reference for a history of coupling and a facsimile of Doeblin's paper. See also [dH]. Section 4.1.2 is based on [Per, Section 6] and Section 4.1.4 is based on [vdH17, Section 5.3].

Section 4.2 Strassen's theorem is due to Strassen [Str65]. Harris' inequality is due to Harris [Har60]. The FKG inequality is due to Fortuin, Kasteleyn, and Ginibre [FKG71]. A "four-function" version of Holley's inequality, which also extends to distributive lattices, was proved by Ahlswede and Daykin [AD78]. See for example [AS11, Section 6.1]. An exposition of submodularity and its connections to convexity can be found in [Lov83]. For more on Markov random fields, see for example [RAS15]. Section 4.2.4 follows [AS11, Sections 8.1, 8.2, 10.1]. Janson's inequality is due to Janson [Jan90]. Boppana and Spencer [BS89] gave the proof presented here. For more on Janson's inequality, see [JLR11, Section 2.2]. The presentation in Section 4.2.5 follows closely [BR06b, Sections 3 and 4]. See also [BR06a, Chapter 3]. Broadbent and Hammersley [BH57, Ham57] initiated the study of the critical value of percolation. Harris' theorem was proved by Harris [Har60] and Kesten's theorem was proved two decades later by Kesten [Kes80], confirming non-rigorous work of Sykes and Essam [SE64]. The RSW lemma was obtained independently by Russo [Rus78] and Seymour and Welsh [SW78]. The proof we gave here is due to Bollobás and Riordan [BR06b]. Another short proof of a version of the RSW lemma for critical site percolation on a triangular lattice was given by Smirnov; see for example [Ste]. The type of "scale invariance" seen in the RSW lemma plays a key role in the contemporary theory of critical two-dimensional percolation and of two-dimensional lattice models more generally. See for example [Law05, Gri10a].

Section 4.3 The material in Section 4.3 borrows heavily from [LPW06, Chapters 5, 14, 15] and [AF, Chapter 12]. Aldous [Ald83] was the first author to make explicit use of coupling to bound total variation distance to stationarity of finite Markov chains. The link between couplings of Markov chains and total variation distance was also used by Griffeath [Gri75] and Pitman [Pit76]. Example 4.3.3 is based on [Str14] and [JH01]. For a more general treatment, see [MT09, Chapter 16]. The proof of Claim 4.3 .7 is partly based on [LPW06, Proposition 7.13]. See also $\left[\mathrm{DGG}^{+} 00\right]$ and $[\mathrm{HS} 07]$ for alternative proofs. Path coupling is due to Bubley and Dyer [BD97]. The optimal transport perspective on the path coupling method in Exercise 4.17 is from [LPW06, Chapter 14]. For more on optimal transport,
see for example [Vil09]. The main result in Section 4.3 .4 is taken from [LPW06, Theorem 15.1]. For more background on the so-called "critical slowdown" of the Glauber dynamics of Ising and Potts models on various graphs, see $\left[\mathrm{CDL}^{+} 12\right.$, LS12].

Section 4.4 The Chen-Stein method was introduced by Chen in [Che75] as an adaptation of the Stein method [Ste72] to the Poisson distribution. The presentation in Section 4.4 is inspired heavily by [Dey] and [vH16]. Example 4.4.9 is taken from [AGG89]. Further applications of the Chen-Stein and Stein methods to random graphs can be found in [JLR11, Chapter 6].

## Chapter 5

## Spectral methods

In this chapter, we develop spectral techniques. We highlight some applications to Markov chain mixing and network analysis. The main tools are the spectral theorem and the variational characterization of eigenvalues, which we review in Section 5.1 together with some related results. We also give a brief introduction to spectral graph theory and detail an application to community recovery. In Section 5.2 we apply the spectral theorem to reversible Markov chains. In particular we define the spectral gap and establish its close relationship to the mixing time. Roughly speaking, we show through an eigendecomposition of the transition matrix that the gap between the eigenvalue 1 (which is the largest in absolute value) and the rest of the spectrum drives how fast $P^{t}$ converges to the stationary distribution. We give several examples. We then show in Section 5.3 that the spectral gap can be bounded using certain isoperimetric properties of the underlying network. We prove Cheeger's inequality, which quantifies this relationship, and introduce expander graphs, an important family of graphs with good "expansion." Applications to mixing times are also discussed. One specific technique is the "canonical paths method," which bounds the spectral graph by formalizing a notion of congestion in the network.

### 5.1 Background

We first review some important concepts from linear algebra. In particular, we recall the spectral theorem as well as the variational characterization of eigenvalues. We also derive a few perturbation results. We end this section with an application

[^7]to community recovery in network analysis.

### 5.1.1 Eigenvalues and their variational characterization

When a matrix $A \in \mathbb{R}^{d \times d}$ is symmetric, that is, $a_{i j}=a_{j i}$ for all $i, j$, a remarkable result is that $A$ is similar to a diagonal matrix by an orthogonal transformation. Put differently, there exists an orthonormal basis of $\mathbb{R}^{d}$ made of eigenvectors of $A$. Recall that a matrix $Q \in \mathbb{R}^{d \times d}$ is orthogonal if $Q Q^{T}=I_{d \times d}$ and $Q^{T} Q=I_{d \times d}$, where $I_{d \times d}$ is the $d \times d$ identity matrix. In words, its columns form an orthonormal basis of $\mathbb{R}^{d}$. For a vector $\mathbf{z}=\left(z_{1}, \ldots, z_{d}\right)$, we let $\operatorname{diag}(\mathbf{z})=\operatorname{diag}\left(z_{1}, \ldots, z_{d}\right)$ be the diagonal matrix with diagonal entries $z_{1}, \ldots, z_{d}$. Unless specified otherwise, a vector is by default a "column vector" and its transpose is a "row vector."
Theorem 5.1.1 (Spectral theorem). Let $A \in \mathbb{R}^{d \times d}$ be a symmetric matrix, that is, $A^{T}=A$. Then $A$ has $d$ orthonormal eigenvectors $\mathbf{q}_{1}, \ldots, \mathbf{q}_{d}$ with corresponding (not necessarily distinct) real eigenvalues $\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{d}$. In matrix form, this is written as the matrix factorization

$$
A=Q \Lambda Q^{T}=\sum_{i=1}^{d} \lambda_{i} \mathbf{q}_{i} \mathbf{q}_{i}^{T},
$$

where $Q$ has columns $\mathbf{q}_{1}, \ldots, \mathbf{q}_{d}$ and $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{d}\right)$. We refer to this factorization as a spectral decomposition of $A$.

The proof uses a greedy sequence maximizing the quadratic form $\langle\mathbf{v}, A \mathbf{v}\rangle$. For a hint as to why that might come about, note that for a unit eigenvector $\mathbf{v}$ with eigenvalue $\lambda$ we have $\langle\mathbf{v}, A \mathbf{v}\rangle=\langle\mathbf{v}, \lambda \mathbf{v}\rangle=\lambda$.

We will need the following formula. Consider the block matrices

$$
\binom{\mathbf{y}}{\mathbf{z}} \quad \text { and } \quad\left(\begin{array}{ll}
A & B \\
C & D
\end{array}\right)
$$

where $\mathbf{y} \in \mathbb{R}^{d_{1}}, \mathbf{z} \in \mathbb{R}^{d_{2}}, A \in \mathbb{R}^{d_{1} \times d_{1}}, B \in \mathbb{R}^{d_{1} \times d_{2}}, C \in \mathbb{R}^{d_{2} \times d_{1}}$, and $D \in$ $\mathbb{R}^{d_{2} \times d_{2}}$. Then it follows by direct calculation that

$$
\binom{\mathbf{y}}{\mathbf{z}}^{T}\left(\begin{array}{ll}
A & B  \tag{5.1.1}\\
C & D
\end{array}\right)\binom{\mathbf{y}}{\mathbf{z}}=\mathbf{y}^{T} A \mathbf{y}+\mathbf{y}^{T} B \mathbf{z}+\mathbf{z}^{T} C \mathbf{y}+\mathbf{z}^{T} D \mathbf{z} .
$$

We will also need the following linear algebra fact. Let $\mathbf{v}_{1}, \ldots, \mathbf{v}_{j}$ be orthonormal vectors in $\mathbb{R}^{d}$, with $j<d$. Then they can be completed into an orthonormal basis $\mathbf{v}_{1}, \ldots, \mathbf{v}_{d}$ of $\mathbb{R}^{d}$.

Proof of Theorem 5.1.1. We proceed by induction.

A first eigenvector Let $A_{1}=A$. Maximizing over the objective function $\left\langle\mathbf{v}, A_{1} \mathbf{v}\right\rangle$, we let

$$
\mathbf{v}_{1} \in \arg \max \left\{\left\langle\mathbf{v}, A_{1} \mathbf{v}\right\rangle:\|\mathbf{v}\|_{2}=1\right\},
$$

and

$$
\lambda_{1}=\max \left\{\left\langle\mathbf{v}, A_{1} \mathbf{v}\right\rangle:\|\mathbf{v}\|_{2}=1\right\} .
$$

Complete $\mathbf{v}_{1}$ into an orthonormal basis of $\mathbb{R}^{d}, \mathbf{v}_{1}, \hat{\mathbf{v}}_{2}, \ldots, \hat{\mathbf{v}}_{d}$, and form the block matrix

$$
\hat{W}_{1}:=\left(\begin{array}{ll}
\mathbf{v}_{1} & \hat{V}_{1}
\end{array}\right)
$$

where the columns of $\hat{V}_{1}$ are $\hat{\mathbf{v}}_{2}, \ldots, \hat{\mathbf{v}}_{d}$. Note that $\hat{W}_{1}$ is orthogonal by construction.

Getting one step closer to diagonalization We show next that $\hat{W}_{1}$ gets us one step closer to a diagonal matrix by similarity transformation. Note first that

$$
\hat{W}_{1}^{T} A_{1} \hat{W}_{1}=\left(\begin{array}{cc}
\lambda_{1} & \mathbf{w}_{1}^{T} \\
\mathbf{w}_{1} & A_{2}
\end{array}\right)
$$

where $\mathbf{w}_{1}:=\hat{V}_{1}^{T} A_{1} \mathbf{v}_{1}$ and $A_{2}:=\hat{V}_{1}^{T} A_{1} \hat{V}_{1}$. The key claim is that $\mathbf{w}_{1}=\mathbf{0}$. This follows from an argument by contradiction.

Suppose $\mathbf{w}_{1} \neq \mathbf{0}$ and consider the unit vector

$$
\mathbf{z}:=\hat{W}_{1} \times \frac{1}{\sqrt{1+\delta^{2}\left\|\mathbf{w}_{1}\right\|_{2}^{2}}}\binom{1}{\delta \mathbf{w}_{1}}
$$

which achieves objective value

$$
\begin{aligned}
\mathbf{z}^{T} A_{1} \mathbf{z} & =\frac{1}{1+\delta^{2}\left\|\mathbf{w}_{1}\right\|_{2}^{2}}\binom{1}{\delta \mathbf{w}_{1}}^{T}\left(\begin{array}{cc}
\lambda_{1} & \mathbf{w}_{1}^{T} \\
\mathbf{w}_{1} & A_{2}
\end{array}\right)\binom{1}{\delta \mathbf{w}_{1}} \\
& =\frac{1}{1+\delta^{2}\left\|\mathbf{w}_{1}\right\|_{2}^{2}}\left(\lambda_{1}+2 \delta\left\|\mathbf{w}_{1}\right\|_{2}^{2}+\delta^{2} \mathbf{w}_{1}^{T} A_{2} \mathbf{w}_{1}\right),
\end{aligned}
$$

where we used (5.1.1). By the Taylor expansion,

$$
\frac{1}{1+\epsilon^{2}}=1-\epsilon^{2}+O\left(\epsilon^{4}\right)
$$

for $\delta$ small enough,

$$
\begin{aligned}
\mathbf{z}^{T} A_{1} \mathbf{z} & =\left(\lambda_{1}+2 \delta\left\|\mathbf{w}_{1}\right\|_{2}^{2}+\delta^{2} \mathbf{w}_{1}^{T} A_{2} \mathbf{w}_{1}\right)\left(1-\delta^{2}\left\|\mathbf{w}_{1}\right\|_{2}^{2}+O\left(\delta^{4}\right)\right) \\
& =\lambda_{1}+2 \delta\left\|\mathbf{w}_{1}\right\|_{2}^{2}+O\left(\delta^{2}\right) \\
& >\lambda_{1} .
\end{aligned}
$$

That gives the desired contradiction.
So, letting $W_{1}:=\hat{W}_{1}$,

$$
W_{1}^{T} A_{1} W_{1}=\left(\begin{array}{cc}
\lambda_{1} & \mathbf{0} \\
\mathbf{0} & A_{2}
\end{array}\right) .
$$

Finally note that $A_{2}=\hat{V}_{1}^{T} A_{1} \hat{V}_{1}$ is symmetric since

$$
A_{2}^{T}=\left(\hat{V}_{1}^{T} A_{1} \hat{V}_{1}\right)^{T}=\hat{V}_{1}^{T} A_{1}^{T} \hat{V}_{1}=\hat{V}_{1}^{T} A_{1} \hat{V}_{1}=A_{2},
$$

by the symmetry of $A_{1}$ itself.
Next step of the induction Apply the same argument to the symmetric submatrix $A_{2} \in \mathbb{R}^{(d-1) \times(d-1)}$, let $\hat{W}_{2} \in \mathbb{R}^{(d-1) \times(d-1)}$ be the corresponding orthogonal matrix, and define $\lambda_{2}$ and $A_{3}$ through the equation

$$
\hat{W}_{2}^{T} A_{2} \hat{W}_{2}=\left(\begin{array}{cc}
\lambda_{2} & \mathbf{0} \\
\mathbf{0} & A_{3}
\end{array}\right) .
$$

Define the block matrix

$$
W_{2}=\left(\begin{array}{cc}
1 & \mathbf{0} \\
\mathbf{0} & \hat{W}_{2}
\end{array}\right)
$$

and observe that

$$
\begin{aligned}
W_{2}^{T} W_{1}^{T} A_{1} W_{1} W_{2} & =W_{2}^{T}\left(\begin{array}{cc}
\lambda_{1} & \mathbf{0} \\
\mathbf{0} & A_{2}
\end{array}\right) W_{2} \\
& =\left(\begin{array}{ccc}
\lambda_{1} & \mathbf{0} \\
\mathbf{0} & \hat{W}_{2}^{T} A_{2} \hat{W}_{2}
\end{array}\right) \\
& =\left(\begin{array}{ccc}
\lambda_{1} & 0 & \mathbf{0} \\
0 & \lambda_{2} & \mathbf{0} \\
\mathbf{0} & \mathbf{0} & A_{3}
\end{array}\right) .
\end{aligned}
$$

Proceeding similarly by induction gives the claim, with the final $Q$ being the product of the $W_{i} \mathrm{~s}$ (which is orthogonal as the product of orthogonal matrices).

We derive an important variational characterization inspired by the proof of the spectral theorem. We will need the following quantity.
Definition 5.1.2 (Rayleigh quotient). Let $A \in \mathbb{R}^{d \times d}$ be a symmetric matrix. The Rayleigh quotient of $A$ is defined as

Rayleigh
quotient

$$
\mathcal{R}_{A}(\mathbf{u})=\frac{\langle\mathbf{u}, A \mathbf{u}\rangle}{\langle\mathbf{u}, \mathbf{u}\rangle}
$$

which is defined for any $\mathbf{u} \neq \mathbf{0}$ in $\mathbb{R}^{d}$.

We let the span of a collection of vectors be defined as

$$
\operatorname{span}\left(\boldsymbol{u}_{\mathbf{1}}, \ldots, \boldsymbol{u}_{n}\right):=\left\{\sum_{i=1}^{n} \alpha_{i} \boldsymbol{u}_{i}: \alpha_{1}, \ldots, \alpha_{n} \in \mathbb{R}\right\} .
$$

Theorem 5.1.3 (Courant-Fischer theorem). Let $A \in \mathbb{R}^{d \times d}$ be a symmetric matrix with spectral decomposition $A=\sum_{i=1}^{d} \lambda_{i} \mathbf{v}_{i} \mathbf{v}_{i}^{T}$ where $\lambda_{1} \geq \cdots \geq \lambda_{d}$. For each $k=1, \ldots, d$, define the subspace

$$
\mathcal{V}_{k}=\operatorname{span}\left(\mathbf{v}_{1}, \ldots, \mathbf{v}_{k}\right) \quad \text { and } \quad \mathcal{W}_{d-k+1}=\operatorname{span}\left(\mathbf{v}_{k}, \ldots, \mathbf{v}_{d}\right) .
$$

Then, for all $k=1, \ldots, d$,

$$
\lambda_{k}=\min _{\mathbf{u} \in \mathcal{V}_{k}} \mathcal{R}_{A}(\mathbf{u})=\max _{\mathbf{u} \in \mathcal{W}_{d-k+1}} \mathcal{R}_{A}(\mathbf{u}) .
$$

Furthermore we have the following min-max formulas, which do not depend on the choice of spectral decomposition, for all $k=1, \ldots, d$

$$
\lambda_{k}=\max _{\operatorname{dim}(\mathcal{V})=k} \min _{\mathbf{u} \in \mathcal{V}} \mathcal{R}_{A}(\mathbf{u})=\min _{\operatorname{dim}(\mathcal{W})=d-k+1} \max _{\mathbf{u} \in \mathcal{W}} \mathcal{R}_{A}(\mathbf{u}) .
$$

Note that, in all these formulas, the vector $\mathbf{u}=\mathbf{v}_{k}$ is optimal. To derive the "local" formula, the first ones above, we expand a vector in $\mathcal{V}_{k}$ into the basis $\mathbf{v}_{1}, \ldots, \mathbf{v}_{k}$ and use the fact that $\mathcal{R}_{A}\left(\mathbf{v}_{i}\right)=\lambda_{i}$ and that eigenvalues are in nonincreasing order. The "global" formulas then follow from a dimension argument.

We will need the following dimension-based fact. Let $\mathcal{U}, \mathcal{V} \subseteq \mathbb{R}^{d}$ be linear subspaces such that $\operatorname{dim}(\mathcal{U})+\operatorname{dim}(\mathcal{V})>d$, where $\operatorname{dim}(\mathcal{U})$ denotes the dimension of $\mathcal{U}$. Then there exists a nonzero vector in the intersection $\mathcal{U} \cap \mathcal{V}$. That is,

$$
\begin{equation*}
\operatorname{dim}(\mathcal{U})+\operatorname{dim}(\mathcal{V})>d \Longrightarrow(\mathcal{U} \cap \mathcal{V}) \backslash\{\mathbf{0}\} \neq \emptyset \tag{5.1.2}
\end{equation*}
$$

Proof of Theorem 5.1.3. We first prove the local formulas, that is, the ones involving a specific decomposition.

Local formulas Since $\mathbf{v}_{1}, \ldots, \mathbf{v}_{k}$ form an orthonormal basis of $\mathcal{V}_{k}$, any nonzero vector $\mathbf{u} \in \mathcal{V}_{k}$ can be written as $\mathbf{u}=\sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle \mathbf{v}_{i}$ and it follows that

$$
\begin{gathered}
\langle\mathbf{u}, \mathbf{u}\rangle=\sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2} \\
\langle\mathbf{u}, A \mathbf{u}\rangle=\left\langle\mathbf{u}, \sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle \lambda_{i} \mathbf{v}_{i}\right\rangle=\sum_{i=1}^{k} \lambda_{i}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2} .
\end{gathered}
$$

Thus,

$$
\mathcal{R}_{A}(\mathbf{u})=\frac{\langle\mathbf{u}, A \mathbf{u}\rangle}{\langle\mathbf{u}, \mathbf{u}\rangle}=\frac{\sum_{i=1}^{k} \lambda_{i}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2}}{\sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2}} \geq \lambda_{k} \frac{\sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2}}{\sum_{i=1}^{k}\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2}}=\lambda_{k}
$$

where we used $\lambda_{1} \geq \cdots \geq \lambda_{k}$ and the fact that $\left\langle\mathbf{u}, \mathbf{v}_{i}\right\rangle^{2} \geq 0$. Moreover $\mathcal{R}_{A}\left(\mathbf{v}_{k}\right)=$ $\lambda_{k}$. So we have established

$$
\lambda_{k}=\min _{\mathbf{u} \in \mathcal{V}_{k}} \mathcal{R}_{A}(\mathbf{u}) .
$$

The expression in terms of $\mathcal{W}_{d-k+1}$ is proved similarly.
Global formulas Since $\mathcal{V}_{k}$ has dimension $k$, it follows from the local formula that

$$
\lambda_{k}=\min _{\mathbf{u} \in \mathcal{V}_{k}} \mathcal{R}_{A}(\mathbf{u}) \leq \max _{\operatorname{dim}(\mathcal{V})=k} \min _{\mathbf{u} \in \mathcal{V}} \mathcal{R}_{A}(\mathbf{u}) .
$$

Let $\mathcal{V}$ be any subspace with dimension $k$. Because $\mathcal{W}_{d-k+1}$ has dimension $d-k+1$, we have that $\operatorname{dim}(\mathcal{V})+\operatorname{dim}\left(\mathcal{W}_{d-k+1}\right)>d$ and there must be nonzero vector $\mathbf{u}_{0}$ in the intersection $\mathcal{V} \cap \mathcal{W}_{d-k+1}$ by the dimension-based fact above. We then have by the other local formula that

$$
\lambda_{k}=\max _{\mathbf{u} \in \mathcal{W}_{d-k+1}} \mathcal{R}_{A}(\mathbf{u}) \geq \mathcal{R}_{A}\left(\mathbf{u}_{0}\right) \geq \min _{\mathbf{u} \in \mathcal{V}} \mathcal{R}_{A}(\mathbf{u})
$$

Since this inequality holds for any subspace of dimension $k$, we have

$$
\lambda_{k} \geq \max _{\operatorname{dim}(\mathcal{V})=k} \min _{\mathbf{u} \in \mathcal{V}} \mathcal{R}_{A}(\mathbf{u})
$$

Combining with the inequality in the other direction above gives the claim. The other global formula is proved similarly.

### 5.1.2 Elements of spectral graph theory

We apply the variational characterization of eigenvalues to matrices arising in graph theory. In this section, graphs have no self-loop.

Unweighted graphs As we have previously seen, a convenient way of specifying a graph is through a matrix representation. Assume the undirected graph $G=$ $(V, E)$ has $n=|V|$ vertices. Recall that the adjacency matrix $A$ of $G$ is the $n \times n$ symmetric matrix defined as

$$
A_{x y}= \begin{cases}1 & \text { if }\{x, y\} \in E \\ 0 & \text { otherwise }\end{cases}
$$

Another matrix of interest is the Laplacian matrix. It is related to the Laplace operator we encountered previously. We will show in particular that it contains useful information about the connectedness of the graph. Recall that, given a graph $G=(V, E)$, the quantity $\delta(v)$ denotes the degree of $v \in V$.

Definition 5.1.4 (Graph Laplacian). Let $G=(V, E)$ be a graph with vertices $V=\{1, \ldots, n\}$ and adjacency matrix $A \in \mathbb{R}^{n \times n}$. Let $D=\operatorname{diag}(\delta(1), \ldots, \delta(n))$ be the degree matrix. The graph Laplacian (or Laplacian matrix, or Laplacian for short) associated to $G$ is defined as $L=D-A$. Its entries are

$$
l_{i j}= \begin{cases}\delta(i) & \text { if } i=j, \\ -1 & \text { if }\{i, j\} \in E \\ 0 & \text { otherwise }\end{cases}
$$

Observe that the Laplacian $L$ of a graph $G$ is a symmetric matrix:

$$
L^{T}=(D-A)^{T}=D^{T}-A^{T}=D-A,
$$

where we used that both $D$ and $A$ are themselves symmetric. The associated quadratic form is particularly simple and will play an important role.

Lemma 5.1.5 (Laplacian quadratic form). Let $G=(V, E)$ be a graph with $n=$ $|V|$ vertices. Its Laplacian $L$ is a positive semi-definite matrix and furthermore we have the following formula for the Laplacian quadratic form (or Dirichlet energy)

$$
\mathbf{x}^{T} L \mathbf{x}=\sum_{e=\{i, j\} \in E}\left(x_{i}-x_{j}\right)^{2}
$$

Laplacian
quadratic
form
for any $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}$.
Proof of Lemma 5.1.5. Let $B$ be an oriented incidence matrix of $G$ (see Definition 1.1.16). We claim that $B B^{T}=L$. Indeed, for $i \neq j$, entry $(i, j)$ of $B B^{T}$ is a sum over all edges containing $i$ and $j$ as endvertices, of which there is at most one. When $e=\{i, j\} \in E$, that entry is -1 , since one of $i$ or $j$ has a 1 in the column of $B$ corresponding to $e$ and the other one has a -1 . For $i=j$, letting $b_{x y}$ be entry $(x, y)$ of $B$,

$$
\left(B B^{T}\right)_{i i}=\sum_{e=\{x, y\} \in E: i \in e} b_{x y}^{2}=\delta(i) .
$$

That shows that $B B^{T}=L$ entry-by-entry.
For any $\mathbf{x}$, we have $\left(B^{T} \mathbf{x}\right)_{k}=x_{v}-x_{u}$ if the edge $e_{k}=\{u, v\}$ is oriented as $(u, v)$ under $B$. That implies

$$
\mathbf{x}^{T} L \mathbf{x}=\mathbf{x}^{T} B B^{T} \mathbf{x}=\left\|B^{T} \mathbf{x}\right\|_{2}^{2}=\sum_{e=\{i, j\} \in E}\left(x_{i}-x_{j}\right)^{2} .
$$

Since the latter quantity is always nonnegative, it also implies that $L$ is positive semidefinite.

As a convention, we denote the eigenvalues of a Laplacian matrix $L$ by

$$
0 \leq \mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{n},
$$

and we will refer to them as Laplacian eigenvalues. Here is a simple observation. For any $G=(V, E)$, the constant unit vector

$$
\mathbf{y}_{1}=\frac{1}{\sqrt{n}}(1, \ldots, 1),
$$

is an eigenvector of the Laplacian with eigenvalue 0 . Indeed, let $B$ be an oriented incidence matrix of $G$ and recall from the proof of Lemma 5.1.5 that $L=B B^{T}$. By construction $B^{T} \mathbf{y}_{1}=\mathbf{0}$ since each column of $B$ has exactly one 1 and one -1 . So $L \mathbf{y}_{1}=B B^{T} \mathbf{y}_{1}=\mathbf{0}$ as claimed. In general, the constant vector may not be the only eigenvector with eigenvalue one.

We are now ready to derive connectivity consequences. Recall that, for any graph $G$, the Laplacian eigenvalue $\mu_{1}=0$.

Lemma 5.1.6 (Laplacian and connectivity). If $G$ is connected, then the Laplacian eigenvalue $\mu_{2}>0$.

Proof. Let $G=(V, E)$ with $n=|V|$ and let $L=\sum_{i=1}^{n} \mu_{i} \mathbf{y}_{i} \mathbf{y}_{i}^{T}$ be a spectral decomposition of its Laplacian $L$ with $0=\mu_{1} \leq \cdots \leq \mu_{n}$. Suppose by way of contradiction that $\mu_{2}=0$. Any eigenvector $\mathbf{y}=\left(y_{1}, \ldots, y_{n}\right)$ with 0 eigenvalue satisfies $L \mathbf{y}=\mathbf{0}$ by definition. By Lemma 5.1.5 then

$$
0=\mathbf{y}^{T} L \mathbf{y}=\sum_{e=\{i, j\} \in E}\left(y_{i}-y_{j}\right)^{2}
$$

In order for this to hold, it must be that any two adjacent vertices $i$ and $j$ have $y_{i}=y_{j}$. That is, $\{i, j\} \in E$ implies $y_{i}=y_{j}$. Furthermore, because $G$ is connected, between any two of its vertices $u$ and $v$ (adjacent or not) there is a path $u=w_{0} \sim \cdots \sim w_{k}=v$ along which the $y_{w} \mathrm{~s}$ must be the same. Thus $\mathbf{y}$ is a constant vector.

But that is a contradiction since the eigenvectors $\mathbf{y}_{1}, \ldots, \mathbf{y}_{n}$ are in fact linearly independent, so that $\mathbf{y}_{1}$ and $\mathbf{y}_{2}$ cannot both be a constant vector.

The quantity $\mu_{2}$ is sometimes referred to as the algebraic connectivity of the graph. The corresponding eigenvector, $\mathbf{y}_{2}$, is known as the Fiedler vector.

We will be interested in more quantitative results of this type. Before proceeding, we start with a simple observation. By our proof of Theorem 5.1.1, the largest eigenvalue $\mu_{n}$ of the Laplacian $L$ is the solution to the optimization problem

$$
\mu_{n}=\max \left\{\langle\mathbf{x}, L \mathbf{x}\rangle:\|\mathbf{x}\|_{2}=1\right\}
$$

Such extremal characterization is useful in order to bound the eigenvalue $\mu_{n}$, since any choice of $\mathbf{x}$ with $\|\mathbf{x}\|_{2}=1$ gives a lower bound through the quantity $\langle\mathbf{x}, L \mathbf{x}\rangle$. We give a simple consequence.

Lemma 5.1.7 (Laplacian and degree). Let $G=(V, E)$ be a graph with maximum degree $\bar{\delta}$. Let $\mu_{n}$ be the largest Laplacian eigenvalue. Then

$$
\mu_{n} \geq \bar{\delta}+1
$$

Proof. Let $u \in V$ be a vertex with degree $\bar{\delta}$. Let $\mathbf{z}$ be the vector with entries

$$
z_{i}= \begin{cases}\bar{\delta} & \text { if } i=u \\ -1 & \text { if }\{i, u\} \in E \\ 0 & \text { otherwise }\end{cases}
$$

and let $\mathbf{x}$ be the unit vector $\mathbf{z} /\|\mathbf{z}\|_{2}$. By definition of the degree of $u,\|\mathbf{z}\|_{2}^{2}=$ $\bar{\delta}^{2}+\bar{\delta}(-1)^{2}=\bar{\delta}(\bar{\delta}+1)$.

Using the Lemma 5.1.5,

$$
\begin{aligned}
\langle\mathbf{z}, L \mathbf{z}\rangle & =\sum_{e=\{i, j\} \in E}\left(z_{i}-z_{j}\right)^{2} \\
& \geq \sum_{i:\{i, u\} \in E}\left(z_{i}-z_{u}\right)^{2} \\
& =\sum_{i:\{i, u\} \in E}(-1-\bar{\delta})^{2} \\
& =\bar{\delta}(\bar{\delta}+1)^{2},
\end{aligned}
$$

where we restricted the sum to those edges incident with $u$ and used the fact that all terms in the sum are nonnegative. Finally

$$
\langle\mathbf{x}, L \mathbf{x}\rangle=\left\langle\frac{\mathbf{z}}{\|\mathbf{z}\|_{2}}, L \frac{\mathbf{z}}{\|\mathbf{z}\|_{2}}\right\rangle=\frac{1}{\|\mathbf{z}\|_{2}^{2}}\langle\mathbf{z}, L \mathbf{z}\rangle=\frac{\bar{\delta}(\bar{\delta}+1)^{2}}{\bar{\delta}(\bar{\delta}+1)}=\bar{\delta}+1,
$$

so that

$$
\mu_{n}=\max \left\{\left\langle\mathbf{x}^{\prime}, L \mathbf{x}^{\prime}\right\rangle:\left\|\mathbf{x}^{\prime}\right\|_{2}=1\right\} \geq\langle\mathbf{x}, L \mathbf{x}\rangle=\bar{\delta}+1,
$$

as claimed.

A special case of Courant-Fischer (Theorem 5.1.3) for the Laplacian matrix is the following.

Corollary 5.1.8 (Variational characterization of $\mu_{2}$ ). Let $G=(V, E)$ be a graph with $n=|V|$ vertices. Assume the Laplacian $L$ of $G$ has spectral decomposition $L=\sum_{i=1}^{n} \mu_{i} \mathbf{y}_{i} \mathbf{y}_{i}^{T}$ with $0=\mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{n}$ and $\mathbf{y}_{1}=\frac{1}{\sqrt{n}}(1, \ldots, 1)$. Then

$$
\mu_{2}=\min \left\{\frac{\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}}{\sum_{u=1}^{n} x_{u}^{2}}: \mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \neq \mathbf{0}, \sum_{u=1}^{n} x_{u}=0\right\}
$$

Proof. By Theorem 5.1.3,

$$
\mu_{2}=\min _{\mathbf{x} \in \mathcal{W}_{d-1}} \mathcal{R}_{L}(\mathbf{x}) .
$$

Since $\mathbf{y}_{1}$ is constant and $\mathcal{W}_{d-1}$ is the subspace orthogonal to it, this is equivalent to restrictring the minimization to those nonzero xs such that

$$
0=\left\langle\mathbf{x}, \mathbf{y}_{1}\right\rangle=\frac{1}{\sqrt{n}} \sum_{u=1}^{m} x_{u} .
$$

Moreover, by Lemma 5.1.5,

$$
\langle\mathbf{x}, L \mathbf{x}\rangle=\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}
$$

so the Rayleigh quotient is

$$
\mathcal{R}_{L}(\mathbf{x})=\frac{\langle\mathbf{x}, L \mathbf{x}\rangle}{\langle\mathbf{x}, \mathbf{x}\rangle}=\frac{\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}}{\sum_{u=1}^{n} x_{u}^{2}} .
$$

That proves the claim.
One application of this extremal characterization is a graph drawing heuristic. Consider the entries of the second Laplacian eigenvector $\mathbf{y}_{2}$ normalized to have unit norm. The entries are centered around 0 by the condition $\sum_{u=1}^{n} x_{u}=0$. Because it minimizes the quantity

$$
\frac{\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}}{\sum_{u=1}^{n} x_{u}^{2}},
$$

over all centered unit vectors, $\mathbf{y}_{2}$ tends to assign similar coordinates to adjacent vertices. A similar reasoning applies to the third Laplacian eigenvector, which in addition is orthogonal to the second one. See Figure 5.1 for an illustration.


Figure 5.1: Top: A 3-by-3 grid graph with vertices located at independent uniformly random points in a square. Bottom: The same 3 -by- 3 grid graph with vertices located at the coordinates corresponding to the second and third eigenvectors of the Laplacian matrix. That is, vertex $i$ is located at position $\left(y_{2, i}, y_{3, i}\right)$.

Example 5.1.9 (Two-component graph). Let $G=(V, E)$ be a graph with two connected components $\emptyset \neq V_{1}, V_{2} \subseteq V$. By the properties of connected components, we have $V_{1} \cap V_{2}=\emptyset$ and $V_{1} \cup V_{2}=V$. Assume the Laplacian $L$ of $G$ has spectral decomposition $L=\sum_{i=1}^{n} \mu_{i} \mathbf{y}_{i} \mathbf{y}_{i}^{T}$ with $0=\mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{n}$ and $\mathbf{y}_{1}=\frac{1}{\sqrt{n}}(1, \ldots, 1)$. We claimed earlier that for such a graph $\mu_{2}=0$. We prove this here using Corollary 5.1.8:

$$
\begin{aligned}
\mu_{2}=\min \{ & \sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}: \\
& \left.\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}, \sum_{u=1}^{n} x_{u}=0, \sum_{u=1}^{n} x_{u}^{2}=1\right\} .
\end{aligned}
$$

Based on this characterization, it suffices to find a vector $\mathbf{x}$ satisfying $\sum_{u=1}^{n} x_{u}=0$ and $\sum_{u=1}^{n} x_{u}^{2}=1$ such that $\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}=0$. Indeed, since $\mu_{2} \geq 0$ and any such x gives an upper bound on $\mu_{2}$, we then necessarily have that $\mu_{2}=0$.

For $\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}$ to be 0 , one might be tempted to take a constant vector $\mathbf{x}$. But then we could not satisfy $\sum_{u=1}^{n} x_{u}=0$ and $\sum_{u=1}^{n} x_{u}^{2}=1$. Instead, we modify this guess slightly. Because the graph has two connected components, there is no edge between $V_{1}$ and $V_{2}$. Hence we can assign a different value to each component and still get $\sum_{\{u, v\} \in E}\left(x_{u}-x_{v}\right)^{2}=0$. So we look for a vector $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right)$ of the form

$$
x_{u}= \begin{cases}\alpha & \text { if } u \in V_{1}, \\ \beta & \text { if } u \in V_{2} .\end{cases}
$$

To satisfy the constraints on $\mathbf{x}$, we require

$$
\sum_{u=1}^{n} x_{u}=\sum_{u \in V_{1}} \alpha+\sum_{u \in V_{2}} \beta=\left|V_{1}\right| \alpha+\left|V_{2}\right| \beta=0
$$

and

$$
\sum_{u=1}^{n} x_{u}^{2}=\sum_{u \in V_{1}} \alpha^{2}+\sum_{u \in V_{2}} \beta^{2}=\left|V_{1}\right| \alpha^{2}+\left|V_{2}\right| \beta^{2}=1
$$

Replacing the first equation in the second one, we get

$$
\left|V_{1}\right|\left(\frac{-\left|V_{2}\right| \beta}{\left|V_{1}\right|}\right)^{2}+\left|V_{2}\right| \beta^{2}=\frac{\left|V_{2}\right|^{2} \beta^{2}}{\left|V_{1}\right|}+\left|V_{2}\right| \beta^{2}=1
$$

or

$$
\beta^{2}=\frac{\left|V_{1}\right|}{\left|V_{2}\right|\left(\left|V_{2}\right|+\left|V_{1}\right|\right)}=\frac{\left|V_{1}\right|}{n\left|V_{2}\right|} .
$$

Take

$$
\beta=-\sqrt{\frac{\left|V_{1}\right|}{n\left|V_{2}\right|}}, \quad \alpha=\frac{-\left|V_{2}\right| \beta}{\left|V_{1}\right|}=\sqrt{\frac{\left|V_{2}\right|}{n\left|V_{1}\right|}} .
$$

The vector $\mathbf{x}$ we constructed is in fact an eigenvector of $L$. Indeed, let $B$ be an oriented incidence matrix of $G$. Then, for $e_{k}=\{u, v\},\left(B^{T} \mathbf{x}\right)_{k}$ is either $x_{u}-x_{v}$ or $x_{v}-x_{u}$. In both cases, that is 0 . So $L \mathbf{x}=B B^{T} \mathbf{x}=\mathbf{0}$, that is, $\mathbf{x}$ is an eigenvector of $L$ with eigenvalue 0 .

We have shown that $\mu_{2}=0$ when $G$ has two connected components. A slight modification of this argument shows that $\mu_{2}=0$ whenever $G$ is not connected.

Networks In the case of a network (i.e., edge-weighted graph) $G=(V, E, w)$, the Laplacian can be defined as follows. As usual, we assume that $w: E \rightarrow \mathbb{R}_{+}$is a function that assigns positive real weights to the edges. We write $w_{e}=w_{i j}$ for the weight of edge $e=\{i, j\}$. Recall that the degree of a vertex $i$ is,

$$
\delta(i)=\sum_{j:\{i, j\} \in E} w_{i j},
$$

the adjacency matrix $A$ of $G$ is the $n \times n$ symmetric matrix defined as

$$
A_{i j}= \begin{cases}w_{i j} & \text { if }\{i, j\} \in E, \\ 0 & \text { otherwise }\end{cases}
$$

Definition 5.1.10 (Network Laplacian). Let $G=(V, E, w)$ be a network with $n=|V|$ vertices and adjacency matrix $A$. Let $D=\operatorname{diag}(\delta(1), \ldots, \delta(n))$ be the degree matrix. The network Laplacian (or Laplacian matrix, or Laplacian for short) associated to $G$ is defined as $L=D-A$.

It can be shown (see Exercise 5.2) that the Laplacian quadratic form satisfies in the edge-weighted case

$$
\begin{equation*}
\langle\mathbf{x}, L \mathbf{x}\rangle=\sum_{\{i, j\} \in E} w_{i j}\left(x_{i}-x_{j}\right)^{2}, \tag{5.1.3}
\end{equation*}
$$

for $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}$. (The keen observer will have noticed that we already encountered this quantity as the "Dirichlet energy" in Section 3.3.3; more on this in Section 5.3.) As a positive semidefinite matrix (see again Exercise 5.2), the network Laplacian has an orthonormal basis of eigenvectors with nonnegative eigenvalues that satisfy the variational characterization we derived above. In particular, if we denote the eigenvalues $0=\mu_{1} \leq \mu_{2} \leq \cdots \leq \mu_{n}$, it follows from

Courant-Fischer (Theorem 5.1.3) that

$$
\begin{aligned}
\mu_{2}=\min \{ & \sum_{\{u, v\} \in E} w_{u v}\left(x_{u}-x_{v}\right)^{2}: \\
& \left.\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right)^{T} \in \mathbb{R}^{n}, \sum_{u=1}^{n} x_{u}=0, \sum_{u=1}^{n} x_{u}^{2}=1\right\} .
\end{aligned}
$$

Other variants of the Laplacian are useful. We introduce the normalized Laplacian next.

Definition 5.1.11 (Normalized Laplacian). The normalized Laplacian of $G=(V, E, w)_{\text {normalized }}$ with adjacency matrix $A$ and degree matrix $D$ is defined as

$$
\mathcal{L}=I-D^{-1 / 2} A D^{-1 / 2} .
$$

The entries of $\mathcal{L}$ are

$$
\mathcal{L}_{i, j}= \begin{cases}1 & \text { if } i=j \\ -\frac{w_{i j}}{\sqrt{\delta(i) \delta(j)}} & \text { otherwise }\end{cases}
$$

We also note the following relation to the (unnormalized) Laplacian:

$$
\begin{equation*}
\mathcal{L}=D^{-1 / 2} L D^{-1 / 2} . \tag{5.1.4}
\end{equation*}
$$

We check that the normalized Laplacian is symmetric:

$$
\begin{aligned}
\mathcal{L}^{T} & =I^{T}-\left(D^{-1 / 2} A D^{-1 / 2}\right)^{T} \\
& =I-\left(D^{-1 / 2}\right)^{T} A^{T}\left(D^{-1 / 2}\right)^{T} \\
& =I-D^{-1 / 2} A D^{-1 / 2} \\
& =\mathcal{L} .
\end{aligned}
$$

It is also positive semidefinite. Indeed,

$$
\mathbf{x}^{T} \mathcal{L} \mathbf{x}=\mathbf{x}^{T} D^{-1 / 2} L D^{-1 / 2} \mathbf{x}=\left(D^{-1 / 2} \mathbf{x}\right)^{T} L\left(D^{-1 / 2} \mathbf{x}\right) \geq 0
$$

by the properties of the Laplacian. Hence by the spectral theorem (Theorem 5.1.1), we can write

$$
\mathcal{L}=\sum_{i=1}^{n} \eta_{i} \mathbf{z}_{i} \mathbf{z}_{i}^{T},
$$

where the $\mathbf{z}_{i} \mathrm{~s}$ are orthonormal eigenvectors of $\mathcal{L}$ and the eigenvalues satisfy

$$
0 \leq \eta_{1} \leq \eta_{2} \leq \cdots \leq \eta_{n}
$$

One more observation: because the constant vector is an eigenvector of $L$ with eigenvalue 0 , we get from (5.1.4) that $D^{1 / 2} \mathbf{1}$ is an eigenvector of $\mathcal{L}$ with eigenvalue 0 . So $\eta_{1}=0$ and we set

$$
\left(\mathbf{z}_{1}\right)_{i}=\left(\frac{D^{1 / 2} \mathbf{1}}{\left\|D^{1 / 2} \mathbf{1}\right\|_{2}}\right)_{i}=\sqrt{\frac{\delta(i)}{\sum_{i \in V} \delta(i)}}, \quad \forall i \in[n]
$$

which makes $\mathbf{z}_{1}$ into a unit norm vector. The relationship to the Laplacian implies (see Exercise 5.3) that

$$
\mathbf{x}^{T} \mathcal{L} \mathbf{x}=\sum_{\{i, j\} \in E} w_{i j}\left(\frac{x_{i}}{\sqrt{\delta(i)}}-\frac{x_{j}}{\sqrt{\delta(j)}}\right)^{2}
$$

for $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}$. Through the change of variables

$$
y_{i}=\frac{x_{i}}{\sqrt{\delta(i)}}
$$

Courant-Fischer (Theorem 5.1.3) gives this time

$$
\begin{align*}
\eta_{2}=\min \{ & \sum_{\{u, v\} \in E} w_{u v}\left(y_{u}-y_{v}\right)^{2}: \\
& \left.\mathbf{y}=\left(y_{1}, \ldots, y_{n}\right) \in \mathbb{R}^{n}, \sum_{u=1}^{n} \delta(u) y_{u}=0, \sum_{u=1}^{n} \delta(u) y_{u}^{2}=1\right\} . \tag{5.1.5}
\end{align*}
$$

### 5.1.3 Perturbation results

We will need some perturbation results for eigenvalues and eigenvectors. Recall the following definition. Define $\mathbb{S}^{m-1}=\left\{\mathbf{x} \in \mathbb{R}^{m}:\|\mathrm{x}\|_{2}=1\right\}$. The spectral norm (or induced 2-norm or 2-norm) of a matrix $A \in \mathbb{R}^{n \times m}$ is

$$
\|A\|_{2}:=\max _{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^{m}} \frac{\|A \mathbf{x}\|}{\|\mathbf{x}\|}=\max _{\mathbf{x} \in \mathbb{S}^{m-1}}\|A \mathbf{x}\| .
$$

The induced 2-norm of a matrix has many other useful properties.

Lemma 5.1.12 (Properties of the induced norm). Let $A, B \in \mathbb{R}^{n \times m}$ and $\alpha \in \mathbb{R}$. The following hold:
(i) $\|A \mathrm{x}\|_{2} \leq\|A\|_{2}\|\mathbf{x}\|_{2}, \forall \mathbf{x} \in \mathbb{R}^{m}$
(ii) $\|A\|_{2} \geq 0$
(iii) $\|A\|_{2}=0$ if and only if $A=0$
(iv) $\|\alpha A\|_{2}=|\alpha|\|A\|_{2}$
(v) $\|A+B\|_{2} \leq\|A\|_{2}+\|B\|_{2}$
(vi) $\|A B\|_{2} \leq\|A\|_{2}\|B\|_{2}$.

Proof. These properties all follow from the definition of the induced norm and the corresponding properties for the vector norm:

- Claims (i) and (ii) are immediate from the definition.
- For (ii) note that $\|A\|_{2}=0$ implies $\|A \mathbf{x}\|_{2}=0, \forall \mathbf{x} \in \mathbb{S}^{m-1}$, so that $A \mathbf{x}=$ $\mathbf{0}, \forall \mathbf{x} \in \mathbb{S}^{m-1}$. In particular, $A_{i j}=\mathbf{e}_{i}^{T} A \mathbf{e}_{j}=0, \forall i, j$.
- For (iv), (v), (vi), observe that for all $\mathrm{x} \in \mathbb{S}^{m-1}$

$$
\begin{gathered}
\|\alpha A \mathbf{x}\|_{2}=|\alpha|\|A \mathbf{x}\|_{2} \\
\|(A+B) \mathbf{x}\|_{2}=\|A \mathbf{x}+B \mathbf{x}\|_{2} \leq\|A \mathbf{x}\|_{2}+\|B \mathbf{x}\|_{2} \leq\|A\|_{2}+\|B\|_{2} \\
\|(A B) \mathbf{x}\|_{2}=\|A(B \mathbf{x})\|_{2} \leq\|A\|_{2}\|B \mathbf{x}\|_{2} \leq\|A\|_{2}\|B\|_{2} .
\end{gathered}
$$

Perturbations of eigenvalues For a symmetric matrix $C \in \mathbb{R}^{d \times d}$, we let $\lambda_{j}(C)$, $j=1, \ldots, d$, be the eigenvalues of $C$ in nonincreasing order with corresponding orthonormal eigenvectors $\mathbf{v}_{j}(C), j=1, \ldots, d$. As in the Courant-Fischer theorem (Theorem 5.1.3), define the subspaces

$$
\mathcal{V}_{k}(C)=\operatorname{span}\left(\mathbf{v}_{1}(C), \ldots, \mathbf{v}_{k}(C)\right)
$$

and

$$
\mathcal{W}_{d-k+1}(C)=\operatorname{span}\left(\mathbf{v}_{k}(C), \ldots, \mathbf{v}_{d}(C)\right) .
$$

The following lemma is one version of what is known as Weyl's inequality.

Lemma 5.1.13 (Weyl's inequality). Let $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times d}$ be symmetric matrices. Then, for all $j=1, \ldots, d$,

$$
\max _{j \in[d]}\left|\lambda_{j}(B)-\lambda_{j}(A)\right| \leq\|B-A\|_{2} .
$$

Proof. Let $H=B-A$. We prove only one upper bound. The other one follows from interchanging the roles of $A$ and $B$. Because

$$
\operatorname{dim}\left(\mathcal{V}_{j}(B)\right)+\operatorname{dim}\left(\mathcal{W}_{d-j+1}(A)\right)=j+(d-j+1)=d+1>d
$$

it follows from (5.1.2) that $\mathcal{V}_{j}(B) \cap \mathcal{W}_{d-j+1}(A)$ contains a nonzero vector. Let $\mathbf{v}$ be a unit vector in that intersection.

By Theorem 5.1.3,

$$
\lambda_{j}(B) \leq\langle\mathbf{v},(A+H) \mathbf{v}\rangle=\langle\mathbf{v}, A \mathbf{v}\rangle+\langle\mathbf{v}, H \mathbf{v}\rangle \leq \lambda_{j}(A)+\langle\mathbf{v}, H \mathbf{v}\rangle .
$$

Moreover, by Cauchy-Schwarz (Theorem B.4.8), since $\|\mathbf{v}\|_{2}=1$

$$
\langle\mathbf{v}, H \mathbf{v}\rangle \leq\|\mathbf{v}\|_{2}\|H \mathbf{v}\|_{2} \leq\|H\|_{2},
$$

which proves the claim after rearranging.
Perturbations of eigenvectors While Weyl's inequality (Lemma 5.1.13) indicates that the eigenvalues of $A$ and $B$ are close when $\|A-B\|_{2}$ is small, it says nothing about the eigenvectors. The following theorem remediates that. It is traditionally stated in terms of the angle between the eigenvectors (whereby the name). Here we give a version that is more suited to the applications we will encounter. We do not optimize the constants. We use the same notation as in the previous paragraph. Recall Parseval's identity: if $\boldsymbol{u}_{1}, \ldots, \boldsymbol{u}_{d}$ is an orthonormal basis of $\mathbb{R}^{d}$, then $\|\boldsymbol{x}\|^{2}=\sum_{i=1}^{d}\left\langle\boldsymbol{x}, \boldsymbol{u}_{i}\right\rangle^{2}$.

Theorem 5.1.14 (Davis-Kahan $\sin \theta$ theorem). Let $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times d}$ be symmetric matrices. For an $i \in\{1, \ldots, d\}$, assume that

$$
\delta:=\min _{j \neq i}\left|\lambda_{i}(A)-\lambda_{j}(A)\right|>0 .
$$

Then

$$
\min _{s \in\{+1,-1\}}\left\|\mathbf{v}_{i}(A)-s \mathbf{v}_{i}(B)\right\|_{2}^{2} \leq \frac{8\|A-B\|_{2}^{2}}{\delta^{2}}
$$

Proof. Expand $\mathbf{v}_{i}(B)$ in the basis formed by the eigenvectors of $A$, that is,

$$
\mathbf{v}_{i}(B)=\sum_{j=1}^{d}\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{j}(A)\right\rangle \mathbf{v}_{j}(A),
$$

where we used the orthonormality of the $\mathbf{v}_{j}(A)$ s. On the one hand,

$$
\begin{aligned}
& \left\|\left(A-\lambda_{i}(A) I_{d \times d}\right) \mathbf{v}_{i}(B)\right\|_{2}^{2} \\
& =\left\|\sum_{j=1}^{d}\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{j}(A)\right\rangle\left(A-\lambda_{i}(A) I_{d \times d}\right) \mathbf{v}_{j}(A)\right\|_{2}^{2} \\
& =\left\|\sum_{j=1, j \neq i}^{d}\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{j}(A)\right\rangle\left(\lambda_{j}(A)-\lambda_{i}(A)\right) \mathbf{v}_{j}(A)\right\|_{2}^{2} \\
& =\sum_{j=1, j \neq i}^{d}\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{j}(A)\right\rangle^{2}\left(\lambda_{j}(A)-\lambda_{i}(A)\right)^{2} \\
& \geq \delta^{2}\left(1-\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{i}(A)\right\rangle^{2}\right),
\end{aligned}
$$

where, on the last two lines, we used the orthonormality of the $\mathbf{v}_{j}(A) \mathrm{s}$ and $\mathbf{v}_{j}(B) \mathrm{s}$ through Parseval's identity, as well as the definition of $\delta$.

On the other hand, letting $E=A-B$, by the triangle inequality

$$
\begin{aligned}
\left\|\left(A-\lambda_{i}(A) I\right) \mathbf{v}_{i}(B)\right\|_{2} & =\left\|\left(B+E-\lambda_{i}(A) I\right) \mathbf{v}_{i}(B)\right\|_{2} \\
& \leq\left\|\left(B-\lambda_{i}(A) I\right) \mathbf{v}_{i}(B)\right\|_{2}+\left\|E \mathbf{v}_{i}(B)\right\|_{2} \\
& \leq\left|\lambda_{i}(B)-\lambda_{i}(A)\right|\left\|\mathbf{v}_{i}(B)\right\|_{2}+\|E\|_{2}\left\|\mathbf{v}_{i}(B)\right\|_{2} \\
& =2\|E\|_{2},
\end{aligned}
$$

where we used Lemma 5.1.12 and Weyl's inequality.
Combining the last two inequalities gives

$$
\left(1-\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{i}(A)\right\rangle^{2}\right) \leq \frac{4\|E\|_{2}^{2}}{\delta^{2}}
$$

The result follows by noting that, since $\left|\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{i}(A)\right\rangle\right| \leq 1$ by Cauchy-Schwarz (Theorem B.4.8), we have

$$
\begin{aligned}
\min _{s \in\{+1,-1\}}\left\|\mathbf{v}_{i}(A)-s \mathbf{v}_{i}(B)\right\|^{2} & =2-2\left|\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{i}(A)\right\rangle\right| \\
& \leq 2\left(1-\left\langle\mathbf{v}_{i}(B), \mathbf{v}_{i}(A)\right\rangle^{2}\right) \\
& \leq \frac{8\|E\|_{2}^{2}}{\delta^{2}}
\end{aligned}
$$

### 5.1.4 $\triangleright$ Data science: community recovery

A common task in network analysis is to recover hidden community structure. Informally, we seek groups of vertices with more edges within the groups then to the rest of the graph. More rigorously, providing statistical guarantees on the output of a community recovery algorithm requires some underlying random graph model. The standard model for this purpose is the stochastic blockmodel, a generalization of the Erdös-Rényi graph model with a "planted partition."

Stochastic blockmodel and recovery requirement We restrict ourselves to the simple case of two strictly balanced communities. Consider a random graph on $n$ (even) nodes where there are two communities, labeled +1 and -1 , consisting of $n / 2$ nodes. Each vertex $i \in V$ is assigned a community label $X_{i} \in\{1,-1\}$ as follows: a subset of $n / 2$ vertices is chosen uniformly at random among all such subsets to form community +1 , and the rest of the vertices form community -1 . For two nodes $i, j$, the edge $\{i, j\}$ is present with probability $p$ if they belong to the same community, and with probability $q$ otherwise. All edges are independent. The following $2 \times 2$ matrix describes the edge density within and across the two communities:

$$
\left.W=\begin{array}{c} 
\\
+1 \\
-1
\end{array} \begin{array}{cc}
+1 & -1 \\
p & q \\
q & p
\end{array}\right] .
$$

We assume that $p \geq q$, encoding the fact that vertices belonging to the same community are more likely to share an edge. To summarize, we say that $(X, G) \sim$ $\mathrm{SBM}_{n, p, q}$ if:

1. (Communities) The assignment $X=\left(X_{1}, \ldots, X_{n}\right)$ is uniformly random blockmodel over

$$
\Pi_{2}^{n}:=\left\{\mathbf{x} \in\{+1,-1\}^{n}: \mathbf{x}^{T} \mathbf{1}=0\right\}
$$

where $\mathbf{1}=(1, \cdots, 1)$ is the all-one vector.
2. (Graph) Conditioned on $X$, the graph $G=([n], E)$ has independent edges where $\{i, j\}$ is present with probability $W_{X_{i}, X_{j}}$ for $\forall i<j$.

We denote the corresponding measure by $\mathbb{P}_{n, p, q}$. We allow $p$ and $q$ to depend on $n$ (although we do not make that dependence explicit).

Roughly speaking, the community recovery problem is the following: given $G$, output $X$. There are different notions of recovery.

Definition 5.1.15 (Agreement). The agreement between two community assignment vectors $\mathbf{x}, \mathbf{y} \in\{+1,-1\}^{n}$ is the largest fraction of common assignments between x and $\pm \mathbf{y}$, that is,

$$
\alpha(\mathbf{x}, \mathbf{y})=\max _{s \in\{+1,-1\}} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\left\{x_{i}=s y_{i}\right\} .
$$

The role of $s$ in this formula is to account for the fact that the community names are not meaningful.

Now consider the following recovery requirements. These are asymptotic notions, as $n \rightarrow+\infty$.

Definition 5.1.16 (Recovery requirement). Let $(X, G) \sim \mathrm{SBM}_{n, p, q}$. For any estimator $\hat{X}:=\hat{X}(G) \in \Pi_{2}^{n}$, we say that it achieves:

- exact recovery if $\mathbb{P}_{n, p, q}[\alpha(X, \hat{X})=1]=1-o(1)$; or
- almost exact recovery if $\mathbb{P}_{n, p, q}[\alpha(X, \hat{X})=1-o(1)]=1-o(1)$.

Next we establish sufficient conditions for almost exact recovery. First we describe a natural estimator $\hat{X}$.

MAP estimator and spectral clustering A natural starting point is the maximum a posteriori (MAP) estimator. Let $\Omega(X)$ be the balanced partition of $[n]$ corresponding to $X$ and $\hat{\Omega}(G)$ be the one corresponding to $\hat{X}(G)$. The probability of error, that is, the probability of not recovering the true partition, is given by

$$
\begin{equation*}
\mathbb{P}[\Omega(X) \neq \hat{\Omega}(G)]=\sum_{g} \mathbb{P}[\hat{\Omega}(g) \neq \Omega(X) \mid G=g] \mathbb{P}[G=g], \tag{5.1.6}
\end{equation*}
$$

where the sum is over all graphs on $n$ vertices (i.e., all possible subsets of edges present) and we dropped the subscript $n, p, q$ to simplify the notation. The MAP estimator $\hat{\Omega}^{\mathrm{MAP}}(G)$ is obtained by minimizing each term $\mathbb{P}[\hat{\Omega}(g) \neq \Omega(X) \mid G=$ $g]$ individually (note that $\mathbb{P}[G=g]>0$ for all $g$ by definition of the $\mathrm{SBM}_{n, p, q}$, a probability which does not depend on the estimator). Equivalently we choose for each $g$ a partition $\gamma$ that maximizes the posterior probability

$$
\begin{align*}
\mathbb{P}[\Omega(X)=\gamma \mid G=g] & =\frac{\mathbb{P}[G=g \mid \Omega(X)=\gamma] \mathbb{P}[\Omega(X)=\gamma]}{\mathbb{P}[G=g]} \\
& =\mathbb{P}[G=g \mid \Omega(X)=\gamma] \cdot \frac{1}{\left|\Pi_{2}^{n}\right| \mathbb{P}[G=g]}, \tag{5.1.7}
\end{align*}
$$

where we applied Bayes' rule on the first line and the uniformity of the partition $X$ on the second line.

Based on (5.1.7), we seek a partition that maximizes $\mathbb{P}[G=g \mid \Omega(X)=\gamma]$. We compute this last probability explicitly. For fixed $g$, let $M:=M(g)$ be the number of edges in $g$. For any $\gamma$, denote by $M_{\mathrm{in}}:=M_{\mathrm{in}}(g, \gamma)$ and $M_{\text {out }}:=M_{\mathrm{out}}(g, \gamma)$ the number of edges within and across communities respectively, and note that $M_{\mathrm{in}}=M-M_{\text {out }}$. By definition of the $\mathrm{SBM}_{n, p, q}$ model, the probability of a graph $g$ given a partition $\gamma$ is expressed simply as

$$
\begin{aligned}
& \mathbb{P}[G=g \mid \Omega(X)=\gamma] \\
& =q^{M_{\text {out }}}(1-q)^{\left(\frac{n}{2}\right)^{2}-M_{\text {out }}} p^{M_{\text {in }}}(1-p)^{\left.\left\{\begin{array}{c}
n \\
2
\end{array}\right)-\left(\frac{n}{2}\right)^{2}\right\}-M_{\text {in }}} \\
& =q^{M_{\text {out }}}(1-q)^{\left(\frac{n}{2}\right)^{2}-M_{\text {out }}} p^{M-M_{\text {out }}}(1-p)^{\left\{\binom{n}{2}-\left(\frac{n}{2}\right)^{2}\right\}-\left\{M-M_{\text {out }}\right\}} \\
& =\left[\frac{q}{1-q} \cdot \frac{1-p}{p}\right]^{M_{\text {out }}}\left\{(1-q)^{\left(\frac{n}{2}\right)^{2}} p^{M}(1-p)^{\left\{\binom{n}{2}-\left(\frac{n}{2}\right)^{2}\right\}-M}\right\} .
\end{aligned}
$$

The expression in curly brackets does not depend on the partition $\gamma$. Moreover, since we assume that $p \geq q$, we have that $\left[\frac{q}{1-q} \cdot \frac{1-p}{p}\right] \leq 1$ (which can be checked directly by rearranging and cancelling). Therefore, to maximize $\mathbb{P}[G=g \mid \Omega(X)=$ $\gamma$ ] over $\gamma$ for a fixed $g$, we need to choose a partition that results in the smallest possible value of $M_{\text {out }}$, the number of edges across the two communities. This problem is well-known in combinatorial optimization, where it is referred to as the minimum bisection problem. It is unfortunately NP-hard and we consider a relaxation that admits a polynomial-time algorithmic solution.

To see how this comes about, observe that the minimum bisection problem can be reformulated as

$$
\max _{\mathbf{x} \in\{+1,-1\}^{n}, \mathbf{x}^{T} \mathbf{1}=0} \mathbf{x}^{T} A \mathbf{x}
$$

where $A$ is the $n \times n$ adjacency matrix. Replacing the combinatorial constraint $\mathbf{x} \in\{+1,-1\}^{n}$ by $\mathbf{x} \in \mathbb{R}^{n}$ with $\|\mathrm{x}\|_{2}=n$ leads to the relaxation

$$
\begin{aligned}
& \max _{\mathbf{z} \in \mathbb{R}^{n}, \mathbf{z}^{T} \mathbf{1}=0,\|\mathbf{z}\|_{2}=n} \mathbf{z}^{T} A \mathbf{z} \\
& =\max _{\mathbf{0} \neq \mathbf{z} \in \mathbb{R}^{n}, \mathbf{z}^{T} \mathbf{1}=0}\left(n \frac{\mathbf{z}}{\|\mathbf{z}\|_{2}}\right)^{T} A\left(n \frac{\mathbf{z}}{\|\mathbf{z}\|_{2}}\right) \\
& \quad=n^{2} \max _{\mathbf{0} \neq \mathbf{z} \in \mathbb{R}^{n}, \mathbf{z}^{T} \mathbf{1}=0} \frac{\mathbf{z}^{T} A \mathbf{z}}{\mathbf{z}^{T} \mathbf{z}},
\end{aligned}
$$

where we changed the notation from $\mathbf{x}$ to $\mathbf{z}$ to emphasize that the solution no longer encodes a partition. We recognize the Rayleigh quotient of $A$ as the objective function in the final formulation. At this point, it is tempting to use Courant-Fischer
(Theorem 5.1.3) and conclude that the maximum above is achieved at the second eigenvalue of $A$. Note however that the vector 1 (appearing in the orthogonality constraint $\mathbf{z}^{T} \mathbf{1}=0$ ) is not in general an eigenvector of $A$ (unless the graph happens to be regular). To leverage the variational characterization of eigenvalues in a statistically justified way, we instead turn to the expected adjacency matrix and then establish concentration.

Lemma 5.1.17 (Expected adjacency). Let $(X, G) \sim \operatorname{SBM}_{n, p, q}$, let $A$ be the adjacency matrix of $G$ and let $\mathcal{A}_{X}=\mathbb{E}_{n, p, q}[A \mid X]$. Then

$$
\mathcal{A}_{X}=n \frac{p+q}{2} \mathbf{u}_{1} \mathbf{u}_{1}^{T}+n \frac{p-q}{2} \mathbf{u}_{2} \mathbf{u}_{2}^{T}-p I,
$$

where

$$
\mathbf{u}_{1}=\frac{1}{\sqrt{n}} \mathbf{1}, \quad \mathbf{u}_{2}=\frac{1}{\sqrt{n}} X .
$$

Proof. For any distinct pair $i, j$, the term

$$
\left(n \frac{p+q}{2} \mathbf{u}_{1} \mathbf{u}_{1}^{T}\right)_{i, j}=n \frac{p+q}{2}\left(\frac{1}{\sqrt{n}}\right)^{2}=\frac{p+q}{2}
$$

while the term

$$
\left(n \frac{p-q}{2} \mathbf{u}_{2} \mathbf{u}_{2}^{T}\right)_{i, j}=n \frac{p-q}{2}\left(\frac{1}{\sqrt{n}}\right)^{2} X_{i} X_{j}=\frac{p-q}{2} X_{i} X_{j} .
$$

The product $X_{i} X_{j}$ is 1 when $i$ and $j$ belong to the same community and is -1 otherwise. In the former case, summing the two terms indeed gives $p$, while in the latter case it gives $q$. Finally, the term $-p I$ accounts for the fact that $A$ has zeros on the diagonal.

Now condition on $X$ and observe that $\mathbf{u}_{1}$ and $\mathbf{u}_{2}$ in Lemma 5.1.17 are orthogonal by our assumption that $X$ corresponds to a balanced partition (i.e., with two communities of equal size). Hence we deduce that an eigenvector decomposition of $\mathcal{A}_{X}$ is formed of $\mathbf{u}_{1}, \mathbf{u}_{2}$ and any orthonormal basis of the orthogonal complement of the span of $\mathbf{u}_{1}$ and $\mathbf{u}_{2}$, with respective eigenvalues

$$
n \frac{p+q}{2}-p, \quad n \frac{p-q}{2}-p, \quad-p .
$$

So the second largest eigenvalue of $\mathcal{A}_{X}$ is $\lambda_{2}\left(\mathcal{A}_{X}\right)=n \frac{p-q}{2}-p$ (independently of $X$ ), and Courant-Fischer implies

$$
\max _{\mathbf{0} \neq \mathbf{z} \in \mathbb{R}^{n}, \mathbf{z}^{T} \mathbf{1}=0} \frac{\mathbf{z}^{T} \mathcal{A}_{X} \mathbf{z}}{\mathbf{z}^{T} \mathbf{z}}=\lambda_{2}\left(\mathcal{A}_{X}\right) .
$$

The corresponding eigenvector, up to scaling and sign, is precisely what we are trying to recover, namely, the community assignment $X$.

These observations motivate the following spectral clustering approach.

1. Input: graph $G$ with adjacency matrix $A$.
2. Compute an eigenvector decomposition of $A$.
3. Let $\hat{\mathbf{u}}_{2}$ be the eigenvector corresponding to the second largest eigenvalue.
4. Output: $\hat{X}(G)=\operatorname{sgn}\left(\hat{\mathbf{u}}_{2}\right)$.

Here we used the notation

$$
(\operatorname{sgn}(\mathbf{z}))_{i}= \begin{cases}+1 & \text { if } z_{i} \geq 0 \\ -1 & \text { otherwise }\end{cases}
$$

Because we used $A$ rather than $\mathcal{A}_{X}$ (which we do not know), it is not immediate that this approach will work. Below, we use Davis-Kahan (Theorem 5.1.14) to show that, under some conditions, the second eigenvector of $A$ is concentrated around that of $\mathcal{A}_{X}$-and therefore almost exact recovery holds.

Before getting to the analysis, we make a final algorithmic remark. The "clustering" above, specifically taking the sign of the second eigenvector, works in this toy model but is perhaps somewhat naive. More generally, in a spectral clustering method, one uses the top eigenvectors (deciding how many is a bit of an art) of the adjacency matrix (or of another matrix associated to the graph such as the Laplacian or normalized Laplacian) to obtain a low-dimensional representation of the input. Then in a second step, one uses a clustering algorithm, for example, $k$-means clustering, to extract communities in the low-dimensional space.

Almost exact recovery We prove the following. We restrict ourselves to the case where $p$ and $q$ are constants not depending on $n$.

Theorem 5.1.18. Let $(X, G) \sim \operatorname{SBM}_{n, p, q}$ and let $A$ be the adjacency matrix of G. Let $\mu:=\min \left\{q, \frac{p-q}{2}\right\}>0$. Clustering according to the sign of the second eigenvector of $A$ identifies the two communities of $G$ with probability at least 1 -$e^{-n}$, except for $C / \mu^{2}$ misclassified nodes for some constant $C>0$.

There are two key ingredients to the proof: concentration of the adjacency matrix and perturbation arguments.

We start with the former.

Lemma 5.1.19 (Norm of the centered adjacency matrix). Let $(X, G) \sim \operatorname{SBM}_{n, p, q}$, let $A$ be the adjacency matrix of $G$ and let $\mathcal{A}_{X}=\mathbb{E}_{n, p, q}[A \mid X]$. There is a constant $C^{\prime}>0$ such that, conditioned on $X$,

$$
\left\|A-\mathcal{A}_{X}\right\|_{2} \leq C^{\prime} \sqrt{n}
$$

with probability at least $1-e^{-n}$.
Proof. Condition on $X$. We use Theorem 2.4.28 on the random matrix $R:=$ $A-\mathcal{A}_{X}$. The entries of $R$ are centered and independent (conditionally on $X$ ). Moreover they are bounded. Indeed, for $i \neq j, A_{i j} \in\{0,1\}$ while $\left(\mathcal{A}_{X}\right)_{i j} \in$ $\{q, p\}$. So $R_{i j} \leq[-p, 1-q]$. On the diagonal, $R_{i i}=0$. Hence, by Hoeffding's lemma (Lemma 2.4.12), the entries are sub-Gaussian with variance factor

$$
\frac{1}{4}(1-q-(-p))^{2} \leq 1
$$

Taking $t=\sqrt{n}$ in Theorem 2.4.28, there is a constant $C>0$ such that with probability $1-e^{-n}$

$$
\left\|A-\mathcal{A}_{X}\right\|_{2} \leq C \sqrt{1}(\sqrt{n}+\sqrt{n}+\sqrt{n}) .
$$

Adjusting the constant gives the claim.
We are ready to prove the theorem.
Proof of Theorem 5.1.18. Condition on $X$. To apply the Davis-Kahan theorem (Theorem 5.1.14), we need to bound the smallest gap $\delta$ between the second largest eigenvalue of $\mathcal{A}_{X}$ and its other eigenvalues. Recall that the eigenvalues are

$$
n \frac{p+q}{2}-p, \quad n \frac{p-q}{2}-p, \quad-p,
$$

so

$$
\delta=\min \left\{n \frac{p-q}{2}, n q\right\}=n \mu>0
$$

By Davis-Kahan and Lemma 5.1.19, with probability at least $1-e^{-n}$, there is $\theta \in\{+1,-1\}$ such that

$$
\left\|\mathbf{u}_{2}-\theta \hat{\mathbf{u}}_{2}\right\|_{2}^{2} \leq \frac{8\left\|A-\mathcal{A}_{X}\right\|_{2}^{2}}{\delta^{2}} \leq \frac{8\left(C^{\prime} \sqrt{n}\right)^{2}}{(n \mu)^{2}}=\frac{C}{n \mu^{2}}
$$

by adjusting the constant. Note that this bound holds for any $X$.

Rearranging and expanding the norm, we get

$$
\sum_{i}\left|\sqrt{n}\left(\mathbf{u}_{2}\right)_{i}-\sqrt{n} \theta\left(\hat{\mathbf{u}}_{2}\right)_{i}\right|^{2} \leq \frac{C}{\mu^{2}}
$$

If the signs of $\left(\mathbf{u}_{2}\right)_{i}$ and $\theta\left(\hat{\mathbf{u}}_{2}\right)_{i}$ disagree, then the $i$-th term in the sum above is $\geq 1$. So there can be at most $C / \mu^{2}$ such disagreements. That establishes the desired bound on the number of misclassified nodes.

Remark 5.1.20. It was shown in [YP14, MNS15a, AS15] that almost exact recovery in the balanced two-community model $\mathrm{SBM}_{n, p_{n}, q_{n}}$ with $p_{n}=a_{n} / n$ and $q_{n}=b_{n} / n$ is achievable (and computationally efficiently so) if and only if

$$
\frac{\left(a_{n}-b_{n}\right)^{2}}{\left(a_{n}+b_{n}\right)}=\omega(1) .
$$

On the other hand, it was shown in [ABH16, MNS15a] that exact recovery in the $\mathrm{SBM}_{n, p_{n}, q_{n}}$ with $p_{n}=\alpha \log n / n$ and $q_{n}=\beta \log n / n$ is achievable and computationally efficiently so if $\sqrt{\alpha}-\sqrt{\beta}>2$ and not achievable if $\sqrt{\alpha}-\sqrt{\beta}<2$.

### 5.2 Spectral techniques for reversible Markov chains

In this section, we apply the spectral theorem to reversible Markov chains. Throughout $\left(X_{t}\right)$ is an irreducible Markov chain on a state space $V$ with transition matrix $P$ reversible with respect to a positive stationary measure $\pi>0$. Recall that this means that $\pi(x) P(x, y)=\pi(y) P(y, x)$ for all $x, y \in V$. We also assume that $P$ is irreducible.

A Hilbert space It will be convenient to introduce a Hilbert space of functions over $V$. Let $\ell^{2}(V, \pi)$ be the space of functions $f: V \rightarrow \mathbb{R}$ such that $\sum_{x \in V} \pi(x) f(x)^{2}<+\infty$. Equipped with the following inner product, it forms a Hilbert space (i.e., a real inner product space that is also a complete metric space (see Theorem B.4.10) with respect to the induced metric; we will work mostly in finite dimension where it is merely a slight generalization of Euclidean space). For $f, g \in \ell^{2}(V, \pi)$, define

$$
\langle f, g\rangle_{\pi}:=\sum_{x \in V} \pi(x) f(x) g(x),
$$

and

$$
\|f\|_{\pi}^{2}:=\langle f, f\rangle_{\pi} .
$$

The inner product is well-defined since the series is summable by Hölder's inequality (Theorem B.4.8), which implies the Cauchy-Schwarz inequality

$$
\langle f, g\rangle_{\pi} \leq\|f\|_{\pi}\|g\|_{\pi}
$$

Minkowski's inequality (Theorem B.4.9) implies the triangle inequality

$$
\|f+g\|_{\pi} \leq\|f\|_{\pi}+\|g\|_{\pi} .
$$

The integral with respect to $\pi$ (see Appendix B) reduces in this case to a sum

$$
\pi(f):=\sum_{x \in V} \pi(x) f(x),
$$

provided $\pi(|f|)<+\infty$ or $f \geq 0$. Here $|f|$ is defined as $|f|(x):=|f(x)|$ for all $x \in V$. We also write $\pi f=\pi(f)$ to simplify the notation.

We recall some standard Hilbert space facts. The countable collection of functions $\left\{f_{i}\right\}_{i=1}^{\infty}$ in $\ell^{2}(V, \pi)$ is an orthonormal basis if: (i) $\left\langle f_{i}, f_{j}\right\rangle_{\pi}=0$ if $i \neq j$ and $=$ 1 if $i=j$; and (ii) any $f \in \ell^{2}(V, \pi)$ can be written as $\lim _{n \rightarrow+\infty} \sum_{i=1}^{n}\left\langle f_{i}, f\right\rangle_{\pi} f_{i}=$ $f$ where the limit is in the norm. We then have Parseval's identity: for any $g \in \ell^{2}(V, \pi)$

$$
\begin{equation*}
\|g\|_{\pi}^{2}=\sum_{j=1}^{\infty}\left\langle g, f_{j}\right\rangle_{\pi}^{2} \tag{5.2.1}
\end{equation*}
$$

Think of $P$ as an operator on $\ell^{2}(V, \pi)$. That is, let $P f: V \rightarrow \mathbb{R}$ be defined as

$$
(P f)(x):=\sum_{y \in V} P(x, y) f(y),
$$

for $x \in V$. For any $f \in \ell^{2}(V, \pi), P f$ is well-defined and further we have

$$
\begin{equation*}
\|P f\|_{\pi} \leq\|f\|_{\pi} \tag{5.2.2}
\end{equation*}
$$

Indeed by Cauchy-Schwarz, stochasticity, Fubini and stationarity

$$
\begin{align*}
\|P|f|\|_{\pi}^{2} & =\sum_{x} \pi(x)\left[\sum_{y} P(x, y) \mid f(y)\right]^{2} \\
& \leq \sum_{x} \pi(x)\left[\sum_{y} P(x, y)|f(y)|^{2} \sum_{z} P(x, z)\right] \\
& =\sum_{y} \sum_{x} \pi(x) P(x, y) f(y)^{2} \\
& =\sum_{y} \pi(y) f(y)^{2} \\
& =\|f\|_{\pi}^{2}<+\infty \tag{5.2.3}
\end{align*}
$$

This shows that $P f$ is well-defined since $\pi>0$ and hence the series in square brackets on the first line is finite for all $x$. Applying the same argument to $\|P f\|_{\pi}^{2}$ gives the inequality above.

Everything above holds whether or not $P$ is reversible, so long as $\pi$ is a stationary measure. Now we use reversibility. We claim that, when $P$ reversible, then it is self-adjoint, that is,

$$
\begin{equation*}
\langle f, P g\rangle_{\pi}=\langle P f, g\rangle_{\pi} \quad \forall f, g \in \ell^{2}(V, \pi) . \tag{5.2.4}
\end{equation*}
$$

This follows immediately by reversibility

$$
\begin{aligned}
\langle f, P g\rangle_{\pi} & =\sum_{x \in V} \pi(x) f(x) \sum_{y \in V} P(x, y) g(y) \\
& =\sum_{x \in V} \sum_{y \in V} \pi(y) P(y, x) f(x) g(y) \\
& =\sum_{y \in V} \pi(y) g(y) \sum_{x \in V} P(y, x) f(x) \\
& =\langle P f, g\rangle_{\pi},
\end{aligned}
$$

where we argue as in (5.2.3) to justify using Fubini.
Throughout this section, we denote by 0 and 1 the all-zero and all-one functions respectively.

### 5.2.1 Spectral gap

In this subsection, we restrict ourselves to a finite state space $V$. Our goal is to bound the mixing time of $\left(X_{t}\right)$ in terms of the eigenvalues of the transition matrix $P$. We assume that $\pi$ is now the stationary distribution, that is, $\sum_{x \in V} \pi(x)=1$ (which is unique by Theorem 1.1.24 and irreducibility). We also let $n:=|V|<$ $+\infty$.

Spectral decomposition Self-adjointness generalizes the notion of a symmetric matrix, with one consequence being that a version of the spectral theorem applies to $P$ (at least in this finite-dimensional case; see Section 5.2.5 for more discussion on this). For completeness, we derive it from Theorem 5.1.1. It will be convenient to assume without loss of generality that $V=[n]$ and identify functions in $\ell^{2}(V, \pi)$ with vectors in $\mathbb{R}^{n}$.

Theorem 5.2.1 (Reversibility: spectral theorem). There is an orthonormal basis of $\ell^{2}(V, \pi)$ formed of real eigenfunctions $\left\{f_{j}\right\}_{j=1}^{n}$ of $P$ with real eigenvalues $\left\{\lambda_{j}\right\}_{j=1}^{n}$.

Proof. Let $D_{\pi}$ be the diagonal matrix with $\pi$ on the diagonal. By reversibility,

$$
\begin{aligned}
M(x, y) & :=\left(D_{\pi}^{1 / 2} P D_{\pi}^{-1 / 2}\right)_{x, y} \\
& =\sqrt{\frac{\pi(x)}{\pi(y)}} P(x, y) \\
& =\sqrt{\frac{\pi(y)}{\pi(x)}} P(y, x) \\
& =\left(D_{\pi}^{1 / 2} P D_{\pi}^{-1 / 2}\right)_{y, x} \\
& =M(y, x) .
\end{aligned}
$$

So $M=(M(x, y))_{x, y}=D_{\pi}^{1 / 2} P D_{\pi}^{-1 / 2}$ is a symmetric matrix. By the spectral theorem (Theorem 5.1.1), it has real eigenvectors $\left\{\phi_{j}\right\}_{j=1}^{n}$ forming an orthonormal basis of $\mathbb{R}^{n}$ with corresponding real eigenvalues $\left\{\lambda_{j}\right\}_{j=1}^{n}$. Define $f_{j}:=D_{\pi}^{-1 / 2} \phi_{j}$. Then

$$
\begin{aligned}
P f_{j} & =P D_{\pi}^{-1 / 2} \phi_{j} \\
& =D_{\pi}^{-1 / 2} D_{\pi}^{1 / 2} P D_{\pi}^{-1 / 2} \phi_{j} \\
& =D_{\pi}^{-1 / 2} M \phi_{j} \\
& =\lambda_{j} D_{\pi}^{-1 / 2} \phi_{j} \\
& =\lambda_{j} f_{j},
\end{aligned}
$$

and

$$
\begin{aligned}
\left\langle f_{i}, f_{j}\right\rangle_{\pi} & =\left\langle D_{\pi}^{-1 / 2} \phi_{i}, D_{\pi}^{-1 / 2} \phi_{j}\right\rangle_{\pi} \\
& =\sum_{x \in V} \pi(x)\left[\pi(x)^{-1 / 2} \phi_{i}(x)\right]\left[\pi(x)^{-1 / 2} \phi_{j}(x)\right] \\
& =\sum_{x \in V} \phi_{i}(x) \phi_{j}(x) .
\end{aligned}
$$

Because $\left\{\phi_{j}\right\}_{j=1}^{n}$ is an orthonormal basis of $\mathbb{R}^{n}$, we have that $\left\{f_{j}\right\}_{j=1}^{n}$ is an orthonormal basis of $\left(\mathbb{R}^{n},\langle\cdot, \cdot\rangle_{\pi}\right)$.

We collect a few more facts about the eigenbasis. Recall that

$$
\|f\|_{\infty}=\max _{x \in V}|f(x)| .
$$

Lemma 5.2.2. Any eigenvalue $\lambda$ of $P$ satisfies $|\lambda| \leq 1$.

Proof. It holds that

$$
P f=\lambda f \Longrightarrow|\lambda|\|f\|_{\infty}=\|P f\|_{\infty}=\max _{x}\left|\sum_{y} P(x, y) f(y)\right| \leq\|f\|_{\infty}
$$

Rearranging gives the claim.
We order the eigenvalues $1 \geq \lambda_{1} \geq \cdots \geq \lambda_{n} \geq-1$. The second eigenvalue will play an important role below.

Lemma 5.2.3. We have $\lambda_{1}=1$ and $\lambda_{2}<1$. Also we can take $f_{1}=1$.
Proof. Because $P$ is stochastic, the all-one vector is a right eigenvector with eigenvalue 1 . Any eigenfunction with eigenvalue 1 is harmonic with respect to $P$ on $V$ (see (3.3.2)). By Corollary 3.3.3, for a finite, irreducible chain the only harmonic functions are the constant functions. So the eigenspace corresponding to 1 is onedimensional. We must have $\lambda_{2}<1$ by Lemma 5.2.2.

When the chain is aperiodic, it cannot have an eigenvalue -1. Exercise 5.9 asks for a proof.

Lemma 5.2.4. If $P$ has an eigenvalue equal to -1 , then $P$ is not aperiodic.
Lemma 5.2.5. For all $j \neq 1, \pi f_{j}=0$.
Proof. By orthonormality, $\left\langle f_{1}, f_{j}\right\rangle_{\pi}=0$. Now use the fact that $f_{1}=1$.
Let $\delta_{x}(y):=\mathbf{1}_{\{x=y\}}$.
Lemma 5.2.6. For all $x, y$,

$$
\sum_{j=1}^{n} f_{j}(x) f_{j}(y)=\pi(x)^{-1} \delta_{x}(y)
$$

Proof. Using the notation of Theorem 5.2.1, the matrix $\Phi$ whose columns are the $\phi_{j} \mathrm{~s}$ is orthogonal so $\Phi \Phi^{T}=I$. That is,

$$
\sum_{j=1}^{n} \phi_{j}(x) \phi_{j}(y)=\delta_{x}(y),
$$

or

$$
\sum_{j=1}^{n} \sqrt{\pi(x) \pi(y)} f_{j}(x) f_{j}(y)=\delta_{x}(y)
$$

Rearranging gives the result.

Using the eigendecomposition of $P$, we get the following expression for its $t$-th power $P^{t}$.

Theorem 5.2.7 (Spectral decomposition of $P^{t}$ ). Let $\left\{f_{j}\right\}_{j=1}^{n}$ be the eigenfunctions of a reversible and irreducible transition matrix $P$ with corresponding eigenvalues $\left\{\lambda_{j}\right\}_{j=1}^{n}$, as defined previously. Assume $\lambda_{1} \geq \cdots \geq \lambda_{n}$. We have the decomposition

$$
\frac{P^{t}(x, y)}{\pi(y)}=1+\sum_{j=2}^{n} f_{j}(x) f_{j}(y) \lambda_{j}^{t} .
$$

Proof. Let $F$ be the matrix whose columns are the eigenvectors $\left\{f_{j}\right\}_{j=1}^{n}$ and let $D_{\lambda}$ be the diagonal matrix with $\left\{\lambda_{j}\right\}_{j=1}^{n}$ on the diagonal. Using the notation in the proof of Theorem 5.2.1,

$$
D_{\pi}^{1 / 2} P^{t} D_{\pi}^{-1 / 2}=M^{t}=\left(D_{\pi}^{1 / 2} F\right) D_{\lambda}^{t}\left(D_{\pi}^{1 / 2} F\right)^{T},
$$

which after rearranging becomes

$$
P^{t} D_{\pi}^{-1}=F D_{\lambda}^{t} F^{T} .
$$

Expanding and using Lemma 5.2.3 gives the result.
Example 5.2.8 (Two-state chain). Let $V:=\{0,1\}$ and

$$
P:=\left(\begin{array}{cc}
1-\alpha & \alpha \\
\beta & 1-\beta
\end{array}\right),
$$

for $\alpha, \beta \in(0,1)$. Observe that $P$ is reversible with respect to the stationary distribution

$$
\pi:=\left(\frac{\beta}{\alpha+\beta}, \frac{\alpha}{\alpha+\beta}\right) .
$$

We know that $f_{1}=1$ is an eigenfunction with eigenvalue 1 . As can be checked by direct computation, the other eigenfunction (in vector form) is

$$
f_{2}:=\left(\sqrt{\frac{\alpha}{\beta}},-\sqrt{\frac{\beta}{\alpha}}\right),
$$

with eigenvalue $\lambda_{2}:=1-\alpha-\beta$. We normalized $f_{2}$ so that $\left\|f_{2}\right\|_{\pi}^{2}=1$.
By Theorem 5.2.7, the spectral decomposition at time $t$ is therefore

$$
P^{t} D_{\pi}^{-1}=\left(\begin{array}{ll}
1 & 1 \\
1 & 1
\end{array}\right)+(1-\alpha-\beta)^{t}\left(\begin{array}{cc}
\frac{\alpha}{\beta} & -1 \\
-1 & \frac{\beta}{\alpha}
\end{array}\right)
$$

Or, rearranging,

$$
P^{t}=\left(\begin{array}{cc}
\frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta} \\
\frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta}
\end{array}\right)+(1-\alpha-\beta)^{t}\left(\begin{array}{cc}
\frac{\alpha}{\alpha+\beta} & -\frac{\alpha}{\alpha+\beta} \\
-\frac{\beta}{\alpha+\beta} & \frac{\beta}{\alpha+\beta}
\end{array}\right) .
$$

Note for instance that the case $\alpha+\beta=1$ corresponds to a rank-one $P$, which immediately converges to stationarity.

Assume $\beta \geq \alpha$. Then, by (1.1.6) and Lemma 4.1.9,

$$
d(t)=\max _{x} \frac{1}{2} \sum_{y}\left|P^{t}(x, y)-\pi(y)\right|=\frac{\beta}{\alpha+\beta}|1-\alpha-\beta|^{t} .
$$

As a result,

$$
\mathrm{t}_{\text {mix }}(\varepsilon)=\left\lceil\frac{\log \left(\varepsilon \frac{\alpha+\beta}{\beta}\right)}{\log |1-\alpha-\beta|}\right\rceil=\left\lceil\frac{\log \varepsilon^{-1}-\log \left(\frac{\alpha+\beta}{\beta}\right)}{\log |1-\alpha-\beta|^{-1}}\right\rceil
$$

Spectral gap and mixing Assume further that $P$ is aperiodic. Recall that by the convergence theorem (Theorem 1.1.33), for all $x, y, P^{t}(x, y) \rightarrow \pi(y)$ as $t \rightarrow+\infty$, and that the mixing time (Definition 1.1.35) is

$$
\operatorname{t}_{\text {mix }}(\varepsilon):=\min \{t \geq 0: d(t) \leq \varepsilon\},
$$

where $d(t):=\max _{x \in V}\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}}$. It will be convenient to work with a different notion of distance.
Definition 5.2.9 (Separation distance). The separation distance is defined as

$$
s_{x}(t):=\max _{y \in V}\left[1-\frac{P^{t}(x, y)}{\pi(y)}\right],
$$

and we let $s(t):=\max _{x \in V} s_{x}(t)$.
Lemma 5.2.10 (Separation distance and total variation distance).

$$
d(t) \leq s(t)
$$

Proof. By Lemma 4.1.15,

$$
\begin{aligned}
\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}} & =\sum_{y: P^{t}(x, y)<\pi(y)}\left[\pi(y)-P^{t}(x, y)\right] \\
& =\sum_{y: P^{t}(x, y)<\pi(y)} \pi(y)\left[1-\frac{P^{t}(x, y)}{\pi(y)}\right] \\
& \leq s_{x}(t) .
\end{aligned}
$$

Since this holds for any $x$, the claim follows.

It follows that, from the spectral decomposition (Theorem 5.2.7), the speed of convergence of $P^{t}(x, y)$ to $\pi(y)$ is dominated by the largest eigenvalue of $P$ not equal to 1 .

Definition 5.2.11 (Spectral gap). The absolute spectral gap is $\gamma_{*}:=1-\lambda_{*}$ where $\lambda_{*}:=\left|\lambda_{2}\right| \vee\left|\lambda_{n}\right|$. The spectral gap is $\gamma:=1-\lambda_{2}$.
absolute spectral
gap Note that the eigenvalues of the lazy version $\frac{1}{2} P+\frac{1}{2} I$ of $P$ are $\left\{\frac{1}{2}\left(\lambda_{j}+1\right)\right\}_{j=1}^{n}$ which are all nonnegative. So, there, $\gamma_{*}=\gamma$.

Definition 5.2.12 (Relaxation time). The relaxation time is defined as

$$
\mathrm{t}_{\mathrm{rel}}:=\gamma_{*}^{-1} .
$$

Example 5.2.13 (Two-state chain (continued)). Returning to Example 5.2.8, there are two cases:

- $\alpha+\beta \leq 1$ : In that case the (absolute) spectral gap is $\gamma_{*}=\gamma=\alpha+\beta$ and the relaxation time is $\mathrm{t}_{\mathrm{rel}}=1 /(\alpha+\beta)$.
- $\alpha+\beta>1$ : In that case the absolute spectral gap is $\gamma_{*}=2-\alpha-\beta$ and the relaxation time is $\mathrm{t}_{\text {rel }}=1 /(2-\alpha-\beta)$.

The following result clarifies the relationship between the mixing and relaxation times. Let $\pi_{\text {min }}=\min _{x} \pi(x)$.

Theorem 5.2.14 (Mixing time and relaxation time). Let $P$ be reversible, irreducible, and aperiodic with positive stationary distribution $\pi$. For all $\varepsilon>0$,

$$
\left(\mathrm{t}_{\text {rel }}-1\right) \log \left(\frac{1}{2 \varepsilon}\right) \leq \mathrm{t}_{\text {mix }}(\varepsilon) \leq \log \left(\frac{1}{\varepsilon \pi_{\text {min }}}\right) \mathrm{t}_{\text {rel }} .
$$

Proof. We start with the upper bound. By Lemma 5.2.10, it suffices to find $t$ such that $s(t) \leq \varepsilon$. By the spectral decomposition and Cauchy-Schwarz,

$$
\left|\frac{P^{t}(x, y)}{\pi(y)}-1\right| \leq \lambda_{*}^{t} \sum_{j=2}^{n}\left|f_{j}(x) f_{j}(y)\right| \leq \lambda_{*}^{t} \sqrt{\sum_{j=2}^{n} f_{j}(x)^{2} \sum_{j=2}^{n} f_{j}(y)^{2}} .
$$

By Lemma 5.2.6, $\sum_{j=2}^{n} f_{j}(x)^{2} \leq \pi(x)^{-1}$. Plugging this back above, we get

$$
\begin{equation*}
\left|\frac{P^{t}(x, y)}{\pi(y)}-1\right| \leq \lambda_{*}^{t} \sqrt{\pi(x)^{-1} \pi(y)^{-1}} \leq \frac{\lambda_{*}^{t}}{\pi_{\min }}=\frac{\left(1-\gamma_{*}\right)^{t}}{\pi_{\min }} \leq \frac{e^{-\gamma_{*} t}}{\pi_{\min }} \tag{5.2.5}
\end{equation*}
$$

where we used that $1-z \leq e^{-z}$ for all $z \in \mathbb{R}$ (see Exercise 1.16). Observe that the right-hand side is less than $\varepsilon$ when $t \geq \log \left(\frac{1}{\varepsilon \pi_{\text {min }}}\right) \mathrm{t}_{\text {rel }}$.

For the lower bound, let $f_{*}$ be an eigenfunction associated with an eigenvalue achieving $\lambda_{*}:=\left|\lambda_{2}\right| \vee\left|\lambda_{n}\right|$. Let $z$ be such that $\left|f_{*}(z)\right|=\left\|f_{*}\right\|_{\infty}$. By Lemma 5.2.5, $\pi f_{*}=0$. Hence

$$
\begin{aligned}
\lambda_{*}^{t}\left|f_{*}(z)\right| & =\left|P^{t} f_{*}(z)\right| \\
& =\left|\sum_{y}\left[P^{t}(z, y) f_{*}(y)-\pi(y) f_{*}(y)\right]\right| \\
& \leq\left\|f_{*}\right\|_{\infty} \sum_{y}\left|P^{t}(z, y)-\pi(y)\right| \leq\left\|f_{*}\right\|_{\infty} 2 d(t),
\end{aligned}
$$

so $d(t) \geq \frac{1}{2} \lambda_{*}^{t}$. When $t=\mathrm{t}_{\text {mix }}(\varepsilon), \varepsilon \geq d\left(\mathrm{t}_{\text {mix }}(\varepsilon)\right) \geq \frac{1}{2} \lambda_{*}^{\mathrm{t}_{\text {mix }}(\varepsilon)}$. Therefore, rearranging and taking a logarithm, we get

$$
\mathrm{t}_{\text {mix }}(\varepsilon)\left(\frac{1}{\lambda_{*}}-1\right) \geq \mathrm{t}_{\text {mix }}(\varepsilon) \log \left(\frac{1}{\lambda_{*}}\right) \geq \log \left(\frac{1}{2 \varepsilon}\right)
$$

where we used $z=1-\lambda_{*}^{-1}$ in $1-z \leq e^{-z}$ to get the first inequality. The result follows from $\left(\frac{1}{\lambda_{*}}-1\right)^{-1}=\left(\frac{1-\lambda_{*}}{\lambda_{*}}\right)^{-1}=\left(\frac{\gamma_{*}}{1-\gamma_{*}}\right)^{-1}=\mathrm{t}_{\text {rel }}-1$.

### 5.2.2 $\triangleright$ Random walks: a spectral look at cycles and hypercubes

We illustrate the results in the previous subsection to random walk on cycles and hypercubes.

## Random walk on a cycle

Consider simple random walk on the $n$-cycle (see Example 1.1.17). That is, $V:=$ $\{0,1, \ldots, n-1\}$ and $P(x, y)=1 / 2$ if and only if $|x-y|=1 \bmod n$. We assume that $n$ is odd to avoid periodicity issues. Let $\pi \equiv n^{-1}$ be the stationary distribution (by symmetry and $|V|=n$ ). We showed in Section 4.3.2 that (for the lazy version of the chain) the mixing time is $\mathrm{t}_{\text {mix }}(\varepsilon)=\Theta\left(n^{2}\right)$.

Here we use spectral techniques. We first compute the eigendecomposition, which in this case can be determined explicitly.

Lemma 5.2.15 (Cycle: eigenbasis). For $j=1, \ldots, n-1$, the function

$$
g_{j}(x):=\sqrt{2} \cos \left(\frac{2 \pi j x}{n}\right), \quad x=0,1, \ldots, n-1,
$$

is an eigenfunction of $P$ with eigenvalue

$$
\mu_{j}:=\cos \left(\frac{2 \pi j}{n}\right),
$$

and $g_{0}=1$ is an eigenfunction with eigenvalue 1 . Moreover the $g_{j}$ s are orthonormal in $\ell^{2}(V, \pi)$.
Proof. We know from Lemma 5.2.3 that 1 is an eigenfunction with eigenvalue 1. Let $j \in\{1, \ldots, n-1\}$. Note that, for all $x$, switching momentarily to the complex representation (where we use $i$ for the imaginary unit)

$$
\begin{aligned}
\sum_{y} P(x, y) g_{j}(y) & =\frac{1}{2}\left[\sqrt{2} \cos \left(\frac{2 \pi j(x-1)}{n}\right)+\sqrt{2} \cos \left(\frac{2 \pi j(x+1)}{n}\right)\right] \\
& =\frac{\sqrt{2}}{2}\left[\frac{e^{i \frac{2 \pi j(x-1)}{n}}+e^{-i \frac{2 \pi j(x-1)}{n}}}{2}+\frac{e^{i \frac{2 \pi j(x+1)}{n}}+e^{-i \frac{2 \pi j(x+1)}{n}}}{2}\right] \\
& =\sqrt{2}\left[\frac{e^{i \frac{2 \pi j x}{n}}+e^{-i \frac{2 \pi j x}{n}}}{2}\right]\left[\frac{e^{i \frac{2 \pi j}{n}}+e^{-i \frac{2 \pi j}{n}}}{2}\right] \\
& =\left[\sqrt{2} \cos \left(\frac{2 \pi j x}{n}\right)\right]\left[\cos \left(\frac{2 \pi j}{n}\right)\right] \\
& =\cos \left(\frac{2 \pi j}{n}\right) g_{j}(x) .
\end{aligned}
$$

The orthonormality follows from standard trigonometric identities. We prove only that the $g_{j}$ s have unit norm. We use the Dirichtlet kernel (see Exercise 5.8)

$$
1+2 \sum_{k=1}^{n} \cos k \theta=\frac{\sin ((n+1 / 2) \theta)}{\sin (\theta / 2)}
$$

for $\theta \neq 0$, and the identity $\cos ^{2}(\theta)=\frac{1}{2}(1+\cos (2 \theta))$. For $j=0, g_{j}=1$ and the norm squared is $\sum_{x} \pi(x)=1$. For $j \neq 0$, we have $\left\|g_{j}\right\|_{\pi}^{2}$ is

$$
\begin{aligned}
\sum_{x \in V} \pi(x) g_{j}(x)^{2} & =\frac{1}{n} \sum_{x=0}^{n-1} 2 \cos ^{2}\left(\frac{2 \pi j x}{n}\right) \\
& =\frac{1}{n} \sum_{x=0}^{n-1}\left(1+\cos \left(\frac{4 \pi j x}{n}\right)\right) \\
& =1+\frac{1}{n} \sum_{k=1}^{n} \cos \left(k \frac{4 \pi j}{n}\right) \\
& =1+\frac{1}{2}\left[\frac{\sin ((n+1 / 2)(4 \pi j / n))}{\sin ((4 \pi j / n) / 2)}-1\right],
\end{aligned}
$$

which is indeed 1.
From the eigenvalues, we derive the relaxation time (Definition 5.2.12) analytically.

Theorem 5.2.16 (Cycle: relaxation time). The relaxation time for lazy simple random walk on the $n$-cycle is

$$
\mathrm{t}_{\mathrm{rel}}=\frac{1}{1-\cos \left(\frac{2 \pi}{n}\right)}=\Theta\left(n^{2}\right) .
$$

Proof. By Lemma 5.2.15, the absolute spectral gap (Definition 5.2.11) is 1 $\cos \left(\frac{2 \pi}{n}\right)$, using that $n$ is odd. By a Taylor expansion,

$$
1-\cos \left(\frac{2 \pi}{n}\right)=\frac{4 \pi^{2}}{n^{2}}+O\left(n^{-4}\right)
$$

Since $\pi_{\text {min }}=1 / n$, we get $\mathrm{t}_{\text {mix }}(\varepsilon)=O\left(n^{2} \log n\right)$ and $\mathrm{t}_{\text {mix }}(\varepsilon)=\Omega\left(n^{2}\right)$ by Theorem 5.2.14.

It turns out our upper bound is off by a logarithmic factor. A sharper bound on the mixing time can be obtained by working directly with the spectral decomposition. By Lemma 4.1.9 and Cauchy-Schwarz (Theorem B.4.8), for any $x \in V$,

$$
\begin{aligned}
4\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}}^{2} & =\left\{\sum_{y} \pi(y)\left|\frac{P^{t}(x, y)}{\pi(y)}-1\right|\right\}^{2} \\
& \leq \sum_{y} \pi(y)\left(\frac{P^{t}(x, y)}{\pi(y)}-1\right)^{2} \\
& =\left\|\sum_{j=1}^{n-1} \mu_{j}^{t} g_{j}(x) g_{j}\right\|_{\pi}^{2} \\
& =\sum_{j=1}^{n-1} \mu_{j}^{2 t} g_{j}(x)^{2}
\end{aligned}
$$

where we used the spectral decomposition of $P^{t}$ (Theorem 5.2.7) on the third line and Parseval's identity (i.e., (5.2.1)) on the fourth line.

Here comes the trick: the total variation distance does not depend on the starting point $x$ by symmetry. Multiplying by $\pi(x)$ and summing over $x$-on the right-
hand side only—gives

$$
\begin{aligned}
4\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}}^{2} & \leq \sum_{x} \pi(x) \sum_{j=1}^{n-1} \mu_{j}^{2 t} g_{j}(x)^{2} \\
& =\sum_{j=1}^{n-1} \mu_{j}^{2 t} \sum_{x} \pi(x) g_{j}(x)^{2} \\
& =\sum_{j=1}^{n-1} \mu_{j}^{2 t},
\end{aligned}
$$

where we used that $\left\|g_{j}\right\|_{\pi}^{2}=1$.
We get

$$
4 d(t)^{2} \leq \sum_{j=1}^{n-1} \cos ^{2 t}\left(\frac{2 \pi j}{n}\right)=2 \sum_{j=1}^{(n-1) / 2} \cos ^{2 t}\left(\frac{2 \pi j}{n}\right)
$$

For $x \in[0, \pi / 2), \cos x \leq e^{-x^{2} / 2}$ (see Exercise 1.16). Then

$$
\begin{aligned}
4 d(t)^{2} & \leq 2 \sum_{j=1}^{(n-1) / 2} \exp \left(-\frac{4 \pi^{2} j^{2}}{n^{2}} t\right) \\
& \leq 2 \exp \left(-\frac{4 \pi^{2}}{n^{2}} t\right) \sum_{j=1}^{\infty} \exp \left(-\frac{4 \pi^{2}\left(j^{2}-1\right)}{n^{2}} t\right) \\
& \leq 2 \exp \left(-\frac{4 \pi^{2}}{n^{2}} t\right) \sum_{\ell=0}^{\infty} \exp \left(-\frac{4 \pi^{2} t}{n^{2}} \ell\right) \\
& =\frac{2 \exp \left(-\frac{4 \pi^{2}}{n^{2}} t\right)}{1-\exp \left(-\frac{4 \pi^{2}}{n^{2}} t\right)}
\end{aligned}
$$

So $\mathrm{t}_{\text {mix }}(\varepsilon)=O\left(n^{2}\right)$.

## Random walk on the hypercube

Consider simple random walk on the hypercube $V:=\{-1,+1\}^{n}$ where $x \sim y$ if they differ at exactly one coordinate. We consider the lazy version to avoid issues of periodicity (see Example 1.1.31). Let $P$ be the transition matrix and let $\pi \equiv 2^{-n}$ be the stationary distribution (by symmetry and $|V|=2^{n}$ ). We showed in Section 4.3.2 that $\mathrm{t}_{\text {mix }}(\varepsilon)=\Theta(n \log n)$. Here we use spectral techniques.

For $J \subseteq[n]$, we let

$$
\chi_{J}(x)=\prod_{j \in J} x_{j}, \quad x \in V
$$

These are called parity functions. We show that the parity functions form an eigenbasis of the transition matrix.

Lemma 5.2.17 (Hypercube: eigenbasis). For all $J \subseteq[n]$, the function $\chi_{J}$ is an eigenfunction of $P$ with eigenvalue

$$
\mu_{J}:=\frac{n-|J|}{n} .
$$

Moreover the $\chi_{J} S$ are orthonormal in $\ell^{2}(V, \pi)$.
Proof. For $x \in V$ and $i \in[n]$, let $x^{[i]}$ be $x$ where coordinate $i$ is flipped. Note that, for all $J, x$,

$$
\begin{aligned}
\sum_{y} P(x, y) \chi_{J}(y) & =\frac{1}{2} \chi_{J}(x)+\frac{1}{2} \sum_{i=1}^{n} \frac{1}{n} \chi_{J}\left(x^{[i]}\right) \\
& =\left\{\frac{1}{2}+\frac{1}{2} \frac{n-|J|}{n}\right\} \chi_{J}(x)-\frac{1}{2} \frac{|J|}{n} \chi_{J}(x) \\
& =\frac{n-|J|}{n} \chi_{J}(x) .
\end{aligned}
$$

For the orthonormality, note that

$$
\sum_{x \in V} \pi(x) \chi_{J}(x)^{2}=\sum_{x \in V} \frac{1}{2^{n}} \prod_{j \in J} x_{j}^{2}=1 .
$$

For $J \neq J^{\prime} \subseteq[n]$,

$$
\begin{aligned}
& \sum_{x \in V} \pi(x) \chi_{J}(x) \chi_{J^{\prime}}(x) \\
& =\sum_{x \in V} \frac{1}{2^{n}} \prod_{j \in J \cap J^{\prime}} x_{j}^{2} \prod_{j \in J \backslash J^{\prime}} x_{j} \prod_{j \in J^{\prime} \backslash J} x_{j} \\
& =\frac{2^{\left|J \cap J^{\prime}\right|}}{2^{n}} \prod_{j \in J \backslash J^{\prime}}\left(\sum_{x_{j} \in\{-1,+1\}} x_{j}\right) \prod_{j \in J^{\prime} \backslash J}\left(\sum_{x_{j} \in\{-1,+1\}} x_{j}\right) \\
& =0,
\end{aligned}
$$

since at least one of $J \backslash J^{\prime}$ or $J^{\prime} \backslash J$ is nonempty.

From the eigenvalues, we obtain the relaxation time.
Theorem 5.2.18 (Hypercube: relaxation time). The relaxation time for lazy simple random walk on the $n$-dimensional hypercube is

$$
\mathrm{t}_{\mathrm{rel}}=n .
$$

Proof. From Lemma 5.2.17, the absolute spectral gap is

$$
\gamma_{*}=\gamma=1-\frac{n-1}{n}=\frac{1}{n} .
$$

Note that $\pi_{\min }=1 / 2^{n}$. Hence, by Theorem 5.2.14, we have $\mathrm{t}_{\text {mix }}(\varepsilon)=O\left(n^{2}\right)$ and $\mathrm{t}_{\text {mix }}(\varepsilon)=\Omega(n)$. Those bounds, it turns out, are both off.

As we did for the cycle, we obtain a sharper upper bound by working directly with the spectral decomposition. By the same argument we used there,

$$
4 d(t)^{2} \leq \sum_{J \neq \emptyset} \mu_{J}^{2 t}
$$

Then

$$
\begin{aligned}
4 d(t)^{2} & \leq \sum_{J \neq \emptyset}\left(\frac{n-|J|}{n}\right)^{2 t} \\
& =\sum_{\ell=1}^{n}\binom{n}{\ell}\left(1-\frac{\ell}{n}\right)^{2 t} \\
& \leq \sum_{\ell=1}^{n}\binom{n}{\ell} \exp \left(-\frac{2 t \ell}{n}\right) \\
& =\left(1+\exp \left(-\frac{2 t}{n}\right)\right)^{n}-1,
\end{aligned}
$$

where we used that $1-x \leq e^{-x}$ for all $x$ (see Exercise 1.16). So, by definition, $\mathrm{t}_{\text {mix }}(\varepsilon) \leq \frac{1}{2} n \log n+O(n)$.

Remark 5.2.19. In fact, lazy simple random walk on the $n$-dimensional hypercube has a "cutoff" at $(1 / 2) n \log n$. Roughly speaking, within a time window of size $O(n)$, the total variation distance to the stationary distribution goes from near 1 to near 0 . See, e.g., [LPW06, Section 18.2.2].

### 5.2.3 $\triangleright$ Markov chains: Varopoulos-Carne and diameter-based bounds on the mixing time

If $\left(S_{t}\right)$ is simple random walk on $\mathbb{Z}$, then Lemma 2.4.3 guarantees that for any $x, y \in \mathbb{Z}$

$$
\begin{equation*}
P^{t}(x, y) \leq e^{-|x-y|^{2} / 2 t} \tag{5.2.6}
\end{equation*}
$$

where $P$ is the transition matrix of $\left(S_{t}\right)$. Interestingly a similar bound holds for any reversible Markov chain-and Lemma 2.4.3 plays an unexpected role in its proof. An application to mixing times is discussed below.

## Varopoulos-Carne bound

Our main bound is the following. Recall that a reversible Markov chain is equivalent to a random walk on the network corresponding to its positive transition probabilities (see Definition 1.2.7 and the discussion following it).

Theorem 5.2.20 (Varopoulos-Carne bound). Let $P$ be the transition matrix of an irreducible Markov chain $\left(X_{t}\right)$ on the countable state space $V$. Assume further that $P$ is reversible with respect to the stationary measure $\pi$ and that the corresponding network $\mathcal{N}$ is locally finite. Then the following hold

$$
\forall x, y \in V, \forall t \in \mathbb{N}, \quad P^{t}(x, y) \leq 2 \sqrt{\frac{\pi(y)}{\pi(x)}} e^{-\rho(x, y)^{2} / 2 t},
$$

where $\rho(x, y)$ is the graph distance between $x$ and $y$ on $\mathcal{N}$.
As a sanity check before proving the theorem, note that if the chain is aperiodic and $\pi$ is the stationary distribution then by the convergence theorem (Theorem 1.1.33)

$$
P^{t}(x, y) \rightarrow \pi(y) \leq 2 \sqrt{\frac{\pi(y)}{\pi(x)}}, \quad \text { as } t \rightarrow+\infty
$$

since $\pi(x), \pi(y) \leq 1$.
Proof of Theorem 5.2.20. The idea of the proof is to show that

$$
P^{t}(x, y) \leq 2 \sqrt{\frac{\pi(y)}{\pi(x)}} \mathbb{P}\left[S_{t} \geq \rho(x, y)\right]
$$

where again $\left(S_{t}\right)$ is simple random walk on $\mathbb{Z}$ started at 0 , and then use the Chernoff bound (Lemma 2.4.3).

By the local finiteness assumption, only a finite number of states can be reached by time $t$. Hence we can reduce the problem to a finite state space. More precisely, let $\tilde{V}=\{z \in V: \rho(x, z) \leq t\}$ and for $z, w \in \tilde{V}$

$$
\tilde{P}(z, w)= \begin{cases}P(z, w) & \text { if } z \neq w \\ P(z, z)+P(z, V \backslash \tilde{V}) & \text { otherwise }\end{cases}
$$

By construction $\tilde{P}$ is reversible with respect to $\tilde{\pi}=\pi / \pi(\tilde{V})$ on $\tilde{V}$. Because within time $t$ one never reaches a state $z$ where $P(z, V \backslash \tilde{V})>0$, by ChapmanKolmogorov (Theorem 1.1.20) and using the fact that $\tilde{\pi}(y) / \tilde{\pi}(x)=\pi(y) / \pi(x)$, it suffices to prove the result for $\tilde{P}$. Hence we assume without loss of generality that $V$ is finite with $|V|=n$.

To relate $\left(X_{t}\right)$ to simple random walk on $\mathbb{Z}$, we use a special representation of $P^{t}$ based on Chebyshev polynomials. For $\xi=\cos \theta \in[-1,1]$,

$$
T_{k}(\xi)=\cos k \theta
$$

is a Chebyshev polynomial of the first kind. Note that $\left|T_{k}(\xi)\right| \leq 1$ on $[-1,1]$ by definition. The classical trigonometric identity (to see this, write it in complex form)

$$
\cos ((k+1) \theta)+\cos ((k-1) \theta)=2 \cos \theta \cos (k \theta),
$$

implies the recursion

$$
T_{k+1}(\xi)+T_{k-1}(\xi)=2 \xi T_{k}(\xi),
$$

which in turn implies that $T_{k}$ is indeed a polynomial. It has degree $k$ from induction and the fact that $T_{0}(\xi)=1$ and $T_{1}(\xi)=\xi$. The connection to simple random walk on $\mathbb{Z}$ comes from the following somewhat miraculous representation (which does not rely on reversibility). Let $T_{k}(P)$ denote the polynomial $T_{k}$ evaluated at $P$ as a matrix polynomial.
Lemma 5.2.21.

$$
P^{t}=\sum_{k=-t}^{t} \mathbb{P}\left[S_{t}=k\right] T_{|k|}(P)
$$

Proof. It suffices to prove

$$
\xi^{t}=\sum_{k=-t}^{t} \mathbb{P}\left[S_{t}=k\right] T_{|k|}(\xi)
$$

as an identity of polynomials. By the binomial theorem (Appendix A),

$$
\xi^{t}=\left(\frac{e^{i \theta}+e^{-i \theta}}{2}\right)^{t}=\sum_{\ell=0}^{t} 2^{-t}\binom{t}{\ell}\left(e^{i \theta}\right)^{\ell}\left(e^{-i \theta}\right)^{t-\ell}=\sum_{k=-t}^{t} \mathbb{P}\left[S_{t}=k\right] e^{i k \theta}
$$

where we used that the probability that

$$
S_{t}=-t+2 \ell=(+1) \ell+(-1)(t-\ell),
$$

is the event of making $\ell$ steps to the right and $t-\ell$ steps to the left. Now take real parts on both sides and use that $\cos (k \theta)=\cos (-k \theta)$ to get the claim. (Put differently, $(\cos \theta)^{t}$ is the characteristic function $\mathbb{E}\left[e^{i \theta S_{t}}\right]$ of $S_{t}$.)

We bound $T_{k}(P)(x, y)$ as follows.
Lemma 5.2.22. It holds that

$$
T_{k}(P)(x, y)=0, \quad \forall k<\rho(x, y),
$$

and

$$
T_{k}(P)(x, y) \leq \sqrt{\frac{\pi(y)}{\pi(x)}}, \quad \forall k \geq \rho(x, y)
$$

Proof. Note that $T_{k}(P)(x, y)=0$ when $k<\rho(x, y)$ because $T_{k}(P)(x, y)$ is a function of the entries $P^{\ell}(x, y)$ for $\ell \leq k$, all of which are 0 .

We work on $\ell^{2}(V, \pi)$. Let $f_{1}, \ldots, f_{n}$ be an eigendecomposition of $P$ orthonormal with respect to the inner product $\langle\cdot, \cdot\rangle_{\pi}$ with eigenvalues $\lambda_{1}, \ldots, \lambda_{n} \in[-1,1]$. Such a decomposition exists by Theorem 5.2.1. Then $f_{1}, \ldots, f_{n}$ is also an eigendecomposition of the polynomial $T_{k}(P)$ with eigenvalues

$$
T_{k}\left(\lambda_{1}\right), \ldots, T_{k}\left(\lambda_{n}\right) \in[-1,1],
$$

by the definition of the Chebyshev polynomials. By decomposing any function $f=\sum_{i=1}^{n} \alpha_{i} f_{i}$ over this eigenbasis, that implies that

$$
\begin{align*}
\left\|T_{k}(P) f\right\|_{\pi}^{2} & =\left\|\sum_{i=1}^{n} \alpha_{i} T_{k}\left(\lambda_{i}\right) f_{i}\right\|_{\pi}^{2} \\
& =\sum_{i=1}^{n} \alpha_{i}^{2} T_{k}\left(\lambda_{i}\right)^{2} \\
& \leq \sum_{i=1}^{n} \alpha_{i}^{2} \\
& =\|f\|_{\pi}^{2} \tag{5.2.7}
\end{align*}
$$

where we used Parseval's identity (5.2.1) twice and the fact that $T_{k}\left(\lambda_{i}\right)^{2} \in[0,1]$.
Let $\delta_{z}$ denote the point mass at $z$. By Cauchy-Schwarz (Theorem B.4.8) and (5.2.7),

$$
T_{k}(P)(x, y)=\frac{\left\langle\delta_{x}, T_{k}(P) \delta_{y}\right\rangle_{\pi}}{\pi(x)} \leq \frac{\left\|\delta_{x}\right\|_{\pi}\left\|\delta_{y}\right\|_{\pi}}{\pi(x)}=\frac{\sqrt{\pi(x)} \sqrt{\pi(y)}}{\pi(x)}=\sqrt{\frac{\pi(y)}{\pi(x)}},
$$

for any $k$ (in particular for $k \geq \rho(x, y)$ ) and we have proved the claim.
Combining the two lemmas gives the result.
Remark 5.2.23. The local finiteness assumption is made for simplicity only. The result holds for any countable-space, reversible chain. See [LP16, Section 13.2].

Lower bound on mixing Let $\left(X_{t}\right)$ be an irreducible aperiodic (for now not necessarily reversible) Markov chain with finite state space $V$ and stationary distribution $\pi$. Recall that, for a fixed $0<\varepsilon<1 / 2$, the mixing time is

$$
\mathrm{t}_{\text {mix }}(\varepsilon)=\min \{t: d(t) \leq \varepsilon\},
$$

where

$$
d(t)=\max _{x \in V}\left\|P^{t}(x, \cdot)-\pi\right\|_{\mathrm{TV}} .
$$

It is intuitively clear that $\mathrm{t}_{\text {mix }}(\varepsilon)$ is at least of the order of the "diameter" of the transition graph of $P$. For $x, y \in V$, let $\rho(x, y)$ be the graph distance between $x$ and $y$ on the undirected version of the transition graph, that is, ignoring the orientation of the edges. With this definition, a shortest directed path from $x$ to $y$ contains at least $\rho(x, y)$ edges. Here we define the diameter of the transition graph as $\Delta:=$ $\max _{x, y \in V} \rho(x, y)$. Let $x_{0}, y_{0}$ be a pair of vertices achieving the diameter. Then we claim that $P^{\lfloor(\Delta-1) / 2\rfloor}\left(x_{0}, \cdot\right)$ and $P^{\lfloor(\Delta-1) / 2\rfloor}\left(y_{0}, \cdot\right)$ are supported on disjoint sets. To see this let

$$
A=\left\{z \in V: \rho\left(x_{0}, z\right)<\rho\left(y_{0}, z\right)\right\},
$$

be the set of states closer to $x_{0}$ than $y_{0}$. See Figure 5.2. By the triangle inequality for $\rho$, any $z$ such that $\rho\left(x_{0}, z\right) \leq\lfloor(\Delta-1) / 2\rfloor$ is in $A$, otherwise we would have $\rho\left(y_{0}, z\right) \leq \rho\left(x_{0}, z\right) \leq\lfloor(\Delta-1) / 2\rfloor$ and hence $\rho\left(x_{0}, y_{0}\right) \leq \rho\left(x_{0}, z\right)+\rho\left(y_{0}, z\right) \leq$ $2\lfloor(\Delta-1) / 2\rfloor<\Delta$, a contradiction. Similarly, if $\rho\left(y_{0}, z\right) \leq\lfloor(\Delta-1) / 2\rfloor$, then $z \in A^{c}$. By the triangle inequality for the total variation distance,

$$
\begin{align*}
d(\lfloor(\Delta-1) / 2\rfloor) & \geq \frac{1}{2}\left\|P^{\lfloor(\Delta-1) / 2\rfloor}\left(x_{0}, \cdot\right)-P^{\lfloor(\Delta-1) / 2\rfloor}\left(y_{0}, \cdot\right)\right\|_{\mathrm{TV}} \\
& \geq \frac{1}{2}\left\{P^{\lfloor(\Delta-1) / 2\rfloor}\left(x_{0}, A\right)-P^{\lfloor(\Delta-1) / 2\rfloor}\left(y_{0}, A\right)\right\} \\
& =\frac{1}{2}\{1-0\}=\frac{1}{2}, \tag{5.2.8}
\end{align*}
$$



Figure 5.2: The supports of $P^{\lfloor(\Delta-1) / 2\rfloor}\left(x_{0}, \cdot\right)$ and $P^{\lfloor(\Delta-1) / 2\rfloor}\left(y_{0}, \cdot\right)$ are contained in $A$ and $A^{c}$ respectively.
where we used (1.1.4) on the second line, so that:

## Claim 5.2.24.

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{\Delta}{2} .
$$

This bound is often far from the truth. Consider for instance simple random walk on a cycle of size $n$. The diameter is $\Delta=n / 2$. But Lemma 2.4.3 suggests that it takes time of order $\Delta^{2}$ to even reach the antipode of the starting point, let alone achieve stationarity. More generally, when $P$ is reversible, the "diffusive behavior" captured by the Varopoulos-Carne bound (Theorem 5.2.20) implies that the mixing time does indeed scale at least as the square of the diameter.

Assume that $P$ is reversible with respect to $\pi$ and has diameter $\Delta$. Letting $n=|V|$ and $\pi_{\text {min }}=\min _{x \in V} \pi(x)$, we then have the following.

Claim 5.2.25. The following lower bound holds

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{\Delta^{2}}{12 \log n+4\left|\log \pi_{\min }\right|}
$$

provided $n \geq \frac{16}{(1-2 \varepsilon)^{2}}$.

Proof. The proof is based on the same argument we used to derive our first diame-ter-based bound, except that the Varopoulos-Carne bound gives a better dependence on the diameter. Namely, let $x_{0}, y_{0}$, and $A$ be as above. By the VaropoulosCarne bound,

$$
P^{t}\left(x_{0}, A^{c}\right)=\sum_{z \in A^{c}} P^{t}\left(x_{0}, z\right) \leq \sum_{z \in A^{c}} 2 \sqrt{\frac{\pi(z)}{\pi\left(x_{0}\right)}} e^{-\frac{\rho^{2}\left(x_{0}, z\right)}{2 t}} \leq 2 n \pi_{\min }^{-1 / 2} e^{-\frac{\Delta^{2}}{8 t}},
$$

where we used that $\left|A^{c}\right| \leq n$ and $\rho\left(x_{0}, z\right) \geq \frac{\Delta}{2}$ for $z \in A^{c}$. For any

$$
\begin{equation*}
t<\frac{\Delta^{2}}{12 \log n+4\left|\log \pi_{\min }\right|}, \tag{5.2.9}
\end{equation*}
$$

we get that

$$
P^{t}\left(x_{0}, A^{c}\right) \leq 2 n \pi_{\min }^{-1 / 2} \exp \left(-\frac{3 \log n+\left|\log \pi_{\min }\right|}{2}\right)=\frac{2}{\sqrt{n}},
$$

or $P^{t}\left(x_{0}, A\right) \geq 1-\frac{2}{\sqrt{n}}$. Similarly, $P^{t}\left(y_{0}, A\right) \leq \frac{2}{\sqrt{n}}$ so that arguing as in (5.2.8)

$$
d(t) \geq \frac{1}{2}\left\{1-\frac{2}{\sqrt{n}}-\frac{2}{\sqrt{n}}\right\}=\frac{1}{2}-\frac{2}{\sqrt{n}} \geq \varepsilon
$$

for $t$ as in (5.2.9) and $n$ as in the statement.
Remark 5.2.26. The dependence on $\Delta$ and $\pi_{\min }$ in Claim 5.2 .25 cannot be improved. See [LP16, Section 13.3].

### 5.2.4 $\triangleright$ Randomized algorithms: Markov chain Monte Carlo and a quantitative ergodic theorem

In Markov chain Monte Carlo methods, one generates samples from a probability distribution of interest $\pi$ over some state space $V$ in order to estimate some of its properties, for example, its mean, by designing and then running a Markov chain with stationary distribution $\pi$. The Metropolis algorithm from Example 1.1.30 is a standard way of constructing such a chain. These techniques play a central role in Bayesian statistics in particular where $\pi$ is the so-called posterior distribution given the data.

We restrict ourselves here to finite $V$ and, without loss of generality, we assume that $V=[n]$. Let $P$ be an irreducible chain reversible with respect to a stationary
distribution $\pi=\left(\pi_{x}\right)_{x \in V}$. As previously, we work on $\ell^{2}(V, \pi)$. Let $f: V \rightarrow \mathbb{R}$ be a function in $\ell^{2}(V, \pi)$. Recall that

$$
\pi f=\sum_{x \in V} \pi_{x} f(x)
$$

Our goal is to estimate $\pi f$ from the sample path of the Markov chain $\left(X_{t}\right)_{t \geq 0}$ with transition matrix $P$. Indeed the ergodic theorem guarantees that

$$
\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right) \rightarrow \pi f
$$

almost surely as $T \rightarrow+\infty$ for any starting point. We derive a simple, quantitative version of this statement that provides insights into how long the chain needs to be run to get an accurate estimate in terms of the spectral gap.

Theorem 5.2.27 (Ergodic theorem: reversible case). Let $P=\left(P_{x, y}\right)_{x, y \in V}$ be an irreducible aperiodic transition matrix over a finite state space $V$ reversible with respect to the stationary distribution $\pi=\left(\pi_{x}\right)_{x \in V}$. Let $f: V \rightarrow \mathbb{R}$ be a function in $\ell^{2}(V, \pi)$. Then for any initial distribution $\mu=\left(\mu_{x}\right)_{x \in V}$

$$
\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right) \rightarrow \pi f
$$

in probability as $T \rightarrow+\infty$. Moreover, for any $\varepsilon>0$,

$$
\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\pi f\right| \geq \varepsilon\right] \leq \frac{9 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \gamma_{*}^{-1} \frac{1}{T}}{\left(\varepsilon-\pi_{\min }^{-1}\|f\|_{\infty} \gamma_{*}^{-1} \frac{1}{T}\right)^{2}},
$$

as $T \rightarrow+\infty$, where $\gamma_{*}>0$ is the absolute spectral gap of $P$.
Recall that, by Lemmas 5.2.3 and 5.2.4, we have $\gamma_{*}>0$ since $P$ is irreducible and aperiodic. We will first need the following lemma.

Lemma 5.2.28 (Convergence of the expectation). For any initial distribution $\mu=$ $\left(\mu_{x}\right)_{x \in V}$ and any $t$

$$
\left|\mathbb{E}\left[f\left(X_{t}\right)\right]-\pi f\right| \leq\left(1-\gamma_{*}\right)^{t} \pi_{\min }^{-1}\|f\|_{\infty}
$$

Proof. We have

$$
\left|\mathbb{E}\left[f\left(X_{t}\right)\right]-\pi f\right|=\left|\sum_{x} \sum_{y} \mu_{x} P_{x, y}^{t} f(y)-\sum_{y} \pi_{y} f(y)\right| .
$$

Because $\sum_{x} \mu_{x}=1$, the right-hand side is

$$
\begin{aligned}
& =\left|\sum_{x} \sum_{y} \mu_{x} P_{x, y}^{t} f(y)-\sum_{x} \sum_{y} \mu_{x} \pi_{y} f(y)\right| \\
& \leq \sum_{x} \mu_{x} \sum_{y}\left|P_{x, y}^{t}-\pi_{y}\right||f(y)|
\end{aligned}
$$

by the triangle inequality.
Now by (5.2.5) this is

$$
\begin{aligned}
& \leq \sum_{x} \mu_{x} \sum_{y}\left(1-\gamma_{*}\right)^{t} \frac{\pi_{y}}{\pi_{\min }}|f(y)| \\
& =\left(1-\gamma_{*}\right)^{t} \frac{1}{\pi_{\min }} \sum_{x} \mu_{x} \sum_{y} \pi_{y}|f(y)| \\
& \leq\left(1-\gamma_{*}\right)^{t} \pi_{\min }^{-1}\|f\|_{\infty} .
\end{aligned}
$$

That proves the claim.
Proof of Theorem 5.2.27. We use Chebyshev's inequality (Theorem 2.1.2), similarly to the proof of the $L^{2}$ weak law of large numbers (Theorem 2.1.6). In particular, we note that the $X_{t} \mathrm{~s}$ are not independent.

By Lemma 5.2.28, the expectation of the time average can be bounded as follows

$$
\begin{aligned}
\left|\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]-\pi f\right| & \leq \frac{1}{T} \sum_{t=1}^{T}\left|\mathbb{E}\left[f\left(X_{t}\right)\right]-\pi f\right| \\
& \leq \frac{1}{T} \sum_{t=1}^{T}\left(1-\gamma_{*}\right)^{t} \pi_{\min }^{-1}\|f\|_{\infty} \\
& \leq \pi_{\min }^{-1}\|f\|_{\infty} \frac{1}{T} \sum_{t=0}^{+\infty}\left(1-\gamma_{*}\right)^{t} \\
& =\pi_{\min }^{-1}\|f\|_{\infty} \gamma_{*}^{-1} \frac{1}{T} \rightarrow 0
\end{aligned}
$$

as $T \rightarrow+\infty$, since $\gamma_{*}>0$.
Next we bound the variance of the sum. We have
$\operatorname{Var}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]=\frac{1}{T^{2}} \sum_{t=1}^{T} \operatorname{Var}\left[f\left(X_{t}\right)\right]+\frac{2}{T^{2}} \sum_{1 \leq s<t \leq T} \operatorname{Cov}\left[f\left(X_{s}\right), f\left(X_{t}\right)\right]$.

We bound the variance and covariance terms separately.
To obtain convergence, a trivial bound on the variance suffices

$$
0 \leq \operatorname{Var}\left[f\left(X_{t}\right)\right] \leq \mathbb{E}\left[f\left(X_{t}\right)^{2}\right] \leq\|f\|_{\infty}^{2}
$$

Hence,

$$
0 \leq \frac{1}{T^{2}} \sum_{t=1}^{T} \operatorname{Var}\left[f\left(X_{t}\right)\right] \leq \frac{T\|f\|_{\infty}^{2}}{T^{2}} \rightarrow 0
$$

as $T \rightarrow+\infty$.
Bounding the covariance requires a more delicate argument. Fix $1 \leq s<t \leq$ $T$. The trick is to condition on $X_{s}$ and use the Markov Property (Theorem 1.1.18). By definition of the covariance, the tower property (Lemma B.6.16) and taking out what is known (Lemma B.6.13),

$$
\begin{aligned}
& \operatorname{Cov}\left[f\left(X_{s}\right), f\left(X_{t}\right)\right] \\
& =\mathbb{E}\left[\left(f\left(X_{s}\right)-\mathbb{E}\left[f\left(X_{s}\right)\right]\right)\left(f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right]\right)\right] \\
& =\sum_{x} \mathbb{E}\left[\left(f\left(X_{s}\right)-\mathbb{E}\left[f\left(X_{s}\right)\right]\right)\left(f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right]\right) \mid X_{s}=x\right] \mathbb{P}\left[X_{s}=x\right] \\
& =\sum_{x} \mathbb{E}\left[f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right] \mid X_{s}=x\right]\left(f(x)-\mathbb{E}\left[f\left(X_{s}\right)\right]\right) \mathbb{P}\left[X_{s}=x\right]
\end{aligned}
$$

We now use the time homogeneity of the chain to note that

$$
\begin{aligned}
& \mathbb{E}\left[f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right] \mid X_{s}=x\right] \\
& =\mathbb{E}\left[f\left(X_{t}\right) \mid X_{s}=x\right]-\mathbb{E}\left[f\left(X_{t}\right)\right] \\
& =\mathbb{E}\left[f\left(X_{t-s}\right) \mid X_{0}=x\right]-\mathbb{E}\left[f\left(X_{t}\right)\right]
\end{aligned}
$$

By Lemma 5.2.28,

$$
\begin{aligned}
& \left|\mathbb{E}\left[f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right] \mid X_{s}=x\right]\right| \\
& =\left|\mathbb{E}\left[f\left(X_{t-s}\right) \mid X_{0}=x\right]-\mathbb{E}\left[f\left(X_{t}\right)\right]\right| \\
& =\left|\left(\mathbb{E}\left[f\left(X_{t-s}\right) \mid X_{0}=x\right]-\pi f\right)-\left(\mathbb{E}\left[f\left(X_{t}\right)\right]-\pi f\right)\right| \\
& \leq\left|\mathbb{E}\left[f\left(X_{t-s}\right) \mid X_{0}=x\right]-\pi f\right|+\left|\mathbb{E}\left[f\left(X_{t}\right)\right]-\pi f\right| \\
& \leq\left(1-\gamma_{*}\right)^{t-s} \pi_{\min }^{-1}\|f\|_{\infty}+\left(1-\gamma_{*}\right)^{t} \pi_{\min }^{-1}\|f\|_{\infty} \\
& \leq 2\left(1-\gamma_{*}\right)^{t-s} \pi_{\text {min }}^{-1}\|f\|_{\infty},
\end{aligned}
$$

which does not depend on $x$. Plugging back above,

$$
\begin{aligned}
& \left|\operatorname{Cov}\left[f\left(X_{s}\right), f\left(X_{t}\right)\right]\right| \\
& \leq \sum_{x}\left|\mathbb{E}\left[f\left(X_{t}\right)-\mathbb{E}\left[f\left(X_{t}\right)\right] \mid X_{s}=x\right]\right|\left|f(x)-\mathbb{E}\left[f\left(X_{s}\right)\right]\right| \mathbb{P}\left[X_{s}=x\right] \\
& \leq 2\left(1-\gamma_{*}\right)^{t-s} \pi_{\min }^{-1}\|f\|_{\infty} \sum_{x}\left|f(x)-\mathbb{E}\left[f\left(X_{s}\right)\right]\right| \mathbb{P}\left[X_{s}=x\right] \\
& \leq 2\left(1-\gamma_{*}\right)^{t-s} \pi_{\min }^{-1}\|f\|_{\infty} \sum_{x} 2\|f\|_{\infty} \mathbb{P}\left[X_{s}=x\right] \\
& \leq 4\left(1-\gamma_{*}\right)^{t-s} \pi_{\min }^{-1}\|f\|_{\infty}^{2} .
\end{aligned}
$$

Returning to the sum over the covariances, the previous bound gives

$$
\begin{aligned}
& \left|\frac{2}{T^{2}} \sum_{1 \leq s<t \leq T} \operatorname{Cov}\left[f\left(X_{s}\right), f\left(X_{t}\right)\right]\right| \\
& \leq \frac{2}{T^{2}} \sum_{1 \leq s<t \leq T}\left|\operatorname{Cov}\left[f\left(X_{s}\right), f\left(X_{t}\right)\right]\right| \\
& \leq \frac{2}{T^{2}} \sum_{1 \leq s<t \leq T} 4\left(1-\gamma_{\star}\right)^{t-s} \pi_{\min }^{-1}\|f\|_{\infty}^{2} .
\end{aligned}
$$

To evaluate the sum we make the change of variable $h=t-s$ to get that the previous expression is

$$
\begin{aligned}
& \leq 4 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \frac{2}{T^{2}} \sum_{1 \leq s \leq T} \sum_{h=1}^{T-s}\left(1-\gamma_{*}\right)^{h} \\
& \leq 4 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \frac{2}{T^{2}} \sum_{1 \leq s \leq T} \sum_{h=0}^{+\infty}\left(1-\gamma_{*}\right)^{h} \\
& =4 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \frac{2}{T^{2}} \sum_{1 \leq s \leq T} \frac{1}{\gamma_{*}} \\
& =8 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \gamma_{*}^{-1} \frac{1}{T} \rightarrow 0,
\end{aligned}
$$

as $T \rightarrow+\infty$.
Combining the variance and covariance bounds, we have shown that

$$
\operatorname{Var}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right] \leq\|f\|_{\infty}^{2} \frac{1}{T}+8 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \gamma_{*}^{-1} \frac{1}{T} \leq 9 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \gamma_{*}^{-1} \frac{1}{T}
$$

For any $\varepsilon>0$

$$
\begin{aligned}
& \mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\pi f\right| \geq \varepsilon\right] \\
& =\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]+\left(\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]-\pi f\right)\right| \geq \varepsilon\right] \\
& \leq \mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]\right|+\left|\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]-\pi f\right| \geq \varepsilon\right] \\
& \leq \mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\mathbb{E}\left[\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)\right]\right| \geq \varepsilon-\pi_{\min }^{-1}\|f\|_{\infty} \gamma_{*}^{-1} \frac{1}{T}\right] .
\end{aligned}
$$

We can now apply Chebyshev's inequality to get

$$
\mathbb{P}\left[\left|\frac{1}{T} \sum_{t=1}^{T} f\left(X_{t}\right)-\pi f\right| \geq \varepsilon\right] \leq \frac{9 \pi_{\min }^{-1}\|f\|_{\infty}^{2} \gamma_{*}^{-1} \frac{1}{T}}{\left(\varepsilon-\pi_{\min }^{-1}\|f\|_{\infty} \gamma_{*}^{-1} \frac{1}{T}\right)^{2}} \rightarrow 0,
$$

as $T \rightarrow+\infty$.

### 5.2.5 Spectral radius

The results in this section have so far concerned finite state spaces. The countably infinite case presents a number of complications. We start with a few observations:

- Suppose $P$ is irreducible, aperiodic and positive recurrent. Then we know from the convergence theorem (Theorem 1.1.33) that, if $\pi$ is the stationary distribution, then for all $x$

$$
\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}} \rightarrow 0,
$$

as $t \rightarrow+\infty$. The convergence rate depends on the starting point $x$. In the infinite state space case, one typically needs to make that dependence explicit to get meaningful results. In particular the mixing time-as we have defined it-may not be a useful concept.

- In the transient and null recurrent cases, there is no stationary distribution to converge to by Theorem 3.1.20. Instead, we have the following by Theorem 3.1.21: if $P$ is an irreducible chain which is either transient or null recurrent, then we have that

$$
\lim _{t} P^{t}(x, y)=0
$$

for all $x, y \in V$.

- Conditions stronger than reversibility are needed for the spectral theoremin a form similar to what we used-to apply. Specifically, one needs that $P$ is a compact operator: whenever $\left(f_{n}\right)_{n} \in \ell^{2}(V, \pi)$ is a bounded sequence, then there exists a subsequence $\left(f_{n_{k}}\right)_{k}$ such that $\left(P f_{n_{k}}\right)$ converges in the norm. Unfortunately that is often not the case, as the next example illustrates, even in the reversible, positive recurrent setting.

Example 5.2.29 (A positive recurrent chain whose $P$ is not compact). For $p<$ $1 / 2$, let $\left(X_{t}\right)$ be the birth-death chain with $V:=\{0,1,2, \ldots\}, P(0,0):=1-p$, $P(0,1)=p, P(x, x+1):=p$ and $P(x, x-1):=1-p$ for all $x \geq 1$, and $P(x, y):=0$ if $|x-y|>1$. As can be checked by direct computation, $P$ is reversible with respect to the stationary distribution $\pi(x)=(1-\gamma) \gamma^{x}$ for $x \geq 0$ where $\gamma:=\frac{p}{1-p}$. For $j \geq 1$, define $g_{j}(x):=\pi(j)^{-1 / 2} \mathbf{1}_{\{x=j\}}$. Then $\left\|g_{j}\right\|_{\pi}^{2}=1$ for all $j$ so $\left\{g_{j}\right\}_{j}$ is bounded in $\ell^{2}(V, \pi)$. On the other hand,

$$
P g_{j}(x)=p \pi(j)^{-1 / 2} \mathbf{1}_{\{x=j-1\}}+(1-p) \pi(j)^{-1 / 2} \mathbf{1}_{\{x=j+1\}} .
$$

So

$$
\begin{aligned}
\left\|P g_{j}\right\|_{\pi}^{2} & =p^{2} \pi(j)^{-1} \pi(j-1)+(1-p)^{2} \pi(j)^{-1} \pi(j+1) \\
& =p^{2} \frac{1-p}{p}+(1-p)^{2} \frac{p}{1-p} \\
& =2 p(1-p) .
\end{aligned}
$$

Hence $\left\{P g_{j}\right\}_{j}$ is also bounded. However, for $j>\ell$

$$
\begin{aligned}
\left\|P g_{j}-P g_{\ell}\right\|_{\pi}^{2} & \geq(1-p)^{2} \pi(j)^{-1} \pi(j+1)+p^{2} \pi(\ell)^{-1} \pi(\ell-1) \\
& =2 p(1-p) .
\end{aligned}
$$

So $\left\{P g_{j}\right\}_{j}$ does not have a converging subsequence.
We will not say much about the spectral theory of infinite networks. In this subsection, we establish a relationship between the operator norm of $P$-which is related to its spectrum-and the decay of $P^{t}(x, y)$.

Let $\ell_{0}(V)$ be the set of real-valued functions on $V$ with finite support. It is dense in $\ell^{2}(V, \pi)$. Indeed let $v_{1}, v_{2}, \ldots$ be an enumeration of $V$ and, for $f \in$ $\ell^{2}(V, \pi)$, define $\left.f\right|_{n}\left(v_{i}\right):=f\left(v_{i}\right) \mathbf{1}_{i \leq n}$ to be $f$ restricted to $v_{1}, \ldots, v_{n}$. Then

$$
\begin{equation*}
\left\|f-\left.f\right|_{n}\right\|_{\pi}^{2}=\sum_{i=n+1}^{\infty} \pi\left(v_{i}\right) f\left(v_{i}\right)^{2} \rightarrow 0 \tag{5.2.10}
\end{equation*}
$$

as $n \rightarrow \infty$, since $\|f\|_{\pi}^{2}=\sum_{x} \pi(x) f(x)^{2}<+\infty$. We will also need the following

$$
\begin{equation*}
\left\|P f-P\left(\left.f\right|_{n}\right)\right\|_{\pi}^{2}=\left\|P\left(f-\left.f\right|_{n}\right)\right\|_{\pi}^{2} \leq\left\|f-\left.f\right|_{n}\right\|_{\pi}^{2} \rightarrow 0 \tag{5.2.11}
\end{equation*}
$$

where we used (5.2.2).
Definition 5.2.30 (Operator norm). The operator norm of $P$ is

$$
\|P\|_{\pi}=\sup \left\{\frac{\|P f\|_{\pi}}{\|f\|_{\pi}}: f \in \ell_{0}(V), f \neq 0\right\}
$$

By definition, for any $f \in \ell_{0}(V)$,

$$
\begin{equation*}
\|P f\|_{\pi} \leq\|P\|_{\pi}\|f\|_{\pi} \tag{5.2.12}
\end{equation*}
$$

The same can be seen to hold for any $f \in \ell^{2}(V, \pi)$ by considering the sequence $\left(\left.f\right|_{n}\right)_{n}$ and noting that $\left\|\left.f\right|_{n}\right\|_{\pi} \rightarrow\|f\|_{\pi}$ and $\left\|P\left(\left.f\right|_{n}\right)\right\|_{\pi} \rightarrow\|P f\|_{\pi}$ as $n \rightarrow \infty$ by (5.2.10), (5.2.11) and the triangle inequality. This latter observation explains why it suffices to restrict the supremum to $\ell_{0}$ in the definition of the norm.

Note that, by (5.2.2), $\|P\|_{\pi} \leq 1$. Note further that, if $V$ is finite or more generally if $\pi$ is summable, then we have in fact $\|P\|_{\pi}=1$ by taking $f \equiv 1$ above. When $P$ is self-adjoint, the norm $\|P\|_{\pi}$ is also equal to what is known as the spectral radius, that is, the radius of the smallest disk centered at 0 in the complex plane that contains the spectrum of $P$. We will not need to define what that means formally here. (But Exercise 5.5 asks for a proof in the setting of symmetric matrices.)

Our main result is the following.
Theorem 5.2.31 (Spectral radius). Let $P$ be irreducible and reversible with respect to $\pi>0$. Then

$$
\rho(P):=\limsup _{t} P^{t}(x, y)^{1 / t}=\|P\|_{\pi}
$$

In particular, the limit does not depend on $x, y$. Moreover, for all $t$,

$$
P^{t}(x, y) \leq \sqrt{\frac{\pi(y)}{\pi(x)}}\|P\|_{\pi}^{t}
$$

In the positive recurrent case (for instance if the chain is finite), we have $P^{t}(x, y) \rightarrow$ $\pi(y)>0$ and so $\rho(P)=1=\|P\|_{\pi}$. The theorem says that the equality between $\rho(P)$ and $\|P\|_{\pi}$ holds in general for reversible chains.

Proof of Theorem 5.2.31. To see that the limit does not depend on $x, y$, let $u, v, x, y \in$ $V$ and $k, m \geq 0$ such that $P^{m}(u, x)>0$ and $P^{k}(y, v)>0$. Then

$$
\begin{aligned}
& P^{t+m+k}(u, v)^{1 /(t+m+k)} \\
& \quad \geq\left(P^{m}(u, x) P^{t}(x, y) P^{k}(y, v)\right)^{1 /(t+m+k)} \\
& \quad \geq P^{m}(u, x)^{1 /(t+m+k)} P^{t}(x, y)^{1 / t} P^{k}(y, v)^{1 /(t+m+k)}
\end{aligned}
$$

which shows that $\limsup _{t} P^{t}(u, v)^{1 / t} \geq \lim \sup _{t} P^{t}(x, y)^{1 / t}$ for all $u, v, x, y$.
We first show that $\rho(P) \leq\|P\|_{\pi}$. Observe that applying (5.2.4) and (5.2.12) repeatedly gives that $P^{t}$ is self-adjoint and satisfies the inequality $\left\|P^{t}\right\|_{\pi} \leq\|P\|_{\pi}^{t}$. Because $\left\|\delta_{z}\right\|_{\pi}^{2}=\pi(z) \leq 1$, by Cauchy-Schwarz

$$
\pi(x) P^{t}(x, y)=\left\langle\delta_{x}, P^{t} \delta_{y}\right\rangle_{\pi} \leq\|P\|_{\pi}^{t}\left\|\delta_{x}\right\|_{\pi}\left\|\delta_{y}\right\|_{\pi}=\|P\|_{\pi}^{t} \sqrt{\pi(x) \pi(y)} .
$$

Hence $P^{t}(x, y) \leq \sqrt{\frac{\pi(y)}{\pi(x)}}\|P\|_{\pi}^{t}$ and

$$
\begin{aligned}
\rho(P) & =\underset{t}{\lim \sup P^{t}(x, y)^{1 / t}} \\
& \leq \limsup _{t}\left(\sqrt{\frac{\pi(y)}{\pi(x)}}\|P\|_{\pi}^{t}\right)^{1 / t} \\
& =\|P\|_{\pi}
\end{aligned}
$$

To establish the inequality in the other direction, we make a series of observations. Fix a nonzero $f \in \ell_{0}(V)$.

- By self-adjointness and Cauchy-Schwarz,

$$
\left\|P^{t+1} f\right\|_{\pi}^{2}=\left\langle P^{t+1} f, P^{t+1} f\right\rangle_{\pi}=\left\langle P^{t+2} f, P^{t} f\right\rangle_{\pi} \leq\left\|P^{t+2} f\right\|_{\pi}\left\|P^{t} f\right\|_{\pi}
$$

or

$$
\frac{\left\|P^{t+1} f\right\|_{\pi}}{\left\|P^{t} f\right\|_{\pi}} \leq \frac{\left\|P^{t+2} f\right\|_{\pi}}{\left\|P^{t+1} f\right\|_{\pi}} .
$$

So $\frac{\left\|P^{t+1} f\right\|_{\pi}}{\left\|P^{t} f\right\|_{\pi}}$ is non-decreasing and therefore has a limit $L \leq+\infty$. Moreover, for $t=0$, we get

$$
\begin{equation*}
\frac{\|P f\|_{\pi}}{\|f\|_{\pi}} \leq L \tag{5.2.13}
\end{equation*}
$$

so it suffices to prove $L \leq \rho(P)$.

- Observe that

$$
\left(\frac{\left\|P^{t} f\right\|_{\pi}}{\|f\|_{\pi}}\right)^{1 / t}=\left(\frac{\|P f\|_{\pi}}{\|f\|_{\pi}} \times \cdots \times \frac{\left\|P^{t} f\right\|_{\pi}}{\left\|P^{t-1} f\right\|_{\pi}}\right)^{1 / t} \rightarrow L
$$

so in fact

$$
L=\lim _{t}\left\|P^{t} f\right\|_{\pi}^{1 / t}
$$

- By self-adjointness again

$$
\left\|P^{t} f\right\|_{\pi}^{2}=\left\langle f, P^{2 t} f\right\rangle_{\pi}=\sum_{x} \pi(x) f(x) \sum_{y} f(y) P^{2 t}(x, y) .
$$

By definition of $\rho(P)$, for any $\varepsilon>0$, there is a $t$ large enough that

$$
P^{2 t}(x, y) \leq(\rho(P)+\varepsilon)^{2 t}
$$

for all $x, y$ in the support of $f$. For such a $t$, plugging back into the previous display

$$
\left\|P^{t} f\right\|_{\pi}^{1 / t} \leq(\rho(P)+\varepsilon)\left(\sum_{x} \pi(x)|f(x)| \sum_{y}|f(y)|\right)^{1 / 2 t} .
$$

The expression in parentheses on the right-hand side is finite because $f$ has finite support. Since $\varepsilon$ is arbitrary, we get

$$
\begin{equation*}
L=\lim _{t}\left\|P^{t} f\right\|_{\pi}^{1 / t} \leq \rho(P) . \tag{5.2.14}
\end{equation*}
$$

So, combining (5.2.13) and (5.2.14), we have shown that $\|P\|_{\pi} \leq \rho(P)$ and that concludes the proof.

Corollary 5.2.32. Let P be irreducible and reversible with respect to $\pi$. If $\|P\|_{\pi}<$ 1 , then $P$ is transient.

Proof. By Theorem 5.2.31, $P^{t}(x, x) \leq\|P\|_{\pi}^{t}$ so

$$
\sum_{t} P^{t}(x, x) \leq \sum_{t}\|P\|_{\pi}^{t}<+\infty .
$$

Let $\left(X_{t}\right)$ be a chain with transition matrix $P$. Because

$$
\sum_{t} P^{t}(x, x)=\mathbb{E}_{x}\left[\sum_{t} \mathbf{1}_{\left\{X_{t}=x\right\}}\right],
$$

we have that $\sum_{t} \mathbf{1}_{\left\{X_{t}=x\right\}}<+\infty, \mathbb{P}_{x}$-a.s., and $\left(X_{t}\right)$ is transient.

This is not an if and only if. Random walk on $\mathbb{Z}^{3}$ is transient, yet $P^{2 t}(0,0)=$ $\Theta\left(t^{-3 / 2}\right)$ so there $\|P\|_{\pi}=\rho(P)=1$.

In the non-reversible case, our definition of $\|P\|_{\pi}$ still makes sense with respect to any stationary measure $\pi$ (although $P$ is not self-adjoint). But the equality in Theorem 5.2.31 no longer holds in general.

Example 5.2.33 (Counter-example). Let $\left(X_{t}\right)$ be asymmetric random walk on $\mathbb{Z}$ with probability $p \in(1 / 2,1)$ of going to the right. Then both $\pi_{0}(x):=\left(\frac{p}{1-p}\right)^{x}$ and $\pi_{1}(x):=1$ define stationary measures, but the transition matrix $P$ is only reversible with respect to $\pi_{0}$.

Under $\pi_{1}$, we have $\|P\|_{\pi_{1}}=1$. Indeed, let $g_{n}(x):=\mathbf{1}_{\{|x| \leq n\}}$ and note that

$$
\left(P g_{n}\right)(x)=\mathbf{1}_{\{|x| \leq n-1\}}+p \mathbf{1}_{\{x=-n-1 \text { or }-n\}}+(1-p) \mathbf{1}_{\{x=n \text { or } n+1\}},
$$

so $\left\|g_{n}\right\|_{\pi_{1}}^{2}=2 n+1$ and $\left\|P g_{n}\right\|_{\pi_{1}}^{2} \geq 2(n-1)+1$. Hence

$$
\limsup _{n} \frac{\left\|P g_{n}\right\|_{\pi_{1}}}{\left\|g_{n}\right\|_{\pi_{1}}} \geq 1
$$

and $\|P\|_{\pi_{1}} \geq 1$. But we already showed that $\|P\|_{\pi_{1}} \leq 1$ in (5.2.2), so the claim follows.

On the other hand, $\mathbb{E}_{0}\left[X_{t}\right]=(2 p-1) t$. So the martingale $Z_{t}:=X_{t}-(2 p-1) t$ (see Example 3.1.29), as a sum of $t$ independent centered random variables in $\{-1-(2 p-1), 1-(2 p-1)\}$, satisfies the assumptions of the Azuma-Hoeffding inequality (Theorem 3.2.1) with increment bound $c_{t}:=2$. So

$$
\begin{aligned}
P^{t}(0,0)^{1 / t} & \leq \mathbb{P}_{0}\left[X_{t} \leq 0\right]^{1 / t} \\
& =\mathbb{P}_{0}\left[X_{t}-(2 p-1) t \leq-(2 p-1) t\right]^{1 / t} \\
& \leq e^{-\frac{2(2 p-1)^{2} t^{2}}{2^{2} t} \frac{1}{t}} .
\end{aligned}
$$

Therefore

$$
\limsup _{t} P^{t}(0,0)^{1 / t} \leq e^{-(2 p-1)^{2} / 2}<1
$$

### 5.3 Geometric bounds

The goal of this section is to relate the spectral gap to certain geometric properties of the underlying network, more specifically isoperimetric properties, that is,
relationships between the "volume" of sets and their "circumference." The classical isoperimetric inequality states that the area enclosed by any rectifiable simple closed curve in the plane is at most the length of the curve squared divided by $4 \pi$. Moreover equality is achieved if and only if the curve is a circle.

Remark 5.3.1. Here is an easy proof in the smooth case. Suppose $\boldsymbol{r}(s)=(x(s), y(s))$, $s \in[0,2 \pi]$, is the parametrization of a positively oriented, smooth, simple closed curve in the plane centered at the origin with arc-length $2 \pi$, where $\left\|\boldsymbol{r}^{\prime}(s)\right\|_{2}=1$ for all $s$, $\int_{0}^{2 \pi} \boldsymbol{r}(s) \mathrm{d} s=\mathbf{0}$ and $x(0)=x(2 \pi)=0$. By Green's theorem, the area enclosed by the curve is

$$
A=\int_{0}^{2 \pi} x(s) y^{\prime}(s) \mathrm{d} s=\frac{1}{2} \int_{0}^{2 \pi}\left[x(s)^{2}+y^{\prime}(s)^{2}-\left(x(s)-y^{\prime}(s)\right)^{2}\right] \mathrm{d} s,
$$

where we used that $2 a b=a^{2}+b^{2}-(a-b)^{2}$. By the one-dimensional Poincaré inequality (Remark 3.2.7),

$$
A \leq \frac{1}{2} \int_{0}^{2 \pi}\left[x(s)^{2}+y^{\prime}(s)^{2}\right] \mathrm{d} s \leq \frac{1}{2} \int_{0}^{2 \pi}\left[x^{\prime}(s)^{2}+y^{\prime}(s)^{2}\right] \mathrm{d} s=\pi,
$$

which is indeed the area of a circle of circumference $2 \pi$.
Edge expansion We define our isoperimetric quantity of interest. Let $\left(X_{t}\right)$ be a finite, irreducible Markov chain on $V$ reversible with respect to a stationary measure $\pi>0$. (In this section, we do not necessarily assume that $\pi$ is a probability distribution.) Let $P$ be its transition matrix. We think of $\left(X_{t}\right)$ as a random walk on the network $\mathcal{N}=(G, c)$ where $G$ is the transition graph and $c(x, y):=$ $\pi(x) P(x, y)=\pi(y) P(y, x)$.

For a subset $S \subseteq V$, we let the edge boundary of $S$ be
isoperimetric inequality

$$
\partial_{\mathrm{E}} S:=\left\{e=\{x, y\} \in E: x \in S, y \in S^{c}\right\}
$$

Let $g: E \rightarrow \mathbb{R}_{+}$be an edge weight function. For $F \subseteq E$ and $W \subseteq V$, we define

$$
|F|_{g}:=\sum_{e \in F} g(e) .
$$

and

$$
|W|_{h}:=\sum_{v \in W} h(v) .
$$

Finally, for $S \subseteq V$, we let

$$
\Phi_{\mathrm{E}}(S ; g, h):=\frac{\left|\partial_{\mathrm{E}} S\right|_{g}}{|S|_{h}} .
$$

Roughly speaking, this is the ratio of the "size of the boundary" of a set to its "volume."

Our main definition, the edge expansion constant, quantifies the worst such ratio. First, one last piece of notation: for disjoint subsets $S_{0}, S_{1} \subseteq V$, we let

$$
c\left(S_{0}, S_{1}\right):=\sum_{x_{0} \in S_{0}} \sum_{x_{1} \in S_{1}} c\left(x_{0}, x_{1}\right) .
$$

Definition 5.3.2 (Edge expansion). For a subset of states $S \subseteq V$, the edge expansion constant (or bottleneck ratio) of $S$ is

$$
\Phi_{\mathrm{E}}(S ; c, \pi)=\frac{\left|\partial_{\mathrm{E}} S\right|_{c}}{|S|_{\pi}}=\frac{c\left(S, S^{c}\right)}{\pi(S)} .
$$

We refer to ( $S, S^{c}$ ) as a cut. The edge expansion constant (or bottleneck ratio or Cheeger number or isoperimetric constant*) of $\mathcal{N}$ is

$$
\Phi_{*}:=\min \left\{\Phi_{\mathrm{E}}(S ; c, \pi): S \subseteq V, 0<\pi(S) \leq \frac{1}{2}\right\} .
$$

Intuitively, a small value of $\Phi_{*}$ suggests the existence of a "bottleneck" in $\mathcal{N}$. Conversely, a large value of $\Phi_{*}$ indicates that all sets "expand out." See Figure 5.3. Note that the quantity $\Phi_{\mathrm{E}}(S ; c, \pi)$ has a natural probabilistic interpretation: pick a stationary state and make one step according to the transition matrix; then $\Phi_{\mathrm{E}}(S ; c, \pi)$ is the conditional probability that, given that the first state is in $S$, the next one is in $S^{c}$.

Equivalently, the edge expansion constant can be expressed as

$$
\Phi_{*}:=\min \left\{\frac{c\left(S, S^{c}\right)}{\pi(S) \wedge \pi\left(S^{c}\right)}: S \subseteq V, 0<\pi(S)<1\right\} .
$$

Example 5.3.3 (Edge expansion: complete graph). Let $G=K_{n}$ be the complete graph on $n$ vertices. Let $c(x, y)=1 / n^{2}$ for all $x, y$ (corresponding to picking any vertex uniformly at random at the next step) and $\pi(x)=1 / n$ for all $x$. For simplicity, take $n$ even. Then for a subset $S$ of size $|S|=k$,

$$
\Phi_{\mathrm{E}}(S ; c, \pi)=\frac{\left|\partial_{\mathrm{E}} S\right|_{c}}{|S|_{\pi}}=\frac{k(n-k) / n^{2}}{k / n}=\frac{n-k}{n} .
$$

Thus, the minimum is achieved for $k=n / 2$ and

$$
\Phi_{*}=\frac{n-n / 2}{n}=\frac{1}{2} .
$$

[^8]

Figure 5.3: A bottleneck.

Dirichlet form, Rayleigh quotient, and normalized Laplacian We relate the edge expansion constant of $\mathcal{N}$ to the spectral gap of $P$. Recall that we denote by $\lambda_{1}, \ldots, \lambda_{n}$ the eigenvalues of $P$ in decreasing order. First, we adapt the variational characterization of Theorem 5.1.3 to the network setting.

The Dirichlet form is defined over $\ell^{2}(V, \pi)$ as the bilinear form

$$
\mathscr{D}(f, g):=\langle f,(I-P) g\rangle_{\pi} .
$$

The associated quadratic form, also known as Dirichlet energy, is $\mathscr{D}(f):=\mathscr{D}(f, f)$.

Note that, using stochasticity and reversibility,

$$
\begin{aligned}
\langle f,(I-P) f\rangle_{\pi}= & \langle f, f\rangle_{\pi}-\langle f, P f\rangle_{\pi} \\
= & \frac{1}{2} \sum_{x} \pi(x) f(x)^{2}+\frac{1}{2} \sum_{y} \pi(y) f(y)^{2}-\sum_{x, y} \pi(x) f(x) f(y) P(x, y) \\
= & \frac{1}{2} \sum_{x} f(x)^{2} \pi(x) \sum_{y} P(x, y) \\
& \quad+\frac{1}{2} \sum_{y} f(y)^{2} \pi(y) \sum_{x} P(y, x)-\sum_{x, y} \pi(x) f(x) f(y) P(x, y) \\
= & \frac{1}{2} \sum_{x, y} \pi(x) P(x, y) f(x)^{2} \\
& +\frac{1}{2} \sum_{x, y} \pi(x) P(x, y) f(y)^{2}-\sum_{x, y} \pi(x) P(x, y) f(x) f(y) \\
= & \frac{1}{2} \sum_{x, y} c(x, y)[f(x)-f(y)]^{2},
\end{aligned}
$$

which is indeed consistent with the expression we encountered previously in Theorem 3.3.25.

The Rayleigh quotient for $I-P$ over $\ell^{2}(V, \pi)$ is then
Rayleigh

$$
\begin{aligned}
\frac{\langle f,(I-P) f\rangle_{\pi}}{\langle f, f\rangle_{\pi}} & =\frac{\frac{1}{2} \sum_{x, y} c(x, y)[f(x)-f(y)]^{2}}{\sum_{x} \pi(x) f(x)^{2}} \\
& =\frac{\mathbf{z}^{T} \mathcal{L}_{\mathbf{z}}}{\mathbf{z}^{T} \mathbf{z}}
\end{aligned}
$$

where $\mathcal{L}$ is the normalized Laplacian of the network $\mathcal{N}$ and we defined the vector $\boldsymbol{z}=\left(z_{x}\right)_{x \in V}$ with $z_{x}:=\sqrt{\pi(x)} f(x)$. Consequently, the Courant-Fischer theorem (Theorem 5.1.3) in the form (5.1.5) gives the following. Here $\eta_{2}=1-\lambda_{2}=\gamma$ is the spectral gap of $P$, which can also be seen as the second smallest eigenvalue of $I-P$ (which has the same eigenfunctions as $P$ itself). We have

$$
\gamma=\inf \left\{\frac{\langle f,(I-P) f\rangle_{\pi}}{\langle f, f\rangle_{\pi}}: \pi f=0, f \neq 0\right\} .
$$

The infimum is achieved by the eigenfunction $f_{2}$ of $P$ corresponding to its second largest eigenvalue $\lambda_{2}$. (Recall from Lemma 5.2.5 that $\pi f_{2}=0$.)

We note further that if $\pi f=0$ then

$$
\langle f, f\rangle_{\pi}=\langle f-\pi f, f-\pi f\rangle_{\pi}=\operatorname{Var}_{\pi}[f]
$$

where the last expression denotes the variance under $\pi$. So the variational characterization of $\gamma$ implies that

$$
\operatorname{Var}_{\pi}[f] \leq \gamma^{-1} \mathscr{D}(f),
$$

for all $f$ such that $\pi f=0$. In fact, it holds for any $f$ by considering $f-\pi f$ and noticing that both sides are unaffected by subtracting a constant to $f$.

We have shown:
Theorem 5.3.4 (Poincaré inequality for $\mathcal{N}$ ). Let $P$ be finite, irreducible and reversible with respect to $\pi$. Then

$$
\begin{equation*}
\operatorname{Var}_{\pi}[f] \leq \gamma^{-1} \mathscr{D}(f), \tag{5.3.1}
\end{equation*}
$$

for all $f \in \ell^{2}(V \pi)$. Equality is achieved by the eigenfunction $f_{2}$ of $P$ corresponding to the second largest eigenvalue $\lambda_{2}$.

An inequality of the type

$$
\begin{equation*}
\operatorname{Var}_{\pi}[f] \leq C \mathscr{D}(f), \quad \forall f, \tag{5.3.2}
\end{equation*}
$$

is known as a Poincaré inequality, a simple version of which we encountered previously in Remark 3.2.7. To see the connection with that one-dimensional version, it will be convenient to work with directed edges. Let $\vec{E}$ be an orientation of $E$, that is, for each $e \in\{x, y\}, \vec{E}$ includes either $(x, y)$ or $(y, x)$ with associated weight $c(\vec{e}):=c(e)>0$ For a function $f: V \rightarrow \mathbb{R}$ and an edge $\vec{e}=(x, y) \in \vec{E}$, we define the "discrete gradient"

$$
\nabla f(\vec{e})=f(y)-f(x)
$$

With this notation, we can rewrite the Dirichlet energy as

$$
\begin{equation*}
\mathscr{D}(f)=\frac{1}{2} \sum_{x, y} c(x, y)[f(x)-f(y)]^{2}=\sum_{\vec{e}} c(\vec{e})[\nabla f(\vec{e})]^{2}, \tag{5.3.3}
\end{equation*}
$$

hence (5.3.1) is a network analogue of (3.2.7).

### 5.3.1 Cheeger's inequality

The edge expansion constant and the spectral gap are related through the following isoperimetric inequalities. The lower bound is known as Cheeger's inequality.

Theorem 5.3.5. Let P be a finite, irreducible, reversible Markov transition matrix and let $\gamma=1-\lambda_{2}$ be the spectral gap of $P$. Then

$$
\frac{\Phi_{*}^{2}}{2} \leq \gamma \leq 2 \Phi_{*}
$$

In terms of the relaxation time $\mathrm{t}_{\text {rel }}=\gamma^{-1}$, these inequalities have an intuitive meaning: the presence or absence of a bottleneck in the state space leads to slow or fast mixing respectively. We detail some applications to mixing times in the next subsections.

Before giving a proof of the theorem, we start with a trivial—yet insightfulexample.

Example 5.3.6 (Two-state chain). Let $V:=\{0,1\}$ and, for $\alpha, \beta \in(0,1)$,

$$
P:=\left(\begin{array}{cc}
1-\alpha & \alpha \\
\beta & 1-\beta
\end{array}\right)
$$

which has stationary distribution

$$
\pi:=\left(\frac{\beta}{\alpha+\beta}, \frac{\alpha}{\alpha+\beta}\right)
$$

Recall from Example 5.2.8 that the second right eigenvector is

$$
f_{2}:=\left(\sqrt{\frac{\alpha}{\beta}},-\sqrt{\frac{\beta}{\alpha}}\right)=\left(\sqrt{\frac{\pi_{1}}{\pi_{0}}},-\sqrt{\frac{\pi_{0}}{\pi_{1}}}\right),
$$

with eigenvalue $\lambda_{2}:=1-\alpha-\beta$, so the spectral gap is $\alpha+\beta$. Assume that $\beta \leq \alpha$. Then the bottleneck ratio is

$$
\Phi_{*}=\frac{c(0,1)}{\pi(0)}=P(0,1)=\alpha .
$$

Then Theorem 5.3.5 reads

$$
\frac{\alpha^{2}}{2} \leq \alpha+\beta \leq 2 \alpha,
$$

which is indeed satisfied for all $0<\beta \leq \alpha<1$. Note that the upper bound is tight when $\alpha=\beta$.

Proof of Theorem 5.3.5. We start with the upper bound. In view of the Poincaré inequality for $\mathcal{N}$ (Theorem 5.3.4), to get an upper bound on the spectral gap, it
suffices to plug in a well-chosen function $f$ in (5.3.1). Taking a hint from Example 5.3.6, for $S \subseteq V$ with $\pi(S) \in(0,1 / 2]$, we let

$$
f_{S}(x):= \begin{cases}-\sqrt{\frac{\pi\left(S^{c}\right)}{\pi(S)}} & x \in S, \\ \sqrt{\frac{\pi(S)}{\pi\left(S^{c}\right)}} & x \in S^{c} .\end{cases}
$$

Then

$$
\sum_{x} \pi(x) f_{S}(x)=\pi(S)\left[-\sqrt{\frac{\pi\left(S^{c}\right)}{\pi(S)}}\right]+\pi\left(S^{c}\right)\left[\sqrt{\frac{\pi(S)}{\pi\left(S^{c}\right)}}\right]=0
$$

and

$$
\sum_{x} \pi(x) f_{S}(x)^{2}=\pi(S)\left[-\sqrt{\frac{\pi\left(S^{c}\right)}{\pi(S)}}\right]^{2}+\pi\left(S^{c}\right)\left[\sqrt{\frac{\pi(S)}{\pi\left(S^{c}\right)}}\right]^{2}=1 .
$$

So $\operatorname{Var}_{\pi}\left[f_{S}\right]=1$. Hence, from Theorem 5.3.4,

$$
\begin{aligned}
\gamma & \leq \frac{\mathscr{D}\left(f_{S}\right)}{\operatorname{Var}_{\pi}\left[f_{S}\right]} \\
& =\frac{1}{2} \sum_{x, y} c(x, y)\left[f_{S}(x)-f_{S}(y)\right]^{2} \\
& =\sum_{x \in S, y \in S^{c}} c(x, y)\left[-\sqrt{\frac{\pi\left(S^{c}\right)}{\pi(S)}}-\sqrt{\frac{\pi(S)}{\pi\left(S^{c}\right)}}\right]^{2} \\
& =\sum_{x \in S, y \in S^{c}} c(x, y)\left[-\frac{\pi\left(S^{c}\right)+\pi(S)}{\sqrt{\pi(S) \pi\left(S^{c}\right)}}\right]^{2} \\
& =\frac{c\left(S, S^{c}\right)}{\pi(S) \pi\left(S^{c}\right)} \\
& \leq 2 \frac{c\left(S, S^{c}\right)}{\pi(S)},
\end{aligned}
$$

as claimed.
The other direction is trickier. Because we seek an upper bound on the edge expansion constant $\Phi_{*}$, our goal is to find a cut $\left(S, S^{c}\right)$ such that

$$
\begin{equation*}
\frac{c\left(S, S^{c}\right)}{\pi(S) \wedge \pi\left(S^{c}\right)} \leq \sqrt{2 \gamma} \tag{5.3.4}
\end{equation*}
$$

Because the eigenfunction $f_{2}$ achieves $\gamma$ in Theorem 5.3.4, it is natural to look to it for "good cuts." Thinking of $f_{2}$ as a one-dimensional embedding of the network, it turns out to be enough to consider only "sweep cuts" of the form $S:=\{v$ : $\left.f_{2}(v) \leq \theta\right\}$ for a threshold $\theta$. How to pick the right threshold is less obvious.

Here we use a probabilistic argument, that is, we construct a random cut $\left(Z, Z^{c}\right)$. Observe that it suffices that

$$
\begin{equation*}
\mathbb{E}\left[c\left(Z, Z^{c}\right)\right] \leq \sqrt{2 \gamma} \mathbb{E}\left[\pi(Z) \wedge \pi\left(Z^{c}\right)\right], \tag{5.3.5}
\end{equation*}
$$

since then $\mathbb{E}\left[\sqrt{2 \gamma} \pi(Z) \wedge \pi\left(Z^{c}\right)-c\left(Z, Z^{c}\right)\right] \geq 0$, which in turn implies that we have $\mathbb{P}\left[\sqrt{2 \gamma} \pi(Z) \wedge \pi\left(Z^{c}\right)-c\left(Z, Z^{c}\right) \geq 0\right]>0$ by the first moment principle (Theorem 2.2.1); in other words, there exists a cut satisfying (5.3.4).

We now describe the random cut $\left(Z, Z^{c}\right)$ :

1. (Cuts from $f_{2}$ ) Let again $f_{2}$ be an eigenfunction corresponding to the eigenvalue $\lambda_{2}$ of $P$ with $\left\|f_{2}\right\|_{\pi}^{2}=1$. Order the vertices $V:=\left\{v_{1}, \ldots, v_{n}\right\}$ in such a way that

$$
f_{2}\left(v_{i}\right) \leq f_{2}\left(v_{i+1}\right), \quad \forall i=1, \ldots, n-1 .
$$

As we described above, the function $f_{2}$ naturally produces a series of cuts ( $S_{i}, S_{i}^{c}$ ) where $S_{i}:=\left\{v_{1}, \ldots, v_{i}\right\}$. By definition of the bottleneck ratio,

$$
\begin{equation*}
\Phi_{*} \leq \frac{c\left(S_{i}, S_{i}^{c}\right)}{\pi\left(S_{i}\right) \wedge \pi\left(S_{i}^{c}\right)} . \tag{5.3.6}
\end{equation*}
$$

2. (Normalization) Let

$$
m:=\min \left\{i: \pi\left(S_{i}\right)>1 / 2\right\}
$$

and define the translated function

$$
f:=f_{2}-f_{2}\left(v_{m}\right) .
$$

We further set $g:=\alpha f$ where $\alpha>0$ is chosen so that

$$
g\left(v_{1}\right)^{2}+g\left(v_{n}\right)^{2}=1 .
$$

Note that, by construction, $g\left(v_{m}\right)=0$ and $g\left(v_{1}\right) \leq \cdots \leq g\left(v_{m}\right)=0 \leq$ $g\left(v_{m+1}\right) \leq \cdots \leq g\left(v_{n}\right)$. The function $g$ is related to $\gamma$ as follows:

Lemma 5.3.7.

$$
\frac{1}{2} \sum_{x, y} c(x, y)(g(x)-g(y))^{2} \leq \gamma \sum_{x} \pi(x) g(x)^{2} .
$$

Proof. By Theorem 5.3.4,

$$
\gamma=\frac{\mathscr{D}\left(f_{2}\right)}{\operatorname{Var}_{\pi}\left[f_{2}\right]} .
$$

Because neither the numerator nor the denominator is affected by adding a constant, we have also

$$
\gamma=\frac{\mathscr{D}(f)}{\operatorname{Var}_{\pi}[f]} .
$$

Furthermore, notice that a constant multiplying $f$ cancels out in the ratio so

$$
\gamma=\frac{\mathscr{D}(g)}{\operatorname{Var}_{\pi}[g]} .
$$

Now use the fact that $\operatorname{Var}_{\pi}[g] \leq \sum_{x} \pi(x) g(x)^{2}$.
3. (Random cut) Pick $\Theta$ in $\left[g\left(v_{1}\right), g\left(v_{n}\right)\right]$ with density $2|\theta|$. Note that

$$
\int_{g\left(v_{1}\right)}^{g\left(v_{n}\right)} 2|\theta| \mathrm{d} \theta=g\left(v_{1}\right)^{2}+g\left(v_{n}\right)^{2}=1 .
$$

Finally define

$$
Z:=\left\{v_{i}: g\left(v_{i}\right)<\Theta\right\} .
$$

The rest of the proof is calculations. We bound the expectations on both sides of (5.3.5).

## Lemma 5.3.8. The following hold:

(i)

$$
\mathbb{E}\left[\pi(Z) \wedge \pi\left(Z^{c}\right)\right]=\sum_{x} \pi(x) g(x)^{2}
$$

(ii)

$$
\mathbb{E}\left[c\left(Z, Z^{c}\right)\right] \leq\left(\frac{1}{2} \sum_{x, y} c(x, y)(g(x)-g(y))^{2}\right)^{1 / 2}\left(2 \sum_{x} \pi(x) g(x)^{2}\right)^{1 / 2}
$$

Lemmas 5.3.7 and 5.3.8 immediately imply (5.3.5) and that concludes the proof of Theorem 5.3.5. So it remains to prove this last lemma.

Proof of Lemma 5.3.8. We start with (i). By definition of $g, \Theta \leq 0$ implies that $\pi(Z) \wedge \pi\left(Z^{c}\right)=\pi(Z)$ and vice versa. Thus

$$
\begin{align*}
\mathbb{E}\left[\pi(Z) \wedge \pi\left(Z^{c}\right)\right] & =\mathbb{E}\left[\sum_{\ell<m} \pi\left(v_{\ell}\right) \mathbf{1}_{\left\{v_{\ell} \in Z\right\}} \mathbf{1}_{\{\Theta \leq 0\}}+\sum_{\ell \geq m} \pi\left(v_{\ell}\right) \mathbf{1}_{\left\{v_{\ell} \in Z^{c}\right\}} \mathbf{1}_{\{\Theta>0\}}\right] \\
& =\mathbb{E}\left[\sum_{\ell<m} \pi\left(v_{\ell}\right) \mathbf{1}_{\left\{g\left(v_{\ell}\right)<\Theta \leq 0\right\}}+\sum_{\ell \geq m} \pi\left(v_{\ell}\right) \mathbf{1}_{\left\{0<\Theta \leq g\left(v_{\ell}\right)\right\}}\right] \\
& =\sum_{\ell<m} \pi\left(v_{\ell}\right) \mathbb{P}\left[g\left(v_{\ell}\right)<\Theta \leq 0\right]+\sum_{\ell \geq m} \pi\left(v_{\ell}\right) \mathbb{P}\left[0<\Theta \leq g\left(v_{\ell}\right)\right] \\
& =\sum_{\ell<m} \pi\left(v_{\ell}\right) g\left(v_{\ell}\right)^{2}+\sum_{\ell \geq m} \pi\left(v_{\ell}\right) g\left(v_{\ell}\right)^{2} \\
& =\sum_{x} \pi(x) g(x)^{2}, \tag{5.3.7}
\end{align*}
$$

where we integrated over the density of $\Theta$ to obtain the fourth line.
We move on to (ii). To compute $\mathbb{E}\left[c\left(Z, Z^{c}\right)\right]$, we note that $x_{k} \in Z$ and $x_{\ell} \in Z^{c}$ if and only if $g\left(v_{k}\right)<\Theta \leq g\left(v_{\ell}\right)$. The probability of that event depends on the signs of $g\left(v_{k}\right)$ and $g\left(v_{\ell}\right)$. If $g\left(v_{k}\right) g\left(v_{\ell}\right) \geq 0$,

$$
\begin{aligned}
\mathbb{P}\left[g\left(v_{k}\right)<\Theta \leq g\left(v_{\ell}\right)\right] & =\left|g\left(v_{k}\right)^{2}-g\left(v_{\ell}\right)^{2}\right| \\
& =\left|g\left(v_{k}\right)-g\left(v_{\ell}\right)\right|\left|g\left(v_{k}\right)+g\left(v_{\ell}\right)\right| \\
& =\left|g\left(v_{k}\right)-g\left(v_{\ell}\right)\right|\left(\left|g\left(v_{k}\right)\right|+\left|g\left(v_{\ell}\right)\right|\right) .
\end{aligned}
$$

If $g\left(v_{k}\right) g\left(v_{\ell}\right)<0$,

$$
\begin{aligned}
\mathbb{P}\left[g\left(v_{k}\right)<\Theta \leq g\left(v_{\ell}\right)\right] & =g\left(v_{k}\right)^{2}+g\left(v_{\ell}\right)^{2} \\
& \leq g\left(v_{k}\right)^{2}+g\left(v_{\ell}\right)^{2}-2 g\left(v_{k}\right) g\left(v_{\ell}\right) \\
& =\left(g\left(v_{k}\right)-g\left(v_{\ell}\right)\right)^{2} \\
& =\left|g\left(v_{k}\right)-g\left(v_{\ell}\right)\right|\left(\left|g\left(v_{k}\right)\right|+\left|g\left(v_{\ell}\right)\right|\right) .
\end{aligned}
$$

We apply Cauchy-Schwarz to get

$$
\begin{aligned}
\mathbb{E}\left[c\left(Z, Z^{c}\right)\right]= & \sum_{k<\ell} c\left(v_{k}, v_{\ell}\right) \mathbb{P}\left[g\left(v_{k}\right)<\Theta \leq g\left(v_{\ell}\right)\right] \\
\leq & \sum_{k<\ell} c\left(v_{k}, v_{\ell}\right)\left|g\left(v_{k}\right)-g\left(v_{\ell}\right)\right|\left(\left|g\left(v_{k}\right)\right|+\left|g\left(v_{\ell}\right)\right|\right) \\
\leq & \left(\sum_{k<\ell} c\left(v_{k}, v_{\ell}\right)\left(g\left(v_{k}\right)-g\left(v_{\ell}\right)\right)^{2}\right)^{1 / 2} \\
& \times\left(\sum_{k<\ell} c\left(v_{k}, v_{\ell}\right)\left(\left|g\left(v_{k}\right)\right|+\left|g\left(v_{\ell}\right)\right|\right)^{2}\right)^{1 / 2}
\end{aligned}
$$

The expression in the first parentheses is equal to $\frac{1}{2} \sum_{x, y} c(x, y)(g(x)-g(y))^{2}$. So it remain to bound the expression in the second parentheses.

Note that

$$
(|g(x)|+|g(y)|)^{2}=2 g(x)^{2}+2 g(y)^{2}-(|g(x)|-|g(y)|)^{2} \leq 2 g(x)^{2}+2 g(y)^{2} .
$$

Therefore, since $\sum_{y} c(x, y)=\sum_{y} c(y, x)=\pi(x)$,

$$
\begin{aligned}
\sum_{k<\ell} c\left(v_{k}, v_{\ell}\right)\left(\left|g\left(v_{k}\right)\right|+\left|g\left(v_{\ell}\right)\right|\right)^{2} & \leq \frac{1}{2} \sum_{x, y} c(x, y)(|g(x)|+|g(y)|)^{2} \\
& \leq \sum_{x} \pi(x) g(x)^{2}+\sum_{y} \pi(y) g(y)^{2} \\
& =2 \sum_{x} \pi(x) g(x)^{2} .
\end{aligned}
$$

That concludes the proof.

### 5.3.2 $\triangleright$ Random walks: trees, cycles, and hypercubes revisited

We use the techniques of the previous subsection to bound the mixing time of random walk on some simple graphs. In particular we revisit the examples of Section 4.3.2.
$b$-ary tree Let $\left(Z_{t}\right)$ be lazy simple random walk on the $\ell$-level rooted $b$-ary tree, $\widehat{\mathbb{T}}_{b}^{\ell}$. The root, 0 , is on level 0 and the leaves, $L$, are on level $\ell$. All vertices have
degree $b+1$, except for the root which has degree $b$ and the leaves which have degree 1 . Recall that the stationary distribution is

$$
\begin{equation*}
\pi(x):=\frac{\delta(x)}{2(n-1)}, \tag{5.3.8}
\end{equation*}
$$

where $n$ is the number of vertices and $\delta(x)$ is the degree of $x$. We take $b=2$ to simplify.

It is intuitively clear that each edge of this graph constitutes a bottleneck, with the root being the most "balanced" one. Let $x_{0}$ be a leaf of $\widehat{\mathbb{T}}_{b}^{\ell}$ and let $A$ be the set of vertices "on the other side of the root (inclusively)," that is, vertices whose graph distance from $x_{0}$ is at least $\ell$. See Figure 4.7. Let $S$ be the remaining vertices. Then by symmetry $\pi(S) \leq 1 / 2$. Note that there is a single edge connecting $S$ and $S^{c}=$ $A$, namely, the edge linking 0 and the root of the subtree $T_{S}$ formed by the vertices in $S$. More precisely, let $v_{S}$ be the root of $T_{S}$. From (5.3.8), $P\left(v_{S}, 0\right)=\frac{1}{2} \cdot \frac{1}{3}=\frac{1}{6}$ (where the $1 / 2$ accounts for the laziness), $\pi\left(v_{S}\right)=\frac{3}{2 n-2}$, and, by symmetry,

$$
\pi(S)=\frac{(2 n-2-2) / 2}{2 n-2}=\frac{n-2}{2 n-2},
$$

where in the numerator we subtracted the degree of the root before dividing the sum of the remaining degrees by 2 . Hence

$$
\Phi_{*} \leq \frac{\frac{1}{6}\left(\frac{3}{2 n-2}\right)}{\frac{n-2}{2 n-2}}=\frac{1}{2(n-2)}
$$

By Theorem 5.3.5,

$$
\gamma \leq 2 \Phi_{*} \leq \frac{1}{n-2} \quad \text { and } \quad \mathrm{t}_{\text {rel }}=\gamma^{-1} \geq n-2
$$

Thus by Theorem 5.2.14 and the fact that the chain is lazy

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq\left(\mathrm{t}_{\text {rel }}-1\right) \log \left(\frac{1}{2 \varepsilon}\right)=\Omega(n)
$$

We showed in Section 4.3.2, using other techniques, that $\mathrm{t}_{\text {mix }}(\varepsilon)=\Theta(n)$.
Cycle Let $\left(Z_{t}\right)$ be lazy simple random walk on the cycle of size $n, \mathbb{Z}_{n}:=$ $\{0,1 \ldots, n-1\}$, where $i \sim j$ if $|j-i|=1(\bmod n)$. Assume $n$ is even.

Consider a subset of vertices $S$. Note that by symmetry $\pi(S)=\frac{|S|}{n}$. Moreover, for all $i \sim j, c(i, j)=\pi(i) P(i, j)=\frac{1}{n} \cdot \frac{1}{2} \cdot \frac{1}{2}=\frac{1}{4 n}$. Among all sets of size $|S|$, consecutive vertices minimize the size of the boundary. So

$$
\Phi_{*} \leq \frac{2 \frac{1}{4 n}}{\frac{\ell}{n}}=\frac{1}{2 \ell},
$$

for all $\ell \leq n / 2$. This expression is minimized for $\ell=n / 2$ so

$$
\Phi_{*}=\frac{1}{n} .
$$

By Theorem 5.3.5,

$$
\frac{1}{2 n^{2}}=\frac{\Phi_{*}^{2}}{2} \leq \gamma \leq 2 \Phi_{*}=\frac{2}{n}
$$

and

$$
\frac{n}{2} \leq \mathrm{t}_{\mathrm{rel}}=\gamma^{-1} \leq 2 n^{2}
$$

Thus by Theorem 5.2.14

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq\left(\mathrm{t}_{\text {rel }}-1\right) \log \left(\frac{1}{2 \varepsilon}\right)=\Omega(n)
$$

and

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \log \left(\frac{1}{\varepsilon \pi_{\min }}\right) \mathrm{t}_{\mathrm{rel}}=O\left(n^{2} \log n\right)
$$

We know from exact eigenvalue computations (see Section 5.2.2 where technically we considered the non-lazy chain; laziness only affects the relaxation time by a factor of 2) that in fact $\gamma=\frac{2 \pi^{2}}{n^{2}}+O\left(n^{-4}\right)$. We also showed in that section that $\mathrm{t}_{\text {mix }}(\varepsilon)=O\left(n^{2}\right)$. (Exercise 4.15 shows this is tight up to a constant factor.)

Hypercube Let $\left(Z_{t}\right)$ be lazy simple random walk on the $n$-dimensional hypercube $\mathbb{Z}_{2}^{n}:=\{0,1\}^{n}$ where $i \sim j$ if $\|i-j\|_{1}=1$.

To get a bound on the edge expansion constant, consider the set $S=\left\{x \in \mathbb{Z}_{2}^{n}\right.$ : $\left.x_{1}=0\right\}$. By symmetry $\pi(S)=\frac{1}{2}$. For each $i \sim j, c(i, j)=\frac{1}{2^{n}} \cdot \frac{1}{2} \cdot \frac{1}{n}=\frac{1}{n 2^{n+1}}$. Hence

$$
\Phi_{*} \leq \frac{2^{n-1} \frac{1}{n 2^{n+1}}}{\frac{1}{2}}=\frac{1}{2 n},
$$

where in the numerator we used that $|S|=2^{n-1}$. By Theorem 5.3.5,

$$
\gamma \leq 2 \Phi_{*} \leq \frac{1}{n}
$$

Thus by Theorem 5.2.14

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq\left(\mathrm{t}_{\text {rel }}-1\right) \log \left(\frac{1}{2 \varepsilon}\right)=\Omega(n)
$$

We know from exact eigenvalue computations (Section 5.2.2) that in fact $\gamma=\frac{1}{n}$.
We also showed in Section 4.3.2 that $\mathrm{t}_{\text {mix }}(\varepsilon)=\Theta(n \log n)$.

### 5.3.3 $\triangleright$ Random graphs: existence of an expander family and application to mixing

In many applications, it is useful to construct "bottleneck-free" graphs. In particular, random walks mix rapidly on such graphs. Formally:

Definition 5.3.9 (Expander family). Let $\left\{G_{n}\right\}_{n}$ be a collection of finite d-regular graphs with $\lim _{n}\left|V\left(G_{n}\right)\right|=+\infty$, where $V\left(G_{n}\right)$ is the vertex set of $G_{n}$. Let

$$
\Phi_{*}\left(G_{n}\right):=\min \left\{\frac{\left|\partial_{\mathrm{E}} S\right|}{d|S|}: S \subseteq V\left(G_{n}\right), 0<|S| \leq \frac{\left|V\left(G_{n}\right)\right|}{2}\right\}
$$

denote the edge expansion constant of $G_{n}$ with unit conductances. ${ }^{\dagger}$ Let $\alpha>0$. We say that $\left\{G_{n}\right\}_{n}$ is a $(d, \alpha)$-expander family iffor all $n$

$$
\Phi_{*}\left(G_{n}\right) \geq \alpha .
$$

The key point of the definition is that the edge expansion constant of all graphs in an expander family is bounded away from 0 uniformly in $n$. Note that it is trivial to construct such a family if we drop the bounded degree assumption: the edge expansion constant of the complete graph $K_{n}$ is $1 / 2$ by Example 5.3.3. On the other hand, it is far from obvious that one can construct a family of sparse graphs (i.e., such that $\left.\left|E\left(G_{n}\right)\right|=O\left(\left|V\left(G_{n}\right)\right|\right)\right)$ with an edge expansion constant uniformly bounded away from 0 . It turns out that a simple probabilistic construction does the trick.

We will need the following definition. For a subset $S \subseteq V$, we let the vertex boundary of $S$ be

$$
\partial_{\mathrm{V}} S:=\left\{y \in S^{c}: \exists x \in S \text { s.t. } x \sim y\right\} .
$$

[^9]

Figure 5.4: A draw from Pinsker's model.

Existence of expander graphs For simplicity, we allow multigraphs (i.e., $E$ is a multiset; or, put differently, there can be multiple edges between the same two vertices) and consider the case $d=3$. We construct a random bipartite multigraph $G_{n}=\left(L_{n}, R_{n}, E_{n}\right)$ on $2 n$ vertices known as Pinsker's model. Denote the vertices by $L_{n}=\left\{\ell_{1}, \ldots, \ell_{n}\right\}$ and $R_{n}=\left\{r_{1}, \ldots, r_{n}\right\}$. Let $\sigma_{n}^{1}$ and $\sigma_{n}^{2}$ be independent uniform random permutations of $[n]$. The edge set of $G_{n}$ is given by

$$
E_{n}:=\left\{\left(\ell_{i}, r_{i}\right): i \in[n]\right\} \cup\left\{\left(\ell_{i}, r_{\sigma_{n}^{1}(i)}\right): i \in[n]\right\} \cup\left\{\left(\ell_{i}, r_{\sigma_{n}^{2}(i)}\right): i \in[n]\right\} .
$$

In words, $G_{n}$ is a union of three independent uniform perfect matchings (and its vertices are labeled so that one of the matchings is $\left.\left\{\left(\ell_{i}, r_{i}\right)\right\}_{i}\right)$. See Figure 5.4. Observe that, as a multigraph, all vertices of $G_{n}$ have degree 3. We show that there exists $\alpha>0$ such that, for all $n$ large enough, with positive (in fact, high) probability $G_{n}$ has an edge expansion constant bounded below by $\alpha$. In particular, such a $G_{n}$ exists for all $n$ large enough and, thus, there exists a ( $3, \alpha$ )-expander family.

Claim 5.3.10 (Pinsker's model: edge expansion constant). There exists $\alpha>0$ such that

$$
\lim _{n} \mathbb{P}\left[\Phi_{*}\left(G_{n}\right) \geq \alpha\right]=1
$$

Proof. For convenience, assume $n$ is even. We need to show that with probability going to 1 , for any $S$ with $|S| \leq\left|V\left(G_{n}\right)\right| / 2=n$, we have $\left|\partial_{\mathrm{E}} S\right| \geq \alpha d|S|$ for some $\alpha>0$. We first reduce the proof to a statement about sets of vertices lying on one side of $G_{n}$.

Lemma 5.3.11. There is $\beta>0$ such that

$$
\lim _{n} \mathbb{P}\left[\left|\partial_{\mathrm{V}} K\right| \geq(1+\beta)|K|, \forall K \subseteq L,|K| \leq n / 2\right]=1
$$

The same holds for $R$.
Before proving Lemma 5.3.11, we argue that it implies Claim 5.3.10. Note that the lemma concerns the vertex boundary of $K$. To relate the latter to the edge boundary, let $S$ with $|S| \leq n$, and let $S_{L}:=S \cap L$ and $S_{R}:=S \cap R$. For any subset $K \subseteq S_{L}$, the size of the edge boundary of $S$ can be bounded below as follows

$$
\begin{equation*}
\left|\partial_{\mathrm{E}} S\right| \geq\left|\partial_{\mathrm{V}} K\right|-\left|S_{R}\right|, \tag{5.3.9}
\end{equation*}
$$

where we took into account that the vertices of $\partial_{\mathrm{V}} K$ in $S_{R}$ do not contribute to the edge boundary, but the others do as they are incident with at least one edge in $\partial_{\mathrm{E}} S$. It remains to find a good $K$.

Assume without loss of generality that $\left|S_{L}\right| \geq\left|S_{R}\right|$ (in the other case, just interchange the roles of $L$ and $R$ ), and suppose that the event in the lemma holds. In particular, $\left|S_{R}\right| \leq|S| / 2$. We claim that there is a subset $K$ of $S_{L}$ such that

$$
\begin{equation*}
\left|S_{R}\right| \leq|S| / 2 \leq|K| \leq n / 2 . \tag{5.3.10}
\end{equation*}
$$

There are two cases:

- If $\left|S_{L}\right|<n / 2$, then take $K=S_{L}$. It follows that $|K|=\left|S_{L}\right| \geq|S| / 2$.
- If $\left|S_{L}\right| \geq n / 2$, then let $K$ be any subset of $S_{L}$ of size $n / 2$. Since $|S| \leq n$, it follows that $|K|=n / 2 \geq|S| / 2$.

Under the event in the lemma, $\left|\partial_{\mathrm{V}} K\right| \geq(1+\beta)|K|$.
Going back to (5.3.9), using the lower bound on $\left|\partial_{\mathrm{V}} K\right|$ and (5.3.10), we get

$$
\left|\partial_{\mathrm{E}} S\right| \geq(1+\beta)|K|-|K|=\beta|K| \geq \frac{\beta}{2}|S|=\alpha|S|,
$$

where we set $\alpha=\beta / 2$. Since this holds for any set $S$ with $|S| \leq n$, we have proved Claim 5.3.10.

It remains to prove the lemma.


Figure 5.5: Illustration of the main step in proof of the lemma.

Proof of Lemma 5.3.11. Let $K \subseteq L$ with $k:=|K| \leq n / 2$. Without loss of generality assume $K=\left\{\ell_{1}, \ldots, \ell_{k}\right\}$. Observe that, by construction, $\partial_{\mathrm{V}} K \supseteq K^{\prime}$ where $K^{\prime}=\left\{r_{1}, \ldots, r_{k}\right\}$. We analyze the "bad event"

$$
\mathcal{B}_{K}:=\left\{\left|\partial_{\mathrm{V}} K\right| \leq k+\lfloor\beta k\rfloor\right\},
$$

by considering all subsets of $\left\{r_{k+1}, \ldots, r_{n}\right\}$ of size $\lfloor\beta k\rfloor$ and bounding the probability that all edges out of $K$ fall into one of them and $K^{\prime}$. Note that there are $\binom{n-k}{[\beta k\rfloor}$ such subsets. See Figure 5.5.

Since $\sigma_{n}^{1}$ and $\sigma_{n}^{2}$ are uniform and independent, they each match $K$ to a uniformly chosen subset of the same size in $R$ and we have by a union bound

$$
\mathbb{P}\left[\mathcal{B}_{K}\right] \leq\binom{ n-k}{\lfloor\beta k\rfloor}\left[\frac{\binom{k+\lfloor\beta k\rfloor}{ k}}{\binom{n}{k}}\right]^{2} \leq\binom{ n}{\lfloor\beta k\rfloor} \frac{\binom{k+\lfloor\beta k\rfloor}{\lfloor\beta k\rfloor}^{2}}{\binom{n}{k}^{2}},
$$

where we used that $\binom{n}{s}=\binom{n}{n-s}$.

Taking a union bound again, this time over $K$ s, we have

$$
\begin{align*}
\mathbb{P}\left[\exists K \subseteq L,|K| \leq n / 2,\left|\partial_{\mathrm{V}} K\right| \leq\right. & (1+\beta)|K|] \\
& \leq \sum_{K \subseteq L,|K| \leq n / 2} \mathbb{P}\left[\mathcal{B}_{K}\right] \\
& \leq \sum_{k=1}^{n / 2}\binom{n}{k}\binom{n}{\lfloor\beta k\rfloor} \frac{\binom{k+\lfloor\beta k\rfloor}{\lfloor\beta k\rfloor}^{2}}{\binom{n}{k}^{2}} \tag{5.3.11}
\end{align*}
$$

We use the bound $\frac{n^{s}}{s^{s}} \leq\binom{ n}{s} \leq \frac{e^{s} n^{s}}{s^{s}} \leq \frac{e^{t} n^{t}}{t^{t}}$ for $s \leq t<n$ (see Appendix A; to see the last inequality, note that $\frac{\mathrm{d}}{\mathrm{d} t} \log \left(\frac{e^{t} n^{t}}{t^{t}}\right)=\log \left(\frac{n}{t}\right)>0$ for $\left.0<t<n\right)$. We obtain that the sum in the last display is bounded as

$$
\begin{align*}
\sum_{k=1}^{n / 2}\binom{n}{k}\binom{n}{\lfloor\beta k\rfloor} \frac{\binom{k+\lfloor\beta k\rfloor}{\lfloor\beta k\rfloor}^{2}}{\binom{n}{k}^{2}} & =\sum_{k=1}^{n / 2}\binom{n}{\lfloor\beta k\rfloor} \frac{\binom{k+\lfloor\beta k\rfloor}{\lfloor\beta k\rfloor}^{2}}{\binom{n}{k}} \\
& \leq \sum_{k=1}^{n / 2} \frac{e^{\beta k} n^{\beta k}}{(\beta k)^{\beta k}} \frac{\left(\frac{e^{\beta k}(k+\beta k)^{\beta k}}{(\beta k)^{\beta k}}\right)^{2}}{\frac{n^{k}}{k^{k}}} \\
& \leq \sum_{k=1}^{n / 2}\left(\frac{k}{n}\right)^{k(1-\beta)}\left(\frac{e^{3}(1+\beta)^{2}}{\beta^{3}}\right)^{\beta k} \\
& =\sum_{k=1}^{\infty} f_{n}(k) \tag{5.3.12}
\end{align*}
$$

where we defined

$$
f_{n}(k):=\mathbf{1}_{\{k \leq n / 2\}}\left(\frac{k}{n}\right)^{k(1-\beta)}\left(\frac{e^{3}(1+\beta)^{2}}{\beta^{3}}\right)^{\beta k}
$$

Let also

$$
g(k):=\left[\left(\frac{1}{2}\right)^{1-\beta}\left(\frac{e^{3}(1+\beta)^{2}}{\beta^{3}}\right)^{\beta}\right]^{k}
$$

and notice that for $\beta$ small enough

$$
\left|f_{n}(k)\right| \leq g(k), \quad \forall k
$$

since $\frac{k}{n} \leq \frac{1}{2}$ for $k \leq \frac{n}{2}$ and

$$
\gamma_{\beta}:=\left(\frac{1}{2}\right)^{1-\beta}\left(\frac{e^{3}(1+\beta)^{2}}{\beta^{3}}\right)^{\beta}<1
$$

using that $\beta^{\beta} \rightarrow 1$ as $\beta \rightarrow 0$. Moreover, for each $k$,

$$
f_{n}(k) \rightarrow 0
$$

as $n \rightarrow+\infty$, and

$$
\sum_{k=1}^{\infty} g(k) \leq \frac{1}{1-\gamma_{\beta}}<+\infty
$$

Hence, by the dominated convergence theorem (Theorem B.4.7), combining (5.3.11) and (5.3.12) we get

$$
\mathbb{P}\left[\exists K \subseteq L,|K| \leq n / 2,\left|\partial_{\mathrm{V}} K\right| \leq(1+\beta)|K|\right]=\sum_{k=1}^{\infty} f_{n}(k) \rightarrow 0 .
$$

That concludes the proof.
That concludes the proof of Claim 5.3.10.
Claim 5.3.10 implies:
Theorem 5.3.12 (Existence of expander family). For $\alpha>0$ small enough, there exists a $(3, \alpha)$-expander (multigraph) family.

Proof. By Claim 5.3.10, for all $n$ large enough, there exists $G_{n}$ with $\Phi_{*}\left(G_{n}\right) \geq \alpha$ for some fixed $\alpha>0$.

Fast mixing on expander graphs As we mentioned above, an important property of an expander graph is that random walk on such a graph mixes rapidly. We make this precise.

Claim 5.3.13 (Mixing on expanders). Let $\left\{G_{n}\right\}$ be a $(d, \alpha)$-expander family. Then $\mathrm{t}_{\text {mix }}(\varepsilon)=\Theta\left(\log \left|V\left(G_{n}\right)\right|\right)$, where the constant depends on $\varepsilon$ and $\alpha$.

Proof. Because of the degree assumption, random walk on $G_{n}$ is reversible with respect to the uniform distribution (see Example 1.1.29). So, by Theorems 5.2.14 and 5.3.5, the mixing time is upper bounded by

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \log \left(\frac{1}{\varepsilon \pi_{\text {min }}}\right) \mathrm{t}_{\text {rel }} \leq \log \left(\frac{\left|V\left(G_{n}\right)\right|}{\varepsilon}\right) 2 \alpha^{-2}=O\left(\log \left|V\left(G_{n}\right)\right|\right) .
$$

By the diameter-based lower bound on the mixing time for reversible chains (Claim 5.2.25), for $n$ large enough

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{\Delta^{2}}{12 \log \left|V\left(G_{n}\right)\right|+4\left|\log \pi_{\min }\right|}
$$

where $\Delta$ is the diameter of $G_{n}$. For a $d$-regular graph $G_{n}$, the diameter is at least $\log \left|V\left(G_{n}\right)\right|$. Indeed, by induction, the number of vertices within graph distance $k$ of any vertex is at most $d^{k}$. For $d^{k}$ to be greater than $\left|V\left(G_{n}\right)\right|$, we need $k \geq$ $\log _{d}\left|V\left(G_{n}\right)\right|$. Finally,

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq \frac{\left(\log _{d}\left|V\left(G_{n}\right)\right|\right)^{2}}{16 \log \left|V\left(G_{n}\right)\right|}=\Omega\left(\log \left|V\left(G_{n}\right)\right|\right)
$$

That concludes the proof.

### 5.3.4 $\triangleright$ Ising model: Glauber dynamics on complete graphs and expanders

Let $G=(V, E)$ be a finite, connected graph with maximal degree $\bar{\delta}$. Define $\mathcal{X}:=$ $\{-1,+1\}^{V}$. Recall from Example 1.2.5 that the (ferromagnetic) Ising model on $V$ with inverse temperature $\beta$ is the probability distribution over spin configurations $\sigma \in \mathcal{X}$ given by

$$
\mu_{\beta}(\sigma):=\frac{1}{\mathcal{Z}(\beta)} e^{-\beta \mathcal{H}(\sigma)},
$$

where

$$
\mathcal{H}(\sigma):=-\sum_{i \sim j} \sigma_{i} \sigma_{j},
$$

is the Hamiltonian and

$$
\mathcal{Z}(\beta):=\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}(\sigma)},
$$

is the partition function. In this context, recall that vertices are often referred to as sites. The single-site Glauber dynamics of the Ising model (Definition 1.2.8) is the Markov chain on $\mathcal{X}$ which, at each time, selects a site $i \in V$ uniformly at random and updates the spin $\sigma_{i}$ according to $\mu_{\beta}(\sigma)$ conditioned on agreeing with $\sigma$ at all sites in $V \backslash\{i\}$. Specifically, for $\gamma \in\{-1,+1\}, i \in V$, and $\sigma \in \mathcal{X}$, let $\sigma^{i, \gamma}$ be the configuration $\sigma$ with the state at $i$ being set to $\gamma$. Then, letting $n=|V|$, the transition matrix of the Glauber dynamics is

$$
Q_{\beta}\left(\sigma, \sigma^{i, \gamma}\right):=\frac{1}{n} \cdot \frac{e^{\gamma \beta S_{i}(\sigma)}}{e^{-\beta S_{i}(\sigma)}+e^{\beta S_{i}(\sigma)}}=\frac{1}{n}\left\{\frac{1}{2}+\frac{1}{2} \tanh \left(\gamma \beta S_{i}(\sigma)\right)\right\},
$$

where

$$
S_{i}(\sigma):=\sum_{j \sim i} \sigma_{j} .
$$

All other transitions have probability 0 . Recall that this chain is irreducible and reversible with respect to $\mu_{\beta}$. In particular $\mu_{\beta}$ is the stationary distribution of $Q_{\beta}$.

We showed in Claim 4.3.15 that the Glauber dynamics is fast mixing at high temperature. More precisely we proved that $\mathrm{t}_{\text {mix }}(\varepsilon)=O(n \log n)$ when $\beta<\bar{\delta}^{-1}$. Here we prove a converse: at low temperature, graphs with good enough expansion properties produce exponentially slow mixing of the Glauber dynamics.

## Curie-Weiss model

Let $G=K_{n}$ be the complete graph on $n$ vertices. In this case, the Ising model is often referred to as the Curie-Weiss model. It is natural to scale $\beta$ with $n$. We define $\alpha:=\beta(n-1)$. Since $\bar{\delta}=n-1$, we have that, when $\alpha<1, \beta=\frac{\alpha}{n-1}<\bar{\delta}^{-1}$ so $\mathrm{t}_{\text {mix }}(\varepsilon)=O(n \log n)$. In the other direction, we prove:

Claim 5.3.14 (Curie-Weiss model: slow mixing at low temperature). For $\alpha>1$, $\mathrm{t}_{\text {mix }}(\varepsilon)=\Omega(\exp (r(\alpha) n))$ for some function $r(\alpha)>0$ not depending on $n$.

Proof. We first prove exponential mixing when $\alpha$ is large enough, an argument which will be useful in the generalization to expander graphs below.

The idea of the proof is to bound the edge expansion constant and use Theorem 5.3.5. To simplify the proof, assume $n$ is odd. We denote the edge expansion constant of the chain by $\Phi_{*}^{\mathcal{X}}$ to avoid confusion with that of the base graph $G$. Intuitively, because the spins tend to align strongly at low temperature, it takes a considerable amount of time to travel from a configuration with a majority of -1 s to a configuration with a majority of +1 s . Because the model tends to prefer agreeing spins but does not favor any particular spin, a natural place to look for a bottleneck is the set

$$
\mathcal{M}:=\left\{\sigma \in \mathcal{X}: \sum_{i} \sigma_{i}<0\right\},
$$

where the quantity $m(\sigma):=\sum_{i} \sigma_{i}$ is called the magnetization. Note that the magnetization is positive if and only if a majority of spins are +1 and that it forms a Markov chain by itself. So the boundary of the set $\mathcal{M}$ must be crossed to travel from configurations with mostly -1 spins to configurations with mostly -1 spins.

Observe further that $\mu_{\beta}(\mathcal{M})=1 / 2$. The edge expansion constant is hence bounded by

$$
\begin{equation*}
\Phi_{*}^{\mathcal{X}} \leq \frac{\sum_{\sigma \in \mathcal{M}, \sigma^{\prime} \notin \mathcal{M}} \mu_{\beta}(\sigma) Q_{\beta}\left(\sigma, \sigma^{\prime}\right)}{\mu_{\beta}(\mathcal{M})}=2 \sum_{\sigma \in \mathcal{M}, \sigma^{\prime} \notin \mathcal{M}} \mu_{\beta}(\sigma) Q_{\beta}\left(\sigma, \sigma^{\prime}\right) . \tag{5.3.13}
\end{equation*}
$$

Because the Glauber dynamics changes a single spin at a time, in order for $\sigma \in \mathcal{M}$ to be adjacent to a configuration $\sigma^{\prime} \notin \mathcal{M}$, it must be that

$$
\sigma \in \mathcal{M}_{-1}:=\{\sigma \in \mathcal{X}: m(\sigma)=-1\}
$$

and that $\sigma^{\prime}=\sigma^{j,+}$ for some site $j$ such that

$$
j \in \mathcal{J}_{\sigma}:=\left\{j \in V: \sigma_{j}=-1\right\} .
$$

Because the number of such sites is $(n+1) / 2$ on $\mathcal{M}_{-1}$, that is, $\left|\mathcal{J}_{\sigma}\right|=(n+1) / 2$ for all $\sigma \in \mathcal{M}_{-1}$, and the Glauber dynamics picks a site uniformly at random, it follows that for $\sigma \in \mathcal{M}_{-1}$

$$
\begin{align*}
\sum_{\sigma \in \mathcal{M}, \sigma^{\prime} \notin \mathcal{M}} \mu_{\beta}(\sigma) Q_{\beta}\left(\sigma, \sigma^{\prime}\right) & =\sum_{\sigma \in \mathcal{M}_{-1}} \mu_{\beta}(\sigma) \sum_{j \in \mathcal{J}_{\sigma}} Q_{\beta}\left(\sigma, \sigma^{j,+}\right) \\
& \leq \sum_{\sigma \in \mathcal{M}_{-1}} \mu_{\beta}(\sigma) \frac{(n+1) / 2}{n}  \tag{5.3.14}\\
& =\frac{1}{2}\left(1+\frac{1}{n}\right) \mu_{\beta}\left(\mathcal{M}_{-1}\right) \tag{5.3.15}
\end{align*}
$$

Thus plugging this back in (5.3.13) gives

$$
\begin{align*}
\Phi_{*}^{\mathcal{X}} & \leq\left(1+\frac{1}{n}\right) \mu_{\beta}\left(\mathcal{M}_{-1}\right) \\
& =(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \frac{e^{-\beta \mathcal{H}(\sigma)}}{\mathcal{Z}(\beta)}  \tag{5.3.16}\\
& =(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \frac{\exp \left(\frac{\alpha}{n-1}\left[\binom{\left|\mathcal{J}_{\sigma}\right|}{2}+\binom{\left|\mathcal{J}_{\sigma}^{c}\right|}{2}-\left|\mathcal{J}_{\sigma}\right|\left|\mathcal{J}_{\sigma}^{c}\right|\right]\right.}{\mathcal{Z}(\beta)} .
\end{align*}
$$

We bound the partition function $\mathcal{Z}(\beta)=\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}(\sigma)}$ with the term for the all-( -1 ) configuration, leading to

$$
\begin{align*}
& =(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \exp \left(-\frac{2 \alpha}{n-1}\left|\mathcal{J}_{\sigma}\right|\left|\mathcal{J}_{\sigma}^{c}\right|\right)  \tag{5.3.17}\\
& =(1+o(1))\binom{n}{(n+1) / 2} \exp \left(-\frac{2 \alpha}{n-1}\left[\frac{n+1}{2}\right]\left[\frac{n-1}{2}\right]\right) \\
& =(1+o(1)) \sqrt{\frac{2}{\pi n}} 2^{n}(1+o(1)) \exp \left(-\frac{\alpha(n+1)}{2}\right) \\
& \leq C_{\alpha} \sqrt{\frac{2}{\pi n}} \exp \left(-n\left[\frac{\alpha}{2}-\log 2\right]\right),
\end{align*}
$$

for some constant $C_{\alpha}>0$ depending on $\alpha$, where we used Stirling's formula (see Appendix A). Hence, by Theorems 5.2.14 and 5.3.5, for $\alpha>2 \log 2$ there is $r(\alpha)>0$

$$
\mathrm{t}_{\text {mix }}(\varepsilon) \geq\left(\mathrm{t}_{\text {rel }}-1\right) \log \left(\frac{1}{2 \varepsilon}\right) \geq \exp (r(\alpha) n) \log \left(\frac{1}{2 \varepsilon}\right)
$$

That proves the weaker result.
We now show that $\alpha>1$ in fact suffices. For this, we need to improve our bound on the partition function in (5.3.17). Writing

$$
\begin{aligned}
\mathcal{Z}(\beta) & =\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}(\sigma)} \\
& =\sum_{k=0}^{n}\binom{n}{k} \exp \left(\frac{\alpha}{n-1}\left[\binom{k}{2}+\binom{n-k}{2}-k(n-k)\right]\right) \\
& =2 \sum_{k=0}^{(n-1) / 2}\binom{n}{k} \exp \left(\frac{\alpha}{n-1}\left[\binom{k}{2}+\binom{n-k}{2}-k(n-k)\right]\right) \\
& =2 \sum_{k=0}^{(n-1) / 2} \mathcal{Y}_{\alpha, k},
\end{aligned}
$$

we see that the partition function is a sum of $O(n)$ exponentially large terms and is therefore dominated by the term corresponding to the largest exponent. Using Stirling's formula,

$$
\log \binom{n}{k}=(1+o(1)) n H(k / n)
$$

where $H(p)=-p \log p-(1-p) \log (1-p)$ is the entropy, and therefore

$$
\log \mathcal{Y}_{\alpha, k}=(1+o(1)) n \underbrace{\left[H(k / n)+\alpha \frac{(k / n)^{2}+(1-k / n)^{2}-2(k / n)(1-k / n)}{2}\right]}_{\mathcal{K}_{\alpha}(k / n)} .
$$

where, for $p \in[0,1]$, we let

$$
\mathcal{K}_{\alpha}(p):=H(p)+\alpha \frac{(1-2 p)^{2}}{2}
$$

Note that the first term in $\mathcal{K}_{\alpha}(p)$ is increasing on $[0,1 / 2]$ while the second term is decreasing on $[0,1 / 2]$. In a sense, we are looking at the tradeoff between the contribution from the entropy (i.e., how many ways are there to have $k$ spins with
value -1 ) and that from the Hamiltonian (i.e., how much such a configuration is favored). We seek to maximize $\mathcal{K}_{\alpha}(p)$ to determine the leading term in the partition function.

By a straightforward computation,

$$
\mathcal{K}_{\alpha}^{\prime}(p)=\log \left(\frac{1-p}{p}\right)-2 \alpha(1-2 p),
$$

and

$$
\mathcal{K}_{\alpha}^{\prime \prime}(p)=-\frac{1}{p(1-p)}+4 \alpha
$$

Observe first that, when $\alpha<1$ (i.e., at high temperature), $\mathcal{K}_{\alpha}^{\prime}(1 / 2)=0$ and $\mathcal{K}_{\alpha}^{\prime \prime}(p)<0$ for all $p \in[0,1]$ since $p(1-p) \leq 1 / 4$. Hence, in that case, $\mathcal{K}_{\alpha}$ is maximized at $p=1 / 2$.

In our case of interest, on the other hand, that is, when $\alpha>1, \mathcal{K}_{\alpha}^{\prime \prime}(p)>0$ in an interval around $1 / 2$ so there is $p_{*}<1 / 2$ with $\mathcal{K}_{\alpha}\left(p_{*}\right)>\mathcal{K}_{\alpha}(1 / 2)=1$. So the distribution significantly favors "unbalanced" configurations and crossing $\mathcal{M}_{-1}$ becomes a bottleneck for the Glauber dynamics. Going back to (5.3.17) and bounding $\mathcal{Z}(\beta) \geq 2 \mathcal{Y}_{\alpha,\left\lfloor p_{*} n\right\rfloor}$, we get

$$
\Phi_{*}^{\mathcal{X}}=O\left(\exp \left(-n\left[\mathcal{K}_{\alpha}\left(p_{*}\right)-\mathcal{K}_{\alpha}(1 / 2)\right]\right)\right) .
$$

Applying Theorems 5.2.14 and 5.3.5 concludes the proof.

## Expander graphs

In the proof of Claim 5.3.14, the bottleneck slowing down the chain arises as a result of the fact that, when $m(\sigma)=-1$, there is a large number of edges in the base graph $K_{n}$ connecting $\mathcal{J}_{\sigma}$ and $\mathcal{J}_{\sigma}^{c}$. That produces a low probability for such configurations under the ferromagnetic Ising model, where agreeing spins are favored. The same argument easily extends to expander graphs. In words, we prove something that-at first-may seem a bit counter-intuitive: good expansion properties in the base graph produces a bottleneck in the Glauber dynamics at low temperature.

Claim 5.3.15 (Ising model on expander graphs: slow mixing of the Glauber dynamics). Let $\left\{G_{n}\right\}_{n}$ be a $(d, \gamma)$-expander family. For large enough inverse temperature $\beta>0$, the Glauber dynamics of the Ising model on $G_{n}$ satisfies $\mathrm{t}_{\text {mix }}(\varepsilon)=$ $\Omega\left(\exp \left(r(\beta)\left|V\left(G_{n}\right)\right|\right)\right)$ for some function $r(\beta)>0$ not depending on $n$.

Proof. Let $\mu_{\beta}$ be the probability distribution over spin configurations under the Ising model over $G_{n}=(V, E)$ with inverse temperature $\beta$. Let $Q_{\beta}$ be the transition
matrix of the Glauber dynamics. For not necessarily disjoint subsets of vertices $W_{0}, W_{1} \subseteq V$ in the base graph $G_{n}$, let

$$
E\left(W_{0}, W_{1}\right):=\left\{\{u, v\}: u \in W_{0}, v \in W_{1},\{u, v\} \in E\right\}
$$

be the set of edges with one endpoint in $W_{0}$ and one endpoint in $W_{1}$. Let $N=$ $\left|V\left(G_{n}\right)\right|$ and assume it is odd for simplicity. We use the notation in the proof of Claim 5.3.14. Following the argument in that proof, we observe that (5.3.15) and (5.3.16) still hold. Thus

$$
\Phi_{*}^{\mathcal{X}} \leq(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \frac{\exp \left(\beta\left[\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}\right)\right|+\left|E\left(\mathcal{J}_{\sigma}^{c}, \mathcal{J}_{\sigma}^{c}\right)\right|-\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}^{c}\right)\right|\right]\right)}{\mathcal{Z}(\beta)} .
$$

As we did in (5.3.17), we bound the partition function $\mathcal{Z}(\beta)=\sum_{\sigma \in \mathcal{X}} e^{-\beta \mathcal{H}(\sigma)}$ with the term for the all- $(-1)$ configuration, leading to

$$
\begin{aligned}
\Phi_{*}^{\mathcal{X}} & \leq(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \frac{\exp \left(\beta\left[\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}\right)\right|+\left|E\left(\mathcal{J}_{\sigma}^{c}, \mathcal{J}_{\sigma}^{c}\right)\right|-\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}^{c}\right)\right|\right]\right)}{\exp \left(\beta\left[\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}\right)\right|+\left|E\left(\mathcal{J}_{\sigma}^{c}, \mathcal{J}_{\sigma}^{c}\right)\right|+\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}^{c}\right)\right|\right]\right)} \\
& =(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \exp \left(-2 \beta\left|E\left(\mathcal{J}_{\sigma}, \mathcal{J}_{\sigma}^{c}\right)\right|\right) \\
& =(1+o(1)) \sum_{\sigma \in \mathcal{M}_{-1}} \exp \left(-2 \beta\left|\partial_{\mathrm{E}} \mathcal{J}_{\sigma}^{c}\right|\right) \\
& \leq(1+o(1))\binom{N}{(N+1) / 2} \exp \left(-2 \beta \gamma d\left|\mathcal{J}_{\sigma}^{c}\right|\right) \\
& =(1+o(1)) \sqrt{\frac{2}{\pi N}} 2^{N}(1+o(1)) \exp (-\beta \gamma d(N-1)) \\
& \leq C_{\beta, \gamma, d} \sqrt{\frac{2}{\pi N}} \exp (-N[\beta \gamma d-\log 2])
\end{aligned}
$$

for some constant $C_{\beta, \gamma, d}>0$. We used the definition of an expander family (Definition 5.3.9) on the fourth line above. Taking $\beta>0$ large enough gives the result.

### 5.3.5 Congestion ratio

Recall from (5.3.2) that an inequality of the type

$$
\begin{equation*}
\operatorname{Var}_{\pi}[f] \leq C \mathscr{D}(f), \tag{5.3.18}
\end{equation*}
$$

holding for all $f$ is known as a Poincaré inequality. By Theorem 5.3.4, it implies the lower bound $\gamma \geq C^{-1}$ on the spectral gap $\gamma=1-\lambda_{2}$. In this section, we derive such an inequality using a formal measure of "congestion" in the network.

Let $\mathcal{N}=(G, c)$ be a finite, connected network with $G=(V, E)$. We assume that $c(x, y)=\pi(x) P(x, y)$ and therefore $c(x)=\sum_{y \sim x} c(x, y)=\pi(x)$, where $\pi$ is the stationary distribution of random walk on $\mathcal{N}$. To state the bound, it will be convenient to work with directed edges-this time in both directions. Let $\widetilde{E}$ contain all edges from $E$ with both orientations, that is, for each $e \in\{x, y\}, \widetilde{E}$ includes $(x, y)$ and $(y, x)$ with associated weight $c(x, y)=c(y, x)=c(e)>0$. For a function $f \in \ell^{2}(V, \pi)$ and an edge $\vec{e}=(x, y) \in \widetilde{E}$, we define as before

$$
\nabla f(\vec{e})=f(y)-f(x)
$$

With this notation, we can rewrite the Dirichlet energy as

$$
\begin{equation*}
\mathscr{D}(f)=\frac{1}{2} \sum_{x, y} c(x, y)[f(x)-f(y)]^{2}=\frac{1}{2} \sum_{\vec{e} \in \widetilde{E}} c(\vec{e})[\nabla f(\vec{e})]^{2} . \tag{5.3.19}
\end{equation*}
$$

For each pair of vertices $x, y$, let $\nu_{x, y}$ be a directed path between $x$ and $y$ in the digraph $\widetilde{G}=(V, \widetilde{E})$, as a collection of directed edges. Let $\left|\nu_{x, y}\right|$ be the number of edges in the path. The congestion ratio associated with the paths $\boldsymbol{\nu}=\left\{\nu_{x, y}\right\}_{x, y \in V}$ is

$$
C_{\boldsymbol{\nu}}=\max _{\vec{e} \in \vec{E}} \frac{1}{c(\vec{e})} \sum_{x, y: \vec{e} \in \nu_{x, y}}\left|\nu_{x, y}\right| \pi(x) \pi(y) .
$$

Note that $C_{\nu}$ tends to be large when many selected paths, called canonical paths, go through the same "congested" edge. To get a good bound in the theorem below, one must choose canonical paths that are well "spread out."
congestion
ratio

Theorem 5.3.16 (Canonical paths method). For any choice of paths $\boldsymbol{\nu}$ as above, we have the following bound on the spectral gap

$$
\gamma \geq \frac{1}{C_{\nu}}
$$

Proof. We establish a Poincaré inequality (5.3.18) with $C:=C_{\nu}$. The proof strategy is to start with the variance and manipulate it to bring out canonical paths.

For any $f \in \ell^{2}(V, \pi)$, it can be checked by expanding that

$$
\begin{equation*}
\operatorname{Var}_{\pi}[f]=\frac{1}{2} \sum_{x, y} \pi(x) \pi(y)(f(x)-f(y))^{2} . \tag{5.3.20}
\end{equation*}
$$

To bring out terms similar to those in (5.3.19), we write $f(x)-f(y)$ as a telescoping sum over the canonical path between $x$ and $y$. That is, letting $\vec{e}_{1}, \ldots, \vec{e}_{\left|\nu_{x, y}\right|}$ be the edges in $\nu_{x, y}$, observe that

$$
f(y)-f(x)=\sum_{i=1}^{\left|\nu_{x, y}\right|} \nabla f\left(\vec{e}_{i}\right) .
$$

By Cauchy-Schwarz (Theorem B.4.8),

$$
\begin{aligned}
(f(y)-f(x))^{2} & =\left(\sum_{i=1}^{\left|\nu_{x, y}\right|} \nabla f\left(\vec{e}_{i}\right)\right)^{2} \\
& \leq\left(\sum_{i=1}^{\left|\nu_{x, y}\right|} 1^{2}\right)\left(\sum_{i=1}^{\left|\nu_{x, y}\right|} \nabla f\left(\vec{e}_{i}\right)^{2}\right) \\
& =\left|\nu_{x, y}\right| \sum_{\vec{e} \in \nu_{x, y}} \nabla f(\vec{e})^{2} .
\end{aligned}
$$

Combining the last display with (5.3.20) and rearranging, we arrive at

$$
\begin{aligned}
\operatorname{Var}_{\pi}[f] & \leq \frac{1}{2} \sum_{x, y} \pi(x) \pi(y)\left|\nu_{x, y}\right| \sum_{\vec{e} \in \nu_{x, y}} \nabla f(\vec{e})^{2} \\
& =\frac{1}{2} \sum_{\vec{e} \in \tilde{E}} \nabla f(\vec{e})^{2} \sum_{x, y: \vec{e} \in \nu_{x, y}}\left|\nu_{x, y}\right| \pi(x) \pi(y) \\
& =\frac{1}{2} \sum_{\vec{e} \in \widetilde{E}} c(\vec{e}) \nabla f(\vec{e})^{2}\left(\frac{1}{c(\vec{e})} \sum_{x, y: \vec{e} \in \nu_{x, y}}\left|\nu_{x, y}\right| \pi(x) \pi(y)\right) \\
& \leq C_{\nu} \mathscr{D}(f) .
\end{aligned}
$$

That concludes the proof.
We give an example next.
Example 5.3.17 (Random walk inside a box). Consider random walk on the following $d$-dimensional box with sides of length $n$ :

$$
\begin{gathered}
V=[n]^{d}=\{1, \ldots, n\}^{d}, \\
E=\left\{x, y \in[n]^{d}:\|x-y\|_{1}=1\right\},
\end{gathered}
$$

$$
\begin{gathered}
P(x, y)=\frac{1}{|\{z: z \sim x\}|}, \quad \forall x, y \in[n]^{d}, x \sim y, \\
\pi(x)=\frac{|\{z: z \sim x\}|}{2|E|},
\end{gathered}
$$

and

$$
c(e)=\frac{1}{2|E|}, \forall e \in E .
$$

We define $\widetilde{E}$ as before.
We use Theorem 5.3.16 to bound the spectral gap. For $x=\left(x_{1}, \ldots, x_{d}\right), y=$ $\left(y_{1}, \ldots, y_{d}\right) \in[n]^{d}$, we construct $\nu_{x, y}$ by matching each coordinate in turn. That is, for two vertices $w, z \in[n]_{\widetilde{d}}^{d}$ with a single distinct coordinate, let $[w, z]$ be the directed path from $w$ to $z$ in $\widetilde{G}=(V, \widetilde{E})$ corresponding to a straight line (or the empty path if $w=z$ ). Then

$$
\begin{equation*}
\nu_{x, y}=\bigcup_{i=1}^{d}\left[\left(y_{1}, \ldots, y_{i-1}, x_{i}, x_{i+1}, \ldots, x_{d}\right),\left(y_{1}, \ldots, y_{i-1}, y_{i}, x_{i+1}, \ldots, x_{d}\right)\right] \tag{5.3.21}
\end{equation*}
$$

It remains to bound

$$
C_{\nu}=\max _{\vec{e} \in \widetilde{E}} \frac{1}{c(\vec{e})} \sum_{x, y: \vec{e} \in \nu_{x, y}}\left|\nu_{x, y}\right| \pi(x) \pi(y),
$$

from above.
Each term in the union defining $\nu_{x, y}$ contains at most $n$ edges, and therefore

$$
\left|\nu_{x, y}\right| \leq d n, \forall x, y .
$$

Not attempting to get the best constant factors, the edge weights (i.e., conductances) satisfy

$$
c(\vec{e})=\frac{1}{2|E|} \geq \frac{1}{2 \cdot 2 d n^{d}}=\frac{1}{4 d n^{d}},
$$

for all $\vec{e}$, since there are $n^{d}$ vertices and each has at most $2 d$ incident edges. Likewise, for any $x$,

$$
\pi(x)=\frac{|\{z: z \sim x\}|}{2|E|} \leq \frac{2 d}{2 \cdot\left(d n^{d}\right) / 2}=\frac{2}{n^{d}},
$$

where we divided by two in the denominator to account for the double-counting of edges. Hence we get

$$
\begin{aligned}
C_{\nu} & \leq \max _{\vec{e} \in \widetilde{E}} \frac{1}{1 /\left(4 d n^{d}\right)} \sum_{x, y: \vec{e} \in \nu_{x, y}}(d n)\left(2 / n^{d}\right)\left(2 / n^{d}\right) \\
& =\frac{16 d^{2}}{n^{d-1}} \max _{\vec{e} \in \vec{E}}\left|\left\{x, y: \vec{e} \in \nu_{x, y}\right\}\right| .
\end{aligned}
$$

To bound the cardinality of the set on the last line, we note that any edge $\vec{e} \in \widetilde{E}$ is of the form

$$
\vec{e}=\left(\left(z_{1}, \ldots, z_{i-1}, z_{i}, z_{i+1}, \ldots, z_{d}\right),\left(z_{1}, \ldots, z_{i-1}, z_{i} \pm 1, z_{i+1}, \ldots, z_{d}\right)\right)
$$

that is, the endvertices differ by exactly one unit along a single coordinate. By the construction of the path $\nu_{x, y}$ in (5.3.21), if $\vec{e} \in \nu_{x, y}$ then it must lie in the subpath

$$
\begin{aligned}
& \left(\left(z_{1}, \ldots, z_{i-1}, z_{i}, z_{i+1}, \ldots, z_{d}\right),\left(z_{1}, \ldots, z_{i-1}, z_{i} \pm 1, z_{i+1}, \ldots, z_{d}\right)\right) \\
& \quad \in\left[\left(y_{1}, \ldots, y_{i-1}, x_{i}, x_{i+1}, \ldots, x_{d}\right),\left(y_{1}, \ldots, y_{i-1}, y_{i}, x_{i+1}, \ldots, x_{d}\right)\right] .
\end{aligned}
$$

But that imposes constraints on $x$ and $y$. Namely, we must have

$$
y_{1}=z_{1}, \ldots, y_{i-1}=z_{i-1}, x_{i+1}=z_{i+1}, \ldots, x_{d}=z_{d}
$$

The remaining components of $x$ and $y$ (of which there are $i$ of the former and $d-(i-1)$ of the latter) each has at most $n$ possible values (although not all of them are allowed), so that

$$
\left|\left\{x, y: \vec{e} \in \nu_{x, y}\right\}\right| \leq n^{i} n^{d-(i-1)}=n^{d+1} .
$$

This upper bound is valid for any $\vec{e}$.
Putting everything together, we get the bound

$$
C_{\nu} \leq \frac{16 d^{2}}{n^{d-1}} n^{d+1}=16 d^{2} n^{2}
$$

so that

$$
\gamma \geq \frac{1}{16 d^{2} n^{2}}
$$

Observe that this lower bound on the spectral gap depends only mildly (i.e., polynomially) in the dimension.

One advantage of the canonical paths method is that it is somewhat robust to modifying the underlying network through comparison arguments. See Exercise 5.17 for a simple illustration.

## Exercises

Exercise 5.1. Let $A$ be an $n \times n$ symmetric random matrix. We assume that the entries on and above the diagonal, $A_{i, j}, i \leq j$, are independent uniform in $\{+1,-1\}$ (and each entry below the diagonal is equal to the corresponding entry above). Use Talagrand's inequality (Theorem 3.2.32) to prove concentration of the largest eigenvalue of $A$ around its mean (which you do not need to compute).

Exercise 5.2. Let $G=(V, E, w)$ be a network.
(i) Prove formula (5.1.3) for the Laplacian quadratic form. (Hint: For an orientation $G^{\sigma}=\left(V, E^{\sigma}\right)$ of $G$ (that is, give an arbitrary direction to each edge to turn it into a digraph), consider the matrix $B^{\sigma} \in \mathbb{R}^{n \times m}$ where the column corresponding to arc $(i, j)$ has $-\sqrt{w_{i j}}$ in row $i$ and $\sqrt{w_{i j}}$ in row $j$, and every other entry is 0 .)
(i) Show that the network Laplacian is positive semidefinite.

Exercise 5.3. Let $G=(V, E, w)$ be a weighted graph with normalized Laplacian $\mathcal{L}$. Show that

$$
\mathbf{x}^{T} \mathcal{L} \mathbf{x}=\sum_{\{i, j\} \in E} w_{i j}\left(\frac{x_{i}}{\sqrt{\delta(i)}}-\frac{x_{j}}{\sqrt{\delta(j)}}\right)^{2}
$$

for $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}$.
Exercise 5.4 (2-norm). Prove that

$$
\sup _{\mathbf{x} \in \mathbb{S}^{n-1}}\|A \mathbf{x}\|_{2}=\sup _{\substack{\mathbf{x} \in \mathbb{S}^{n-1} \\ \mathbf{y} \in \mathbb{S}^{m-1}}}\langle A \mathbf{x}, \mathbf{y}\rangle
$$

[Hint: Use Cauchy-Schwarz (Theorem B.4.8) for one direction, and set $\mathbf{y}=$ $A \mathbf{x} /\|A \mathbf{x}\|_{2}$ for the other one.]

Exercise 5.5 (Spectral radius of a symmetric matrix). Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix. The set $\sigma(A)$ of eigenvalues of $A$ is called the spectrum of $A$ and

$$
\rho(A)=\max \{|\lambda|: \lambda \in \sigma(A)\},
$$

is its spectral radius. Prove that
spectral
radius
where recall that

$$
\|A\|_{2}=\max _{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^{m}} \frac{\|A \mathbf{x}\|}{\|\mathbf{x}\|}
$$

Exercise 5.6 (Community recovery in sparse networks). Assume without proof the following theorem.

Theorem 5.3.18 (Remark 3.13 of [BH16]). Consider a symmetric matrix $\mathbf{Z}=$ $\left[Z_{i, j}\right] \in \mathbb{R}^{n \times n}$ whose entries are independent and obey, $\mathbb{E} Z_{i, j}=0$ and $Z_{i, j} \leq B$, $\forall 1 \leq i, j \leq n, \mathbb{E} Z_{i, j}^{2} \leq \sigma^{2}$ then with high probability we have $\|\mathbf{Z}\| \lesssim \sigma \sqrt{n}+$ $B \sqrt{\log n}$.

Let $(X, G) \sim \mathrm{SBM}_{n, p_{n}, q_{n}}$. Show that, under the conditions $p_{n} \gtrsim \frac{\operatorname{logn}}{n}$ and $\sqrt{\frac{p_{n}}{n}}=$ $o\left(p_{n}-q_{n}\right)$, spectral clustering achieves almost exact recovery.

Exercise 5.7 (Parseval's identity). Prove Parseval's identity (i.e., (5.2.1)) in the finite-dimensional case.

Exercise 5.8 (Dirichlet kernel). Prove that for $\theta \neq 0$

$$
1+2 \sum_{k=1}^{n} \cos k \theta=\frac{\sin ((n+1 / 2) \theta)}{\sin (\theta / 2)}
$$

[Hint: Switch to the complex representation and use the formula for a geometric series.]

Exercise 5.9 (Eigenvalues and periodicity). Let $P$ be a finite irreducible transition matrix reversible with respect to $\pi$ over $V$. Show that if $P$ has a nonzero eigenfunction $f$ with eigenvalue -1 , then $P$ is not aperiodic. [Hint: Look at $x$ achieving $\|f\|_{\infty}$.]

Exercise 5.10 (Mixing time: necessary condition for cutoff). Consider a sequence of Markov chains indexed by $n=1,2, \ldots$. Assume that each chain has a finite state space and is irreducible, aperiodic, and reversible. Let $\mathrm{t}_{\text {mix }}^{(n)}(\varepsilon)$ and $\mathrm{t}_{\text {rel }}^{(n)}$ be respectively the mixing time and relaxation time of the $n$-th chain. The sequence is said to have pre-cutoff if

$$
\sup _{0<\varepsilon<1 / 2} \limsup _{n \rightarrow+\infty} \frac{\mathrm{t}_{\text {mix }}^{(n)}(\varepsilon)}{\mathrm{t}_{\text {mix }}^{(n)}(1-\varepsilon)}<+\infty .
$$

Show that if for some $\varepsilon>0$

$$
\sup _{n \geq 1} \frac{\mathrm{t}_{\text {mix }}^{(n)}(\varepsilon)}{\mathrm{t}_{\mathrm{rel}}^{(n)}}<+\infty,
$$

then there is no pre-cutoff. In particular, there is no cutoff, as defined in Remark 4.3.8.

Exercise 5.11 (Relaxation time and variance). Let $P$ be a finite irreducible transition matrix reversible with respect to $\pi$ over $V$. Define

$$
\operatorname{Var}_{\pi}[g]=\sum_{x \in V} \pi(x)[g(x)-\pi g]^{2}
$$

Let $\gamma_{*}$ be the absolute spectral gap of $P$. Show that

$$
\operatorname{Var}_{\pi}\left[P^{t} f\right] \leq\left(1-\gamma_{*}\right)^{2 t} \operatorname{Var}_{\pi}[f] .
$$

Exercise 5.12 (Lumping). Let $\left(X_{t}\right)$ be a Markov chain on a finite state space $V$ with transition $P$. Suppose there is an equivalence relation $\sim$ on $V$ with equivalence classes $V^{\sharp}$, denoting by $[x]$ the equivalence class of $x$, such that $\left[X_{t}\right]$ is a Markov chain with transition matrix $P^{\sharp}([x],[y])=P(x,[y])$.
(i) Let $f: V \rightarrow \mathbb{R}$ be an eigenfunction of $P$ with eigenvalue $\lambda$ and assume that $f$ is constant on each equivalence class. Prove that $f^{\sharp}([x]):=f(x)$ defines an eigenfunction of $P^{\sharp}$. What is its eigenvalue?
(ii) Suppose $g: V^{\sharp} \rightarrow \mathbb{R}$ is eigenfunction of $P^{\sharp}$ with eigenvalue $\lambda$. Prove that $g^{b}: V \rightarrow \mathbb{R}$ defined by $g^{b}(x):=g([x])$ is eigenfunction of $P$. What is its eigenvalue?

Exercise 5.13 (Random walk on path with reflecting boundaries). Let $n$ be an even positive integer. Let $\left(X_{t}\right)$ be simple random walk on the path $\{1, \ldots, n\}$ with reflecting boundaries, that is, the transition matrix $P$ is defined by $P(x, x-1)=$ $P(x, x+1)=1 / 2$ for $x \in\{2, \ldots, n-1\}$, and $P(1,2)=P(n, n-1)=1$. Use Exercise 5.12 to compute the eigenfunctions of $P$. [Hint: Use the results of Section 5.2.2.]

Exercise 5.14 (Product chain). For $j=1, \ldots, d$, let $P_{j}$ be a transition matrix on the finite state space $V_{j}$ reversible with respect to the stationary distribution $\pi_{j}$. Let $w=\left(w_{j}\right)_{j \in[d]}$ be a probability distribution over [d]. Consider the following Markov chain $\left(X_{t}\right)$ on $V:=V_{1} \times \cdots \times V_{d}$ : at each step, pick $j$ according to $w$, then take one step along the $j$-th coordinate according to $P_{j}$.
(i) Compute the transition matrix $P$ and stationary distribution $\pi$ of the chain $\left(X_{t}\right)$. Show that $P$ is reversible with respect to $\pi$.
(ii) Construct an orthonormal basis of $\ell^{2}(V, \pi)$ made of eigenfunctions of $P$ in terms of eigenfunctions of the $P_{j} \mathrm{~s}$. What are the corresponding eigenvalues?
(iii) Compute the spectral gap $\gamma$ of $P$ in terms of the spectral gaps $\gamma_{j}$ of the $P_{j}$ s.

Exercise 5.15 (Hypercube revisited). Use Exercise 5.14 to recover Lemma 5.2.17.

Exercise 5.16 (Norm and Rayleigh quotient). Let $P$ be irreducible and reversible with respect to $\pi>0$.
(i) Prove the polarization identity

$$
\langle P f, g\rangle_{\pi}=\frac{1}{4}\left[\langle P(f+g), f+g\rangle_{\pi}-\langle P(f-g), f-g\rangle_{\pi}\right] .
$$

(ii) Show that

$$
\|P\|_{\pi}=\sup \left\{\frac{\langle f, P f\rangle_{\pi}}{\langle f, f\rangle_{\pi}}: f \in \ell_{0}(V), f \neq 0\right\} .
$$

Exercise 5.17 (Random walk on a box with holes). Consider the random walk in Example 5.3.17 with $d=2$. Suppose we remove from the network an arbitrary collection of horizontal edges at even heights. Use the canonical paths method to derive a lower bound on the spectral gap of the form $\gamma \geq 1 /\left(C n^{2}\right)$. [Hint: Modify the argument in Example 5.3.17 and relate the congestion ratio before and after the removal.]

## Bibliographic Remarks

Section 5.1 General references on the spectral theorem, the Courant-Fischer and perturbation results include the classics [HJ13, Ste98]. Much more on spectral graph theory can be gleaned from [Chu97, Nic18]. Section 5.1.4 is based largely on [Abb18], which gives a broad survey of theoretical results for community recovery, and [Ver18, Section 4.5] as well as on scribe notes by Joowon Lee, Aidan Howells, Govind Gopakumar, and Shuyao Li for "MATH 888: Topics in Mathematical Data Science" taught at the University of Wisconsin-Madison in Fall 2021.

Section 5.2 For a great introduction to Hilbert space theory and its applications (including to the Dirichlet problem), consult [SS05, Chapters 4,5]. Section 5.2.1 borrows from [LP17, Chapter 12]. A representation-theoretic approach to computing eigenvalues and eigenfunctions, greatly generalizing the calculations in 5.2.2, is presented in [Dia88]. The presentation in Section 5.2.3 follows [KP, Section 3] and [LP16, Section 13.3]. The Varopoulos-Carne bound is due to Carne [Car85] and Varopoulos [Var85]. For a probabilistic approach to the Varopoulos-Carne bound see Peyre's proof [Pey08]. The application to mixing times is from [LP16]. There are many textbooks dedicated to Markov chain Monte Carlo (MCMC) and its uses in data analysis, for example, [RC04, GL06, GCS ${ }^{+}$14]. See also [Dia09]. A good overview of the techniques developed in the statistics literature to bound the rate of convergence of MCMC methods (a combination of coupling and Lyapounov arguments) is [JH01]. A deeper treatment of these ideas is developed in [MT09]. A formal definition of the spectral radius and its relationship to the operator norm can be found, for instance, in [Rud73, Part III].

Section 5.3 This section follows partly the presentation in [LP16, Section 6.4], [LPW06, Section 13.6], and [Spi12]. Various proofs of the isoperimetric inequality can be found in [SS03, SS05]. Theorem 5.3.5 is due to [SJ89, LS88]. The approach to its proof used here is due to Luca Trevisan The original Cheeger inequality was proved, in the context of manifolds, in [Che70]. For a fascinating introduction to expander graphs and their applications, see [HLW06]. A detailed account of the Curie-Weiss model can be found in [FV18]. Section 5.3.5 is based partly on [Ber14, Sections 3 and 4]. The method of canonical paths, and some related comparison techniques, were developed in [JS89, DS91, DSC93b, DSC93a]. For more advanced functional techniques for bounding the mixing time, see e.g. [MT06].

## Chapter 6

## Branching processes

Branching processes, which are the focus of this chapter, arise naturally in the study of stochastic processes on trees and locally tree-like graphs. Similarly to martingales, finding a hidden (or not-so-hidden) branching process within a probabilistic model can lead to useful bounds and insights into asymptotic behavior. After a review of the basic extinction theory of branching processes in Section 6.1 and of a fruitful random-walk perspective in Section 6.2, we give a couple examples of applications in discrete probability in Section 6.3. In particular we analyze the height of a binary search tree, a standard data structure in computer science. We also give an introduction to phylogenetics, where a "multitype" variant of the Galton-Watson branching process plays an important role; we use the techniques derived in this chapter to establish a phase transition in the reconstruction of ancestral molecular sequences. We end this chapter in Section 6.4 with a detailed look into the phase transition of the Erdös-Rényi graph model. The random-walk perspective mentioned above allows one to analyze the "exploration" of a largest connected component, leading to information about the "evolution" of its size as edge density increases. Tools from all chapters come to bear on this final, marquee application.

### 6.1 Background

We begin with a review of the theory of Galton-Watson branching processes, a standard stochastic model for population growth. In particular we discuss extinction theory. We also briefly introduce a multitype variant, where branching process

[^10]and Markov chain aspects interact to produce interesting new behavior.

### 6.1.1 Basic definitions

Recall the definition of a Galton-Watson process.
Definition 6.1.1. A Galton-Watson branching process is a Markov chain of the following form:

- Let $Z_{0}:=1$.
- Let $X(i, t), i \geq 1, t \geq 1$, be an array of i.i.d. $\mathbb{Z}_{+}$-valued random variables with finite mean $m=\mathbb{E}[X(1,1)]<+\infty$, and define inductively,

$$
Z_{t}:=\sum_{1 \leq i \leq Z_{t-1}} X(i, t) .
$$

We denote by $\left\{p_{k}\right\}_{k \geq 0}$ the law of $X(1,1)$. We also let $f(s):=\mathbb{E}\left[s^{X(1,1)}\right]$ be the corresponding probability generating function. To avoid trivialities we assume $\mathbb{P}[X(1,1)=i]<1$ for all $i \geq 0$. We further assume that $p_{0}>0$.

In words, $Z_{t}$ models the size of a population at time (or generation) $t$. The random variable $X(i, t)$ corresponds to the number of offspring of the $i$-th individual (if there is one) in generation $t-1$. Generation $t$ is formed of all offspring of the individuals in generation $t-1$.

By tracking genealogical relationships, that is, who is whose child, we obtain a tree $T$ rooted at the single individual in generation 0 with a vertex for each individual in the progeny and an edge for each parent-child relationship. We refer to $T$ as a Galton-Watson tree.

A basic observation about Galton-Watson processes is that their growth (or decay) is exponential in $t$.

Lemma 6.1.2 (Exponential growth I). Let

$$
\begin{equation*}
W_{t}:=m^{-t} Z_{t} . \tag{6.1.1}
\end{equation*}
$$

Then $\left(W_{t}\right)$ is a nonnegative martingale with respect to the filtration

$$
\mathcal{F}_{t}=\sigma\left(Z_{0}, \ldots, Z_{t}\right) .
$$

In particular, $\mathbb{E}\left[Z_{t}\right]=m^{t}$.

Proof. We use Lemma B.6.17. Observe that on $\left\{Z_{t-1}=k\right\}$

$$
\mathbb{E}\left[Z_{t} \mid \mathcal{F}_{t-1}\right]=\mathbb{E}\left[\sum_{1 \leq j \leq k} X(j, t) \mid \mathcal{F}_{t-1}\right]=m k=m Z_{t-1}
$$

This is true for all $k$. Rearranging shows that $\left(W_{t}\right)$ is a martingale. For the second claim, note that $\mathbb{E}\left[W_{t}\right]=\mathbb{E}\left[W_{0}\right]=1$.

In fact, the martingale convergence theorem (Theorem 3.1.47) gives the following.
Lemma 6.1.3 (Exponential growth II). We have $W_{t} \rightarrow W_{\infty}<+\infty$ almost surely for some nonnegative random variable $W_{\infty} \in \sigma\left(\cup_{t} \mathcal{F}_{t}\right)$ with $\mathbb{E}\left[W_{\infty}\right] \leq 1$.

Proof. This follows immediately from the martingale convergence theorem for nonnegative martingales (Corollary 3.1.48).

### 6.1.2 Extinction

Observe that 0 is a fixed point of the process. The event

$$
\left\{Z_{t} \rightarrow 0\right\}=\left\{\exists t: Z_{t}=0\right\},
$$

is called extinction. Establishing when extinction occurs is a central question in branching process theory. We let $\eta$ be the probability of extinction. Recall that, to avoid trivialities, we assume that $p_{0}>0$ and $p_{1}<1$. Here is a first observation about extinction.

Lemma 6.1.4. Almost surely either $Z_{t} \rightarrow 0$ or $Z_{t} \rightarrow+\infty$.
Proof. The process $\left(Z_{t}\right)$ is integer-valued and 0 is the only fixed point of the process under the assumption that $p_{1}<1$. From any state $k$, the probability of never coming back to $k>0$ is at least $p_{0}^{k}>0$, so every state $k>0$ is transient. So the only possibilities left are $Z_{t} \rightarrow 0$ and $Z_{t} \rightarrow+\infty$, and the claim follows.

In the critical case, that immediately implies almost sure extinction.
Theorem 6.1.5 (Extinction: critical case). Assume $m=1$. Then $Z_{t} \rightarrow 0$ almost surely, that is, $\eta=1$.

Proof. When $m=1,\left(Z_{t}\right)$ itself is a martingale. Hence $\left(Z_{t}\right)$ must converge to 0 by Lemma 6.1.3.

We address the general case using probability generating functions. Let $f_{t}(s)=$ $\mathbb{E}\left[s^{Z_{t}}\right]$, where by convention we set $f_{t}(0):=\mathbb{P}\left[Z_{t}=0\right]$. Note that, by monotonicity,

$$
\begin{equation*}
\eta=\mathbb{P}\left[\exists t \geq 0: Z_{t}=0\right]=\lim _{t \rightarrow+\infty} \mathbb{P}\left[Z_{t}=0\right]=\lim _{t \rightarrow+\infty} f_{t}(0) \tag{6.1.2}
\end{equation*}
$$

Moreover, by the tower property (Lemma B.6.16) and the Markov property (Theorem 1.1.18), $f_{t}$ has a natural recursive form

$$
\begin{align*}
f_{t}(s) & =\mathbb{E}\left[s^{Z_{t}}\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[s^{Z_{t}} \mid \mathcal{F}_{t-1}\right]\right] \\
& =\mathbb{E}\left[f(s)^{Z_{t-1}}\right] \\
& =f_{t-1}(f(s))=\cdots=f^{(t)}(s) \tag{6.1.3}
\end{align*}
$$

where $f^{(t)}$ is the $t$-th iterate of $f$. The subcritical case below has an easier proof (see Exercise 6.1).

Theorem 6.1.6 (Extinction: subcriticial and supercritical cases). The probability of extinction $\eta$ is given by the smallest fixed point of $f$ in $[0,1]$. Moreover:
(i) (Subcritical regime) If $m<1$ then $\eta=1$.
(ii) (Supercritical regime) If $m>1$ then $\eta<1$.

Proof. The case $p_{0}+p_{1}=1$ is straightforward: the process dies almost surely after a geometrically distributed time. So we assume $p_{0}+p_{1}<1$ for the rest of the proof.

We first summarize without proof some properties of $f$ which follow from standard power series facts.
Lemma 6.1.7. On $[0,1]$, the function $f$ satisfies:
(i) $f(0)=p_{0}, f(1)=1$;
(ii) $f$ is infinitely differentiable on $[0,1)$;
(iii) $f$ is strictly convex and increasing; and
(iv) $\lim _{s \uparrow 1} f^{\prime}(s)=m<+\infty$.

We first characterize the fixed points of $f$. See Figure 6.1 for an illustration.
Lemma 6.1.8. We have the following.
(i) If $m>1$ then $f$ has a unique fixed point $\eta_{0} \in[0,1)$.


Figure 6.1: Fixed points of $f$ in subcritical (left) and supercritical (right) cases.
(ii) If $m<1$ then $f(t)>t$ for $t \in[0,1)$. Let $\eta_{0}:=1$ in that case.

Proof. Assume $m>1$. Since $f^{\prime}(1)=m>1$, there is $\delta>0$ such that $f(1-\delta)<$ $1-\delta$. On the other hand $f(0)=p_{0}>0$ so by continuity of $f$ there must be a fixed point in $(0,1-\delta)$. Moreover, by strict convexity and the fact that $f(1)=1$, if $x \in(0,1)$ is a fixed point then $f(y)<y$ for $y \in(x, 1)$, proving uniqueness.

The second part follows by strict convexity and monotonicity.
It remains to prove convergence of the iterates to the appropriate fixed point. See Figure 6.2 for an illustration.

Lemma 6.1.9. We have the following.
(i) If $x \in\left[0, \eta_{0}\right)$, then $f^{(t)}(x) \uparrow \eta_{0}$.
(ii) If $x \in\left(\eta_{0}, 1\right)$ then $f^{(t)}(x) \downarrow \eta_{0}$.

Proof. We only prove (i). The argument for (ii) is similar. By monotonicity, for $x \in\left[0, \eta_{0}\right)$, we have $x<f(x)<f\left(\eta_{0}\right)=\eta_{0}$. Iterating

$$
x<f^{(1)}(x)<\cdots<f^{(t)}(x)<f^{(t)}\left(\eta_{0}\right)=\eta_{0} .
$$



Figure 6.2: Convergence of iterates to a fixed point.

So $f^{(t)}(x) \uparrow L \leq \eta_{0}$ as $t \rightarrow \infty$. By continuity of $f$, we can take the limit $t \rightarrow \infty$ inside of $f$ on the right-hand side of the equality

$$
f^{(t)}(x)=f\left(f^{(t-1)}(x)\right),
$$

to get $L=f(L)$. So by definition of $\eta_{0}$ we must have $L=\eta_{0}$.
The result then follows from the above lemmas together with Equations (6.1.2) and (6.1.3).

Example 6.1.10 (Poisson branching process). Consider the offspring distribution $X(1,1) \sim \operatorname{Poi}(\lambda)$ with mean $\lambda>0$. We refer to this case as the Poisson branching process. Then

$$
f(s)=\mathbb{E}\left[s^{X(1,1)}\right]=\sum_{i \geq 0} e^{-\lambda} \frac{\lambda^{i}}{i!} s^{i}=e^{\lambda(s-1)} .
$$

So the process goes extinct with probability 1 when $\lambda \leq 1$. For $\lambda>1$, the probability of extinction $\eta$ is the smallest solution in $[0,1]$ to the equation

$$
e^{-\lambda(1-x)}=x .
$$

The survival probability $\zeta_{\lambda}:=1-\eta$ satisfies $1-e^{-\lambda \zeta_{\lambda}}=\zeta_{\lambda}$.
We can use these extinction results to obtain more information on the limit in Lemma 6.1.3. Recall the definition of $\left(W_{t}\right)$ in (6.1.1). Of course, conditioned on extinction, $W_{\infty}=0$ almost surely. On the other hand:

Lemma 6.1.11 (Exponential growth III). Conditioned on nonextinction, either $W_{\infty}=0$ almost surely or $W_{\infty}>0$ almost surely.

As a result, $\mathbb{P}\left[W_{\infty}=0\right] \in\{\eta, 1\}$.
Proof of Lemma 6.1.11. A property of rooted trees is said to be inherited if all finite trees satisfy the property and whenever a tree satisfies the property then so do all subtrees rooted at the children of the root. The property $\left\{W_{\infty}=0\right\}$, as a property of the Galton-Watson tree $T$, is inherited, seeing that $Z_{t}$ is a sum over the children of the root of the number of descendants at the corresponding generation $t-1$. The result then follows from the following 0-1 law.

Lemma 6.1.12 (0-1 law for inherited properties). For a Galton-Watson tree T, an inherited property $A$ has, conditioned on nonextinction, probability 0 or 1 .

Proof. Let $T^{(1)}, \ldots, T^{\left(Z_{1}\right)}$ be the descendant subtrees of the children of the root. We use the notation $T \in A$ to mean that tree $T$ satisfies $A$. By the tower property, the definition of inherited, and conditional independence,

$$
\begin{aligned}
\mathbb{P}[A] & =\mathbb{E}\left[\mathbb{P}\left[T \in A \mid Z_{1}\right]\right] \\
& \leq \mathbb{E}\left[\mathbb{P}\left[T^{(i)} \in A, \forall i \leq Z_{1} \mid Z_{1}\right]\right] \\
& =\mathbb{E}\left[\mathbb{P}[A]^{Z_{1}}\right] \\
& =f(\mathbb{P}[A]) .
\end{aligned}
$$

So $\mathbb{P}[A] \in[0, \eta] \cup\{1\}$ by the proof of Lemma 6.1.8.
Moreover since $A$ holds for finite trees, we have $\mathbb{P}[A] \geq \eta$, where recall that $\eta$ is the probability of extinction. Hence, in fact, $\mathbb{P}[A] \in\{\eta, 1\}$. Conditioning on nonextinction gives the claim.

That concludes the proof.
A further moment assumption provides a more detailed picture.
Lemma 6.1.13 (Exponential growth IV). Let $\left(Z_{t}\right)$ be a Galton-Watson branching process with $m=\mathbb{E}[X(1,1)]>1$ and $\sigma^{2}=\operatorname{Var}[X(1,1)]<+\infty$. Then, $\left(W_{t}\right)$ converges in $L^{2}$ and, in particular, $\mathbb{E}\left[W_{\infty}\right]=1$. Further, $\mathbb{P}\left[W_{\infty}=0\right]=\eta$.

Proof. We bound $\mathbb{E}\left[W_{t}^{2}\right]$ by computing it explicitly by induction. From the orthogonality of increments (Lemma 3.1.50), it holds that

$$
\mathbb{E}\left[W_{t}^{2}\right]=\mathbb{E}\left[W_{t-1}^{2}\right]+\mathbb{E}\left[\left(W_{t}-W_{t-1}\right)^{2}\right] .
$$

Since $\mathbb{E}\left[W_{t} \mid \mathcal{F}_{t-1}\right]=W_{t-1}$ by the martingale property,

$$
\begin{aligned}
\mathbb{E}\left[\left(W_{t}-W_{t-1}\right)^{2} \mid \mathcal{F}_{t-1}\right] & =\operatorname{Var}\left[W_{t} \mid \mathcal{F}_{t-1}\right] \\
& =m^{-2 t} \operatorname{Var}\left[Z_{t} \mid \mathcal{F}_{t-1}\right] \\
& =m^{-2 t} \operatorname{Var}\left[\sum_{i=1}^{Z_{t-1}} X(i, t) \mid \mathcal{F}_{t-1}\right] \\
& =m^{-2 t} Z_{t-1} \sigma^{2} .
\end{aligned}
$$

Hence, taking expectations and using Lemma 6.1.2, we get

$$
\mathbb{E}\left[W_{t}^{2}\right]=\mathbb{E}\left[W_{t-1}^{2}\right]+m^{-t-1} \sigma^{2} .
$$

Since $\mathbb{E}\left[W_{0}^{2}\right]=1$, induction gives

$$
\mathbb{E}\left[W_{t}^{2}\right]=1+\sigma^{2} \sum_{i=2}^{t+1} m^{-i}
$$

which is uniformly bounded from above when $m>1$.
By the convergence theorem for martingales bounded in $L^{2}$ (Theorem 3.1.51), $\left(W_{t}\right)$ converges almost surely and in $L^{2}$ to a finite limit $W_{\infty}$ and

$$
1=\mathbb{E}\left[W_{t}\right] \rightarrow \mathbb{E}\left[W_{\infty}\right] .
$$

The last statement follows from Lemma 6.1.11.
Remark 6.1.14. A theorem of Kesten and Stigum gives a necessary and sufficient condition for $\mathbb{E}\left[W_{\infty}\right]=1$ to hold [KS66b]. See, e.g., [LP16, Chapter 12].

### 6.1.3 $\triangleright$ Percolation: Galton-Watson trees

Let $T$ be the Galton-Watson tree for an offspring distribution with mean $m>1$. Now perform bond percolation on $T$ with density $p$ (see Definition 1.2.1). Let $\mathcal{C}_{0}$ be the open cluster of the root in $T$. Recall from Section 2.3.3 that the critical value is

$$
p_{c}(T)=\sup \{p \in[0,1]: \theta(p)=0\},
$$

where the percolation function (conditioned on $T$ ) is $\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty \mid T\right]$.

Theorem 6.1.15 (Bond percolation on Galton-Watson trees). Assume $m>1$. Conditioned on nonextinction of $T$,

$$
p_{\mathrm{c}}(T)=\frac{1}{m},
$$

almost surely.
Proof. We can think of $\mathcal{C}_{0}$ (or more precisely, its size on each level) as being itself generated by a Galton-Watson branching process, where this time the offspring distribution is the law of $\sum_{i=1}^{X(1,1)} I_{i}$ where the $I_{i}$ s are i.i.d. $\operatorname{Ber}(p)$ and $X(1,1)$ is distributed according to the offspring distribution of $T$. In words, we are "thinning" $T$. By conditioning on $X(1,1)$ and then using the tower property (Lemma B.6.16), the offspring mean under the process generating $\mathcal{C}_{0}$ is $m p$.

If $m p \leq 1$ then by the extinction theory (Theorems 6.1.5 and 6.1.6)

$$
1=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty\right]=\mathbb{E}\left[\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty \mid T\right]\right],
$$

and we must have $\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty \mid T\right]=1$ almost surely. Taking $p=1 / m$, we get $p_{\mathrm{c}}(T) \geq \frac{1}{m}$ almost surely. That holds, in particular, on the nonextinction of $T$ which happens with positive probability.

For the other direction, fix $p$ such that $m p>1$. The property of trees $\left\{\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<\right.\right.$ $+\infty \mid T]=1\}$ is inherited. So by Lemma 6.1.12, conditioned on nonextinction of $T$, it has probability 0 or 1 . That probability is of course 1 on extinction. By Theorem 6.1.6,

$$
1>\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty\right]=\mathbb{E}\left[\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty \mid T\right]\right],
$$

and, conditioned on nonextinction of $T$, we must have $\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty \mid T\right]=0$ i.e., $p_{\mathrm{c}}(T)<p$-almost surely. Repeating this argument for a sequence $p_{n} \downarrow$ $1 / m$ simultaneously (i.e., on the same $T$ ) and using the monotonicity of $\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<\right.$ $+\infty \mid T]$, we get that $p_{\mathrm{c}}(T) \leq 1 / m$ almost surely conditioned on nonextinction of $T$. That proves the claim.

### 6.1.4 Multitype branching processes

Multitype branching processes are a useful generalization of Galton-Watson processes (Definition 6.1.1). Their behavior combines aspects of branching processes (exponential growth, extinction, etc.) and Markov chains (reducibility, mixing, etc.). We will not develop the full theory here. In this section, we define this class of processes and hint (largely without proofs) at their properties. In Section 6.3.2, we illustrate some of the more intricate interplay between the driving phenomena involved in a special example of practical importance.
multitype branching processes

Definition In a multitype branching process, each individual has one of $\tau$ types, which we will denote in this section by $1, \ldots, \tau$ for simplicity. Each type $\alpha \in[\tau]=$ $\{1, \ldots, \tau\}$ has its own offspring distribution $\left\{p_{\mathbf{k}}^{(\alpha)}: \mathbf{k} \in \mathbb{Z}_{+}^{\tau}\right\}$, which specifies the distribution of the number of offspring of each type it has. Just to emphasize, this is a collection of (typically distinct) multivariate distributions.

For reasons that will become clear below, it will be convenient to work with row vectors. For each $\alpha \in[\tau]$, let

$$
\mathbf{X}^{(\alpha)}(i, t)=\left(X_{1}^{(\alpha)}(i, t), \ldots, X_{\tau}^{(\alpha)}(i, t)\right), \quad \forall i, t \geq 1
$$

be an array of i.i.d. $\mathbb{Z}_{+}^{\tau}$-valued random row vectors with distribution $\left\{p_{\mathbf{k}}^{(\alpha)}\right\}$. Let

$$
\mathbf{Z}_{0}=\mathbf{k}_{0} \in \mathbb{Z}_{+}^{\tau},
$$

be the initial population at time 0 , again as a row vector. Recursively, the population vector

$$
\mathbf{Z}_{t}=\left(Z_{t, 1}, \ldots, Z_{t, \tau}\right) \in \mathbb{Z}_{+}^{\tau},
$$

at time $t \geq 1$ is set to

$$
\begin{equation*}
\mathbf{Z}_{t}:=\sum_{\alpha=1}^{\tau} \sum_{i=1}^{Z_{t-1, \alpha}} \mathbf{X}^{(\alpha)}(i, t) \tag{6.1.4}
\end{equation*}
$$

In words, the $i$-th individual of type $\alpha$ at generation $t-1$ produces $X_{\beta}^{(\alpha)}(i, t)$ individuals of type $\beta$ at generation $t$ (before itself dying). Let $\mathcal{F}_{t}=\sigma\left(\mathbf{Z}_{0}, \ldots, \mathbf{Z}_{t}\right)$ be the corresponding filtration. We assume throughout that $\mathbb{P}\left[\left\|\mathbf{X}^{(\alpha)}(1,1)\right\|_{1}=\right.$ $1]<1$ for at least one $\alpha$ (which is referred to as the nonsingular case); otherwise the process reduces to a simple finite Markov chain.

Martingales As in the single-type case, the means of the offspring distributions play a key role in the theory. This time however they form a matrix, the so-called mean matrix $M=\left(m_{\alpha, \beta}\right)$ with entries

$$
m_{\alpha, \beta}=\mathbb{E}\left[X_{\beta}^{(\alpha)}(1,1)\right], \quad \forall \alpha, \beta \in[\tau] .
$$

That is, $m_{\alpha, \beta}$ is the expected number of offspring of type $\beta$ of an individual of type $\alpha$. We assume throughout that $m_{\alpha, \beta}<+\infty$ for all $\alpha, \beta$.

To see how $M$ drives the growth of the process, we generalize the proof of Lemma 6.1.2. By the recursive formula (6.1.4),

$$
\begin{align*}
\mathbb{E}\left[\mathbf{Z}_{t} \mid \mathcal{F}_{t-1}\right] & =\mathbb{E}\left[\sum_{\alpha=1}^{\tau} \sum_{i=1}^{Z_{t-1, \alpha}} \mathbf{X}^{(\alpha)}(i, t) \mid \mathcal{F}_{t-1}\right] \\
& =\sum_{\alpha=1}^{\tau} \sum_{i=1}^{Z_{t-1, \alpha}} \mathbb{E}\left[\mathbf{X}^{(\alpha)}(i, t) \mid \mathcal{F}_{t-1}\right] \\
& =\sum_{\alpha=1}^{\tau} Z_{t-1, \alpha} \mathbb{E}\left[\mathbf{X}^{(\alpha)}(1,1)\right] \\
& =\mathbf{Z}_{t-1} M, \tag{6.1.5}
\end{align*}
$$

where recall that $\mathbf{Z}_{t-1}$ and $\mathbf{Z}_{t}$ are row vectors. Inductively,

$$
\begin{equation*}
\mathbb{E}\left[\mathbf{Z}_{t} \mid \mathbf{Z}_{0}\right]=\mathbf{Z}_{0} M^{t} \tag{6.1.6}
\end{equation*}
$$

Moreover any real right eigenvector $\mathbf{u}$ (as a column vector) of $M$ with real eigenvalue $\lambda \neq 0$ gives rise to a martingale

$$
\begin{equation*}
U_{t}:=\lambda^{-t} \mathbf{Z}_{t} \mathbf{u}, \quad t \geq 0 \tag{6.1.7}
\end{equation*}
$$

since

$$
\begin{aligned}
\mathbb{E}\left[U_{t} \mid \mathcal{F}_{t-1}\right] & =\mathbb{E}\left[\lambda^{-t} \mathbf{Z}_{t} \mathbf{u} \mid \mathcal{F}_{t-1}\right] \\
& =\lambda^{-t} \mathbb{E}\left[\mathbf{Z}_{t} \mid \mathcal{F}_{t-1}\right] \mathbf{u} \\
& =\lambda^{-t} \mathbf{Z}_{t-1} M \mathbf{u} \\
& =\lambda^{-t} \mathbf{Z}_{t-1} \lambda \mathbf{u} \\
& =U_{t-1} .
\end{aligned}
$$

Extinction The classical Perron-Frobenius Theorem characterizes the direction of largest growth of the matrix $M$. We state a version of it without proof in the case where all entries of $M$ are strictly positive, which is referred to as the positive regular case. Note that, unlike the case of simple finite Markov chains, the matrix $M$ is not in general stochastic, as it also reflects the "growth" of the population in addition to the "transitions" between types. We encountered the following concept in Section 5.2.5 and Exercise 5.5.

Definition 6.1.16. The spectral radius $\rho(A)$ of a matrix $A$ is the maximum of the eigenvalues of $A$ in absolute value.
positive regular case
spectral

Theorem 6.1.17 (Perron-Frobenius theorem: positive regular case). Let $M$ be a strictly positive, square matrix. Then $\rho:=\rho(M)$ is an eigenvalue of $M$ with algebraic and geometric multiplicities 1. It is also the only eigenvalue with absolute value $\rho$. The corresponding left and right eigenvectors, denoted by $\mathbf{v}$ (as a row vector) and w (as a column vector) respectively, are positive vectors. They are referred to as left and right Perron vector. We assume that they are normalized so that $1 \mathbf{w}=1$ and $\mathrm{vw}=1$. Here $\mathbf{1}$ is the all-one row vector.

Because $w$ is positive, the martingale

$$
W_{t}:=\rho^{-t} \mathbf{Z}_{t} \mathbf{w}, \quad t \geq 0
$$

is nonnegative. Therefore it converges almost surely to a random limit with a finite mean by Corollary 3.1.48. When $\rho<1$, an argument based on Markov's inequality (Theorem 2.1.1) implies that the process goes extinct almost surely. Formally, let $q^{(\alpha)}$ be the probability of extinction when started with a single individual of type $\alpha$, that is,

$$
q^{(\alpha)}:=\mathbb{P}\left[\mathbf{Z}_{t}=\mathbf{0} \text { for some } t \mid \mathbf{Z}_{0}=\mathbf{e}_{\alpha}\right]
$$

where $\mathbf{e}_{\alpha} \in \mathbb{Z}_{+}^{\tau}$ is the standard basis row vector with a one in the $\alpha$-th coordinate, and let $\mathbf{q}:=\left(q^{(1)}, \ldots, q^{(\tau)}\right)$. Then

$$
\begin{equation*}
\rho<1 \Longrightarrow \mathbf{q}=\mathbf{1} \tag{6.1.8}
\end{equation*}
$$

Exercise 6.1 asks for the proof. We state the following more general result without proof. We use the notation of Theorem 6.1.17. We will also refer to the generating functions

$$
f^{(\alpha)}(\mathbf{s}):=\mathbb{E}\left[\prod_{\beta=1}^{\tau} s_{\beta}^{X_{\beta}^{(\alpha)}(1,1)}\right], \quad \mathbf{s} \in[0,1]^{\tau}
$$

with $\mathbf{f}=\left(f^{(1)}, \ldots, f^{(\tau)}\right)$.
Theorem 6.1.18 (Extinction: multitype case). Let $(\mathbf{Z})_{t}$ be a positive regular, nonsingular multitype branching process with a finite mean matrix $M$.
(i) If $\rho \leq 1$ then $\mathbf{q}=\mathbf{1}$.
(ii) If $\rho>1$ then:
$a$ - It holds that $\mathbf{q}<\mathbf{1}$.
$b$ - The unique solution to $\mathbf{f}(\mathbf{s})=\mathbf{s}$ in $[0,1)^{\tau}$ is $\mathbf{q}$.
c- Almost surely

$$
\lim _{t \rightarrow+\infty} \rho^{-t} \mathbf{Z}_{\mathbf{t}}=\mathbf{v} W_{\infty},
$$

where $W_{\infty}$ is a nonnegative random variable.
$d$ - If in addition $\operatorname{Var}\left[X_{\beta}^{(\alpha)}(1,1)\right]<+\infty$ for all $\alpha, \beta$ then

$$
\mathbb{E}\left[W_{\infty} \mid \mathbf{Z}_{0}=\mathbf{e}_{\alpha}\right]=w_{\alpha},
$$

and

$$
q^{(\alpha)}=\mathbb{P}\left[W_{\infty}=0 \mid \mathbf{Z}_{0}=\mathbf{e}_{\alpha}\right],
$$

for all $\alpha \in[\tau]$.
Remark 6.1.19. As in the single-type case, a theorem of Kesten and Stigum gives a necessary and sufficient condition for the last claim of Theorem 6.1.18 (ii) to hold [KS66b].

Linear functionals Theorem 6.1.18 also characterizes the limit behavior of linear functionals of the form $\mathbf{Z}_{\mathbf{t}} \mathbf{u}$ for any vector that is not orthogonal to $\mathbf{v}$. In contrast, interesting new behavior arises when $\mathbf{u}$ is orthogonal to $\mathbf{v}$. We will not derive the general theory here. We only show through a second moment calculation that a phase transition takes place.

We restrict ourselves to the supercritical case $\rho>1$ and to $\mathbf{u}=\left(u_{1}, \ldots, u_{\tau}\right)$ being a real right eigenvector of $M$ with a real eigenvalue $\lambda \notin\{0, \rho\}$. Let $U_{t}$ be the corresponding martingale from (6.1.7). The vector $\mathbf{u}$ is necessarily orthogonal to $\mathbf{v}$. Indeed $\mathbf{v} M \mathbf{u}$ is equal to both $\rho \mathbf{v u}$ and $\lambda \mathbf{v u}$. Because $\rho \neq \lambda$ by assumption, this is only possible if all three expressions are 0 . That implies $\mathbf{v u}=0$ since we also have $\rho \neq 0$ by assumption.

To compute the second moment of $U_{t}$, we mimic the computations in the proof of Lemma 6.1.13. We have

$$
\mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right]=\mathbb{E}\left[U_{t-1}^{2} \mid \mathbf{Z}_{0}\right]+\mathbb{E}\left[\left(U_{t}-U_{t-1}\right)^{2} \mid \mathbf{Z}_{0}\right],
$$

by the orthogonality of increments (Lemma 3.1.50). Since $\mathbb{E}\left[U_{t} \mid \mathcal{F}_{t-1}\right]=U_{t-1}$ by the martingale property, we get

$$
\begin{aligned}
\mathbb{E}\left[\left(U_{t}-U_{t-1}\right)^{2} \mid \mathcal{F}_{t-1}\right] & =\operatorname{Var}\left[U_{t} \mid \mathcal{F}_{t-1}\right] \\
& =\operatorname{Var}\left[\lambda^{-t} \mathbf{Z}_{t} \mathbf{u} \mid \mathcal{F}_{t-1}\right] \\
& =\lambda^{-2 t} \operatorname{Var}\left[\left(\sum_{\alpha=1}^{\tau} \sum_{i=1}^{Z_{t-1, \alpha}} \mathbf{X}^{(\alpha)}(i, t)\right) \mathbf{u} \mid \mathcal{F}_{t-1}\right] \\
& =\lambda^{-2 t} \sum_{\alpha=1}^{\tau} Z_{t-1, \alpha} \operatorname{Var}\left[\mathbf{X}^{(\alpha)}(1,1) \mathbf{u}\right] \\
& =\lambda^{-2 t} \mathbf{Z}_{t-1} \mathbf{S}^{(\mathbf{u})}
\end{aligned}
$$

where $\mathbf{S}^{(\mathbf{u})}=\left(\operatorname{Var}\left[\mathbf{X}^{(1)}(1,1) \mathbf{u}\right], \ldots, \operatorname{Var}\left[\mathbf{X}^{(\tau)}(1,1) \mathbf{u}\right]\right)$ as a column vector. In the last display, we used (6.1.4) on the third line and the independence of the random vectors $\mathbf{X}^{(\alpha)}(i, t)$ on the fourth line. Hence, taking expectations and using (6.1.6), we get

$$
\mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right]=\mathbb{E}\left[U_{t-1}^{2} \mid \mathbf{Z}_{0}\right]+\lambda^{-2 t} \mathbf{Z}_{0} M^{t-1} \mathbf{S}^{(\mathbf{u})}
$$

and finally

$$
\begin{equation*}
\mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right]=\left(\mathbf{Z}_{0} \mathbf{u}\right)^{2}+\sum_{s=1}^{t} \lambda^{-2 s} \mathbf{Z}_{0} M^{s-1} \mathbf{S}^{(\mathbf{u})} \tag{6.1.9}
\end{equation*}
$$

The case $\mathbf{S}^{(\mathbf{u})}=\mathbf{0}$ is trivial (see Exercise 6.4), so we exclude it from the following lemma.
Lemma 6.1.20 (Second moment of $U_{t}$ ). Assume $\mathbf{S}^{(\mathbf{u})} \neq \mathbf{0}$ and $\mathbf{Z}_{0} \neq \mathbf{0}$. The sequence $\mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right], t=0,1,2, \ldots$, is non-decreasing and satisfies

$$
\sup _{t \geq 0} \mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right] \begin{cases}<+\infty & \text { if } \rho<\lambda^{2} \\ =+\infty & \text { otherwise } .\end{cases}
$$

Proof. Because $\mathbf{S}^{(\mathbf{u})} \neq \mathbf{0}$ and nonnegative and the matrix $M$ is strictly positive by assumption, we have that

$$
\widetilde{\mathbf{S}}^{(\mathbf{u})}:=M \mathbf{S}^{(\mathbf{u})}>\mathbf{0} .
$$

Since $\mathbf{w}$ is also strictly positive, there is $0<C^{-} \leq C^{+}<+\infty$ such that

$$
C^{-} \mathbf{w} \leq \widetilde{\mathbf{S}}^{(\mathbf{u})} \leq C^{+} \mathbf{w} .
$$

Moreover, since $M$ is positive, each inequality is preserved when multiplying on both sides by $M$, that is, for any $s \geq 1$

$$
\begin{equation*}
C^{-} \rho^{s} \mathbf{w} \leq M^{s} \widetilde{\mathbf{S}}^{(\mathbf{u})} \leq C^{+} \rho^{s} \mathbf{w} . \tag{6.1.10}
\end{equation*}
$$

Now rewrite (6.1.9) as

$$
\mathbb{E}\left[U_{t}^{2} \mid \mathbf{Z}_{0}\right]=\left(\mathbf{Z}_{0} \mathbf{u}\right)^{2}+\lambda^{-2} \mathbf{Z}_{0} \mathbf{S}^{(\mathbf{u})}+\lambda^{-4} \sum_{s=2}^{t} \mathbf{Z}_{0}\left(1 / \lambda^{2}\right)^{s-2} M^{s-2} \widetilde{\mathbf{S}}^{(\mathbf{u})}
$$

There are two cases:

- When $\rho<\lambda^{2}$, using (6.1.10), the sum on the right-hand side can be bounded above by

$$
C^{+} \mathbf{Z}_{0} \mathbf{w} \sum_{s=2}^{t}\left(\frac{\rho}{\lambda^{2}}\right)^{s-2} \leq C^{+} \mathbf{Z}_{0} \mathbf{w} \frac{1}{1-\left(\rho / \lambda^{2}\right)}<+\infty
$$

uniformly in $t$.

- When $\rho \geq \lambda^{2}$, the same sum can be bounded from below by

$$
C^{-} \mathbf{Z}_{0} \mathbf{w} \sum_{s=2}^{t}\left(\frac{\rho}{\lambda^{2}}\right)^{s-2} \rightarrow+\infty
$$

as $t \rightarrow+\infty$. Indeed, $\mathbf{Z}_{0} \neq \mathbf{0}$ implies that the inner product $\mathbf{Z}_{0} \mathbf{w}$ is strictly positive.

That proves the claim.
In the case $\rho<\lambda^{2}$, the martingale $\left(U_{t}\right)$ is bounded in $L^{2}$ and therefore converges almost surely to a limit $U_{\infty}$ with $\mathbb{E}\left[U_{\infty} \mid \mathbf{Z}_{0}\right]=\mathbf{Z}_{0} \mathbf{u}$ by Theorem 3.1.51. On the other hand, when $\rho \geq \lambda^{2}$, it can be shown (we will not do this here) that $\mathbf{Z}_{t} \mathbf{u} / \sqrt{\mathbf{Z}_{t} \mathbf{w}}$ satisfies a central limit theorem with a limit independent of $\mathbf{Z}_{0}$. Implications of these claims are illustrated in Section 6.3.2.

### 6.2 Random-walk representation

In this section, we develop a random-walk representation of the Galton-Watson process. We give two applications: a characterization of the Galton-Watson process conditioned on extinction in terms of a dual branching process; and a formula for the size of the total progeny. We illustrate both in Section 6.2.4, where we revisit percolation on the infinite $b$-ary tree.

### 6.2.1 Exploration process

We introduce an exploration process where a random-walk perspective will naturally arise.

## Exploration of a graph

Because this will be useful again later, we describe it first in the context of a locally finite graph $G=(V, E)$. The exploration process starts at an arbitrary vertex $v \in V$ and has 3 types of vertices:

- $\mathcal{A}_{t}$ : active vertices,
- $\mathcal{E}_{t}:$ explored vertices,
- $\mathcal{N}_{t}:$ neutral vertices.


Figure 6.3: Exploration process for $\mathcal{C}_{v}$.
At the beginning, we have $\mathcal{A}_{0}:=\{v\}, \mathcal{E}_{0}:=\emptyset$, and $\mathcal{N}_{0}$ contains all other vertices in $G$. At time $t$, if $\mathcal{A}_{t-1}=\emptyset$ (i.e., there are no active vertices) we let $\left(\mathcal{A}_{t}, \mathcal{E}_{t}, \mathcal{N}_{t}\right):=$ $\left(\mathcal{A}_{t-1}, \mathcal{E}_{t-1}, \mathcal{N}_{t-1}\right)$. Otherwise, we pick an element, $a_{t}$, from $\mathcal{A}_{t-1}$ (say in firstcome, first-served basis to be explicit) and set:

- $\mathcal{A}_{t}:=\left(\mathcal{A}_{t-1} \backslash\left\{a_{t}\right\}\right) \cup\left\{x \in \mathcal{N}_{t-1}:\left\{x, a_{t}\right\} \in E\right\}$
- $\mathcal{E}_{t}:=\mathcal{E}_{t-1} \cup\left\{a_{t}\right\}$
- $\mathcal{N}_{t}:=\mathcal{N}_{t-1} \backslash\left\{x \in \mathcal{N}_{t-1}:\left\{x, a_{t}\right\} \in E\right\}$

We imagine revealing the edges of $G$ as they are encountered in the exploration process. In words, starting with $v$, the connected component $\mathcal{C}_{v}$ of $v$ is progressively grown by adding to it at each time a vertex adjacent to one of the previously explored vertices and uncovering its remaining neighbors in $G$. In this process, $\mathcal{E}_{t}$ is the set of previously explored vertices and $\mathcal{A}_{t}$-the frontier of the process-is the set of vertices who are known to belong to $\mathcal{C}_{v}$ but whose full neighborhood is waiting to be uncovered. The rest of the vertices form the set $\mathcal{N}_{t}$. See Figure 6.3.

Let $A_{t}:=\left|\mathcal{A}_{t}\right|, E_{t}:=\left|\mathcal{E}_{t}\right|$, and $N_{t}:=\left|\mathcal{N}_{t}\right|$. Note that $\left(E_{t}\right)$-not to be confused with the edge set-is non-decreasing while $\left(N_{t}\right)$ is non-increasing. Let

$$
\tau_{0}:=\inf \left\{t \geq 0: A_{t}=0\right\}
$$

be the first time $A_{t}$ is 0 (which by convention is $+\infty$ if there is no such $t$ ). The process is fixed for all $t>\tau_{0}$. Notice that $E_{t}=t$ for all $t \leq \tau_{0}$, as exactly one
vertex is explored at each time until the set of active vertices is empty. The size of the connected component of $v$ can be characterized as follows.

Lemma 6.2.1.

$$
\tau_{0}=\left|\mathcal{C}_{v}\right|
$$

Proof. Indeed a single vertex of $\mathcal{C}_{v}$ is explored at each time until all of $\mathcal{C}_{v}$ has been visited. At that point, $\mathcal{A}_{t}$ is empty.

## Random-walk representation of a Galton-Watson tree

Let $\left(Z_{i}\right)_{i \geq 0}$ be a Galton-Watson branching process and let $T$ be the corresponding Galton-Watson tree. We run the exploration process above on $T$ started at the root 0 . We will refer to the index $i$ in $Z_{i}$ as a "generation," and to the index $t$ in the exploration process as "time"-they are not the same. Let $\left(\mathcal{A}_{t}, \mathcal{E}_{t}, \mathcal{N}_{t}\right)$ and $A_{t}:=\left|\mathcal{A}_{t}\right|, E_{t}:=\left|\mathcal{E}_{t}\right|$, and $N_{t}:=\left|\mathcal{N}_{t}\right|$ be as above. Let $\left(\mathcal{F}_{t}\right)$ be the corresponding filtration. Because we explore the vertices on first-come, first-served basis, we exhaust all vertices in generation $i$ before considering vertices in generation $i+1$ (i.e., we perform breadth-first search).

The random-walk representation is the following. Observe that the process $\left(A_{t}\right)$ admits a simple recursive form. We start with $A_{0}:=1$. Then, conditioning on $\mathcal{F}_{t-1}$ :

- If $A_{t-1}=0$, the exploration process has finished its course and $A_{t}=0$.
- Otherwise, (a) one active vertex becomes an explored vertex and (b) its offspring become active vertices. That is,

$$
A_{t}= \begin{cases}A_{t-1}+(\underbrace{-1}_{\text {(a) }}+\underbrace{X_{t}}_{\text {(b) }}) & \text { if } t-1<\tau_{0}, \\ 0 & \text { otherwise }\end{cases}
$$

where $X_{t}$ is distributed according to the offspring distribution.
We let $Y_{t}:=X_{t}-1$ and

$$
S_{t}:=1+\sum_{s=1}^{t} Y_{s}
$$

with $S_{0}:=1$. Then

$$
\tau_{0}=\inf \left\{t \geq 0: S_{t}=0\right\}
$$

and

$$
\left(A_{t}\right)=\left(S_{t \wedge \tau_{0}}\right),
$$

is a random walk started at 1 with i.i.d. increments $\left(Y_{t}\right)$ stopped when it hits 0 for the first time.

We refer to

$$
H=\left(X_{1}, \ldots, X_{\tau_{0}}\right),
$$

as the history of the process $\left(Z_{i}\right)$. Observe that, under breadth-first search, the process $\left(Z_{i}\right)$ can be reconstructed from $H: Z_{0}=1, Z_{1}=X_{1}, Z_{2}=X_{2}+\ldots+$ $X_{Z_{1}+1}$, and so forth. (Exercise 6.5 asks for a general formula.) As a result, $\left(Z_{i}\right)$ can be recovered from $\left(S_{t}\right)$ as well. We call $\left(x_{1}, \ldots, x_{t}\right)$ a valid history if

$$
1+\left(x_{1}-1\right)+\cdots+\left(x_{s}-1\right)>0,
$$

for all $s<t$ and

$$
1+\left(x_{1}-1\right)+\cdots+\left(x_{t}-1\right)=0 .
$$

Note that a valid history may have probability 0 under the offspring distribution.

### 6.2.2 Duality principle

The random-walk representation above is useful to prove the following duality principle.

Theorem 6.2.2 (Duality principle). Let $\left(Z_{i}\right)$ be a branching process with offspring distribution $\left\{p_{k}\right\}_{k \geq 0}$ and extinction probability $\eta<1$. Let $\left(Z_{i}^{\prime}\right)$ be a branching process with offspring distribution $\left\{p_{k}^{\prime}\right\}_{k \geq 0}$ where

$$
p_{k}^{\prime}=\eta^{k-1} p_{k} .
$$

Then $\left(Z_{i}\right)$ conditioned on extinction has the same distribution as $\left(Z_{i}^{\prime}\right)$, which is referred to as the dual branching process.

Let $f$ be the probability generating function of the offspring distribution of $\left(Z_{i}\right)$. process Note that

$$
\sum_{k \geq 0} p_{k}^{\prime}=\sum_{k \geq 0} \eta^{k-1} p_{k}=\eta^{-1} f(\eta)=1,
$$

because $\eta$ is a fixed point of $f$ by Theorem 6.1.6. So $\left\{p_{k}^{\prime}\right\}_{k \geq 0}$ is indeed a probability distribution. Note further that its expectation is

$$
\sum_{k \geq 0} k p_{k}^{\prime}=\sum_{k \geq 0} k \eta^{k-1} p_{k}=f^{\prime}(\eta)<1,
$$

since, by Lemma 6.1.7, $f^{\prime}$ is strictly increasing, $f(\eta)=\eta<1$ and $f(1)=1$ (which would not be possible if $f^{\prime}(\eta)$ were greater or equal to 1 ; see Figure 6.1 for an illustration). So the dual branching process is subcritical.

Proof of Theorem 6.2.2. We use the random-walk representation. Let $H=\left(X_{1}, \ldots, X_{\tau_{0}}\right)$ and $H^{\prime}=\left(X_{1}^{\prime}, \ldots, X_{\tau_{0}^{\prime}}^{\prime}\right)$ be the histories of $\left(Z_{i}\right)$ and $\left(Z_{i}^{\prime}\right)$ respectively. In the case of extinction of $\left(Z_{i}\right)$, the history $H$ has finite length.

By definition of the conditional probability, for a valid history $\left(x_{1}, \ldots, x_{t}\right)$ with a finite $t$,

$$
\mathbb{P}\left[H=\left(x_{1}, \ldots, x_{t}\right) \mid \tau_{0}<+\infty\right]=\frac{\mathbb{P}\left[H=\left(x_{1}, \ldots, x_{t}\right)\right]}{\mathbb{P}\left[\tau_{0}<+\infty\right]}=\eta^{-1} \prod_{s=1}^{t} p_{x_{s}}
$$

Because $\left(x_{1}-1\right)+\cdots+\left(x_{t}-1\right)=-1$,

$$
\eta^{-1} \prod_{s=1}^{t} p_{x_{s}}=\eta^{-1} \prod_{s=1}^{t} \eta^{1-x_{s}} p_{x_{s}}^{\prime}=\prod_{s=1}^{t} p_{x_{s}}^{\prime}=\mathbb{P}\left[H^{\prime}=\left(x_{1}, \ldots, x_{t}\right)\right] .
$$

Since this is true for all valid histories and the processes can be recovered from their histories, we have proved the claim.

Example 6.2.3 (Poisson branching process). Let $\left(Z_{i}\right)$ be a Galton-Watson branching process with offspring distribution $\operatorname{Poi}(\lambda)$ where $\lambda>1$. Then the dual probability distribution is given by

$$
p_{k}^{\prime}=\eta^{k-1} p_{k}=\eta^{k-1} e^{-\lambda} \frac{\lambda^{k}}{k!}=\eta^{-1} e^{-\lambda} \frac{(\lambda \eta)^{k}}{k!},
$$

where recall from Example 6.1.10 that $e^{-\lambda(1-\eta)}=\eta$, so

$$
p_{k}^{\prime}=e^{\lambda(1-\eta)} e^{-\lambda} \frac{(\lambda \eta)^{k}}{k!}=e^{-\lambda \eta} \frac{(\lambda \eta)^{k}}{k!} .
$$

That is, the dual branching process has offspring distribution $\operatorname{Poi}(\lambda \eta)$.

### 6.2.3 Hitting-time theorem

The random-walk representation also gives a formula for the distribution of the size of the progeny.

Law of total progeny The key is the following claim.
Lemma 6.2.4 (Total progeny and random-walk representation). Let $W$ be the total progeny of the Galton-Watson branching process $\left(Z_{i}\right)$. Then

$$
W=\tau_{0} .
$$

Proof. Recall that

$$
\tau_{0}:=\inf \left\{t \geq 0: A_{t}=0\right\}
$$

If the process does not go extinct, then $\tau_{0}=+\infty$ as there are always more vertices to explore.

Suppose the process goes extinct and that $W=n$. Notice that $E_{t}=t$ for all $t \leq \tau_{0}$, as exactly one vertex is explored at each time until the set of active vertices is empty. Moreover, for all $t,\left(\mathcal{A}_{t}, \mathcal{E}_{t}, \mathcal{N}_{t}\right)$ forms a partition of $[n]$ so

$$
A_{t}+t+N_{t}=n, \quad \forall t \leq \tau_{0} .
$$

At $t=\tau_{0}, A_{t}=N_{t}=0$ and we get

$$
\tau_{0}=n
$$

That proves the claim.
To compute the distribution of $W=\tau_{0}$, we use the following hitting-time theorem, which is proved later in this subsection.

Theorem 6.2.5 (Hitting-time theorem). Let $\left(R_{t}\right)$ be a random walk started at 0 with i.i.d. increments $\left(U_{t}\right)$ satisfying

$$
\mathbb{P}\left[U_{t} \leq 1\right]=1
$$

Fix a positive integer $\ell$. Let $\sigma_{\ell}$ be the first time $t$ such that $R_{t}=\ell$. Then

$$
\mathbb{P}\left[\sigma_{\ell}=t\right]=\frac{\ell}{t} \mathbb{P}\left[R_{t}=\ell\right] .
$$

Finally we get:
Theorem 6.2.6 (Law of total progeny). Let $\left(Z_{t}\right)$ be a Galton-Watson branching process with total progeny $W$. In the random-walk representation of $\left(Z_{t}\right)$,

$$
\mathbb{P}[W=t]=\frac{1}{t} \mathbb{P}\left[X_{1}+\cdots+X_{t}=t-1\right]
$$

for all $t \geq 1$.

Proof. Recall that $Y_{t}:=X_{t}-1 \geq-1$ and

$$
S_{t}=1+\sum_{s=1}^{t} Y_{s}
$$

with $S_{0}=1$, and that

$$
\begin{aligned}
\tau_{0} & =\inf \left\{t \geq 0: S_{t}=0\right\} \\
& =\inf \left\{t \geq 0: 1+\left(X_{1}-1\right)+\cdots+\left(X_{t}-1\right)=0\right\} \\
& =\inf \left\{t \geq 0: X_{1}+\cdots+X_{t}=t-1\right\} .
\end{aligned}
$$

Define $R_{t}:=1-S_{t}$ and $U_{t}:=-Y_{t}$ for all $t$. Then $R_{0}:=0$,

$$
\left\{X_{1}+\cdots+X_{t}=t-1\right\}=\left\{R_{t}=1\right\},
$$

and

$$
\tau_{0}=\inf \left\{t \geq 0: R_{t}=1\right\}
$$

The process $\left(R_{t}\right)$ satisfies the assumptions of the hitting-time theorem (Theorem 6.2.5) with $\ell=1$ and $\sigma_{\ell}=\tau_{0}=W$. Applying the theorem gives the claim.

Example 6.2.7 (Poisson branching process (continued)). Let ( $Z_{i}$ ) be a GaltonWatson branching process with offspring distribution $\operatorname{Poi}(\lambda)$ where $\lambda>0$. Let $W$ be its total progeny. By the hitting-time theorem, for $t \geq 1$,

$$
\begin{aligned}
\mathbb{P}[W=t] & =\frac{1}{t} \mathbb{P}\left[X_{1}+\cdots+X_{t}=t-1\right] \\
& =\frac{1}{t} e^{-\lambda t} \frac{(\lambda t)^{t-1}}{(t-1)!} \\
& =e^{-\lambda t} \frac{(\lambda t)^{t-1}}{t!}
\end{aligned}
$$

where we used that a sum of independent Poisson is Poisson.

Spitzer's combinatorial lemma Before proving the hitting-time theorem, we begin with a combinatorial lemma of independent interest. Let $u_{1}, \ldots, u_{t} \in \mathbb{R}$ and define $r_{0}:=0$ and $r_{j}:=u_{1}+\cdots+u_{j}$ for $1 \leq j \leq t$. We say that $j$ is a ladder index if $r_{j}>r_{0} \vee \cdots \vee r_{j-1}$. Consider the cyclic permutations of $\boldsymbol{u}=\left(u_{1}, \ldots, u_{t}\right)$, that is, $\boldsymbol{u}^{(0)}=\boldsymbol{u}, \boldsymbol{u}^{(1)}=\left(u_{2}, \ldots, u_{t}, u_{1}\right), \ldots, \boldsymbol{u}^{(t-1)}=\left(u_{t}, u_{1}, \ldots, u_{t-1}\right)$. Define the corresponding partial sums $r_{j}^{(\beta)}:=u_{1}^{(\beta)}+\cdots+u_{j}^{(\beta)}$ for $j=1, \ldots, t$ and $\beta=0, \ldots, t-1$.

Lemma 6.2.8 (Spitzer's combinatorial lemma). Assume $r_{t}>0$. Let $\ell$ be the number of cyclic permutations such that $t$ is a ladder index. Then $\ell \geq 1$ and each such cyclic permutation has exactly $\ell$ ladder indices.

Proof. We will need the following observation

$$
\begin{align*}
\left(r_{1}^{(\beta)},\right. & \left.\ldots, r_{t}^{(\beta)}\right) \\
= & \left(r_{\beta+1}-r_{\beta}, r_{\beta+2}-r_{\beta}, \ldots, r_{t}-r_{\beta},\right. \\
& \left.\quad\left[r_{t}-r_{\beta}\right]+r_{1},\left[r_{t}-r_{\beta}\right]+r_{2}, \ldots,\left[r_{t}-r_{\beta}\right]+r_{\beta}\right) \\
= & \left(r_{\beta+1}-r_{\beta}, r_{\beta+2}-r_{\beta}, \ldots, r_{t}-r_{\beta},\right. \\
& \left.\quad r_{t}-\left[r_{\beta}-r_{1}\right], r_{t}-\left[r_{\beta}-r_{2}\right], \ldots, r_{t}-\left[r_{\beta}-r_{\beta-1}\right], r_{t}\right) . \tag{6.2.1}
\end{align*}
$$

We first show that $\ell \geq 1$, that is, there is at least one cyclic permutation where $t$ is a ladder index. Let $\beta \geq 1$ be the smallest index achieving the maximum of $r_{1}, \ldots, r_{t}$, that is,

$$
r_{\beta}>r_{1} \vee \cdots \vee r_{\beta-1} \quad \text { and } \quad r_{\beta} \geq r_{\beta+1} \vee \cdots \vee r_{t} .
$$

Moreover, $r_{t}>0=r_{0}$ by assumption. Hence,

$$
r_{\beta+j}-r_{\beta} \leq 0<r_{t}, \quad \forall j=1, \ldots, t-\beta,
$$

and

$$
r_{t}-\left[r_{\beta}-r_{j}\right]<r_{t}, \quad \forall j=1, \ldots, \beta-1 .
$$

From (6.2.1), in $\boldsymbol{u}^{(\beta)}, t$ is a ladder index.
For the second claim, since $\ell \geq 1$, we can assume without loss of generality that $\boldsymbol{u}$ is such that $t$ is a ladder index. (Note that $r_{t}^{(\beta)}=r_{t}$ for all $\beta$.) We show that $\beta$ is a ladder index in $\boldsymbol{u}$ if and only if $t$ is a ladder index in $\boldsymbol{u}^{(\beta)}$. That does indeed imply the claim as there are $\ell$ cyclic permutations where $t$ is a ladder index by assumption. We use (6.2.1) again. Observe that $\beta$ is a ladder index in $\boldsymbol{u}$ if and only if

$$
r_{\beta}>r_{0} \vee \cdots \vee r_{\beta-1},
$$

which holds if and only if

$$
\begin{equation*}
r_{\beta}>r_{0}=0 \quad \text { and } \quad r_{t}-\left[r_{\beta}-r_{j}\right]<r_{t}, \forall j=1, \ldots, \beta-1 . \tag{6.2.2}
\end{equation*}
$$

Moreover, because $r_{t}>r_{j}$ for all $j$ by the assumption that $t$ is ladder index, the last display holds if and only if

$$
\begin{equation*}
r_{\beta+j}-r_{\beta}<r_{t}, \forall j=1, \ldots, t-\beta, \tag{6.2.3}
\end{equation*}
$$

and

$$
\begin{equation*}
r_{t}-\left[r_{\beta}-r_{j}\right]<r_{t}, \forall j=1, \ldots, \beta-1, \tag{6.2.4}
\end{equation*}
$$

that is, if and only if $t$ is a ladder index in $\boldsymbol{u}^{(\beta)}$ by (6.2.1). Indeed, the second condition (i.e., (6.2.4)) is intact from (6.2.2), while the first one (i.e., (6.2.3)) can be rewritten as $r_{\beta}>-\left(r_{t}-r_{\beta+j}\right)$ where the right-hand side is $<0$ for $j=$ $1, \ldots, t-\beta-1$ and $=0$ for $j=t-\beta$.

Proof of hitting-time theorem We are now ready to prove the hitting-time theorem. We only handle the case $\ell=1$ (which is the one we used for the law of the total progeny). Exercise 6.6 asks for the full proof.

Proof of Theorem 6.2.5. Recall that $R_{t}=\sum_{s=1}^{t} U_{s}$ and $\sigma_{1}=\inf \left\{j \geq 0: R_{j}=\right.$ $1\}$. By the assumption that $U_{s} \leq 1$ almost surely for all $s$,

$$
\left\{\sigma_{1}=t\right\}=\left\{t \text { is the first ladder index in } R_{1}, \ldots, R_{t}\right\} .
$$

By symmetry, for all $\beta=0, \ldots, t-1$,

$$
\begin{aligned}
& \mathbb{P}\left[t \text { is the first ladder index in } R_{1}, \ldots, R_{t}\right] \\
& \quad=\mathbb{P}\left[t \text { is the first ladder index in } R_{1}^{(\beta)}, \ldots, R_{t}^{(\beta)}\right] .
\end{aligned}
$$

Let $\mathcal{E}_{\beta}$ be the event on the last line. Then,

$$
\mathbb{P}\left[\sigma_{1}=t\right]=\mathbb{E}\left[\mathbf{1}_{\mathcal{E}_{0}}\right]=\frac{1}{t} \mathbb{E}\left[\sum_{\beta=0}^{t-1} \mathbf{1}_{\mathcal{E}_{\beta}}\right]
$$

By Spitzer's combinatorial lemma (Lemma 6.2.8), there is at most one cyclic permutation where $t$ is the first ladder index. (There is at least one cyclic permutation where $t$ is a ladder index-but it may not be the first one, that is, there may be multiple ladder indices.) In particular, $\sum_{\beta=0}^{t-1} \mathbf{1}_{\mathcal{E}_{\beta}} \in\{0,1\}$. So, by the previous display,

$$
\mathbb{P}\left[\sigma_{1}=t\right]=\frac{1}{t} \mathbb{P}\left[\cup_{\beta=0}^{t-1} \mathcal{E}_{\beta}\right] .
$$

Finally we claim that $\left\{R_{t}=1\right\}=\cup_{\beta=0}^{t-1} \mathcal{E}_{\beta}$. Indeed, because $R_{0}=0$ and $U_{s} \leq 1$ for all $s$, the partial sum at the $j$-th ladder index must take value $j$. So the event $\cup_{\beta=0}^{t-1} \mathcal{E}_{\beta}$ implies $\left\{R_{t}=1\right\}$ since the last partial sum of all cyclic permutations is $R_{t}$. Similarly, because there is at least one cyclic permutation such that $t$ is a ladder index, the event $\left\{R_{t}=1\right\}$ implies that $t$ is in fact the first ladder index in that cyclic permutation, and therefore it implies $\cup_{\beta=1}^{t} \mathcal{E}_{\beta}$. Hence,

$$
\mathbb{P}\left[\sigma_{1}=t\right]=\frac{1}{t} \mathbb{P}\left[R_{t}=1\right],
$$

which concludes the proof (for the case $\ell=1$ ).

### 6.2.4 $\triangleright$ Percolation: critical exponents on the infinite b-ary tree

In this section, we use branching processes to study bond percolation (Definition 1.2.1) on the infinite $b$-ary tree $\widehat{\mathbb{T}}_{b}$ and derive explicit expressions for quantities of interest. Close to the critical value, we prove the existence of "critical exponents." We illustrate the use of both the duality principle (Theorem 6.2.2) and the hitting-time theorem (Theorem 6.2.6).

Critical value We denote the root by 0 . Similarly to what we did in Section 6.1.3, we think of the open cluster of the root, $\mathcal{C}_{0}$, as the progeny of a branching process as follows. Denote by $\partial_{n}$ the $n$-th level of $\widehat{\mathbb{T}}_{b}$, that is, the vertices of $\widehat{\mathbb{T}}_{b}$ at graph distance $n$ from the root. In the branching process interpretation, we think of the immediate descendants in $\mathcal{C}_{0}$ of a vertex $v$ as the offspring of $v$. By construction, $v$ has at most $b$ children, independently of all other vertices in the same generation. In this branching process, the offspring distribution $\left\{q_{k}\right\}_{k=0}^{b}$ is binomial with parameters $b$ and $p ; Z_{n}:=\left|\mathcal{C}_{0} \cap \partial_{n}\right|$ represents the size of the progeny at generation $n$; and $W:=\left|\mathcal{C}_{0}\right|$ is the total progeny of the process. In particular $\left|\mathcal{C}_{0}\right|<+\infty$ if and only if the process goes extinct. Because the mean number of offspring is $b p$, by Theorem 6.1.6, this leads immediately to a second proof of (a rooted variant of) Claim 2.3.9:

## Claim 6.2.9.

$$
p_{c}\left(\widehat{\mathbb{T}}_{b}\right)=\frac{1}{b} .
$$

Percolation function The generating function of the offspring distribution is $\phi(s):=((1-p)+p s)^{b}$. So, by Theorems 6.1.5 and 6.1.6, the percolation function

$$
\theta(p)=\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=+\infty\right],
$$

is 0 on $[0,1 / b]$, while on $(1 / b, 1]$ the quantity $\eta(p):=1-\theta(p)$ is the unique solution in $[0,1)$ of the fixed point equation

$$
\begin{equation*}
s=((1-p)+p s)^{b} . \tag{6.2.5}
\end{equation*}
$$

For $b=2$, for instance, we can compute the fixed point explicitly by noting that

$$
\begin{aligned}
0 & =((1-p)+p s)^{2}-s \\
& =p^{2} s^{2}+[2 p(1-p)-1] s+(1-p)^{2},
\end{aligned}
$$

whose solution for $p \in(1 / 2,1]$ is

$$
\begin{aligned}
s^{*} & =\frac{-[2 p(1-p)-1] \pm \sqrt{[2 p(1-p)-1]^{2}-4 p^{2}(1-p)^{2}}}{2 p^{2}} \\
& =\frac{-[2 p(1-p)-1] \pm \sqrt{1-4 p(1-p)}}{2 p^{2}} \\
& =\frac{-[2 p(1-p)-1] \pm(2 p-1)}{2 p^{2}} \\
& =\frac{2 p^{2}+[(1-2 p) \pm(2 p-1)]}{2 p^{2}} .
\end{aligned}
$$

So, rejecting the fixed point 1 ,

$$
\theta(p)=1-\frac{2 p^{2}+2(1-2 p)}{2 p^{2}}=\frac{2 p-1}{p^{2}} .
$$

We have proved:
Claim 6.2.10. For $b=2$,

$$
\theta(p)= \begin{cases}0, & 0 \leq p \leq \frac{1}{2} \\ \frac{2\left(p-\frac{1}{2}\right)}{p^{2}} & \frac{1}{2}<p \leq 1 .\end{cases}
$$

Since $\eta(p)=(1-\theta(p))$, we have in that case

$$
\eta(p)= \begin{cases}1, & 0 \leq p \leq \frac{1}{2} \\ \frac{(1-p)^{2}}{p^{2}} & \frac{1}{2}<p \leq 1\end{cases}
$$

Conditioning on a finite cluster The expected size of the population at generation $n$ is $(b p)^{n}$ by Lemma 6.1.2, so for $p \in\left[0, \frac{1}{b}\right)$

$$
\begin{equation*}
\mathbb{E}_{p}\left|\mathcal{C}_{0}\right|=\sum_{n \geq 0}(b p)^{n}=\frac{1}{1-b p} . \tag{6.2.6}
\end{equation*}
$$

For $p \in\left(\frac{1}{b}, 1\right)$, the total progeny is infinite with positive probability (and in particular the expectation is infinite), but we can compute the expected cluster size on the event that $\left|\mathcal{C}_{0}\right|<+\infty$. For this purpose we use the duality principle.

Recall that $q_{k}=\binom{b}{k} p^{k}(1-p)^{b-k}, k=0, \ldots, b$, is the offspring distribution. For $0 \leq k \leq b$, we let the dual offspring distribution be

$$
\begin{aligned}
\hat{q}_{k} & :=[\eta(p)]^{k-1} q_{k} \\
& =[\eta(p)]^{k-1}\binom{b}{k} p^{k}(1-p)^{b-k} \\
& =\frac{[\eta(p)]^{k}}{((1-p)+p \eta(p))^{b}}\binom{b}{k} p^{k}(1-p)^{b-k} \\
& =\binom{b}{k}\left(\frac{p \eta(p)}{(1-p)+p \eta(p)}\right)^{k}\left(\frac{1-p}{(1-p)+p \eta(p)}\right)^{b-k} \\
& =:\binom{b}{k} \hat{p}^{k}(1-\hat{p})^{b-k},
\end{aligned}
$$

where we used (6.2.5) and implicitly defined the dual density

$$
\begin{equation*}
\hat{p}:=\frac{p \eta(p)}{(1-p)+p \eta(p)} . \tag{6.2.7}
\end{equation*}
$$

In particular $\left\{\hat{q}_{k}\right\}$ is a probability distribution as expected under Theorem 6.2.2it is in fact binomial with parameters $b$ and $\hat{p}$. Summarizing the implications of Theorem 6.2.2:

Claim 6.2.11. Conditioned on $\left|\mathcal{C}_{0}\right|<+\infty$, (supercritical) percolation on $\widehat{\mathbb{T}}_{b}$ with density $p \in\left(\frac{1}{b}, 1\right)$ has the same distribution as (subcritical) percolation on $\widehat{\mathbb{T}}_{b}$ with density defined by (6.2.7).

Hence, using (6.2.6) with both $p$ and $\hat{p}$ as well as the fact that $\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|<+\infty\right]=$ $\eta(p)$, we have the following.

## Claim 6.2.12.

$$
\chi^{\mathrm{f}}(p):=\mathbb{E}_{p}\left[\left|\mathcal{C}_{0}\right| \mathbf{1}_{\left\{\left|\mathcal{C}_{0}\right|<+\infty\right\}}\right]= \begin{cases}\frac{1}{1-b p}, & p \in\left[0, \frac{1}{b}\right), \\ \frac{\eta(p)}{1-b \bar{p}}, & p \in\left(\frac{1}{b}, 1\right) .\end{cases}
$$

For $b=2, \eta(p)=1-\theta(p)=\left(\frac{1-p}{p}\right)^{2}$ so

$$
\hat{p}=\frac{p\left(\frac{1-p}{p}\right)^{2}}{(1-p)+p\left(\frac{1-p}{p}\right)^{2}}=\frac{(1-p)^{2}}{p(1-p)+(1-p)^{2}}=1-p
$$

and

Claim 6.2.13. For $b=2$,

$$
\chi^{\mathrm{f}}(p)= \begin{cases}\frac{1 / 2}{\frac{1}{2}-p}, & p \in\left[0, \frac{1}{2}\right), \\ \frac{1}{2}\left(\frac{1-p}{p}\right)^{2} \\ p-\frac{1}{2} & \\ \hline & p \in\left(\frac{1}{2}, 1\right) .\end{cases}
$$

Distribution of the open cluster size In fact the hitting-time theorem gives an explicit formula for the distribution of $\left|\mathcal{C}_{0}\right|$. Namely, recall that $\left|\mathcal{C}_{0}\right| \stackrel{\text { d }}{=} \tau_{0}$ where

$$
\tau_{0}=\inf \left\{t \geq 0: S_{t}=0\right\},
$$

for $S_{t}=\sum_{\ell \leq t} X_{\ell}-(t-1)$ where $S_{0}=1$ and the $X_{\ell}$ s are i.i.d. binomial with parameters $b$ and $p$. By Theorem 6.2.6,

$$
\mathbb{P}\left[\tau_{0}=t\right]=\frac{1}{t} \mathbb{P}\left[S_{t}=0\right],
$$

and we have

$$
\begin{equation*}
\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=\ell\right]=\frac{1}{\ell} \mathbb{P}\left[\sum_{i \leq \ell} X_{\ell}=\ell-1\right]=\frac{1}{\ell}\binom{b \ell}{\ell-1} p^{\ell-1}(1-p)^{b \ell-(\ell-1)}, \tag{6.2.8}
\end{equation*}
$$

where we used that a sum of independent binomials with the same $p$ is itself binomial. In particular at criticality (where $\left|\mathcal{C}_{0}\right|<+\infty$ almost surely; see Claim 3.1.52), using Stirling's formula (see Appendix A) it can be checked that

$$
\left.\mathbb{P}_{p_{\mathrm{c}}}| | \mathcal{C}_{0} \mid=\ell\right] \sim \frac{1}{\ell} \frac{1}{\sqrt{2 \pi p_{\mathrm{c}}\left(1-p_{\mathrm{c}}\right) b \ell}}=\frac{1}{\sqrt{2 \pi\left(1-p_{\mathrm{c}}\right) \ell^{3}}}
$$

as $\ell \rightarrow+\infty$.

Critical exponents Close to criticality, physicists predict that many quantities behave according to power laws of the form $\left|p-p_{\mathrm{c}}\right|^{\beta}$, where the exponent is referred to as a critical exponent. The critical exponents are believed to satisfy certain "universality" properties. But even proving the existence of such exponents in general remains a major open problem. On trees, though, we can simply read off the critical exponents from the above formulas. For $b=2$, Claims 6.2.10 and 6.2.13 imply for instance that, as $p \rightarrow p_{\mathrm{c}}$,

$$
\theta(p) \sim 8\left(p-p_{\mathrm{c}}\right) \mathbf{1}_{\{p>1 / 2\}},
$$

and

$$
\chi^{\mathrm{f}}(p) \sim \frac{1}{2}\left|p-p_{\mathrm{c}}\right|^{-1} .
$$

In fact, as can be seen from Claim 6.2.12, the critical exponent of $\chi^{\mathrm{f}}(p)$ does not depend on $b$. The same holds for $\theta(p)$ (see Exercise 6.9). Using (6.2.8), the higher moments of $\left|\mathcal{C}_{0}\right|$ can also be studied around criticality (see Exercise 6.10).

### 6.3 Applications

We develop two applications of branching processes in discrete probability. First, we prove a result about the height of a random binary search tree. Then we describe a phase transition in an Ising model on a tree with applications to evolutionary biology. In the next section, we also use branching processes to study the phase transition of Erdős-Rényi random graph model.

### 6.3.1 $\triangleright$ Probabilistic analysis of algorithms: binary search tree

A binary search tree (BST) is a commonly used data structure in computer science. It consists of a rooted binary tree $T_{n}=\left(V_{n}, E_{n}\right)$. Each vertex has a "left" and "right" subtree (possibly empty) and a "key" from an input sequence $x_{1}, \ldots, x_{n} \in$ $\mathbb{R}$ (which we assume are distinct) that satisfies the BST property: the key at vertex $v \in V$ is greater than all keys in the left subtree below it and less than all keys in the right subtree below it. Such a data structure can be used for a variety of algorithmic tasks, such as searching for keys or sorting them.

The tree is constructed recursively as follows. Assume that the keys $x_{1}, \ldots, x_{i}$ have already been inserted and that the current tree $T_{i}$ satisfies the BST property. To insert $x_{i+1}$ :

- start at the root;
- if the root's key is strictly larger than $x_{i+1}$, then move to its left descendant, otherwise move to its right descendant;
- if such a descendant does not exist then create it and assign it $x_{i+1}$ as its key;
- otherwise repeat.

Inserting keys (and other operations such as deleting keys, which we do not describe) takes time proportional to the height $H_{n}$ of the tree $T_{n}$, that is, the length of the longest path from the root to a leaf. While, in general, the height can be as large as $n$ (if keys are inserted in order for instance), the typical behavior can be much smaller.

Indeed, here we study the case of $n$ keys $X_{1}, \ldots, X_{n}$ i.i.d. from a continuous distribution on $\mathbb{R}$ and establish a much better behavior for the random height. Let $\gamma$ be the unique solution greater than 1 of

$$
\begin{equation*}
\left(\frac{1}{e}\right)\left(\frac{2 e}{\gamma}\right)^{\gamma}=1 . \tag{6.3.1}
\end{equation*}
$$

See Exercise 6.14 for a proof that $\gamma$ is well-defined and that the left-hand side is strictly decreasing at $\gamma$. We show:

Claim 6.3.1. $H_{n} / \log n \rightarrow_{\mathrm{p}} \gamma$ as $n \rightarrow \infty$.

## Alternative representation of the height

The main idea of the proof is to relate the height $H_{n}$ of the tree $T_{n}$ to a product of independent uniform random variables. We make a series of observations about the structure of the tree. First:

Observation 1. Keys affect the construction of the binary search tree only through their ordering. Let $\sigma$ be the corresponding (random) permutation, that is,

$$
X_{\sigma(1)}<X_{\sigma(2)}<\cdots<X_{\sigma(n)} .
$$

Let $t[\sigma]$ be the binary search tree generated by the permutation $\sigma$.
Second, by symmetry:
Observation 2. The permutation $\sigma$ is uniformly distributed.
Denote by $S_{v}$ the size of the subtree rooted at $v$ (including $v$ itself) in $t[\sigma]$. At the root $\rho$, we have $S_{\rho}=n$. What is the size of the subtree rooted at the left descendant $\rho^{\prime}$ of $\rho$ ? Eventually all keys with a rank lower than $\sigma^{-1}(1)$, that is, those keys with indices in $\left\{\sigma(i): i<\sigma^{-1}(1)\right\}$, find their way into the left subtree of the root. In other words,

$$
S_{\rho^{\prime}}=\sigma^{-1}(1)-1 .
$$

Similarly, denoting by $\rho^{\prime \prime}$ the right descendant of $\rho$, we see that

$$
S_{\rho^{\prime \prime}}=n-\sigma^{-1}(1) .
$$

We refer to $\sigma^{-1}(1)$ as the rank of the root. By Observation 2:
Observation 3. The rank $\sigma^{-1}(1)$ of the root is uniformly distributed in $[n]$. Moreover it is identically distributed to $\left\lfloor S_{\rho} W_{\rho}\right\rfloor+1$, where $W_{\rho}$ is uniform in $[0,1]$.

The second part of this last observation can be checked by direct computation. Rename $X_{1}^{\prime}, \ldots, X_{S_{\rho^{\prime}}}^{\prime}$ the keys in the subtree rooted at $\rho^{\prime}$ in the order that they are inserted and let $\sigma^{\prime}$ be the (random) permutation corresponding to their ordering, that is,

$$
X_{\sigma^{\prime}(1)}^{\prime}<X_{\sigma^{\prime}(2)}^{\prime}<\cdots<X_{\sigma^{\prime}\left(S_{\rho^{\prime}}\right)}^{\prime}
$$

Define $\sigma^{\prime \prime}$ similarly for $\rho^{\prime \prime}$. Again by symmetry:
Observation 4. Conditioned on $\sigma^{-1}(1)$ (and therefore on $S_{\rho^{\prime}}$ and $S_{\rho^{\prime \prime}}$ ), the permutations $\sigma^{\prime}$ and $\sigma^{\prime \prime}$ are independent and uniformly distributed.

Finally, recursively:
Observation 5. The binary search tree $t[\sigma]$ is obtained by appending the left subtree $t\left[\sigma^{\prime}\right]$ and right subtree $t\left[\sigma^{\prime \prime}\right]$ to the root $\rho$.

If $S_{\rho^{\prime}}=0$, then $t\left[\sigma^{\prime}\right]=\emptyset$ (and there is in fact no $\rho^{\prime}$ ); while, if $S_{\rho^{\prime}}=1$, the tree $t\left[\sigma^{\prime}\right]$ is comprised of the single vertex $\rho^{\prime}$. Similarly for $\sigma^{\prime \prime}$. Hence this recursive process stops whenever we reach a vertex $v$ with $S_{v} \in\{0,1\}$. But it will be convenient to extend it indefinitely to produce an infinite binary tree $\mathcal{T}=\widehat{\mathbb{T}}_{2}$, where all additional vertices $v$ are assigned $S_{v}=0$.

The upshot of all these observations is that we obtain the following alternative characterization of the height $H_{n}$ :

- assign an independent $U[0,1]$ (i.e., uniform in $[0,1]$ ) random variable $W_{v}$ to each vertex $v$ in the infinite binary tree $\mathcal{T}$;
- at the root $\rho$, set

$$
S_{\rho}=n ;
$$

- then recursively from the root down, set

$$
\begin{equation*}
S_{v^{\prime}}:=\left\lfloor S_{v} W_{v}\right\rfloor \quad \text { and } \quad S_{v^{\prime \prime}}:=\left\lfloor S_{v}\left(1-W_{v}\right)\right\rfloor, \tag{6.3.2}
\end{equation*}
$$

where $v^{\prime}$ and $v^{\prime \prime}$ are the left and right descendants of $v$ in $\mathcal{T}$.
It can be checked that $S_{v^{\prime}}+S_{v^{\prime \prime}}=S_{v}-1$ almost surely, provided $S_{v} \geq 1$ (see Exercise 6.15). Moreover notice that, when $S_{v}=1$, then $S_{v^{\prime}}=S_{v^{\prime \prime}}=0$ almost surely; while, if $S_{v}=0$, then $S_{v^{\prime}}=S_{v^{\prime \prime}}=0$. Finally the height $H_{n}$ is the highest level containing a vertex with subtree size at least 1 , that is,

$$
\begin{equation*}
H_{n}=\sup \left\{h: \exists v \in \mathcal{L}_{h}, S_{v} \geq 1\right\}, \tag{6.3.3}
\end{equation*}
$$

where $\mathcal{L}_{h}$ is the set of vertices of $\mathcal{T}$ at graph distance $h$ from the root.

## Key technical bound

Because $W \sim U[0,1]$ implies also that $(1-W) \sim U[0,1]$, we immediately get from (6.3.2) that:

Lemma 6.3.2 (Distribution of subtree size). Let v be a vertex at topological distance $\ell$ from the root of $\mathcal{T}$. Let $U_{1}, \ldots, U_{\ell}$ be i.i.d. $U[0,1]$. Then we have the equality in distribution

$$
S_{v} \stackrel{\mathrm{~d}}{=}\left\lfloor\cdots\left\lfloor\left\lfloor n U_{1}\right\rfloor U_{2}\right\rfloor \cdots U_{\ell}\right\rfloor .
$$

From Lemma 6.3.2 and the characterization of the height in (6.3.3), we need to control how fast products of independent uniforms decrease. But that is only half of the story: the number of paths of length $\ell$ from the root grows exponentially with $\ell$. The following lemma, which takes both effects into account, will play a key role in the analysis. It also explains the definition of $\gamma$ in (6.3.1). Note that we ignore-for the time being-the repeated rounding in Lemma 6.3.2; it will turn out to have a minor effect.

Lemma 6.3.3 (Product of uniforms). Let $U_{1}, U_{2}, \ldots$ be i.i.d. $U[0,1]$. Then

$$
\lim _{\ell \rightarrow+\infty} 2^{\ell} \mathbb{P}\left[U_{1} \cdots U_{\ell} \geq e^{-\ell / c}\right]= \begin{cases}+\infty & \text { if } c<\gamma \\ 0 & \text { if } c>\gamma\end{cases}
$$

Proof. Taking logarithms turns the product on the left-hand side into a sum of i.i.d. random variables

$$
\begin{equation*}
2^{\ell} \mathbb{P}\left[U_{1} \cdots U_{\ell} \geq e^{-\ell / c}\right]=2^{\ell} \mathbb{P}\left[\sum_{i=1}^{\ell}\left(-\log U_{i}\right) \leq \ell / c\right] \tag{6.3.4}
\end{equation*}
$$

Now it is elementary to bound the right-hand side.
Lemma 6.3.4 (A tail bound). Let $U_{1}, \ldots, U_{\ell}$ be i.i.d. $U[0,1]$. Then for any $y>0$

$$
\begin{equation*}
\frac{y^{\ell} e^{-y}}{\ell!} \leq \mathbb{P}\left[\sum_{i=1}^{\ell}\left(-\log U_{i}\right) \leq y\right] \leq \frac{y^{\ell} e^{-y}}{\ell!}\left(\frac{1}{1-\frac{y}{\ell+1}}\right) . \tag{6.3.5}
\end{equation*}
$$

Proof. We prove a more general claim, specifically

$$
\mathbb{P}\left[\sum_{i=1}^{\ell}\left(-\log U_{i}\right) \leq y\right]=e^{-y}\left\{\sum_{i=\ell}^{+\infty} \frac{y^{i}}{i!}\right\},
$$

from which (6.3.5) follows: the lower bound is obtained by keeping only the first term in the sum; the upper bound is obtained by factoring out $y^{\ell} e^{-y} / \ell!$ and relating the remaining sum to a geometric series.

So it remains to prove the general claim. First note that $-\log U_{1}$ is exponentially distributed. Indeed, for any $y \geq 0$,

$$
\mathbb{P}\left[-\log U_{1}>y\right]=\mathbb{P}\left[U_{1}<e^{-y}\right]=e^{-y} .
$$

So

$$
\mathbb{P}\left[-\log U_{1} \leq y\right]=1-e^{-y}=e^{-y}\left\{\sum_{i=0}^{+\infty} \frac{y^{i}}{i!}-1\right\}=e^{-y}\left\{\sum_{i=1}^{+\infty} \frac{y^{i}}{i!}\right\},
$$

as claimed in the base case $\ell=1$.
Proceeding by induction, suppose the claim holds up to $\ell-1$. Then

$$
\begin{aligned}
\mathbb{P}\left[\sum_{i=1}^{\ell}\left(-\log U_{i}\right)>y\right] & =\int_{0}^{+\infty} e^{-z} \mathbb{P}\left[\sum_{i=1}^{\ell-1}\left(-\log U_{i}\right)>y-z\right] \mathrm{d} z \\
& =e^{-y}+\int_{0}^{y} e^{-z} \mathbb{P}\left[\sum_{i=1}^{\ell-1}\left(-\log U_{i}\right)>y-z\right] \mathrm{d} z \\
& =e^{-y}+\int_{0}^{y} e^{-z} e^{-(y-z)}\left\{\sum_{i=0}^{\ell-2} \frac{(y-z)^{i}}{i!}\right\} \mathrm{d} z \\
& =e^{-y}+e^{-y} \sum_{i=0}^{\ell-2} \frac{y^{i+1}}{i!(i+1)} \\
& =e^{-y} \sum_{j=0}^{\ell-1} \frac{y^{j}}{j!} .
\end{aligned}
$$

That proves the claim.
We return to the proof of Lemma 6.3.3. Plugging (6.3.5) into (6.3.4), we get

$$
\begin{align*}
\frac{2^{\ell}(\ell / c)^{\ell} e^{-(\ell / c)}}{\ell!} \leq 2^{\ell} \mathbb{P}[ & \left.U_{1} \cdots U_{\ell} \geq e^{-\ell / c}\right] \\
& \leq \frac{2^{\ell}(\ell / c)^{\ell} e^{-(\ell / c)}}{\ell!}\left(\frac{1}{1-\frac{(\ell / c)}{\ell+1}}\right) \tag{6.3.6}
\end{align*}
$$

As $\ell \rightarrow+\infty$,

$$
\begin{equation*}
\frac{1}{1-\frac{(\ell / c)}{\ell+1}} \rightarrow \frac{1}{1-\frac{1}{c}}, \tag{6.3.7}
\end{equation*}
$$

which is positive when $c>1$. We will use the standard bound (see Exercise 1.3 for a proof)

$$
\frac{\ell^{\ell}}{e^{\ell-1}} \leq \ell!\leq \frac{\ell^{\ell+1}}{e^{\ell-1}}
$$

It implies immediately that

$$
\frac{2^{\ell}(\ell / c)^{\ell} e^{-(\ell / c)} e^{\ell-1}}{\ell^{\ell+1}} \leq \frac{2^{\ell}(\ell / c)^{\ell} e^{-(\ell / c)}}{\ell!} \leq \frac{2^{\ell}(\ell / c)^{\ell} e^{-(\ell / c)} e^{\ell-1}}{\ell^{\ell}}
$$

which after simplifying gives

$$
\begin{equation*}
(e \ell)^{-1}\left[\left(\frac{1}{e}\right)\left(\frac{2 e}{c}\right)^{c}\right]^{\ell / c} \leq 2^{\ell} \frac{(\ell / c)^{\ell} e^{-(\ell / c)}}{\ell!} \leq e^{-1}\left[\left(\frac{1}{e}\right)\left(\frac{2 e}{c}\right)^{c}\right]^{\ell / c} . \tag{6.3.8}
\end{equation*}
$$

By (6.3.1) and the remark following it, the expression in square brackets is $>1$ or $<1$ depending on whether $c<\gamma$ or $c>\gamma$. Combining (6.3.6), (6.3.7) and (6.3.8) and taking a limit as $\ell \rightarrow+\infty$ gives the claim.

As an immediate consequence of Lemma 6.3.3, we bound the height from above. Fix any $\varepsilon>0$ and let $h:=(\gamma+\varepsilon) \log n$. We use a union bound as follows

$$
\begin{equation*}
\mathbb{P}\left[H_{n} \geq h\right]=\mathbb{P}\left[\bigcup_{v \in \mathcal{L}_{h}}\left\{S_{v} \geq 1\right\}\right] \leq \sum_{v \in \mathcal{L}_{h}} \mathbb{P}\left[S_{v} \geq 1\right]=2^{h} \mathbb{P}\left[S_{v} \geq 1\right] \tag{6.3.9}
\end{equation*}
$$

for any $v \in \mathcal{L}_{h}$, where the first equality follows from (6.3.3). Since

$$
\left\lfloor\cdots\left\lfloor\left\lfloor n U_{1}\right\rfloor U_{2}\right\rfloor \cdots U_{h}\right\rfloor \leq n U_{1} U_{2} \cdots U_{h},
$$

Lemmas 6.3.2 and 6.3.3 imply that

$$
\begin{align*}
2^{h} \mathbb{P}\left[S_{v} \geq 1\right] & \leq 2^{h} \mathbb{P}\left[n U_{1} U_{2} \cdots U_{h} \geq 1\right] \\
& =2^{h} \mathbb{P}\left[U_{1} U_{2} \cdots U_{h} \geq e^{-h /(\gamma+\varepsilon)}\right] \\
& \rightarrow 0 \tag{6.3.10}
\end{align*}
$$

as $h \rightarrow+\infty$. From (6.3.9) and (6.3.10), we obtain finally that for any $\varepsilon>0$

$$
\mathbb{P}\left[H_{n} / \log n \geq \gamma+\varepsilon\right] \rightarrow 0,
$$

as $n \rightarrow+\infty$, which establishes one direction of Claim 6.3.1.

## Lower bounding the height: a branching process

Establishing the other direction is where branching processes enter the scene. We will need some additional notation. Fix $c<\gamma$ and let $\ell$ be a positive integer that will be set later on. For any pair of vertex $v, w \in \mathcal{T}$ with $w$ a descendant of $v$, let $\mathcal{Q}[v, w]$ be the set of vertices on the path between $v$ and $w$, including $v$ but excluding $w$. Further, recalling (6.3.2), define

$$
\mathcal{U}[v, w]=\prod_{z \in \mathcal{Q}[u, v]} U_{z}^{v, w},
$$

where $U_{z}^{v, w}=W_{z}$ (respectively $1-W_{z}$ ) if the path from $v$ to $w$ takes the left (respectively right) edge upon exiting $z$. Denote by $\mathcal{L}_{\ell}[v]$ the set of descendant vertices of $v$ in $\mathcal{T}$ at graph distance $\ell$ from $v$ and consider the random subset

$$
\mathcal{L}_{\ell}^{*}[v]=\left\{w \in \mathcal{L}_{\ell}[v]: \mathcal{U}[v, w] \geq e^{-\ell / c}\right\} .
$$

Fix a vertex $u \in \mathcal{T}$. We define the following Galton-Watson branching process.

- Initialize $Z_{0}^{u, \ell}:=1$ and $u_{0,1}:=u$.
- For $t \geq 1$, set

$$
Z_{t}^{u, \ell}=\sum_{r=1}^{Z_{t-1}^{u, \ell}}\left|\mathcal{L}_{\ell}^{*}\left[u_{t-1, r}\right]\right|,
$$

and let $u_{t, 1}, \ldots, u_{t, Z_{t}^{u, \ell}}$ be the vertices in $\cup_{r=1}^{Z_{t-1}^{u, \ell}} \mathcal{L}_{\ell}^{*}\left[u_{t-1, r}\right]$ from left to right.
In words, $Z_{1}^{u, \ell}$ counts the number of vertices $\ell$ levels below $u$ whose subtree sizes (ignoring rounding) have not decreased "too much" compared to that of $u$ (in the sense of Lemma 6.3.3). We let such vertices (if any) be $u_{1,1}, \ldots, u_{1, Z_{1}^{u, \ell}}$. Similarly, $Z_{2}^{u, \ell}$ counts the same quantity over all vertices $\ell$ levels below the vertices $u_{1,1}, \ldots, u_{1, Z_{1}^{u, \ell}}$, and so forth.

Because the $W_{v}$ s are i.i.d., this process is indeed a Galton-Watson branching process. The expectation of the offspring distribution (which by symmetry does not depend on the choice of $u$ ) is

$$
m=\mathbb{E}\left[Z_{1}^{u, \ell}\right]=2^{\ell} \mathbb{P}\left[U_{1} \cdots U_{\ell} \geq e^{-\ell / c}\right]
$$

where we used the notation of Lemma 6.3.3. By that lemma, we can choose $\ell$ large enough that $m>1$. Fix such an $\ell$ for the rest of the proof. In that case, by Theorem 6.1.6, the process survives with probability $1-\eta$ for some $0 \leq \eta<1$.

The relevance of this observation can be seen from taking $u=\rho$.

Claim 6.3.5. Let $c^{\prime}<c$. Conditioned on survival of $\left(Z_{t}^{\rho, \ell}\right)$, for $n$ large enough $H_{n} \geq c^{\prime} \log n-\theta_{n} \ell$ almost surely for some $\theta_{n} \in[0,1)$.

Proof. To account for the rounding, we will need the inequality

$$
\begin{equation*}
\left\lfloor\cdots\left\lfloor\left\lfloor n U_{1}\right\rfloor U_{2}\right\rfloor \cdots U_{s}\right\rfloor \geq n U_{1} U_{2} \cdots U_{s}-s, \tag{6.3.11}
\end{equation*}
$$

which holds for all $n, s \geq 1$, as can be checked by induction. Write $s=k \ell$ for some positive integer $k$ to be determined. Conditioned on survival of $\left(Z_{t}^{\rho, \ell}\right)$, the population at generation $k$ satisfies

$$
Z_{k}^{\rho, \ell} \geq 1
$$

which implies that, for some $v^{*} \in \mathcal{L}_{s}[\rho]$, it holds that

$$
n \mathcal{U}\left[\rho, v^{*}\right] \geq n\left(e^{-\ell / c}\right)^{k} .
$$

Now take $s=c^{\prime} \log n-\theta_{n} \ell$ with $c^{\prime}<c$ and $\theta_{n} \in[0,1)$ such that $s$ is a multiple of $\ell$. Then

$$
\begin{align*}
& n\left(e^{-\ell / c}\right)^{k}=n\left(e^{-s / c}\right)=n\left(n^{-c^{\prime} / c} e^{-\theta_{n} \ell / c}\right)=n^{1-c^{\prime} / c} e^{-\theta_{n} \ell / c} \\
& \geq c^{\prime} \log n-\theta_{n} \ell+1=s+1, \tag{6.3.12}
\end{align*}
$$

for all $n$ large enough, where we used that $1-c^{\prime} / c>0, \theta_{n} \in[0,1)$ and $\ell$ is fixed. So, using the characterization of the height in (6.3.2) and (6.3.3) together with inequality (6.3.11), we derive

$$
\begin{equation*}
S_{v^{*}} \geq n \mathcal{U}\left[\rho, v^{*}\right]-s \geq n\left(e^{-\ell / c}\right)^{k}-s \geq 1 \tag{6.3.13}
\end{equation*}
$$

That is, $H_{n} \geq c^{\prime} \log n-\theta_{n} \ell$.
But this is not quite what we want: this last claim holds only conditioned on survival; or put differently, it holds with probability $1-\eta$, a value which could be significantly smaller than 1 in general. To handle this last issue, we consider a large number of independent copies of the Galton-Watson process above in order to "boost" the probability that at least one of them survives to a value arbitrarily close to 1 .

Claim 6.3.6. For any $\delta>0$, there is a $J$ so that $H_{n} \geq c^{\prime} \log n-\theta_{n} \ell+J \ell$ with probability at least $1-\delta$ for all $n$ large enough.

Proof. Let $J \ell$ be a multiple of $\ell$ and let

$$
u_{1}^{*}, \ldots, u_{2 J \ell}^{*}
$$

be the vertices on level $J \ell$ from left to right. Each process

$$
\left(Z_{t}^{u_{i}^{*}, \ell}\right)_{t \geq 0}, \quad i=1, \ldots, 2^{J \ell}
$$

is an independent copy of $\left(Z_{t}^{\rho, \ell}\right)_{t \geq 0}$.
We define two "bad events":

- (No survival) Let $\mathcal{B}_{1}$ be the event that all $\left(Z_{t}^{u_{i}^{*}, \ell}\right)$ s go extinct and choose $J$ large enough that this event has probability $<\delta / 2$, that is,

$$
\mathbb{P}\left[\mathcal{B}_{1}\right]=\eta^{2^{J \ell}}<\delta / 2 .
$$

Under $\mathcal{B}_{1}^{c}$, at least one of the branching processes survives; let $I$ be the lowest index among them.

- (Fast decay at the top) To bound the height, we also need to control the effect of the first $J \ell$ levels on the subtree sizes. Let $\mathcal{B}_{2}$ be the event that at least one of the $W$-values associated with the $2^{J \ell}-1$ vertices ancestral to the $u_{i}^{*} \mathrm{~s}$ is outside the interval $(\alpha, 1-\alpha)$. Choose $\alpha$ small enough that this event has probability $<\delta / 2$, that is,

$$
\mathbb{P}\left[\mathcal{B}_{2}\right] \leq(2 \alpha)\left(2^{J \ell}-1\right)<\delta / 2 .
$$

Under $\mathcal{B}_{2}^{c}$, we have almost surely the lower bound

$$
\begin{equation*}
\mathcal{U}\left[\rho, u_{I}^{*}\right] \geq \alpha^{J \ell} \tag{6.3.14}
\end{equation*}
$$

since it in fact holds for all $u_{i}^{*}$ s simultaneously.
We are now ready to conclude. Assume $\mathcal{B}_{1}^{c}$ and $\mathcal{B}_{2}^{c}$ hold. Taking

$$
s=k \ell=c^{\prime} \log n-\theta_{n} \ell,
$$

as before, we have $Z_{k}^{u_{I}^{*}, \ell} \geq 1$ so there is $v^{*} \in \mathcal{L}_{s}\left[u_{I}^{*}\right]$ such that

$$
n \mathcal{U}\left[\rho, v^{*}\right]=n \mathcal{U}\left[\rho, u_{I}^{*}\right] \mathcal{U}\left[u_{I}^{*}, v^{*}\right] \geq n \alpha^{J \ell}\left(e^{-\ell / c}\right)^{k},
$$

where we used (6.3.14). Observe that (6.3.12) remains valid (for potentially larger $n$ ) even after multiplying all expressions on the left-hand side of the inequality by $\alpha^{J \ell}$. Arguing as in (6.3.13), we get that $H_{n} \geq c^{\prime} \log n-\theta_{n} \ell+J \ell$. This event holds with probability at least

$$
\mathbb{P}\left[\left(\mathcal{B}_{1} \cup \mathcal{B}_{2}\right)^{c}\right]=1-\mathbb{P}\left[\mathcal{B}_{1}\right]-\mathbb{P}\left[\mathcal{B}_{2}\right] \geq 1-\delta .
$$

We have proved the claim.

For any $\varepsilon>0$, we can choose $c^{\prime}=\gamma-\varepsilon$ and $c^{\prime}<c<\gamma$. Further, $\delta$ can be made arbitrarily small (provided $n$ is large enough). Put differently, we have proved that for any $\varepsilon>0$

$$
\mathbb{P}\left[H_{n} / \log n \geq \gamma-\varepsilon\right] \rightarrow 1,
$$

as $n \rightarrow+\infty$, which establishes the other direction of Claim 6.3.1.

### 6.3.2 $\triangleright$ Data science: the reconstruction problem, the Kesten-Stigum bound and a phase transition in phylogenetics

In this section, we explore an application of multitype branching processes in statistical phylogenetics, the reconstruction of evolutionary trees from molecular data. Informally, we consider a ferromagnetic Ising model (Example 1.2.5) on an infinite binary tree and we ask: when do the states at level $h$ "remember" the state at the root? We establish the existence of a phase transition. Before defining the problem formally and explaining its connection to evolutionary biology, we describe an equivalent definition of the model. This alternative "Markov chain on a tree" perspective will make it easier to derive recursions for quantities of interest. Equivalence between the two models is proved in Exercise 6.16.

## The reconstruction problem

Consider a rooted infinite binary tree $\mathcal{T}=\widehat{\mathbb{T}}_{2}$, where the root is denoted by 0 . Fix a parameter $0<p<1 / 2$, which we will refer to as the mutation probability for reasons that will be explained below. We assign a state $\sigma_{v}$ in $\mathcal{C}=\{+1,-1\}$ to each vertex $v$ as follows. At the root 0 , the state $\sigma_{0}$ is picked uniformly at random in $\{+1,-1\}$. Moving away from the root, the state $\sigma_{v}$ at a vertex $v$, conditioned on the state at its immediate ancestor $u$, is equal to $\sigma_{u}$ with probability $1-p$ and to $-\sigma_{u}$ with probability $p$. In the computational biology literature, this model is referred to as the Cavender-Farris-Neyman (CFN) model.

For $h \geq 0$, let $\mathcal{L}_{h}$ be the set of vertices in $\mathcal{T}$ at graph distance $h$ from the root. We denote by $\sigma_{h}=\left(\sigma_{\ell}\right)_{\ell \in \mathcal{L}_{h}}$ the vector of states at level $h$ and we denote by $\mu_{h}$ the distribution of $\boldsymbol{\sigma}_{h}$. The reconstruction problem consists in trying to "guess" the state at the root $\sigma_{0}$ given the states $\boldsymbol{\sigma}_{h}$ at level $h$. We first note that in general we cannot expect an arbitrarily good estimator. Indeed, rewriting the Markov transition matrix along the edges (i.e., the matrix encoding the probability of the state at a vertex given the state at its immediate ancestor) in its random cluster form

$$
P:=\left(\begin{array}{cc}
1-p & p  \tag{6.3.15}\\
p & 1-p
\end{array}\right)=(1-2 p)\left(\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right)+(2 p)\left(\begin{array}{ll}
1 / 2 & 1 / 2 \\
1 / 2 & 1 / 2
\end{array}\right),
$$

CFN model
reconstruction
problem
we see that the states $\sigma_{1}$ at the first level are completely randomized (i.e., independent of $\sigma_{0}$ ) with probability $(2 p)^{2}$-in which case we cannot hope to reconstruct the root state better than a coin flip. Intuitively the reconstruction problem is solvable if we can find an estimator of the root state which outperforms a random coin flip as $h$ grows to $+\infty$. Let $\mu_{h}^{+}$be the distribution $\mu_{h}$ conditioned on the root state $\sigma_{0}$ being +1 , and similarly for $\mu_{h}^{-}$. Observe that $\mu_{h}=\frac{1}{2} \mu_{h}^{+}+\frac{1}{2} \mu_{h}^{-}$. Recall also that

$$
\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}=\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}}\left|\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right|
$$

Definition 6.3.7 (Reconstruction solvability). We say that the reconstruction problem for $0<p<1 / 2$ is solvable if

$$
\lim \inf _{h \rightarrow+\infty}\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}>0
$$

otherwise the problem is unsolvable.
(Exercise 6.17 asks for a proof that $\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}$ is monotone in $h$ and therefore has a limit.)

To see the connection with the description above, consider an arbitrary root estimator $\hat{\sigma}_{0}\left(\mathbf{s}_{h}\right)$. Then the probability of a mistake is

$$
\begin{aligned}
\mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right]= & \frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}^{-}\left(\mathbf{s}_{h}\right) \mathbf{1}\left\{\hat{\sigma}_{0}\left(\mathbf{s}_{h}\right)=+1\right\} \\
& +\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}^{+}\left(\mathbf{s}_{h}\right) \mathbf{1}\left\{\hat{\sigma}_{0}\left(\mathbf{s}_{h}\right)=-1\right\}
\end{aligned}
$$

This expression is minimized by choosing for each $s_{h}$ separately

$$
\hat{\sigma}_{0}\left(\mathbf{s}_{h}\right)= \begin{cases}+1 & \text { if } \mu_{h}^{+}\left(\mathbf{s}_{h}\right) \geq \mu_{h}^{-}\left(\mathbf{s}_{h}\right) \\ -1 & \text { otherwise }\end{cases}
$$

Let $\mu_{h}\left(s_{0} \mid \mathbf{s}_{h}\right)$ be the posterior probability of the root state, that is, the conditional probability of the root state $s_{0}$ given the states $\mathbf{s}_{h}$ at level $h$. By Bayes' rule,

$$
\mu_{h}\left(+1 \mid \mathbf{s}_{h}\right)=\frac{(1 / 2) \mu_{h}^{+}\left(\mathbf{s}_{h}\right)}{\mu_{h}\left(\mathbf{s}_{h}\right)}
$$

and similarly for $\mu_{h}\left(+1 \mid \mathbf{s}_{h}\right)$. Hence the choice above is equivalent to

$$
\hat{\sigma}_{0}\left(\mathbf{s}_{h}\right)= \begin{cases}+1 & \text { if } \mu_{h}\left(+1 \mid \mathbf{s}_{h}\right) \geq \mu_{h}\left(-\mid \mathbf{s}_{h}\right) \\ -1 & \text { otherwise }\end{cases}
$$

which is known as the maximum a posteriori (MAP) estimator. (We encountered it in a different context in Section 5.1.4.) For short, we will denote it by $\hat{\sigma}_{0}^{\text {MAP }}$.

Now note that

$$
\begin{aligned}
& \mathbb{P}\left[\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\boldsymbol{\sigma}_{h}\right)=\sigma_{0}\right]-\mathbb{P}\left[\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right] \\
& =\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}^{+}\left(\mathbf{s}_{h}\right)\left[\mathbf{1}\left\{\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right)=+1\right\}-\mathbf{1}\left\{\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right)=-1\right\}\right] \\
& \quad+\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{2}}} \mu_{h}^{-}\left(\mathbf{s}_{h}\right)\left[\mathbf{1}\left\{\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right)=-1\right\}-\mathbf{1}\left\{\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right)=+1\right\}\right] \\
& =\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}^{+}\left(\mathbf{s}_{h}\right) \hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right) \\
& \quad-\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{2}}} \mu_{h}^{-}\left(\mathbf{s}_{h}\right) \hat{\sigma}_{0}^{\mathrm{MAP}}\left(\mathbf{s}_{h}\right) \\
& =\frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}}\left|\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right| \\
& =\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}},
\end{aligned}
$$

where the third equality comes from

$$
|a-b|=(a-b) \mathbf{1}\{a \geq b\}+(b-a) \mathbf{1}\{a<b\} .
$$

Since $\mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right)=\sigma_{0}\right]+\mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right]=1$, the display above can be rewritten as

$$
\mathbb{P}\left[\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right]=\frac{1}{2}-\frac{1}{2}\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}
$$

Given that $\hat{\sigma}_{0}^{\text {MAP }}$ was chosen to minimize the error probability, we also have that for any root estimator $\hat{\sigma}_{0}$

$$
\mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right] \geq \frac{1}{2}-\frac{1}{2}\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}
$$

Since this last inequality also applies to the estimator $-\hat{\sigma}_{0}$, we have also that

$$
\mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right] \leq \frac{1}{2}
$$

The next lemma summarizes the discussion above.
Lemma 6.3.8 (Probability of erroneous reconstruction). The probability of an erroneous root reconstruction behaves as follows.
(i) If the reconstruction problem is solvable, then

$$
\lim _{h \rightarrow+\infty} \mathbb{P}\left[\hat{\sigma}_{0}^{\mathrm{MAP}}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right]<\frac{1}{2}
$$

(ii) If the reconstruction problem is unsolvable, then for any root estimator $\hat{\sigma}_{0}$

$$
\lim _{h \rightarrow+\infty} \mathbb{P}\left[\hat{\sigma}_{0}\left(\boldsymbol{\sigma}_{h}\right) \neq \sigma_{0}\right]=\frac{1}{2} .
$$

It turns out that the accuracy of the MAP estimator undergoes a phase transition at a critical mutation probability $p_{*}$. Our main theorem is the following.

Theorem 6.3.9 (Solvability). Let $\theta_{*}$ be the unique positive solution to

$$
2 \theta_{*}^{2}=1,
$$

and set $p^{*}=\frac{1-\theta^{*}}{2}$. Then the reconstruction problem is:
(i) solvable if $0<p<p_{*}$;
(ii) unsolvable if $p_{*} \leq p<1 / 2$.

We will prove this theorem in the rest of the section.
But first, what does all of this have to do with evolutionary biology? Truncate $\mathcal{T}$ at level $h$ to obtain a finite tree $\mathcal{T}_{h}$ with leaf set $\mathcal{L}_{h}$. In phylogenetics, one uses such a tree to depict evolutionary relationships between extant species that are represented by its leaves. Each internal branching corresponds to a past speciation event. Extinctions have been pruned from the tree. The genomes of ancestral species, starting from the most recent common ancestor at the root, are posited to have evolved along the (deterministic) tree $\mathcal{T}_{h}$ according to a random process of single-site substitutions. To simplify, each position in the genome is assumed to take one of two values, +1 or -1 , and it evolves independently from all other positions under a CFN model on $\mathcal{T}_{h}$. That is, on each edge of the tree a mutation occurs with probability $p$, changing the state of the immediate descendant species at that position. This is of course only a toy model, but it is not far from what evolutionary biologists actually use in practice with great success. One practical problem of interest is to reconstruct the genome of ancestors given access to contemporary genomes. This is, in a nutshell, the reconstruction problem.

## Kesten-Stigum bound

The condition in Theorem 6.3.9 is referred to as the Kesten-Stigum bound. We explain why next. We showed in Lemma 6.3.8 that the MAP estimator has an error probability bounded away from $1 / 2$ if and only if the reconstruction problem is solvable. Of course, other estimators may also achieve that same desirable outcome. In fact, from the lemma, to establish reconstruction solvability it suffices to exhibit one such "better-than-random" estimator. So, rather than analyzing $\hat{\sigma}_{0}^{\mathrm{MAP}}$, we look at a simpler estimator first and prove half of Theorem 6.3.9. The other half will be proven below using different ideas.

The key is to notice that a multitype branching process (see Section 6.1.4) hides in the background. For $h \geq 0$, consider the random row vector $\mathbf{Z}_{h}=\left(Z_{h,+}, Z_{h,-}\right)$ where the first component records the number of +1 states (which we refer to as belonging to the + type) in $\sigma_{h}$ and, likewise, the second component counts the -1 states (referred to as of - type). Then $\left(\mathbf{Z}_{h}\right)_{h \geq 0}$ is a two-type Galton-Watson process where each individual has exactly two children. Their types depend on the type of the parent. A type + individual has the following offspring distribution:

$$
p_{\mathbf{k}}^{(+)}= \begin{cases}(1-p)^{2} & \text { if } \mathbf{k}=(2,0) \\ 2 p(1-p) & \text { if } \mathbf{k}=(1,1) \\ p^{2} & \text { if } \mathbf{k}=(0,2) \\ 0 & \text { otherwise }\end{cases}
$$

Similar expressions hold for $p_{\mathbf{k}}^{(-)}$. The mean matrix is given by

$$
\begin{aligned}
M & =\left(\begin{array}{cc}
2(1-p)^{2}+2 p(1-p) & 2 p(1-p)+2 p^{2} \\
2 p(1-p)+2 p^{2} & 2(1-p)^{2}+2 p(1-p)
\end{array}\right) \\
& =2\left(\begin{array}{cc}
(1-p)(1-p+p) & p(1-p+p) \\
p(1-p+p) & (1-p)(1-p+p)
\end{array}\right) \\
& =2\left(\begin{array}{cc}
1-p & p \\
p & 1-p
\end{array}\right) \\
& =2 P
\end{aligned}
$$

where (not coincidentally) we have already encountered the matrix $P$ in (6.3.15). As a symmetric matrix, by the spectral theorem (Theorem 5.1.1), $P$ has a real
eigenvector decomposition

$$
\begin{aligned}
P & =\left(\begin{array}{cc}
1-p & p \\
p & 1-p
\end{array}\right) \\
& =\left(\begin{array}{cc}
1 / 2 & 1 / 2 \\
1 / 2 & 1 / 2
\end{array}\right)+(1-2 p)\left(\begin{array}{cc}
1 / 2 & -1 / 2 \\
-1 / 2 & 1 / 2
\end{array}\right) \\
& =\lambda_{1} \mathbf{x}_{1} \mathbf{x}_{1}^{T}+\lambda_{2} \mathbf{x}_{2} \mathbf{x}_{2}^{T},
\end{aligned}
$$

where the eigenvalues and eigenvectors are

$$
\lambda_{1}=1, \quad \lambda_{2}=1-2 p, \quad \mathbf{x}_{1}=\binom{1 / \sqrt{2}}{1 / \sqrt{2}} \quad \mathbf{x}_{2}=\binom{1 / \sqrt{2}}{-1 / \sqrt{2}} .
$$

The eigenvalues of $M$ are twice those of $P$ while the eigenvectors are the same. In particular, using the notation and convention of the Perron-Frobenius Theorem (Theorem 6.1.17), we have

$$
\rho=2, \quad \mathbf{w}=\binom{1 / 2}{1 / 2} .
$$

These should not come entirely as a surprise. In particular, recall from Theorem 6.1.18 that $\rho$ can be interpreted as an "overall rate of growth" of the population, which here is two since each individual has exactly two children (ignoring the types).

Let $\mathbf{u}=(1,-1)$ be a column vector proportional to the second right eigenvector of $M$. We know from Section 6.1.4 that

$$
U_{h}=\left(2 \lambda_{2}\right)^{-h} \mathbf{Z}_{h} \mathbf{u}=\frac{1}{2^{h} \theta^{h}} \sum_{\ell \in \mathcal{L}_{h}} \sigma_{\ell}, \quad h \geq 0,
$$

is a martingale, where we used the notation

$$
\theta:=\lambda_{2}=1-2 p .
$$

Upon looking more closely, the quantity $U_{h}$ has a natural interpretation: its sign is the majority estimator, that is, $\operatorname{sgn}\left(U_{h}\right)=+1$ if a majority of individuals at level $h$ are of type + (breaking ties in favor of + ), and is -1 otherwise. We indicated previously that we only need to find one estimator with an error probability bounded away from $1 / 2$ to establish reconstruction solvability for a given value of $p$. The majority estimator

$$
\hat{\sigma}_{0}^{\mathrm{Maj}}:=\operatorname{sgn}\left(U_{h}\right),
$$

is an obvious one to try. What is less obvious is that it works-all the way to the threshold. This essentially follows from the results of Section 6.1.4, as we detail next.

We begin with an informal discussion. When can $\hat{\sigma}_{0}^{\text {Maj }}$ be expected to work? We will not in fact bound the error probability of $\hat{\sigma}_{0}^{\text {Maj }}$, but instead analyze directly the properties of $\left(U_{h}\right)$. By our modeling assumptions, $\mathbf{Z}_{0}$ is either $(1,0)$ or $(0,1)$ with equal probability. Hence, by the martingale property, we obtain that

$$
\begin{equation*}
\mathbb{E}\left[U_{h} \mid \mathbf{Z}_{0}\right]=\mathbf{Z}_{0} \mathbf{u}=\sigma_{0} \tag{6.3.16}
\end{equation*}
$$

In words, $U_{h}$ is "centered" around the root state. Intuitively, its second moment therefore captures how informative it is about $\sigma_{0}$. Lemma 6.1.20 exhibits a phase transition for $\mathbb{E}\left[U_{h}^{2} \mid \mathbf{Z}_{0}\right]$. The condition for that lemma to hold is

$$
\left(\operatorname{Var}\left[\mathbf{X}^{(+)}(1,1) \mathbf{u}\right], \operatorname{Var}\left[\mathbf{X}^{(-)}(1,1) \mathbf{u}\right]\right) \neq \mathbf{0}
$$

where $\mathbf{X}^{(+)}(1,1) \sim\left\{p_{\mathbf{k}}^{(+)}\right\}$and $\mathbf{X}^{(-)}(1,1) \sim\left\{p_{\mathbf{k}}^{(-)}\right\}$. This is indeed satisfied. The lemma then states that $\mathbb{E}\left[U_{h}^{2} \mid \mathbf{Z}_{0}\right]$ is uniformly bounded if and only if $\rho<$ $\left(2 \lambda_{2}\right)^{2}$, or after rearranging

$$
\begin{equation*}
2 \theta^{2}>1 \tag{6.3.17}
\end{equation*}
$$

Note that this is the condition in Theorem 6.3.9. It arises as a tradeoff between the rate of growth $\rho=2$ and the second largest eigenvalue $\lambda_{2}=\theta$ of the Markov transition matrix $P$. One way to make sense of it is to observe the following:

- On any infinite path out of the root, the process performs a finite Markov chain with transition matrix $P$. We know from Theorem 5.2.14 (see in particular Example 5.2.8) that the chain mixes-and therefore "forgets" its starting state $\sigma_{0}$-at a rate governed by the spectral gap $1-\lambda_{2}$.
- On the other hand, the tree itself is growing at rate $\rho=2$, which produces an exponentially large number of (overlapping) paths out of the root. That growth helps preserve the information about $\sigma_{0}$ down the tree through the duplication of the state (with mutation) at each branching.
- The condition $\rho<\left(2 \lambda_{2}\right)^{2}$ says in essence that when mixing is slow enoughcorresponding to larger values of $\lambda_{2}$-compared to the growth, then the reconstruction problem is solvable. Lemma 6.1 .20 was first proved by Kesten and Stigum, and (6.3.17) is thereby known as the Kesten-Stigum bound.

It remains to turn these observations into a formal proof.
Denote by $\mathbb{E}^{+}$the expectation conditioned on $\sigma_{0}=+1$, and similarly for $\mathbb{E}^{-}$. The following lemma is a consequence of (6.3.16). We give a quick alternative proof.

Lemma 6.3.10 (Unbiasedness of $U_{h}$ ). We have

$$
\mathbb{E}^{+}\left[U_{h}\right]=+1, \quad \mathbb{E}^{-}\left[U_{h}\right]=-1 .
$$

Proof. By applying the Markov transition matrix $P$ on the first level and using the symmetries of the model, for any $\ell \in \mathcal{L}_{h}$ and $\ell^{\prime} \in \mathcal{L}_{h-1}$, we have

$$
\begin{aligned}
\mathbb{E}^{+}\left[\sigma_{\ell}\right] & =(1-p) \mathbb{E}^{+}\left[\sigma_{\ell^{\prime}}\right]+p \mathbb{E}^{-}\left[\sigma_{\ell^{\prime}}\right] \\
& =(1-p) \mathbb{E}^{+}\left[\sigma_{\ell^{\prime}}\right]+p \mathbb{E}^{+}\left[-\sigma_{\ell^{\prime}}\right] \\
& =(1-2 p) \mathbb{E}^{+}\left[\sigma_{\ell^{\prime}}\right] \\
& =\theta \mathbb{E}^{+}\left[\sigma_{\ell^{\prime}}\right] .
\end{aligned}
$$

Iterating, we get $\mathbb{E}^{+}\left[\sigma_{\ell}\right]=\theta^{h}$. The claim follows by linearity of expectation.
Although we do not strictly need it, we also derive an explicit formula for the variance. The proof is typical of how conditional independence properties of this kind of Markov model on trees can be used to derive recursions for quantities of interest.

Lemma 6.3.11 (Variance of $U_{h}$ ). We have

$$
\operatorname{Var}\left[U_{h}\right] \rightarrow \begin{cases}\frac{1 / 2}{1-\left(2 \theta^{2}\right)^{-1}} & \text { if } 2 \theta^{2}>1, \\ +\infty, & \text { otherwise }\end{cases}
$$

Proof. By the conditional variance formula

$$
\begin{align*}
\operatorname{Var}\left[U_{h}\right] & =\operatorname{Var}\left[\mathbb{E}\left[U_{h} \mid \sigma_{0}\right]\right]+\mathbb{E}\left[\operatorname{Var}\left[U_{h} \mid \sigma_{0}\right]\right] \\
& =\operatorname{Var}\left[\sigma_{0}\right]+\mathbb{E}\left[\operatorname{Var}\left[U_{h} \mid \sigma_{0}\right]\right] \\
& =1+\operatorname{Var}^{+}\left[U_{h}\right], \tag{6.3.18}
\end{align*}
$$

where the last line follows from symmetry, with $\mathrm{Var}^{+}$indicating the conditional variance given that the root state $\sigma_{0}$ is +1 . Write $U_{h}=\dot{U}_{h}+\ddot{U}_{h}$ as a sum over the left and right subtrees below the root respectively. Using the conditional independence of those two subtrees given the root state, we get from (6.3.18) that

$$
\begin{align*}
\operatorname{Var}\left[U_{h}\right] & =1+\operatorname{Var}^{+}\left[U_{h}\right] \\
& =1+\operatorname{Var}^{+}\left[\dot{U}_{h}+\ddot{U}_{h}\right] \\
& =1+2 \operatorname{Var}^{+}\left[\dot{U}_{h}\right] \\
& =1+2\left(\mathbb{E}^{+}\left[\dot{U}_{h}^{2}\right]-\mathbb{E}^{+}\left[\dot{U}_{h}\right]^{2}\right) . \tag{6.3.19}
\end{align*}
$$

We now use the Markov transition matrix on the first level to derive a recursion in $h$. Let $\dot{\sigma}_{0}$ be the state at the left child of the root. We use the fact that the random variables $2 \theta \dot{U}_{h}$ conditioned on $\dot{\sigma}_{0}=+1$ and $U_{h-1}$ conditioned on $\sigma_{0}=+1$ are identically distributed. Using $\mathbb{E}^{+}\left[\dot{U}_{h}\right]=1 / 2$ (by Lemma 6.3.10 and symmetry), we get from (6.3.19) that

$$
\begin{align*}
\operatorname{Var}\left[U_{h}\right] & =1-2 \mathbb{E}^{+}\left[\dot{U}_{h}\right]^{2}+2 \mathbb{E}^{+}\left[\dot{U}_{h}^{2}\right] \\
& =1-2(1 / 2)^{2}+2\left[(1-p) \mathbb{E}^{+}\left[(2 \theta)^{-2} U_{h-1}^{2}\right]+p \mathbb{E}^{-}\left[(2 \theta)^{-2} U_{h-1}^{2}\right]\right] \\
& =1 / 2+\left(2 \theta^{2}\right)^{-1} \mathbb{E}^{+}\left[U_{h-1}^{2}\right] \\
& =1 / 2+\left(2 \theta^{2}\right)^{-1} \operatorname{Var}\left[U_{h-1}\right], \tag{6.3.20}
\end{align*}
$$

where we used that

$$
\operatorname{Var}\left[U_{h-1}\right]=\mathbb{E}\left[U_{h-1}^{2}\right]=\mathbb{E}^{+}\left[U_{h-1}^{2}\right]=\mathbb{E}^{-}\left[U_{h-1}^{2}\right]
$$

by symmetry and the fact that $\mathbb{E}\left[U_{h-1}\right]=0$. Solving the affine recursion (6.3.20) gives

$$
\operatorname{Var}\left[U_{h}\right]=\left(2 \theta^{2}\right)^{-h}+(1 / 2) \sum_{i=0}^{h-1}\left(2 \theta^{2}\right)^{-i}
$$

where we used that $\operatorname{Var}\left[U_{0}\right]=\operatorname{Var}\left[\sigma_{0}\right]=1$. The result follows.
We can now prove the first part of Theorem 6.3.9.
Proof of Theorem $6.3 .9(i)$. Let $\bar{\mu}_{h}$ be the distribution of $U_{h}$ and define $\bar{\mu}_{h}^{+}$and $\bar{\mu}_{h}^{-}$ similarly. We give a bound on $\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}$ through a bound on $\left\|\bar{\mu}_{h}^{+}-\bar{\mu}_{h}^{-}\right\|_{\mathrm{TV}}$. Let $\overline{\mathbf{s}}_{h}$ be the $U_{h}$-value associated to $\mathbf{s}_{h}=\left(s_{h, \ell}\right)_{\ell \in \mathcal{L}_{h}} \in\{+1,-1\}^{2^{h}}$, that is,

$$
\overline{\mathbf{s}}_{h}=\frac{1}{2^{h} \theta^{h}} \sum_{\ell \in \mathcal{L}_{h}} s_{h, \ell} .
$$

Then, by marginalizing and the triangle inequality,

$$
\begin{aligned}
\sum_{z}\left|\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)\right| & =\sum_{z}\left|\sum_{\mathbf{s}_{h}: \bar{s}_{h}=z}\left(\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right)\right| \\
& \leq \sum_{z} \sum_{\mathbf{s}_{h}: \overline{\mathbf{s}}_{h}=z}\left|\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right| \\
& =\sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2}}\left|\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right|,
\end{aligned}
$$

where the first sum is over the support of $\bar{\mu}_{h}$. So it suffices to bound from below the left-hand side on the first line.

For that purpose, we apply Cauchy-Schwarz and use the variance bound in Lemma 6.3.11. First note that $\frac{1}{2} \bar{\mu}_{h}^{+}+\frac{1}{2} \bar{\mu}_{h}^{-}=\bar{\mu}_{h}$ so that, by the triangle inequality,

$$
\begin{equation*}
\frac{\left|\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)\right|}{2 \bar{\mu}_{h}(z)} \leq \frac{\bar{\mu}_{h}^{+}(z)+\bar{\mu}_{h}^{-}(z)}{2 \bar{\mu}_{h}(z)}=1 . \tag{6.3.21}
\end{equation*}
$$

Hence, we get

$$
\begin{aligned}
\sum_{z}\left|\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)\right| & =\sum_{z} \frac{\left|\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)\right|}{2 \bar{\mu}_{h}(z)} 2 \bar{\mu}_{h}(z) \\
& \geq 2 \sum_{z}\left(\frac{\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)}{2 \bar{\mu}_{h}(z)}\right)^{2} \bar{\mu}_{h}(z) \\
& \geq 2 \frac{\left(\sum_{z} z\left(\frac{\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)}{2 \bar{\mu}_{h}(z)}\right) \bar{\mu}_{h}(z)\right)^{2}}{\sum_{z} z^{2} \bar{\mu}_{h}(z)} \\
& \geq \frac{1}{2} \frac{\left(\sum_{z} z\left(\bar{\mu}_{h}^{+}(z)-\bar{\mu}_{h}^{-}(z)\right)\right)^{2}}{\sum_{z} z^{2} \bar{\mu}_{h}(z)} \\
& =\frac{1}{2} \frac{\left(\mathbb{E}^{+}\left[U_{h}\right]-\mathbb{E}^{-}\left[U_{h}\right]\right)^{2}}{\operatorname{Var}\left[U_{h}\right]} \\
& \geq 4\left(1-\left(2 \theta^{2}\right)^{-1}\right)>0,
\end{aligned}
$$

where we used (6.3.21) on the second line, Cauchy-Schwarz on the third line (after rearranging), and Lemmas 6.3.10 and 6.3.11 on the last line.

Remark 6.3.12. The proof above and a correlation inequality of [EKPSOO, Theorem 1.4] give a lower bound on the probability of reconstruction of the majority estimator.

## Impossibility of reconstruction

The previous result was based on showing that majority voting, that is, $\hat{\sigma}_{0}^{\text {Maj }}$, produces a good root-state estimator-up to $p=p_{*}$. Here we establish that this result is best possible. Majority is not in fact the best root-state estimator: in general its error probability can be higher than $\hat{\sigma}_{0}^{\mathrm{MAP}}$ as the latter also takes into account the configuration of the states at level $h$. However, perhaps surprisingly, it turns out that the critical threshold for $\hat{\sigma}_{0}^{\mathrm{Maj}}$ coincides with that of $\hat{\sigma}_{0}^{\mathrm{MAP}}$ in the CFN model.

To prove the second part of Theorem 6.3.9 we analyze the MAP estimator. Recall that $\mu_{h}\left(s_{0} \mid \mathbf{s}_{h}\right)$ is the conditional probability of the root state $s_{0}$ given the
states $\mathbf{s}_{h}$ at level $h$. It will be more convenient to work with the following "root magnetization"

$$
R_{h}:=\mu_{h}\left(+1 \mid \boldsymbol{\sigma}_{h}\right)-\mu_{h}\left(-1 \mid \boldsymbol{\sigma}_{h}\right),
$$

which, as a function of $\sigma_{h}$, is a random variable. Note that $\mathbb{E}\left[R_{h}\right]=0$ by symmetry. By Bayes' rule and the fact that $\mu_{h}\left(+1 \mid \boldsymbol{\sigma}_{h}\right)+\mu_{h}\left(-1 \mid \boldsymbol{\sigma}_{h}\right)=1$, we have the following alternative formulas which will prove useful

$$
\begin{align*}
R_{h} & =\frac{1}{2 \mu_{h}\left(\boldsymbol{\sigma}_{h}\right)}\left[\mu_{h}^{+}\left(\boldsymbol{\sigma}_{h}\right)-\mu_{h}^{-}\left(\boldsymbol{\sigma}_{h}\right)\right],  \tag{6.3.22}\\
R_{h} & =2 \mu_{h}\left(+1 \mid \boldsymbol{\sigma}_{h}\right)-1=\frac{\mu_{h}^{+}\left(\boldsymbol{\sigma}_{h}\right)}{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)}-1,  \tag{6.3.23}\\
R_{h} & =1-2 \mu_{h}\left(-1 \mid \boldsymbol{\sigma}_{h}\right)=1-\frac{\mu_{h}^{-}\left(\boldsymbol{\sigma}_{h}\right)}{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)} . \tag{6.3.24}
\end{align*}
$$

It turns out to be enough to prove an upper bound on the variance of $R_{h}$.
Lemma 6.3.13 (Second moment bound). It holds that

$$
\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}} \leq \sqrt{\mathbb{E}\left[R_{h}^{2}\right]} .
$$

Proof. By (6.3.22),

$$
\begin{aligned}
& \frac{1}{2} \sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}}\left|\mu_{h}^{+}\left(\mathbf{s}_{h}\right)-\mu_{h}^{-}\left(\mathbf{s}_{h}\right)\right| \\
& =\sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}\left(\mathbf{s}_{h}\right)\left|\mu_{h}\left(+1 \mid \mathbf{s}_{h}\right)-\mu_{h}\left(-1 \mid \mathbf{s}_{h}\right)\right| \\
& =\mathbb{E}\left|R_{h}\right| \\
& \leq \sqrt{\mathbb{E}\left[R_{h}^{2}\right]}
\end{aligned}
$$

where we used Cauchy-Schwarz on the last line.
Let $\bar{z}_{h}=\mathbb{E}\left[R_{h}^{2}\right]$. In view of Lemma 6.3.13, the proof of Theorem 6.3 .9 (ii) will follow from establishing the limit

$$
\lim _{h \rightarrow+\infty} \bar{z}_{h}=0 .
$$

We apply the same kind of recursive argument we used for the analysis of majority (see in particular Lemma 6.3.11): we condition on the root to exploit conditional independence; we use the Markov transition matrix on the top edges.

We first derive a recursion for $R_{h}$ itself-as a random variable. We proceed in two steps:

- Step 1: we break up the first $h$ levels of the tree into two identical $(h-1)$ level trees with an additional edge at their respective root through conditional independence;
- Step 2: we account for that edge through the Markov transition matrix.

We will need some notation. Let $\dot{\boldsymbol{\sigma}}_{h}$ be the states at level $h$ (from the root) below the left child of the root and let $\dot{\mu}_{h}$ be the distribution of $\dot{\boldsymbol{\sigma}}_{h}$ (and use a superscript + to denote the conditional probability given the root is + , and so on). Define

$$
\dot{Y}_{h}=\dot{\mu}_{h}\left(+1 \mid \dot{\boldsymbol{\sigma}}_{h}\right)-\dot{\mu}_{h}\left(-1 \mid \dot{\boldsymbol{\sigma}}_{h}\right),
$$

where $\dot{\mu}_{h}\left(s_{0} \mid \dot{\mathbf{s}}_{h}\right)$ is the conditional probability that the root is $s_{0}$ given that $\dot{\boldsymbol{\sigma}}_{h}=$ $\dot{\mathbf{s}}_{h}$. Similarly, denote with a double dot the same quantities with respect to the subtree below the right child of the root. Expressions similar to (6.3.22), (6.3.23) and (6.3.24) also hold.

Lemma 6.3.14 (Recursion: Step 1). It holds almost surely that

$$
R_{h}=\frac{\dot{Y}_{h}+\ddot{Y}_{h}}{1+\dot{Y}_{h} \ddot{Y}_{h}} .
$$

Proof. Using $\mu_{h}^{+}\left(\mathbf{s}_{h}\right)=\dot{\mu}_{h}^{+}\left(\dot{\mathbf{s}}_{h}\right) \ddot{\mu}_{h}^{+}\left(\ddot{\mathbf{s}}_{h}\right)$ by conditional independence, (6.3.22) applied to $R_{h}$, and (6.3.23) and (6.3.24) applied to $\dot{Y}_{h}$ and $\ddot{Y}_{h}$, we get

$$
\begin{aligned}
R_{h} & =\frac{1}{2} \sum_{\gamma=+,-} \gamma \frac{\mu_{h}^{\gamma}\left(\boldsymbol{\sigma}_{h}\right)}{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)} \\
& =\frac{1}{2} \frac{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)}{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)} \sum_{\gamma=+,-} \gamma \frac{\dot{\mu}_{h}^{\gamma}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}^{\gamma}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)}{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)} \\
& =\frac{1}{2} \frac{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)}{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)} \sum_{\gamma=+,-} \gamma\left(1+\gamma \dot{Y}_{h}\right)\left(1+\gamma \ddot{Y}_{h}\right) \\
& =\frac{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)}{\mu_{h}\left(\dot{\sigma}_{h}\right)}\left(\dot{Y}_{h}+\ddot{Y}_{h}\right) .
\end{aligned}
$$

The factor in front can be computed as follows

$$
\begin{aligned}
\frac{\mu_{h}\left(\boldsymbol{\sigma}_{h}\right)}{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)} & =\sum_{\gamma=+,-} \frac{1}{2} \frac{\mu_{h}^{\gamma}\left(\boldsymbol{\sigma}_{h}\right)}{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)} \\
& =\sum_{\gamma=+,-} \frac{1}{2} \frac{\dot{\mu}_{h}^{\gamma}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}^{\gamma}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)}{\dot{\mu}_{h}\left(\dot{\boldsymbol{\sigma}}_{h}\right) \ddot{\mu}_{h}\left(\ddot{\boldsymbol{\sigma}}_{h}\right)} \\
& =\frac{1}{2} \sum_{\gamma=+,--}\left(1+\gamma \dot{Y}_{h}\right)\left(1+\gamma \ddot{Y}_{h}\right) \\
& =1+\dot{Y}_{h} \ddot{Y}_{h} .
\end{aligned}
$$

That proves the claim.
For the second step of the recursion, we define

$$
\dot{D}_{h}=\dot{\nu}_{h}\left(+1 \mid \dot{\boldsymbol{\sigma}}_{h}\right)-\dot{\nu}_{h}\left(-1 \mid \dot{\boldsymbol{\sigma}}_{h}\right),
$$

where $\dot{\nu}_{h}\left(\dot{s}_{0} \mid \dot{\mathbf{s}}_{h}\right)$ is the conditional probability that the left child of the root is $\dot{s}_{0}$ given that the states at level $h$ (from the root) below the left child are $\dot{\boldsymbol{\sigma}}_{h}=\dot{\mathbf{s}}_{h}$; and similarly for the right child of the root. Again expressions similar to (6.3.22), (6.3.23) and (6.3.24) hold. The following lemma is left as an exercise (see Exercise 6.18).

Lemma 6.3.15 (Recursion: Step 2). It holds almost surely that

$$
\dot{Y}_{h}=\theta \dot{D}_{h} .
$$

We are now ready to prove the second half of our main theorem.
Proof of Theorem 6.3 .9 (ii). Putting Lemmas 6.3.14 and 6.3.15 together, we get

$$
\begin{equation*}
R_{h}=\frac{\theta\left(\dot{D}_{h}+\ddot{D}_{h}\right)}{1+\theta^{2} \dot{D}_{h} \ddot{D}_{h}} . \tag{6.3.25}
\end{equation*}
$$

We now take expectations. Recall that we seek to compute the second moment of
$R_{h}$. However an important simplification arises from the following observation

$$
\begin{aligned}
\mathbb{E}^{+}\left[R_{h}\right] & =\sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}^{+}\left(\mathbf{s}_{h}\right) R_{h}\left(\mathbf{s}_{h}\right) \\
& =\sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}\left(\mathbf{s}_{h}\right) \frac{\mu_{h}^{+}\left(\mathbf{s}_{h}\right)}{\mu_{h}\left(\mathbf{s}_{h}\right)} R_{h}\left(\mathbf{s}_{h}\right) \\
& =\sum_{\mathbf{s}_{h} \in\{+1,-1\}^{2^{h}}} \mu_{h}\left(\mathbf{s}_{h}\right)\left(1+R_{h}\left(\mathbf{s}_{h}\right)\right) R_{h}\left(\mathbf{s}_{h}\right) \\
& =\mathbb{E}\left[\left(1+R_{h}\right) R_{h}\right] \\
& =\mathbb{E}\left[R_{h}^{2}\right],
\end{aligned}
$$

where we used (6.3.23) on the third line and $\mathbb{E}\left[R_{h}\right]=0$ on the fifth line. So it suffices to compute the conditional first moment.

Using the expansion

$$
\frac{1}{1+r}=1-r+\frac{r^{2}}{1+r},
$$

with $r=\theta^{2} \dot{D}_{h} \ddot{D}_{h}$, we have by (6.3.25) that

$$
\begin{align*}
R_{h} & =\theta\left(\dot{D}_{h}+\ddot{D}_{h}\right)-\theta^{3}\left(\dot{D}_{h}+\ddot{D}_{h}\right) \dot{D}_{h} \ddot{D}_{h}+\theta^{4} \dot{D}_{h}^{2} \ddot{D}_{h}^{2} R_{h} \\
& \leq \theta\left(\dot{D}_{h}+\ddot{D}_{h}\right)-\theta^{3}\left(\dot{D}_{h}+\ddot{D}_{h}\right) \dot{D}_{h} \ddot{D}_{h}+\theta^{4} \dot{D}_{h}^{2} \ddot{D}_{h}^{2}, \tag{6.3.26}
\end{align*}
$$

where we used $\left|R_{h}\right| \leq 1$.
We will need the conditional first and second moments of $\dot{D}_{h}$. For the first moment, note that by symmetry (more precisely, by the fact that $R_{h-1}$ conditioned on $\sigma_{0}=-1$ is equal in distribution to $-R_{h-1}$ conditioned on $\sigma_{0}=+1$ )

$$
\begin{aligned}
\mathbb{E}^{+}\left[\dot{D}_{h}\right] & =(1-p) \mathbb{E}^{+}\left[R_{h-1}\right]+p \mathbb{E}^{-}\left[R_{h-1}\right] \\
& =(1-p) \mathbb{E}^{+}\left[R_{h-1}\right]+p \mathbb{E}^{+}\left[-R_{h-1}\right] \\
& =(1-2 p) \mathbb{E}^{+}\left[R_{h-1}\right] \\
& =\theta \mathbb{E}^{+}\left[R_{h-1}\right] .
\end{aligned}
$$

Similarly, for the second moment, we have

$$
\begin{aligned}
\mathbb{E}^{+}\left[\dot{D}_{h}^{2}\right] & =(1-p) \mathbb{E}^{+}\left[R_{h-1}^{2}\right]+p \mathbb{E}^{-}\left[R_{h-1}^{2}\right] \\
& =\mathbb{E}\left[R_{h-1}^{2}\right] \\
& =\mathbb{E}^{+}\left[R_{h-1}\right]
\end{aligned}
$$

where we used that $\mathbb{E}^{+}\left[R_{h-1}^{2}\right]=\mathbb{E}^{-}\left[R_{h-1}^{2}\right]$ by symmetry so that $\mathbb{E}\left[R_{h-1}^{2}\right]=$ $(1 / 2) \mathbb{E}^{+}\left[R_{h-1}^{2}\right]+(1 / 2) \mathbb{E}^{-}\left[R_{h-1}^{2}\right]=\mathbb{E}^{+}\left[R_{h-1}^{2}\right]$.

Taking expectations in (6.3.26), using conditional independence, and plugging in the formulas for $\mathbb{E}^{+}\left[\dot{D}_{h}\right]$ and $\mathbb{E}^{+}\left[\dot{D}_{h}^{2}\right]$ above, we obtain

$$
\begin{align*}
\bar{z}_{h}= & \mathbb{E}^{+}\left[R_{h}\right] \\
\leq & \theta\left(\mathbb{E}^{+}\left[\dot{D}_{h}\right]+\mathbb{E}^{+}\left[\ddot{D}_{h}\right]\right)-\theta^{3}\left(\mathbb{E}^{+}\left[\dot{D}_{h}^{2}\right] \mathbb{E}^{+}\left[\ddot{D}_{h}\right]+\mathbb{E}^{+}\left[\ddot{D}_{h}^{2}\right] \mathbb{E}^{+}\left[\dot{D}_{h}\right]\right) \\
& \quad+\theta^{4} \mathbb{E}^{+}\left[\dot{D}_{h}^{2}\right] \mathbb{E}^{+}\left[\ddot{D}_{h}^{2}\right] \\
= & 2 \theta^{2} \mathbb{E}^{+}\left[R_{h-1}\right]-2 \theta^{4} \mathbb{E}^{+}\left[R_{h-1}\right]^{2}+\theta^{4} \mathbb{E}^{+}\left[R_{h-1}\right]^{2} \\
= & 2 \theta^{2} \bar{z}_{h-1}-\theta^{4} \bar{z}_{h-1}^{2} . \tag{6.3.27}
\end{align*}
$$

We analyze this recursion next. At $h=0$, we have $\bar{z}_{0}=\mathbb{E}^{+}\left[R_{0}\right]=1$.

- When $2 \theta^{2}<1$, the sequence $\bar{z}_{h}$ decreases to 0 exponentially fast

$$
\bar{z}_{h} \leq\left(2 \theta^{2}\right)^{h}, \quad h \geq 0 .
$$

- When $2 \theta^{2}=1$ on the other hand, convergence to 0 occurs at a slower rate. We show by induction that

$$
\bar{z}_{h} \leq \frac{4}{h}, \quad h \geq 0 .
$$

Note that $\bar{z}_{1} \leq \bar{z}_{0}-\theta^{4} \bar{z}_{0}^{2}=3 / 4 \leq 4$ since $\theta^{4}=1 / 4$, which proves the base case. Assuming the bound holds for $h-1$, we have from (6.3.27) that

$$
\begin{aligned}
\bar{z}_{h} & \leq \bar{z}_{h-1}-\frac{1}{4} \bar{z}_{h-1}^{2} \\
& \leq \frac{4}{h-1}-\frac{4}{(h-1)^{2}} \\
& =4 \frac{h-2}{(h-1)^{2}} \\
& \leq \frac{4}{h},
\end{aligned}
$$

where the last line follows from checking that $h(h-2) \leq(h-1)^{2}$.
Since

$$
\lim _{h \rightarrow+\infty} \bar{z}_{h}=0,
$$

the claim follows from Lemma 6.3.13.
Remark 6.3.16. While Theorem 6.3.9 part (i) can be generalized beyond the CFN model (see, e.g., [MP03]), part (ii) cannot. A striking construction of [Mos01] shows that, under more general models, certain root-state estimators taking into account the configuration of the states at level h can "beat" the Kesten-Stigum bound.

## 6.4 $\triangleright$ Finale: the phase transition of the Erdốs-Rényi model

A compelling way to view an Erdős-Rényi random graph-as its density variesis the following coupling or "evolution." For each pair $\{i, j\}$, let $U_{\{i, j\}}$ be independent uniform random variables in $[0,1]$ and set $\mathcal{G}(p):=([n], \mathcal{E}(p))$ where $\{i, j\} \in \mathcal{E}(p)$ if and only if $U_{\{i, j\}} \leq p$. Then $\mathcal{G}(p)$ is distributed according to $\mathbb{G}_{n, p}$. As $p$ varies from 0 to 1 , we start with an empty graph and progressively add edges until the complete graph is obtained.

We showed in Section 2.3.2 that $\frac{\log n}{n}$ is a threshold function for connectivity. Before connectivity occurs in the evolution of the random graph, a quantity of interest is the size of the largest connected component. As we show in the current section, this quantity itself undergoes a remarkable phase transition: when $p=$ $\frac{\lambda}{n}$ with $\lambda<1$, the largest component has size $\Theta(\log n)$; as $\lambda$ crosses 1 , many components quickly merge to form a so-called "giant component" of size $\Theta(n)$.

This celebrated result is often referred to as "the" phase transition of the ErdősRényi graph model. Although the proof is quite long, it is well worth studying in details. It employs most tools we have seen up to this point: first and second moment methods, Chernoff-Cramér bounds, martingale techniques, coupling and stochastic domination, and branching processes. It is quintessential discrete probability.

### 6.4.1 Statement and proof sketch

Before stating the main theorems, we recall a basic result from Chapter 2.

- (Poisson tail) Let $S_{n}$ be a sum of $n$ i.i.d. Poi $(\lambda)$ variables. Recall from (2.4.10) and (2.4.11) that for $a>\lambda$

$$
\begin{equation*}
-\frac{1}{n} \log \mathbb{P}\left[S_{n} \geq a n\right] \geq a \log \left(\frac{a}{\lambda}\right)-a+\lambda=: I_{\lambda}^{\mathrm{Poi}}(a), \tag{6.4.1}
\end{equation*}
$$

and similarly for $a<\lambda$

$$
\begin{equation*}
-\frac{1}{n} \log \mathbb{P}\left[S_{n} \leq a n\right] \geq I_{\lambda}^{\mathrm{Poi}}(a) . \tag{6.4.2}
\end{equation*}
$$

To simplify the notation, we let

$$
\begin{equation*}
I_{\lambda}:=I_{\lambda}^{\mathrm{Poi}}(1)=\lambda-1-\log \lambda \geq 0, \tag{6.4.3}
\end{equation*}
$$

where the inequality follows from the convexity of $I_{\lambda}$ and the fact that it attains its minimum at $\lambda=1$ where it is 0 .

We let $p=\frac{\lambda}{n}$ and denote by $\mathcal{C}_{\text {max }}$ a largest connected component. In the subcritical case, that is, when $\lambda<1$, we show that the largest connected component has logarithmic size in $n$.

Theorem 6.4.1 (Subcritical case: upper bound on the largest cluster). Let $G_{n} \sim$ $\mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda \in(0,1)$. For all $\kappa>0$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right| \geq(1+\kappa) I_{\lambda}^{-1} \log n\right]=o(1),
$$

where $I_{\lambda}$ is defined in (6.4.3).
We also give a matching logarithmic lower bound on the size of $\mathcal{C}_{\text {max }}$ in Theorem 6.4.11.

In the supercritical case, that is, when $\lambda>1$, we prove the existence of a unique connected component of size linear in $n$, which is referred to as the giant component.

Theorem 6.4.2 (Supercritical regime: giant component). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda>1$. For any $\gamma \in(1 / 2,1)$ and $\delta<2 \gamma-1$,

$$
\mathbb{P}_{n, p_{n}}\left[| | \mathcal{C}_{\max }\left|-\zeta_{\lambda} n\right| \geq n^{\gamma}\right]=O\left(n^{-\delta}\right),
$$

where $\zeta_{\lambda}$ be the unique solution in $(0,1)$ to the fixed point equation

$$
1-e^{-\lambda \zeta}=\zeta .
$$

In fact, with probability $1-O\left(n^{-\delta}\right)$, there is a unique largest component and the second largest connected component has size $O(\log n)$.

See Figure 6.4 for an illustration.
At a high level, the proof goes as follows:

- (Subcritical regime) In the subcritical case, we use an exploration process and a domination argument to approximate the size of the connected components with the progeny of a branching process. The result then follows from the hitting-time theorem and the Poisson tail.
- (Supercritical regime) In the supercritical case, a similar argument gives a bound on the expected size of the giant component, which is related to the survival of the branching process. Chebyshev's inequality gives concentration. The hard part there is to bound the variance.


Figure 6.4: Illustration of the phase transition.

### 6.4.2 Bounding cluster size: domination by branching processes

For a vertex $v \in[n]$, let $\mathcal{C}_{v}$ be the connected component containing $v$, which we also refer to as the cluster of $v$. To analyze the size of $\mathcal{C}_{v}$, we use the exploration process introduced in Section 6.2.1 and show that it is dominated above and below by branching processes.

## Exploration process

Recall that the exploration process started at $v$ has 3 types of vertices: the active vertices $\mathcal{A}_{t}$, the explored vertices $\mathcal{E}_{t}$, and the neutral vertices $\mathcal{N}_{t}$. We start with $\mathcal{A}_{0}:=\{v\}, \mathcal{E}_{0}:=\emptyset$, and $\mathcal{N}_{0}$ contains all other vertices in $G_{n}$. We imagine revealing the edges of $G_{n}$ as they are encountered in this process and we let $\left(\mathcal{F}_{t}\right)$ be the corresponding filtration. In words, starting with $v$, the cluster of $v$ is progressively grown by adding to it at each time a vertex adjacent to one of the previously explored vertices and uncovering its remaining neighbors in $G_{n}$.

Let as before $A_{t}:=\left|\mathcal{A}_{t}\right|, E_{t}:=\left|\mathcal{E}_{t}\right|$, and $N_{t}:=\left|\mathcal{N}_{t}\right|$, and

$$
\tau_{0}:=\inf \left\{t \geq 0: A_{t}=0\right\}=\left|\mathcal{C}_{v}\right|,
$$

where the rightmost equality is from Lemma 6.2.1. Recall that $\left(E_{t}\right)$ is non-decreasing while $\left(N_{t}\right)$ is non-increasing, and that the process is fixed for all $t>\tau_{0}$. Since $E_{t}=t$ for all $t \leq \tau_{0}$ (as exactly one vertex is explored at each time until the set of active vertices is empty) and $\left(\mathcal{A}_{t}, \mathcal{E}_{t}, \mathcal{N}_{t}\right)$ forms a partition of $[n]$ for all $t$, we have

$$
\begin{equation*}
A_{t}+t+N_{t}=n, \quad \forall t \leq \tau_{0} . \tag{6.4.4}
\end{equation*}
$$

Hence, in tracking the size of the exploration process, we can work with $A_{t}$ or $N_{t}$. Moreover at $t=\tau_{0}$ we have

$$
\begin{equation*}
\left|\mathcal{C}_{v}\right|=\tau_{0}=n-N_{\tau_{0}} . \tag{6.4.5}
\end{equation*}
$$

Similarly to the case of a Galton-Watson tree, the processes $\left(A_{t}\right)$ and $\left(N_{t}\right)$ admit a simple recursive form. Conditioning on $\mathcal{F}_{t-1}$ :

- (Active vertices) If $A_{t-1}=0$, the exploration process has finished its course and $A_{t}=0$. Otherwise, (a) one active vertex becomes explored and (b) its neutral neighbors become active vertices. That is,

$$
\begin{equation*}
A_{t}=A_{t-1}+1_{\left\{A_{t-1}>0\right\}}[\underbrace{-1}_{\text {(a) }}+\underbrace{X_{t}}_{\text {(b) }}], \tag{6.4.6}
\end{equation*}
$$

where $X_{t}$ is binomial with parameters $N_{t-1}$ and $p_{n}$. By (6.4.4), $N_{t-1}$ can be written in terms of $A_{t-1}$ as $N_{t-1}=n-(t-1)-A_{t-1}$. For the coupling arguments below, it will be useful to think of $X_{t}$ as a sum of independent Bernoulli variables. That is, let ( $I_{t, j}: t \geq 1, j \geq 1$ ) be an array of independent, identically distributed $\{0,1\}$-variables with $\mathbb{P}\left[I_{1,1}=1\right]=p_{n}$. We write

$$
\begin{equation*}
X_{t}=\sum_{i=1}^{N_{t-1}} I_{t, i} . \tag{6.4.7}
\end{equation*}
$$

- (Neutral vertices) Similarly, if $A_{t-1}>0$, that is, $N_{t-1}<n-(t-1), X_{t}$ neutral vertices become active. That is,

$$
\begin{equation*}
N_{t}=N_{t-1}-\mathbf{1}_{\left\{N_{t-1}<n-(t-1)\right\}} X_{t} . \tag{6.4.8}
\end{equation*}
$$

## Poisson branching process approximation

With these observations, we now relate the size of the cluster of $v$ to the total progeny of a Poisson branching process with an appropriately chosen offspring mean. The intuition is simple: when $p_{n}=\lambda / n$, the number of neighbors of a vertex is well approximated by a Poisson distribution; therefore, exploration of the
cluster of $v$ is similar to that of the corresponding branching process. We will see that this holds long enough to prove accurate results about the subcritical regime (see Lemma 6.4.6). It will also be useful in the supercritical regime, but additional arguments will be required there (see Lemmas 6.4.7 and 6.4.8).

Lemma 6.4.3 (Cluster size: Poisson branching process approximation). Let $G_{n} \sim$ $\mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda>0$ and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. Let $W_{\lambda}$ be the total progeny of a branching process with offspring distribution $\operatorname{Poi}(\lambda)$. Then, for $1 \leq k_{n}=o(\sqrt{n})$,

$$
\mathbb{P}\left[W_{\lambda} \geq k_{n}\right]-O\left(\frac{k_{n}^{2}}{n}\right) \leq \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq k_{n}\right] \leq \mathbb{P}\left[W_{\lambda} \geq k_{n}\right] .
$$

From Example 6.2.7, we have an explicit formula for the distribution of $W_{\lambda}$.
Before proving the lemma, recall the following simple domination results from Chapter 4:

- (Binomial domination) We have

$$
\begin{equation*}
n \geq m \Longrightarrow \operatorname{Bin}(n, p) \succeq \operatorname{Bin}(m, p) . \tag{6.4.9}
\end{equation*}
$$

The binomial distribution is also dominated by the Poisson distribution in the following way:

$$
\begin{equation*}
\lambda \in(0,1) \quad \Longrightarrow \quad \operatorname{Poi}(\lambda) \succeq \operatorname{Bin}\left(n-1, \frac{\lambda}{n}\right) . \tag{6.4.10}
\end{equation*}
$$

For the proofs, see Examples 4.2.4 and 4.2.8.
We use these domination results to relate the size of a connected component to the progeny of a branching process.

Proof of Lemma 6.4.3. We start with the upper bound.
Upper bound: Because $N_{t-1}=n-(t-1)-A_{t-1} \leq n-1$, conditioned on $\mathcal{F}_{t-1}$, the following stochastic domination relations hold

$$
\operatorname{Bin}\left(N_{t-1}, \frac{\lambda}{n}\right) \preceq \operatorname{Bin}\left(n-1, \frac{\lambda}{n}\right) \preceq \operatorname{Poi}(\lambda),
$$

by (6.4.9) and (6.4.10). Observe that the center and rightmost distributions do not depend on $N_{t-1}$. Let $\left(X_{t}^{\succ}\right)$ be a sequence of independent $\operatorname{Poi}(\lambda)$.

Using the coupling in Example 4.2.8, we can couple the processes $\left(I_{t, j}\right)_{j}$ and $\left(X_{t}^{\succ}\right)$ in such way that $X_{t}^{\succ} \geq \sum_{j=1}^{n-1} I_{t, j}$ almost surely for all $t$. Then by induction on $t$,

$$
A_{t} \leq A_{t}^{\succ},
$$

almost surely for all $t$, where we define (recalling (6.4.6))

$$
\begin{equation*}
A_{t}^{\succ}:=A_{t-1}^{\succ}+\mathbf{1}_{\left\{A_{t-1}^{\succ}>0\right\}}\left[-1+X_{t}^{\succ}\right], \tag{6.4.11}
\end{equation*}
$$

with $A_{0}^{\succ}:=1$. In words, $\left(A_{t}^{\succ}\right)$ is the size of the active set of a Galton-Watson branching process with offspring distribution $\operatorname{Poi}(\lambda)$, as defined in Section 6.2.1.

As a result, letting

$$
W_{\lambda}=\tau_{0}^{\succ}:=\inf \left\{t \geq 0: A_{t}^{\succ}=0\right\},
$$

be the total progeny of this branching process, we immediately get

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq k_{n}\right]=\mathbb{P}_{n, p_{n}}\left[\tau_{0} \geq k_{n}\right] \leq \mathbb{P}\left[\tau_{0}^{\succ} \geq k_{n}\right]=\mathbb{P}\left[W_{\lambda} \geq k_{n}\right] .
$$

Lower bound: In the other direction, we proceed in two steps. We first show that, up to a certain time, the process is bounded from below by a branching process with binomial offspring distribution. In a second step, we show that this binomial branching process can be approximated by a Poisson branching process.

1. (Domination from below) Let $A_{t}^{\prec}$ be defined as (again recalling (6.4.6))

$$
\begin{equation*}
A_{t}^{\prec}:=A_{t-1}^{\prec}+\mathbf{1}_{\left\{A_{t-1}^{\prec}>0\right\}}\left[-1+X_{t}^{\prec}\right], \tag{6.4.12}
\end{equation*}
$$

with $A_{0}^{\prec}:=1$, where

$$
\begin{equation*}
X_{t}^{\prec}:=\sum_{i=1}^{n-k_{n}} I_{t, j} . \tag{6.4.13}
\end{equation*}
$$

Note that we use the same $I_{t, j}$ as in the definition of $X_{t}$, that is, we couple the two processes. This time $\left(A_{t}^{\prec}\right)$ is the size of the active set in the exploration process of a Galton-Watson branching process with offspring distribution $\operatorname{Bin}\left(n-k_{n}, p_{n}\right)$. Let

$$
\tau_{0}^{\prec}:=\inf \left\{t \geq 0: A_{t}^{\prec}=0\right\},
$$

be the total progeny of this branching process. We prove the following relationship between $\tau_{0}$ and $\tau_{0}^{\prec}$.

Lemma 6.4.4. We have

$$
\mathbb{P}\left[\tau_{0}^{\prec} \geq k_{n}\right] \leq \mathbb{P}_{n, p_{n}}\left[\tau_{0} \geq k_{n}\right] .
$$

Proof. We claim that $A_{t}$ is bounded from below by $A_{t}^{\prec}$ up to the stopping time

$$
\sigma_{n-k_{n}}:=\inf \left\{t \geq 0: N_{t} \leq n-k_{n}\right\},
$$

which by convention is $+\infty$ if the event is not reached (i.e., if the cluster is "small"; see below). Indeed, $N_{0}=n-1$ and for all $t \leq \sigma_{n-k_{n}}, N_{t-1}>$ $n-k_{n}$ by definition. Hence, by the coupling (6.4.7) and (6.4.13), $X_{t} \geq X_{t}^{\prec}$ for all $t \leq \sigma_{n-k_{n}}$ and as a result, by induction on $t$,

$$
A_{t} \geq A_{t}^{\prec}, \quad \forall t \leq \sigma_{n-k_{n}}
$$

where we used the recursions (6.4.6) and (6.4.13).
Because the inequality between $A_{t}$ and $A_{t}^{\prec}$ holds only up to time $\sigma_{n-k_{n}}$, we cannot compare directly $\tau_{0}$ and $\tau_{0}^{\prec}$. However, we will use the following observation: the size of the cluster of $v$ is at least the total number of active and explored vertices at any time $t$. In particular, when $\sigma_{n-k_{n}}<+\infty$,

$$
\tau_{0}=\left|\mathcal{C}_{v}\right| \geq A_{\sigma_{n-k_{n}}}+E_{\sigma_{n-k_{n}}}=n-N_{\sigma_{n-k_{n}}} \geq k_{n}
$$

On the other hand, when $\sigma_{n-k_{n}}=+\infty, N_{t}>n-k_{n}$ for all $t$-in particular for $t=\tau_{0}$-and therefore $\left|\mathcal{C}_{v}\right|=\tau_{0}=n-N_{\tau_{0}}<k_{n}$ by (6.4.5). Moreover in that case, because $A_{t} \geq A_{t}^{\prec}$ for all $t \leq \sigma_{n-k_{n}}=+\infty$, it holds in addition that $\tau_{0}^{\prec} \leq \tau_{0}<k_{n}$. To sum up, we have proved the implications

$$
\tau_{0}^{\prec} \geq k_{n} \Longrightarrow \sigma_{n-k_{n}}<+\infty \Longrightarrow \tau_{0} \geq k_{n} .
$$

In particular, we have proved the lemma.
2. (Poisson approximation) Our next step is approximate the tail of $\tau_{0}^{\prec}$ by that of $\tau_{0}^{\succ}$.

Lemma 6.4.5. We have

$$
\mathbb{P}\left[\tau_{0}^{\prec} \geq k_{n}\right]=\mathbb{P}\left[\tau_{0}^{\succ} \geq k_{n}\right]+O\left(\frac{k_{n}^{2}}{n}\right) .
$$

Proof. By Theorem 6.2.6,

$$
\begin{equation*}
\mathbb{P}\left[\tau_{0}^{\prec}=t\right]=\frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\prec}=t-1\right], \tag{6.4.14}
\end{equation*}
$$

where the $X_{i}^{\prec}$ s are independent $\operatorname{Bin}\left(n-k_{n}, p_{n}\right)$. Note further that, because the sum of independent binomials with the same success probability is binomial,

$$
\sum_{i=1}^{t} X_{i}^{\prec} \sim \operatorname{Bin}\left(t\left(n-k_{n}\right), p_{n}\right) .
$$

Recall on the other hand that $\left(X_{t}^{\succ}\right)$ is $\operatorname{Poi}(\lambda)$ and, because a sum of independent Poisson is Poisson (see Exercise 6.7), we have

$$
\begin{equation*}
\mathbb{P}\left[\tau_{0}^{\succ}=t\right]=\frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right], \tag{6.4.15}
\end{equation*}
$$

where

$$
\sum_{i=1}^{t} X_{i}^{\succ} \sim \operatorname{Poi}(t \lambda)
$$

We use the Poisson approximation result in Theorem 4.1.18 to compare the probabilities on the right-hand sides of (6.4.14) and (6.4.15). In fact, because the Poisson approximation is in terms of the total variation distance-which bounds any event-one might be tempted to apply it directly to the tails of $\tau_{0}^{\prec}$ and $\tau_{0}^{\succ}$ by summing over $t$. However note that the factor of $1 / t$ in (6.4.14) and (6.4.15) prevents us from doing so.
Instead, we argue for each $t$ separately and use that

$$
\begin{aligned}
\mid \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\prec}=\right. & t-1]-\mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right] \mid \\
& \leq\left\|\operatorname{Bin}\left(t\left(n-k_{n}\right), p_{n}\right)-\operatorname{Poi}(t \lambda)\right\|_{\mathrm{TV}}
\end{aligned}
$$

by the observations in the previous paragraph. Theorem 4.1.18 tells us that

$$
\begin{aligned}
&\left\|\operatorname{Bin}\left(t\left(n-k_{n}\right), p_{n}\right)-\operatorname{Poi}\left(t\left(n-k_{n}\right)\left[-\log \left(1-p_{n}\right)\right]\right)\right\|_{\mathrm{TV}} \\
& \leq \frac{1}{2} t\left(n-k_{n}\right)\left[-\log \left(1-p_{n}\right)\right]^{2} .
\end{aligned}
$$

We must adjust the mean of the Poisson distribution. To do so, we argue as in Example 4.1.12 to get

$$
\begin{aligned}
&\left\|\operatorname{Poi}\left(t\left(n-k_{n}\right)\left[-\log \left(1-p_{n}\right)\right]\right)-\operatorname{Poi}(t \lambda)\right\|_{\mathrm{TV}} \\
& \leq\left|t \lambda-t\left(n-k_{n}\right)\left(-\log \left(1-p_{n}\right)\right)\right| .
\end{aligned}
$$

Finally, recalling that $p_{n}=\lambda / n$, combining the last three displays and using the triangle inequality for the total variation distance,

$$
\begin{aligned}
& \left|\mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\prec}=t-1\right]-\mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right]\right| \\
& \quad \leq \frac{1}{2} t\left(n-k_{n}\right)\left[-\log \left(1-p_{n}\right)\right]^{2}+\left|t \lambda-t\left(n-k_{n}\right)\left(-\log \left(1-p_{n}\right)\right)\right| \\
& \quad \leq \frac{1}{2} t n\left(\frac{\lambda}{n}+O\left(\frac{\lambda^{2}}{n^{2}}\right)\right)^{2}+\left|t \lambda-t\left(n-k_{n}\right)\left(\frac{\lambda}{n}+O\left(\frac{\lambda^{2}}{n^{2}}\right)\right)\right| \\
& \\
& \quad=O\left(\frac{t k_{n}}{n}\right),
\end{aligned}
$$

where we used that $k_{n} \geq 1$ and $\lambda$ is fixed.
So, by (6.4.14) and (6.4.15), dividing by $t$ and then summing over $t<k_{n}$ gives

$$
\left|\mathbb{P}\left[\tau_{0}^{\prec}<k_{n}\right]-\mathbb{P}\left[\tau_{0}^{\succ}<k_{n}\right]\right|=O\left(\frac{k_{n}^{2}}{n}\right) .
$$

Rearranging proves the lemma.
Putting together Lemmas 6.4.4 and 6.4.5 gives

$$
\begin{aligned}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq k_{n}\right] & =\mathbb{P}_{n, p_{n}}\left[\tau_{0} \geq k_{n}\right] \\
& \geq \mathbb{P}\left[\tau_{0}^{\succ} \geq k_{n}\right]-O\left(\frac{k_{n}^{2}}{n}\right) \\
& =\mathbb{P}\left[W_{\lambda} \geq k_{n}\right]-O\left(\frac{k_{n}^{2}}{n}\right),
\end{aligned}
$$

as claimed.

## Subcritical regime: largest cluster

We are now ready to analyze the subcritical regime, that is, the case $\lambda<1$.
Lemma 6.4.6 (Subcritical regime: upper bound on cluster size). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda \in(0,1)$ and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. For all $\kappa>0$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq(1+\kappa) I_{\lambda}^{-1} \log n\right]=O\left(n^{-(1+\kappa)}\right)
$$

Proof. We use the Poisson branching process approximation (Lemma 6.4.3). To apply the lemma we need to bound the tail of the progeny $W_{\lambda}$ of a Poisson branching process. Using the notation of Lemma 6.4.3, by Theorem 6.2.6,

$$
\begin{equation*}
\mathbb{P}\left[W_{\lambda} \geq k_{n}\right]=\mathbb{P}\left[W_{\lambda}=+\infty\right]+\sum_{t \geq k_{n}} \frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right], \tag{6.4.16}
\end{equation*}
$$

where the $X_{i}^{\succ}$ s are i.i.d. $\operatorname{Poi}(\lambda)$. Both terms on the right-hand side depend on whether or not the mean $\lambda$ is smaller or larger than 1 . When $\lambda<1$, the Poisson branching process goes extinct with probability 1 by the extinction theory (Theorem 6.1.6). Hence $\mathbb{P}\left[W_{\lambda}=+\infty\right]=0$.

As to the second term, the sum of the $X_{i}^{\succ}$ s is $\operatorname{Poi}(\lambda t)$. Using the Poisson tail (6.4.1) for $\lambda<1$ and $k_{n}=\omega(1)$,

$$
\begin{align*}
\sum_{t \geq k_{n}} \frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right] & \leq \sum_{t \geq k_{n}} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ} \geq t-1\right] \\
& \leq \sum_{t \geq k_{n}} \exp \left(-t I_{\lambda}^{\mathrm{Poi}}\left(\frac{t-1}{t}\right)\right) \\
& \leq \sum_{t \geq k_{n}} \exp \left(-t\left(I_{\lambda}-O\left(t^{-1}\right)\right)\right) \\
& \leq \sum_{t \geq k_{n}} C \exp \left(-t I_{\lambda}\right) \\
& =O\left(\exp \left(-I_{\lambda} k_{n}\right)\right) \tag{6.4.17}
\end{align*}
$$

for some constant $C>0$.
Let $c=(1+\kappa) I_{\lambda}^{-1}$ for $\kappa>0$. By Lemma 6.4.3,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq c \log n\right] \leq \mathbb{P}\left[W_{\lambda} \geq c \log n\right] .
$$

By (6.4.16) and (6.4.17),

$$
\begin{equation*}
\mathbb{P}\left[W_{\lambda} \geq c \log n\right]=O\left(\exp \left(-I_{\lambda} c \log n\right)\right), \tag{6.4.18}
\end{equation*}
$$

which proves the claim.
As before, let $\mathcal{C}_{\text {max }}$ be a largest connected component of $G_{n}$ (choosing the component containing the lowest label if there is more than one such component). A union bound and the previous lemma immediately imply an upper bound on the size of $\mathcal{C}_{\text {max }}$ in the subcritical case.

Proof of Theorem 6.4.1. Let again $c=(1+\kappa) I_{\lambda}^{-1}$ for $\kappa>0$. By a union bound and symmetry,

$$
\begin{align*}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right| \geq c \log n\right] & =\mathbb{P}_{n, p_{n}}\left[\exists v,\left|\mathcal{C}_{v}\right|>c \log n\right] \\
& \leq n \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{1}\right| \geq c \log n\right] . \tag{6.4.19}
\end{align*}
$$

By Lemma 6.4.6,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right| \geq c \log n\right]=O\left(n \cdot n^{-(1+\kappa)}\right)=O\left(n^{-\kappa}\right) \rightarrow 0
$$

as $n \rightarrow+\infty$.
In fact we prove below that the largest component is indeed of size roughly $I_{\lambda}^{-1} \log n$. But first we turn to the supercritical regime.

## Supercritical regime: two phases

Applying the Poisson branching process approximation in the supercritical regime gives the following.

Lemma 6.4.7 (Supercritical regime: extinction). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda>1$. and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. Let $\zeta_{\lambda}$ be the unique solution in $(0,1)$ to the fixed point equation

$$
1-e^{-\lambda \zeta}=\zeta .
$$

For any $\kappa>0$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq(1+\kappa) I_{\lambda}^{-1} \log n\right]=\zeta_{\lambda}+O\left(\frac{\log ^{2} n}{n}\right)
$$

Note the small-but critical difference-with Lemma 6.4.6: this time the branching process can survive. This happens with probability $\zeta_{\lambda}$ by extinction theory (Theorem 6.1.6). In that case, we will need further arguments to nail down the cluster size. Observe also that the result holds for a fixed vertex $v$-and therefore does not yet tell us about the largest cluster. We come back to the latter in the next subsection.

Proof of Lemma 6.4.7. We adapt the proof of Lemma 6.4.6, beginning with (6.4.16) which recall states

$$
\mathbb{P}\left[W_{\lambda} \geq k_{n}\right]=\mathbb{P}\left[W_{\lambda}=+\infty\right]+\sum_{t \geq k_{n}} \frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right],
$$

where the $X_{i}^{\succ}$ s are i.i.d. $\operatorname{Poi}(\lambda)$. When $\lambda>1, \mathbb{P}\left[W_{\lambda}=+\infty\right]=\zeta_{\lambda}$, where $\zeta_{\lambda}>0$ is the survival probability of the branching process by Example 6.1.10. As to the second term, using (6.4.2) for $\lambda>1$,

$$
\begin{align*}
\sum_{t \geq k_{n}} \frac{1}{t} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ}=t-1\right] & \leq \sum_{t \geq k_{n}} \mathbb{P}\left[\sum_{i=1}^{t} X_{i}^{\succ} \leq t\right] \\
& \leq \sum_{t \geq k_{n}} \exp \left(-t I_{\lambda}\right) \\
& \leq C \exp \left(-I_{\lambda} k_{n}\right) \tag{6.4.20}
\end{align*}
$$

for a constant $C>0$.
Now let $c=(1+\kappa) I_{\lambda}^{-1}$ for $\kappa>0$. By Lemma 6.4.3,

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq c \log n\right]=\mathbb{P}\left[W_{\lambda} \geq c \log n\right]+O\left(\frac{\log ^{2} n}{n}\right) . \tag{6.4.21}
\end{equation*}
$$

By (6.4.16) and (6.4.20),

$$
\begin{align*}
\mathbb{P}\left[W_{\lambda} \geq c \log n\right] & =\zeta_{\lambda}+O\left(\exp \left(-c I_{\lambda} \log n\right)\right) \\
& =\zeta_{\lambda}+O\left(n^{-(1+\kappa)}\right) . \tag{6.4.22}
\end{align*}
$$

Combining (6.4.21) and (6.4.22), for any $\kappa>0$,

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \geq c \log n\right]=\zeta_{\lambda}+O\left(\frac{\log ^{2} n}{n}\right) \tag{6.4.23}
\end{equation*}
$$

as claimed.
Recall that the Poisson branching process approximation was based on the fact that the degree of a vertex is well approximated by a Poisson distribution. When the exploration process goes on for too long however (i.e., when $k_{n}$ is large), this approximation is not as accurate because of a saturation effect: at each step of the exploration, we uncover edges to the neutral vertices (which then become active); and, because an Erdős-Rényi graph has a finite pool of vertices from which to draw these edges, as the number of neutral vertices decreases so does the expected number of uncovered edges. Instead we use the following lemma which explicitly accounts for the dwindling size of $\mathcal{N}_{t}$. Roughly speaking, we model the set of neutral vertices as a process that discards a fraction $p_{n}$ of its current set at each time step (i.e., those neutral vertices with an edge to the current explored vertex).

Lemma 6.4.8. Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda>0$ and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. Let $Y_{t} \sim \operatorname{Bin}\left(n-1,1-\left(1-p_{n}\right)^{t}\right)$. Then, for any $t$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|=t\right] \leq \mathbb{P}\left[Y_{t}=t-1\right] .
$$

Proof. We work with neutral vertices. By (6.4.4) and Lemma 6.2.1, for any $t$,

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|=t\right]=\mathbb{P}_{n, p_{n}}\left[\tau_{0}=t\right] \leq \mathbb{P}_{n, p_{n}}\left[N_{t}=n-t\right] . \tag{6.4.24}
\end{equation*}
$$

Recall that $N_{0}=n-1$ and

$$
N_{t}=N_{t-1}-\mathbf{1}_{\left\{N_{t-1}<n-(t-1)\right\}} \sum_{i=1}^{N_{t-1}} I_{t, i} .
$$

It is easier to consider the process without the indicator as it has a simple distribution. Define $N_{0}^{0}:=n-1$ and

$$
N_{t}^{0}:=N_{t-1}^{0}-\sum_{i=1}^{N_{t-1}^{0}} I_{t, i},
$$

and observe that $N_{t} \geq N_{t}^{0}$ for all $t$, as the two processes agree up to time $\tau_{0}$ at which point $N_{t}$ stays fixed. The interpretation of $N_{t}^{0}$ is straightforward: starting with $n-1$ vertices, at each time each remaining vertex is discarded with probability $p_{n}$. Hence, the number of surviving vertices at time $t$ has distribution

$$
N_{t}^{0} \sim \operatorname{Bin}\left(n-1,\left(1-p_{n}\right)^{t}\right)
$$

by the independence of the steps. Arguing as in (6.4.24),

$$
\begin{aligned}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|=t\right] & \leq \mathbb{P}_{n, p_{n}}\left[N_{t}^{0}=n-t\right] \\
& =\mathbb{P}_{n, p_{n}}\left[(n-1)-N_{t}^{0}=t-1\right] \\
& =\mathbb{P}\left[Y_{t}=t-1\right],
\end{aligned}
$$

which concludes the proof.
The previous lemma gives the following additional bound on the cluster size in the supercritical regime. Together with Lemma 6.4 .7 it shows that, when $\left|\mathcal{C}_{v}\right|>$ $c \log n$, the cluster size is in fact linear in $n$ with high probability. We will have more to say about the largest cluster in the next subsection.

Lemma 6.4.9 (Supercritical regime: saturation). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda>1$ and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. Let $\zeta_{\lambda}$ be the unique solution in $(0,1)$ to the fixed point equation

$$
1-e^{-\lambda \zeta}=\zeta .
$$

For any $\alpha<\zeta_{\lambda}$ and any $\delta>0$, there exists $\kappa_{\delta, \alpha}>0$ large enough so that

$$
\begin{equation*}
\mathbb{P}_{n, p_{n}}\left[\left(1+\kappa_{\delta, \alpha}\right) I_{\lambda}^{-1} \log n \leq\left|\mathcal{C}_{v}\right| \leq \alpha n\right]=O\left(n^{-(1+\delta)}\right) \tag{6.4.25}
\end{equation*}
$$

Proof. By Lemma 6.4.8,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|=t\right] \leq \mathbb{P}\left[Y_{t}=t-1\right] \leq \mathbb{P}\left[Y_{t} \leq t\right],
$$

where $Y_{t} \sim \operatorname{Bin}\left(n-1,1-\left(1-p_{n}\right)^{t}\right)$. Roughly, the right-hand side is negligible until the mean $\mu_{t}:=(n-1)\left(1-(1-\lambda / n)^{t}\right)$ is of the order of $t$. Let $\zeta_{\lambda}$ be as above, and recall that it is a solution to

$$
1-e^{-\lambda \zeta}-\zeta=0
$$

Note in particular that, when $t=\zeta_{\lambda} n$,

$$
\mu_{t}=(n-1)\left(1-(1-\lambda / n)^{\zeta_{\lambda} n}\right) \approx n\left(1-e^{-\lambda \zeta_{\lambda}}\right)=\zeta_{\lambda} n=t .
$$

Let $\alpha<\zeta_{\lambda}$.
For any $t \in[c \log n, \alpha n]$, by the Chernoff bound for Poisson trials (Theorem 2.4.7 (ii)(b)),

$$
\begin{equation*}
\mathbb{P}\left[Y_{t} \leq t\right] \leq \exp \left(-\frac{\mu_{t}}{2}\left(1-\frac{t}{\mu_{t}}\right)^{2}\right) \tag{6.4.26}
\end{equation*}
$$

For $t / n \leq \alpha<\zeta_{\lambda}$, using $1-x \leq e^{-x}$ for $x \in(0,1)$ (see Exercise 1.16), there is $\gamma_{\alpha}>1$ such that

$$
\begin{aligned}
\mu_{t} & \geq(n-1)\left(1-e^{-\lambda(t / n)}\right) \\
& =t\left(\frac{n-1}{n}\right) \frac{1-e^{-\lambda(t / n)}}{t / n} \\
& \geq t\left(\frac{n-1}{n}\right) \frac{1-e^{-\lambda \alpha}}{\alpha} \\
& \geq \gamma_{\alpha} t,
\end{aligned}
$$

for $n$ large enough, where we used that $1-e^{-\lambda x}$ is increasing in $x$ on the third line and that $1-e^{-\lambda x}-x>0$ for $0<x<\zeta_{\lambda}$ on the fourth line (as can be checked by
computing the first and second derivatives). Plugging this back into (6.4.26), we get

$$
\mathbb{P}\left[Y_{t} \leq t\right] \leq \exp \left(-t\left\{\frac{\gamma_{\alpha}}{2}\left(1-\frac{1}{\gamma_{\alpha}}\right)^{2}\right\}\right)
$$

Therefore

$$
\begin{aligned}
\sum_{t=c \log n}^{\alpha n} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|=t\right] & \leq \sum_{t=c \log n}^{\alpha n} \mathbb{P}\left[Y_{t} \leq t\right] \\
& \leq \sum_{t=c \log n}^{+\infty} \exp \left(-t\left\{\frac{\gamma_{\alpha}}{2}\left(1-\frac{1}{\gamma_{\alpha}}\right)^{2}\right\}\right) \\
& =O\left(\exp \left(-c \log n\left\{\frac{\gamma_{\alpha}}{2}\left(1-\frac{1}{\gamma_{\alpha}}\right)^{2}\right\}\right)\right) .
\end{aligned}
$$

Taking $\kappa>0$ large enough proves (6.4.25).

### 6.4.3 Concentration of cluster size: second moment bounds

To characterize the size of the largest cluster in the supercritical case, we use Chebyshev's inequality. We also use a related second moment argument to give a lower bound on the largest cluster in the subcritical regime.

## Supercritical regime: giant component

Assume $\lambda>1$. Our goal is to characterize the size of the largest component. We do this by bounding what is not in it (i.e., intuitively those vertices whose exploration process goes extinct). For $\delta>0$ and $\alpha<\zeta_{\lambda}$, let $\kappa_{\delta, \alpha}$ be as defined in Lemma 6.4.9. Set

$$
\underline{k}_{n}:=\left(1+\kappa_{\delta, \alpha}\right) I_{\lambda}^{-1} \log n \quad \text { and } \quad \bar{k}_{n}:=\alpha n .
$$

We call a vertex $v$ such that $\left|\mathcal{C}_{v}\right| \leq \underline{k}_{n}$ a small vertex.
Let

$$
S_{k}:=\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right| \leq k\right\}} .
$$

It will also be useful to work with

$$
B_{k}=n-S_{k}=\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right|>k\right\}} .
$$

The quantity $S_{\underline{k}_{n}}$ is the number of small vertices. By Lemma 6.4.7, its expectation is

$$
\begin{equation*}
\mathbb{E}_{n, p_{n}}\left[S_{\underline{k}_{n}}\right]=n\left(1-\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|>\underline{k}_{n}\right]\right)=\left(1-\zeta_{\lambda}\right) n+O\left(\log ^{2} n\right) . \tag{6.4.27}
\end{equation*}
$$

Using Chebyshev's inequality (Theorem 2.1.2), we prove that $S_{\underline{k}_{n}}$ is concentrated.
Lemma 6.4.10 (Concentration of $S_{\underline{k}_{n}}$ ). For any $\gamma \in(1 / 2,1)$ and $\delta<2 \gamma-1$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|S_{\underline{k}_{n}}-\left(1-\zeta_{\lambda}\right) n\right| \geq n^{\gamma}\right]=O\left(n^{-\delta}\right) .
$$

Lemma 6.4.10, which is proved below, leads to our main result in the supercritical case: the existence of the giant component, a unique cluster $\mathcal{C}_{\text {max }}$ of size linear in $n$.

Proof of Theorem 6.4.2. Take $\alpha \in\left(\zeta_{\lambda} / 2, \zeta_{\lambda}\right)$ and let $\underline{k}_{n}, \bar{k}_{n}$, and $\gamma$ be as above. Let $\mathcal{B}_{1, n}:=\left\{\left|B_{\underline{k}_{n}}-\zeta_{\lambda} n\right| \geq n^{\gamma}\right\}$. Because $\gamma<1$, the event $\mathcal{B}_{1, n}^{c}$ implies that

$$
\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right|>\underline{k}_{n}\right\}}=B_{\underline{\underline{k}}_{n}}>\zeta_{\lambda} n-n^{\gamma} \geq 1,
$$

for $n$ large enough. That is, there is at least one "large" cluster of size $>\underline{k}_{n}$. In turn, that implies

$$
\left|\mathcal{C}_{\max }\right| \leq B_{{\underline{k_{n}}}}
$$

since there are at most $B_{\underline{k}_{n}}$ vertices in that large cluster.
Let $\mathcal{B}_{2, n}:=\left\{\exists v,\left|\mathcal{C}_{v}\right| \in\left[\underline{k}_{n}, \bar{k}_{n}\right]\right\}$. If $\mathcal{B}_{2, n}^{c}$ holds, in addition to $\mathcal{B}_{1, n}^{c}$, then

$$
\left|\mathcal{C}_{\max }\right| \leq B_{\underline{k}_{n}}=B_{\bar{k}_{n}},
$$

since there is no cluster whose size falls in $\left[\underline{k}_{n}, \bar{k}_{n}\right]$. Moreover there is equality across the last display if there is a unique cluster of size greater than $\bar{k}_{n}$.

This is indeed the case under $\mathcal{B}_{1, n}^{c} \cap \mathcal{B}_{2, n}^{c}$ : if there were two distinct clusters of size $\bar{k}_{n}$, then since $2 \alpha>\zeta_{\lambda}$ we would have for $n$ large enough

$$
B_{\underline{k}_{n}}=B_{\bar{k}_{n}}>2 \bar{k}_{n}=2 \alpha n>\zeta_{\lambda} n+n^{\gamma},
$$

a contradiction. Hence we have proved that under $\mathcal{B}_{1, n}^{c} \cap \mathcal{B}_{2, n}^{c}$

$$
\left|\mathcal{C}_{\max }\right|=B_{\underline{\underline{k}}_{n}}=B_{\bar{k}_{n}} .
$$

Take $\delta<2 \gamma-1$. Applying Lemmas 6.4.9 and 6.4.10

$$
\mathbb{P}\left[\mathcal{B}_{1, n} \cup \mathcal{B}_{2, n}\right] \leq O\left(n^{-\delta}\right)+n \cdot O\left(n^{-(1+\delta)}\right)=O\left(n^{-\delta}\right),
$$

which concludes the proof.

It remains to prove Lemma 6.4.10.
Proof of Lemma 6.4.10. As mentioned above, we use Chebyshev's inequality. Hence our main task is to bound the variance of $S_{\underline{k}_{n}}$.

Our starting point is the following expression for the second moment

$$
\begin{align*}
\mathbb{E}_{n, p_{n}}\left[S_{k}^{2}\right]= & \sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k,\left|\mathcal{C}_{v}\right| \leq k\right] \\
= & \sum_{u, v \in[n]}\left\{\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k,\left|\mathcal{C}_{v}\right| \leq k, u \leftrightarrow v\right]\right. \\
& \left.\quad+\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k,\left|\mathcal{C}_{v}\right| \leq k, u \leftrightarrow v\right]\right\} \tag{6.4.28}
\end{align*}
$$

where $u \leftrightarrow v$ indicates that $u$ and $v$ are in the same connected component.
To bound the first term in (6.4.28), we note that $u \leftrightarrow v$ implies that $\mathcal{C}_{u}=\mathcal{C}_{v}$. Hence,

$$
\begin{align*}
\sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k,\left|\mathcal{C}_{v}\right| \leq k, u \leftrightarrow v\right] & =\sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k, v \in \mathcal{C}_{u}\right] \\
& =\sum_{u, v \in[n]} \mathbb{E}_{n, p_{n}}\left[\mathbf{1}_{\left\{\left|\mathcal{C}_{u}\right| \leq k\right\}} \mathbf{1}_{\left\{v \in \mathcal{C}_{u}\right\}}\right] \\
& =\sum_{u \in[n]} \mathbb{E}_{n, p_{n}}\left[\mathbf{1}_{\left\{\left|\mathcal{C}_{u}\right| \leq k\right\}} \sum_{v \in[n]} \mathbf{1}_{\left\{v \in \mathcal{C}_{u}\right\}}\right] \\
& =\sum_{u \in[n]} \mathbb{E}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \mathbf{1}_{\left\{\left|\mathcal{C}_{u}\right| \leq k\right\}}\right] \\
& =n \mathbb{E}_{n, p_{n}}\left[\left|\mathcal{C}_{1}\right| \mathbf{1}_{\left\{\left|\mathcal{C}_{1}\right| \leq k\right\}}\right] \\
& \leq n k . \tag{6.4.29}
\end{align*}
$$

To bound the second term in (6.4.28), we sum over the size of $\mathcal{C}_{u}$ and note that, conditioned on $\left\{\left|\mathcal{C}_{u}\right|=\ell, u \nleftarrow v\right\}$, the size of $\mathcal{C}_{v}$ has the same distribution as the unconditional size of $\mathcal{C}_{1}$ in a $\mathbb{G}_{n-\ell, p_{n}}$ random graph, that is,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k| | \mathcal{C}_{u} \mid=\ell, u \nleftarrow v\right]=\mathbb{P}_{n-\ell, p_{n}}\left[\left|\mathcal{C}_{1}\right| \leq k\right] .
$$

Observe that the probability on the right-hand side is increasing in $\ell$ (as can be
seen, e.g., by coupling; see below for a related argument). Hence

$$
\begin{aligned}
& \sum_{u, v \in[n]} \sum_{\ell \leq k} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right|=\ell,\left|\mathcal{C}_{v}\right| \leq k, u \nrightarrow v\right] \\
& \quad=\sum_{u, v \in[n]} \sum_{\ell \leq k} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right|=\ell, u \nrightarrow v\right] \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k| | \mathcal{C}_{u} \mid=\ell, u \leftrightarrow v\right] \\
& \quad=\sum_{u, v \in[n]} \sum_{\ell \leq k} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right|=\ell, u \nrightarrow v\right] \mathbb{P}_{n-\ell, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k\right] \\
& \quad \leq \sum_{u, v \in[n]} \sum_{\ell \leq k} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right|=\ell\right] \mathbb{P}_{n-k, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k\right] \\
& =\sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k\right] \mathbb{P}_{n-k, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k\right] .
\end{aligned}
$$

To get a bound on the variance of $S_{k}$, we need to relate this last expression to $\left(\mathbb{E}_{n, p_{n}}\left[S_{k}\right]\right)^{2}$, where we will use that

$$
\begin{equation*}
\mathbb{E}_{n, p_{n}}\left[S_{k}\right]=\mathbb{E}_{n, p_{n}}\left[\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right| \leq k\right\}}\right]=\sum_{v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k\right] . \tag{6.4.30}
\end{equation*}
$$

We define

$$
\Delta_{k}:=\mathbb{P}_{n-k, p_{n}}\left[\left|\mathcal{C}_{1}\right| \leq k\right]-\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{1}\right| \leq k\right] .
$$

Then, plugging this back above, we get

$$
\begin{aligned}
& \sum_{u, v \in[n]} \sum_{\ell \leq k} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right|=\ell,\left|\mathcal{C}_{v}\right| \leq k, u \leftrightarrow v\right] \\
& \quad \leq \sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k\right]\left(\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right| \leq k\right]+\Delta_{k}\right) \\
& \quad \leq\left(\mathbb{E}_{n, p_{n}}\left[S_{k}\right]\right)^{2}+n^{2}\left|\Delta_{k}\right|,
\end{aligned}
$$

by (6.4.30). It remains to bound $\Delta_{k}$.
We use a coupling argument. Let $H \sim \mathbb{G}_{n-k, p_{n}}$ and construct $H^{\prime} \sim \mathbb{G}_{n, p_{n}}$ in the following manner: let $H^{\prime}$ coincide with $H$ on the first $n-k$ vertices then pick the rest the edges independently. Then clearly $\Delta_{k} \geq 0$ since the cluster of 1 in $H^{\prime}$ includes the cluster of 1 in $H$. In fact, $\Delta_{k}$ is the probability that under this coupling the cluster of 1 has at most $k$ vertices in $H$ but not in $H^{\prime}$. That implies in particular that at least one of the vertices in the cluster of 1 in $H$ is connected to a vertex in $\{n-k+1, \ldots, n\}$. Hence by a union bound over those $k^{2}$ potential edges

$$
\Delta_{k} \leq k^{2} p_{n}
$$

and

$$
\begin{equation*}
\sum_{u, v \in[n]} \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{u}\right| \leq k,\left|\mathcal{C}_{v}\right| \leq k, u \leftrightarrow v\right] \leq\left(\mathbb{E}_{n, p_{n}}\left[S_{k}\right]\right)^{2}+\lambda n k^{2} . \tag{6.4.31}
\end{equation*}
$$

Combining (6.4.29) and (6.4.31), we get

$$
\operatorname{Var}\left[S_{k}\right] \leq 2 \lambda n k^{2}
$$

The result follows from (6.4.27) and Chebyshev's inequality

$$
\begin{aligned}
& \mathbb{P}\left[\left|S_{\underline{k}_{n}}-\left(1-\zeta_{\lambda}\right) n\right| \geq n^{\gamma}\right] \\
& \leq \mathbb{P}\left[\left|S_{\underline{k}_{n}}-\mathbb{E}_{n, p_{n}}\left[S_{\underline{k}_{n}}\right]\right| \geq n^{\gamma}-C \log ^{2} n\right] \\
& \leq \frac{2 \lambda n \underline{k}_{n}^{2}}{\left(n^{\gamma}-C \log ^{2} n\right)^{2}} \\
& \leq \frac{2 \lambda n\left(1+\kappa_{\delta, \alpha}\right)^{2} I_{\lambda}^{-2} \log ^{2} n}{C^{\prime} n^{2 \gamma}} \\
& \leq C^{\prime \prime} n^{-\delta},
\end{aligned}
$$

for constants $C, C^{\prime}, C^{\prime \prime}>0$ and $n$ large enough, where we used that $2 \gamma>1$ and $\delta<2 \gamma-1$.

## Subcritical regime: second moment argument

A second moment argument also gives a lower bound on the size of the largest component in the subcritical case. We proved in Theorem 6.4.1 that, when $\lambda<1$, the probability of observing a connected component of size larger than $I_{\lambda}^{-1} \log n$ is vanishingly small. In the other direction, we get:

Theorem 6.4.11 (Subcritical regime: lower bound on the largest cluster). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{\lambda}{n}$ with $\lambda \in(0,1)$. For all $\kappa \in(0,1)$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right| \leq(1-\kappa) I_{\lambda}^{-1} \log n\right]=o(1) .
$$

Proof. Recall that

$$
B_{k}=\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right|>k\right\}} .
$$

It suffices to prove that with probability $1-o(1)$ we have $B_{k}>0$ when $k=$ $(1-\kappa) I_{\lambda}^{-1} \log n$. To apply the second moment method (Theorem 2.3.2), we need an upper bound on the second moment of $B_{k}$ and a lower bound on its first moment. The following lemma is closely related to Lemma 6.4.10. Exercise 6.12 asks for a proof.

Lemma 6.4.12 (Second moment of $X_{k}$ ). Assume $\lambda<1$. There is a constant $C>0$ such that

$$
\mathbb{E}_{n, p_{n}}\left[B_{k}^{2}\right] \leq\left(\mathbb{E}_{n, p_{n}}\left[B_{k}\right]\right)^{2}+C n k e^{-k I_{\lambda}}, \quad \forall k \geq 0
$$

Lemma 6.4.13 (First moment of $X_{k}$ ). Let $k_{n}=(1-\kappa) I_{\lambda}^{-1} \log n$. Then, for any $\beta \in(0, \kappa)$ we have that

$$
\mathbb{E}_{n, p_{n}}\left[B_{k_{n}}\right]=\Omega\left(n^{\beta}\right),
$$

for $n$ large enough.
Proof. By Lemma 6.4.3,

$$
\begin{align*}
\mathbb{E}_{n, p_{n}}\left[B_{k_{n}}\right] & =n \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{1}\right|>k_{n}\right] \\
& \geq n \mathbb{P}\left[W_{\lambda}>k_{n}\right]-O\left(\left\lceil k_{n}\right\rceil^{2}\right) . \tag{6.4.32}
\end{align*}
$$

Once again, we use the random-walk representation of the total progeny of a branching process (Theorem 6.2.6). In contrast to the proof of Lemma 6.4.6, we need a lower bound this time. For this purpose, we use the explicit expression for the law of the total progeny $W_{\lambda}$ from Example 6.2.7

$$
\mathbb{P}\left[W_{\lambda}>k_{n}\right]=\sum_{t>k_{n}} \frac{1}{t} e^{-\lambda t} \frac{(\lambda t)^{t-1}}{(t-1)!}
$$

Using Stirling's formula (see Appendix A) and (6.4.3), we note that

$$
\begin{aligned}
\frac{1}{t} e^{-\lambda t} \frac{(\lambda t)^{t-1}}{(t-1)!} & =e^{-\lambda t} \frac{(\lambda t)^{t-1}}{t!} \\
& =e^{-\lambda t} \frac{(\lambda t)^{t}}{\lambda t(t / e)^{t} \sqrt{2 \pi t}(1+o(1))} \\
& =\frac{1-o(1)}{\lambda \sqrt{2 \pi t^{3}}} \exp (-t \lambda+t \log \lambda+t) \\
& =\frac{1-o(1)}{\lambda \sqrt{2 \pi t^{3}}} \exp \left(-t I_{\lambda}\right) .
\end{aligned}
$$

Hence, for any $\varepsilon>0$,

$$
\begin{aligned}
\mathbb{P}\left[W_{\lambda}>k_{n}\right] & \geq \lambda^{-1} \sum_{t>k_{n}} \exp \left(-t\left(I_{\lambda}+\varepsilon\right)\right) \\
& =\Omega\left(\exp \left(-k_{n}\left(I_{\lambda}+\varepsilon\right)\right)\right),
\end{aligned}
$$

for $n$ large enough. For any $\beta \in(0, \kappa)$, taking $\varepsilon$ small enough we have

$$
\begin{aligned}
n \mathbb{P}\left[W_{\lambda}>k_{n}\right] & =\Omega\left(n \exp \left(-k_{n}\left(I_{\lambda}+\varepsilon\right)\right)\right) \\
& =\Omega\left(\exp \left(\left\{1-(1-\kappa) I_{\lambda}^{-1}\left(I_{\lambda}+\varepsilon\right)\right\} \log n\right)\right) \\
& =\Omega\left(n^{\beta}\right) .
\end{aligned}
$$

Plugging this back into (6.4.32) gives

$$
\mathbb{E}_{n, p_{n}}\left[B_{k_{n}}\right]=\Omega\left(n^{\beta}\right),
$$

which proves the claim.
We return to the proof of Theorem 6.4.11. Let again $k_{n}=(1-\kappa) I_{\lambda}^{-1} \log n$. By the second moment method and Lemmas 6.4.12 and 6.4.13,

$$
\begin{aligned}
\mathbb{P}_{n, p_{n}}\left[B_{k_{n}}>0\right] & \geq \frac{\left(\mathbb{E} B_{k_{n}}\right)^{2}}{\mathbb{E}\left[B_{k_{n}}^{2}\right]} \\
& \geq\left(1+\frac{O\left(n k_{n} e^{-k_{n} I_{\lambda}}\right)}{\Omega\left(n^{2 \beta}\right)}\right)^{-1} \\
& =\left(1+\frac{O\left(n k_{n} e^{(\kappa-1) \log n}\right)}{\Omega\left(n^{2 \beta}\right)}\right)^{-1} \\
& =\left(1+\frac{O\left(k_{n} n^{\kappa}\right)}{\Omega\left(n^{2 \beta}\right)}\right)^{-1} \\
& \rightarrow 1,
\end{aligned}
$$

for $\beta$ close enough to $\kappa$. That proves the claim.

### 6.4.4 Critical case via martingales

It remains to consider the critical case, that is, when $\lambda=1$. As it turns out, the model goes through a "double jump": as $\lambda$ crosses 1 , the largest cluster size goes from order $\log n$ to order $n^{2 / 3}$ to order $n$. Here we use martingale methods to show the following.

Theorem 6.4.14 (Critical case: upper bound on the largest cluster). Let $G_{n} \sim$ $\mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{1}{n}$. For all $\kappa>1$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right|>\kappa n^{2 / 3}\right] \leq \frac{C}{\kappa^{3 / 2}},
$$

for some constant $C>0$.

Remark 6.4.15. One can also derive a lower bound on the probability that $\left|\mathcal{C}_{\max }\right|>$ $\kappa n^{2 / 3}$ for some $\kappa>0$ [ER60]. Exercise 6.20 provides a sketch based on counting tree components; the combinatorial approach has the advantage of giving insights into the structure of the graph (see [Bol01] for more on this). See also [NP10] for a martingale proof of the lower bound as well as a better upper bound.

The key technical bound is the following.
Lemma 6.4.16. Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{1}{n}$ and let $\mathcal{C}_{v}$ be the connected component of $v \in[n]$. There are constants $c, c^{\prime}>0$ such that for all $k \geq c$

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|>k\right] \leq \frac{c^{\prime}}{\sqrt{k}}
$$

Before we establish the lemma, we prove the theorem assuming it.
Proof of Theorem 6.4.14. Recall that

$$
B_{k}=\sum_{v \in[n]} \mathbf{1}_{\left\{\left|\mathcal{C}_{v}\right|>k\right\}}
$$

Take

$$
k_{n}:=\kappa n^{2 / 3}
$$

By Markov's inequality (Theorem 2.1.1) and Lemma 6.4.16,

$$
\begin{aligned}
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{\max }\right|>k_{n}\right] & \leq \mathbb{P}_{n, p_{n}}\left[B_{k_{n}}>k_{n}\right] \\
& \leq \frac{\mathbb{E}_{n, p_{n}}\left[B_{k_{n}}\right]}{k_{n}} \\
& =\frac{n \mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|>k_{n}\right]}{k_{n}} \\
& \leq \frac{n c^{\prime}}{k_{n}^{3 / 2}} \\
& \leq \frac{C}{\kappa^{3 / 2}}
\end{aligned}
$$

for some constant $C>0$.
It remains to prove the lemma.
Proof of Lemma 6.4.16. Once again, we use the exploration process defined in Section 6.4.2 started at $v$. Let $\left(\mathcal{F}_{t}\right)$ be the corresponding filtration and let $A_{t}=\left|\mathcal{A}_{t}\right|$ be the size of the active set.

Domination by a martingale Recalling (6.4.6), we define

$$
\begin{equation*}
M_{t}:=M_{t-1}+\left[-1+\widetilde{X}_{t}\right], \tag{6.4.33}
\end{equation*}
$$

with $M_{0}:=1$ and $\left(\widetilde{X}_{t}\right)$ are i.i.d. $\operatorname{Bin}(n, 1 / n)$. We couple $\left(A_{t}\right)$ and $\left(M_{t}\right)$ through the equation (6.4.7) by letting

$$
\widetilde{X}_{t}=\sum_{i=1}^{n} I_{t, i} .
$$

In particular $M_{t} \geq A_{t}$ for all $t$.
Furthermore, we have

$$
\mathbb{E}\left[M_{t} \mid \mathcal{F}_{t-1}\right]=M_{t-1}-1+n \frac{1}{n}=M_{t-1} .
$$

So $\left(M_{t}\right)$ is a martingale. We define the stopping time

$$
\tilde{\tau}_{0}:=\inf \left\{t \geq 0: M_{t}=0\right\} .
$$

Recalling that

$$
\tau_{0}=\inf \left\{t \geq 0: A_{t}=0\right\}=\left|\mathcal{C}_{v}\right|,
$$

by Lemma 6.2.1, we have $\tilde{\tau}_{0} \geq \tau_{0}=\left|\mathcal{C}_{v}\right|$ almost surely. So

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|>k\right] \leq \mathbb{P}\left[\tilde{\tau}_{0}>k\right] .
$$

The tail of $\tilde{\tau}_{0} \quad$ To bound the tail of $\tilde{\tau}_{0}$, we introduce a modified stopping time. For $h>0$, let

$$
\tau_{h}^{\prime}:=\inf \left\{t \geq 0: M_{t}=0 \text { or } M_{t} \geq h\right\} .
$$

We will use the inequality

$$
\mathbb{P}\left[\tilde{\tau}_{0}>k\right]=\mathbb{P}\left[M_{t}>0, \forall t \leq k\right] \leq \mathbb{P}\left[\tau_{h}^{\prime}>k\right]+\mathbb{P}\left[M_{\tau_{h}^{\prime}} \geq h\right],
$$

and we will choose $h$ below to minimize the rightmost expression (or, more specifically, an upper bound on it). The rest of the analysis is similar to the gambler's ruin problem in Example 3.1.41, with some slight complications arising from the fact that the process is not nearest-neighbor.

We note that by the exponential tail of hitting times on finite state spaces (Lemma 3.1.25), the stopping time $\tau_{h}^{\prime}$ is almost surely finite and, in fact, has a finite expectation. By two applications of Markov's inequality,

$$
\mathbb{P}\left[M_{\tau_{h}^{\prime}} \geq h\right] \leq \frac{\mathbb{E}\left[M_{\tau_{h}^{\prime}}\right]}{h}
$$

and

$$
\mathbb{P}\left[\tau_{h}^{\prime}>k\right] \leq \frac{\mathbb{E} \tau_{h}^{\prime}}{k}
$$

We bound the expectations on the right-hand sides.

Bounding $\mathbb{E} M_{\tau_{h}^{\prime}}$ and $\mathbb{E} \tau_{h}^{\prime} \quad$ To compute $\mathbb{E} M_{\tau_{h}^{\prime}}$, we use the optional stopping theorem in the uniformly bounded case (Theorem 3.1.38 (ii)) to the stopped process ( $M_{t \wedge \tau_{h}^{\prime}}$ ) (which is also a martingale by Lemma 3.1.37) to get that

$$
\mathbb{E}\left[M_{\tau_{h}^{\prime}}\right]=\mathbb{E}\left[M_{0}\right]=1
$$

We conclude that

$$
\begin{equation*}
\mathbb{P}\left[M_{\tau_{h}^{\prime}} \geq h\right] \leq \frac{1}{h} \tag{6.4.34}
\end{equation*}
$$

To compute $\mathbb{E} \tau_{h}^{\prime}$, we use a different martingale (adapted from Example 3.1.31), specifically

$$
L_{t}:=M_{t}^{2}-\sigma^{2} t,
$$

where we let $\sigma^{2}:=n \frac{1}{n}\left(1-\frac{1}{n}\right)=\left(1-\frac{1}{n}\right)$, which is $\geq \frac{1}{2}$ when $n \geq 2$. To see that $\left(L_{t}\right)$ is a martingale, note that by taking out what is known (Lemma B.6.13) and using the fact that $\left(M_{t}\right)$ is itself a martingale

$$
\begin{aligned}
\mathbb{E}\left[L_{t} \mid \mathcal{F}_{t-1}\right] & =\mathbb{E}\left[\left(M_{t-1}+\left(M_{t}-M_{t-1}\right)\right)^{2}-\sigma^{2} t \mid \mathcal{F}_{t-1}\right] \\
& =\mathbb{E}\left[M_{t-1}^{2}+2 M_{t-1}\left(M_{t}-M_{t-1}\right)+\left(M_{t}-M_{t-1}\right)^{2}-\sigma^{2} t \mid \mathcal{F}_{t-1}\right] \\
& =M_{t-1}^{2}+2 M_{t-1} \cdot 0+\sigma^{2}-\sigma^{2} t \\
& =L_{t-1} .
\end{aligned}
$$

By Lemma 3.1.37, the stopped process ( $L_{t \wedge \tau_{h}^{\prime}}$ ) is also a martingale; and it has bounded increments since

$$
\begin{aligned}
\left|L_{(t+1) \wedge \tau_{h}^{\prime}}-L_{t \wedge \tau_{h}^{\prime}}\right| & \leq\left|M_{(t+1) \wedge \tau_{h}^{\prime}}^{2}-M_{t \wedge \tau_{h}^{\prime}}^{2}\right|+\sigma^{2} \\
& \leq\left|\left(-1+\widetilde{X}_{t+1}\right)^{2}+2 h\right|-1+\widetilde{X}_{t+1}| |+\sigma^{2} \\
& \leq n^{2}+2 h n+1 .
\end{aligned}
$$

We use the optional stopping theorem in the bounded increments case (Theorem 3.1.38 (iii)) on ( $L_{t \wedge \tau_{h}^{\prime}}$ ) to get

$$
\mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2}-\sigma^{2} \tau_{h}^{\prime}\right]=\mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2}\right]-\sigma^{2} \mathbb{E} \tau_{h}^{\prime}=1
$$

After rearranging (6.4.35)

$$
\begin{equation*}
\mathbb{E} \tau_{h}^{\prime} \leq \frac{1}{\sigma^{2}} \mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2}\right] \leq 2 \mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2}\right], \tag{6.4.35}
\end{equation*}
$$

where we used the fact that $\sigma^{2} \geq 1 / 2$.
To bound $\mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2}\right]$, we need to control by how much the process "overshoots" $h$. A stochastic domination argument gives the desired bound; Exercise 6.21 asks for a proof.

Lemma 6.4.17 (Overshoot bound). Let $f$ be an increasing function and $W \sim$ $\operatorname{Bin}(n, 1 / n)$. Then

$$
\mathbb{E}\left[f\left(M_{\tau_{h}^{\prime}}-h\right) \mid M_{\tau_{h}^{\prime}} \geq h\right] \leq \mathbb{E}[f(W)] .
$$

The lemma implies that

$$
\begin{aligned}
\mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2} \mid M_{\tau_{h}^{\prime}} \geq h\right] & =\mathbb{E}\left[\left(M_{\tau_{h}^{\prime}}-h\right)^{2}+2\left(M_{\tau_{h}^{\prime}}-h\right) h+h^{2} \mid M_{\tau_{h}^{\prime}} \geq h\right] \\
& \leq\left(\sigma^{2}+1\right)+2 h+h^{2} \\
& \leq 4 h^{2} .
\end{aligned}
$$

Plugging back into (6.4.35) gives

$$
\mathbb{E} \tau_{h}^{\prime} \leq 2\left\{\frac{1}{h} \mathbb{E}\left[M_{\tau_{h}^{\prime}}^{2} \mid M_{\tau_{h}^{\prime}} \geq h\right]\right\} \leq 8 h,
$$

where we used (6.4.34).
Putting everything together Finally take $h:=\sqrt{\frac{k}{8}}$. Putting everything together

$$
\mathbb{P}_{n, p_{n}}\left[\left|\mathcal{C}_{v}\right|>k\right] \leq \mathbb{P}\left[\tilde{\tau}_{0}>k\right] \leq \mathbb{P}\left[\tau_{h}^{\prime}>k\right]+\mathbb{P}\left[M_{\tau_{h}^{\prime}} \geq h\right] \leq \frac{8 h}{k}+\frac{1}{h}=2 \sqrt{\frac{8}{k}}
$$

That concludes the proof.

### 6.4.5 $\triangleright$ Encore: random walk on the Erdös-Rényi graph

So far in this section we have used techniques from all chapters of the book-with the exception of Chapter 5 . Not to be outdone, we discuss one last result that will make use of spectral techniques. We venture a little further down the evolution of the Erdôs-Rényi graph model to the connected regime. Specifically, recall from Section 2.3.2 that $G_{n}=\left(V_{n}, E_{n}\right) \sim \mathbb{G}_{n, p_{n}}$ is connected with probability $1-o(1)$ when $n p_{n}=\omega(\log n)$.

We show in that regime that lazy simple random walk $\left(X_{t}\right)$ on $G_{n}$ "mixes fast." Recall from Example 1.1.29 that, when the graph is connected, the corresponding transition matrix $P$ is reversible with respect to the stationary distribution

$$
\pi(v):=\frac{\delta(v)}{2\left|E_{n}\right|},
$$

where $\delta(v)$ is the degree of $v$. For a fixed $\varepsilon>0$, the mixing time (see Definition 1.1.35) is

$$
\mathrm{t}_{\text {mix }}(\varepsilon)=\inf \{t \geq 0: d(t) \leq \varepsilon\},
$$

where

$$
d(t)=\sup _{x \in V_{n}}\left\|P^{t}(x, \cdot)-\pi(\cdot)\right\|_{\mathrm{TV}}
$$

By convention, we let $\mathrm{t}_{\text {mix }}(\varepsilon)=+\infty$ if the graph is not connected. Our main result is the following.

Theorem 6.4.18 (Mixing on a connected Erdős-Rényi graph). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ with $n p_{n}=\omega(\log n)$. With probability $1-o(1)$, the mixing time is $O(\log n)$.

Edge expansion We use Cheeger's inequality (Theorem 5.3.5) which, recall, states that

$$
\gamma \geq \frac{\Phi_{*}^{2}}{2}
$$

where $\gamma$ is the spectral gap of $P$ (see Definition 5.2.11) and

$$
\Phi_{*}=\min \left\{\Phi_{\mathrm{E}}(S ; c, \pi): S \subseteq V, 0<\pi(S) \leq \frac{1}{2}\right\},
$$

is the edge expansion constant (see Definition 5.3.2), with

$$
\Phi_{\mathrm{E}}(S ; c, \pi)=\frac{c\left(S, S^{c}\right)}{\pi(S)}
$$

for a subset of vertices $S \subseteq V_{n}$. Here, for a pair of vertices $x, y$ connected by an edge,

$$
c(x, y)=\pi(x) P(x, y)=\frac{\delta(x)}{2\left|E_{n}\right|} \frac{1}{\delta(x)}=\frac{1}{2\left|E_{n}\right|} .
$$

Hence

$$
c\left(S, S^{c}\right)=\frac{\left|E\left(S, S^{c}\right)\right|}{2\left|E_{n}\right|}
$$

where $E\left(S, S^{c}\right)$ is the set of edges between $S$ and $S^{c}$. Similarly,

$$
\pi(S)=\frac{\sum_{x \in S} \delta(x)}{2\left|E_{n}\right|}
$$

The numerator is referred to as the volume of $S$ and we use the notation $\operatorname{vol}(S)=$ $\sum_{x \in S} \delta(x)$. So

$$
\frac{c\left(S, S^{c}\right)}{\pi(S)}=\frac{\left|E\left(S, S^{c}\right)\right|}{\operatorname{vol}(S)} .
$$

Because the random walk is lazy, the spectral gap is equal to the absolute spectral gap (see Definition 5.2.11), and as a consequence the relaxation time (see Definition 5.2.12) is

$$
\mathrm{t}_{\mathrm{rel}}=\gamma^{-1}
$$

Using Theorem 5.2.14, we get

$$
\begin{equation*}
\mathrm{t}_{\text {mix }}(\varepsilon) \leq \log \left(\frac{1}{\varepsilon \pi_{\text {min }}}\right) \mathrm{t}_{\text {rel }} \leq \log \left(\frac{1}{\varepsilon \pi_{\text {min }}}\right) \frac{2}{\Phi_{*}^{2}}, \tag{6.4.37}
\end{equation*}
$$

where

$$
\pi_{\min }=\min _{x} \pi(x)=\min _{x} \frac{\delta(x)}{2\left|E_{n}\right|}=\frac{\min _{x} \delta(x)}{\sum_{y} \delta(y)} .
$$

So our main task is to bound $\delta(x)$ and $\left|E\left(S, S^{c}\right)\right|$ with high probability. We do this next.

Bounding the degrees In fact, we have already done half the work. Indeed in Example 2.4.18 we studied the maximum degree of $G_{n}$

$$
D_{n}=\max _{v \in V_{n}} \delta(v),
$$

in the regime $n p_{n}=\omega(\log n)$. We showed that for any $\zeta>0$, as $n \rightarrow+\infty$,

$$
\mathbb{P}\left[\left|D_{n}-n p_{n}\right| \leq 2 \sqrt{(1+\zeta) n p_{n} \log n}\right] \rightarrow 1
$$

The proof of that result actually shows something stronger: all degrees satisfy the inequality simultaneously, that is,

$$
\begin{equation*}
\mathbb{P}\left[\forall v \in V_{n},\left|\delta(v)-n p_{n}\right| \leq 2 \sqrt{(1+\zeta) n p_{n} \log n}\right]=1-o(1) . \tag{6.4.38}
\end{equation*}
$$

We will use the fact that $2 \sqrt{(1+\zeta) n p_{n} \log n}=o\left(n p_{n}\right)$ when $n p_{n}=\omega(\log n)$. In essence, all degrees are roughly $n p_{n}$. That implies the following claims.

Lemma 6.4.19 (Bounds on stationary distribution and volume). The following hold with probability $1-o(1)$.
(i) The smallest stationary probability satisfies

$$
\pi_{\min } \geq \frac{1-o(1)}{n} .
$$

(ii) For any set of vertices $S \subseteq V_{n}$ with $|S|>2 n / 3$, we have

$$
\pi(S)>\frac{1}{2} .
$$

(iii) For any set of vertices $S \subseteq V_{n}$ with $s:=|S|$

$$
\operatorname{vol}(S)=\operatorname{snp_{n}}(1+o(1))
$$

Proof. We assume that the event in (6.4.38) holds.
For (i), that means

$$
\pi_{\min } \geq \frac{n p_{n}-2 \sqrt{(1+\zeta) n p_{n} \log n}}{n\left(n p_{n}+2 \sqrt{\left.(1+\zeta) n p_{n} \log n\right)}\right.}=\frac{1}{n}(1-o(1))
$$

when $n p_{n}=\omega(\log n)$.
For (ii), we get

$$
\pi(S)=\frac{\sum_{x \in S} \delta(x)}{\sum_{x \in V_{n}} \delta(x)} \geq \frac{|S|\left(n p_{n}-2 \sqrt{(1+\zeta) n p_{n} \log n}\right)}{n\left(n p_{n}+2 \sqrt{(1+\zeta) n p_{n} \log n}\right)}>\frac{2}{3}(1-o(1))
$$

Finally (iii) follows similarly.
Bounding the cut size An application of Bernstein's inequality (Theorem 2.4.17) gives the following bound.

Lemma 6.4.20 (Bound on the edge expansion). With probability $1-o(1)$,

$$
\Phi_{*}=\Omega(1) .
$$

Proof. By the definition of $\Phi_{*}$ and Lemma 6.4.19 (ii), we can restrict ourselves to sets $S$ of size at most $2 n / 3$. Let $S$ be such a set with $s=|S|$. Then $\left|E\left(S, S^{c}\right)\right|$ is $\operatorname{Bin}\left(s(n-s), p_{n}\right)$. By Bernstein's inequality with $c=1$ and $\nu_{i}=p_{n}\left(1-p_{n}\right)$,

$$
\mathbb{P}_{n, p_{n}}\left[\left|E\left(S, S^{c}\right)\right| \leq s(n-s) p_{n}-\beta\right] \leq \exp \left(-\frac{\beta^{2}}{4 s(n-s) p_{n}\left(1-p_{n}\right)}\right)
$$

for $\beta \leq s(n-s) p_{n}\left(1-p_{n}\right)$. We take $\beta=\frac{1}{2} s(n-s) p_{n}$ and get

$$
\mathbb{P}_{n, p_{n}}\left[\left|E\left(S, S^{c}\right)\right| \leq \frac{1}{2} s(n-s) p_{n}\right] \leq \exp \left(-\frac{s(n-s) p_{n}}{16\left(1-p_{n}\right)}\right) .
$$

By a union bound over all sets of size $s$ and using the fact that $\binom{n}{s} \leq\left(\frac{n e}{s}\right)^{s}$ (see Appendix A), there is a constant $C>0$ such that

$$
\begin{aligned}
\mathbb{P}_{n, p_{n}}[\exists S,|S|=s, & \left.\left|E\left(S, S^{c}\right)\right| \leq \frac{1}{2} s(n-s) p_{n}\right] \\
& \leq\binom{ n}{s} \exp \left(-\frac{s(n-s) p_{n}}{16\left(1-p_{n}\right)}\right) \\
& \leq \exp \left(-s \frac{n p_{n}}{48}+s \log (n e / s)\right) \\
& \leq \exp \left(-C \sin p_{n}\right)
\end{aligned}
$$

for $n$ large enough, where we also used that $n-s \geq n / 3$ and $n p_{n}=\omega(\log n)$. Summing over $s$ gives, for a constant $C^{\prime}>0$,
$\mathbb{P}_{n, p_{n}}\left[\exists S, 1 \leq|S| \leq 2 n / 3,\left|E\left(S, S^{c}\right)\right| \leq \frac{1}{2}|S|(n-|S|) p_{n}\right] \leq C^{\prime} \exp \left(-C n p_{n}\right)$,
which goes to 0 as $n \rightarrow+\infty$.
Using (6.4.36) and Lemma 6.4.19 (iii), any set $S$ such that $\left|E\left(S, S^{c}\right)\right|>$ $\frac{1}{2}|S|(n-|S|) p_{n}$ has edge expansion

$$
\Phi_{\mathrm{E}}(S ; c, \pi) \geq \frac{\frac{1}{2}|S|(n-|S|) p_{n}}{|S| n p_{n}(1+o(1))} \geq \frac{1}{6}(1-o(1)) .
$$

That proves the claim.
Proof of the theorem Finally, we are ready to prove the main result.
Proof of Theorem 6.4.18. Plugging Lemma 6.4.19 (i) and Lemma 6.4.20 into (6.4.37) gives

$$
\mathrm{t}_{\mathrm{mix}}(\varepsilon) \leq \log \left(\frac{1}{\varepsilon \pi_{\min }}\right) \frac{2}{\Phi_{*}^{2}} \leq C^{\prime \prime} \log \left(\varepsilon^{-1} n(1+o(1))\right)=O(\log n)
$$

for some constant $C^{\prime \prime}>0$.
Remark 6.4.21. A mixing time of $O(\log n)$ in fact holds for lazy simple random walk on $\mathbb{G}_{n, p_{n}}$ when $p_{n}=\frac{\lambda \log n}{n}$ with $\lambda>1$ [CF07]. See also [Dur06, Section 6.5]. Mixing time on the giant component has also been studied. See, e.g., [FR08, BKW14, DKLP11].

## Exercises

Exercise 6.1 (Galton-Watson process: subcritical case). We use Markov's inequality to analyze the subcritical case.
(i) Let $\left(Z_{t}\right)$ be a Galton-Watson process with offspring distribution mean $m<$ 1. Use Markov's inequality (Theorem 2.1.1) to prove that extinction occurs almost surely.
(ii) Prove the equivalent result in the multitype case, that is, prove (6.1.8).

Exercise 6.2 (Galton-Watson process: geometric offspring). Let $\left(Z_{t}\right)$ be a GaltonWatson branching process with geometric offspring distribution (started at 0 ), that is, $p_{k}=p(1-p)^{k}$ for all $k \geq 0$, for some $p \in(0,1)$. Let $q:=1-p$, let $m$ be the mean of the offspring distribution, and let $W_{t}=m^{-t} Z_{t}$.
(i) Compute the probability generating function $f$ of $\left\{p_{k}\right\}_{k \geq 0}$ and the extinction probability $\eta:=\eta_{p}$ as a function of $p$.
(ii) If $G$ is a $2 \times 2$ matrix, define

$$
G(s):=\frac{G_{11} s+G_{12}}{G_{21} s+G_{22}} .
$$

Show that $G(H(s))=(G H)(s)$.
(iii) Assume $m \neq 1$. Use (ii) to derive

$$
f_{t}(s)=\frac{p m^{t}(1-s)+q s-p}{q m^{t}(1-s)+q s-p} .
$$

Deduce that when $m>1$

$$
\mathbb{E}\left[\exp \left(-\lambda W_{\infty}\right)\right]=\eta+(1-\eta) \frac{(1-\eta)}{\lambda+(1-\eta)} .
$$

(iv) Assume $m=1$. Show that

$$
f_{t}(s)=\frac{t-(t-1) s}{t+1-t s}
$$

and deduce that

$$
\mathbb{E}\left[e^{-\lambda Z_{t} / t} \mid Z_{t}>0\right] \rightarrow \frac{1}{1+\lambda}
$$

Exercise 6.3 (Supercritical branching process: infinite line of descent). Let $\left(Z_{t}\right)$ be a supercritical Galton-Watson branching process with offspring distribution $\left\{p_{k}\right\}_{k \geq 0}$. Let $\eta$ be the extinction probability and define $\zeta:=1-\eta$. Let $Z_{t}^{\infty}$ be the number of individuals in the $t$-th generation with an infinite line of descent, i.e., whose descendant subtree is infinite. Denote by $\mathcal{S}$ the event of nonextinction of $\left(Z_{t}\right)$. Define $p_{0}^{\infty}:=0$ and

$$
p_{k}^{\infty}:=\zeta^{-1} \sum_{j \geq k}\binom{j}{k} \eta^{j-k} \zeta^{k} p_{j} .
$$

(i) Show that $\left\{p_{k}^{\infty}\right\}_{k \geq 0}$ is a probability distribution and compute its expectation.
(ii) Show that for any $k \geq 0$

$$
\mathbb{P}\left[Z_{1}^{\infty}=k \mid \mathcal{S}\right]=p_{k}^{\infty} .
$$

[Hint: Condition on $Z_{1}$.]
(iii) Show by induction on $t$ that, conditioned on nonextinction, the process $\left(Z_{t}^{\infty}\right)$ has the same distribution as a Galton-Watson branching process with offspring distribution $\left\{p_{k}^{\infty}\right\}_{k \geq 0}$.

Exercise 6.4 (Multitype branching processes: a special case). Extend Lemma 6.1.20 to the case $\mathbf{S}^{(\mathbf{u})}=\mathbf{0}$. [Hint: Show that $U_{t}=\mathbf{Z}_{0} \mathbf{u}$ for all $t$ almost surely.]

Exercise 6.5 (Galton-Watson: Inverting history). Let

$$
H=\left(X_{1}, \ldots, X_{\tau_{0}}\right),
$$

be the history (see Section 6.2) of the Galton-Watson process $\left(Z_{i}\right)$. Write $Z_{i}$ as a function of $H$, for all $i$.

Exercise 6.6 (Spitzer's lemma). Prove Theorem 6.2.5.
Exercise 6.7 (Sum of Poisson). Let $Q_{1}$ and $Q_{2}$ be independent Poisson random variables with respective means $\lambda_{1}$ and $\lambda_{2}$. Show by direct computation of the convolution that the sum $Q_{1}+Q_{2}$ is Poisson with mean $\lambda_{1}+\lambda_{2}$. [Hint: Recall that $\mathbb{P}\left[Q_{1}=k\right]=e^{-\lambda_{1}} \lambda_{1}^{k} / k!$ for all $k \in \mathbb{Z}_{+}$.]
Exercise 6.8 (Percolation on bounded-degree graphs). Let $G=(V, E)$ be a countable graph such that all vertices have degree bounded by $b+1$ for $b \geq 2$. Let 0 be a distinguished vertex in $G$. For bond percolation on $G$, prove that

$$
p_{\mathrm{c}}(G) \geq p_{\mathrm{c}}\left(\widehat{\mathbb{T}}_{b}\right),
$$

by bounding the expected size of the cluster of 0 . [Hint: Consider self-avoiding paths started at 0.]

Exercise 6.9 (Percolation on $\widehat{\mathbb{T}}_{b}$ : critical exponent of $\theta(p)$ ). Consider bond percolation on the rooted infinite $b$-ary tree $\widehat{\mathbb{T}}_{b}$ with $b>2$. For $\varepsilon \in\left[0,1-\frac{1}{b}\right]$ and $u \in[0,1]$, define

$$
h(\varepsilon, u):=u-\left(\left(1-\frac{1}{b}-\varepsilon\right)(1-u)+\frac{1}{b}+\varepsilon\right)^{b} .
$$

(i) Show that there is a constant $C>0$ not depending on $\varepsilon, u$ such that

$$
\left|h(\varepsilon, u)-b \varepsilon u+\frac{b-1}{2 b} u^{2}\right| \leq C\left(u^{3} \vee \varepsilon u^{2}\right) .
$$

(ii) Use (i) to prove that

$$
\lim _{p \downarrow p_{c}\left(\widehat{\mathbb{T}}_{b}\right)} \frac{\theta(p)}{\left(p-p_{\mathrm{c}}\left(\widehat{\mathbb{T}}_{b}\right)\right)}=\frac{2 b^{2}}{b-1} .
$$

Exercise 6.10 (Percolation on $\widehat{\mathbb{T}}_{2}$ : higher moments of $\left.\left|\mathcal{C}_{0}\right|\right)$. Consider bond percolation on the rooted infinite binary tree $\widehat{\mathbb{T}}_{2}$. For density $p<\frac{1}{2}$, let $Z_{p}$ be an integer-valued random variable with distribution

$$
\mathbb{P}_{p}\left[Z_{p}=\ell\right]=\frac{\ell \mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=\ell\right]}{\mathbb{E}_{p}\left|\mathcal{C}_{0}\right|}, \quad \forall \ell \geq 1
$$

(i) Using the explicit formula for $\mathbb{P}_{p}\left[\left|\mathcal{C}_{0}\right|=\ell\right]$ derived in Section 6.2.4, show that for all $0<a<b<+\infty$

$$
\mathbb{P}_{p}\left[\frac{Z_{p}}{(1 / 4)\left(\frac{1}{2}-p\right)^{-2}} \in[a, b]\right] \rightarrow C \int_{a}^{b} x^{-1 / 2} e^{-x} \mathrm{~d} x,
$$

as $p \uparrow \frac{1}{2}$, for some constant $C>0$.
(ii) Show that for all $k \geq 2$ there is $C_{k}>0$ such that

$$
\lim _{p \uparrow p_{c}\left(\widehat{\mathbb{T}}_{2}\right)} \frac{\mathbb{E}_{p}\left|\mathcal{C}_{0}\right|^{k}}{\left(p_{\mathrm{c}}\left(\widehat{\mathbb{T}}_{2}\right)-p\right)^{-1-2(k-1)}}=C_{k} .
$$

(iii) What happens when $p \downarrow p_{\mathrm{c}}\left(\widehat{\mathbb{T}}_{2}\right)$ ?

Exercise 6.11 (Branching process approximation: improved bound). Let $p_{n}=\frac{\lambda}{n}$ with $\lambda>0$. Let $W_{n, p_{n}}$, respectively $W_{\lambda}$, be the total progeny of a branching process with offspring distribution $\operatorname{Bin}\left(n, p_{n}\right)$, respectively $\operatorname{Poi}(\lambda)$.
(i) Show that

$$
\begin{aligned}
\mid \mathbb{P}\left[W_{n, p_{n}}\right. & \geq k]-\mathbb{P}\left[W_{\lambda} \geq k\right] \mid \\
& \leq \max \left\{\mathbb{P}\left[W_{n, p_{n}} \geq k, W_{\lambda}<k\right], \mathbb{P}\left[W_{n, p_{n}}<k, W_{\lambda} \geq k\right]\right\} .
\end{aligned}
$$

(ii) Couple the two processes step-by-step and use (i) to show that

$$
\left|\mathbb{P}\left[W_{n, p_{n}} \geq k\right]-\mathbb{P}\left[W_{\lambda} \geq k\right]\right| \leq \frac{\lambda^{2}}{n} \sum_{i=1}^{k-1} \mathbb{P}\left[W_{\lambda} \geq i\right] .
$$

Exercise 6.12 (Subcritical Erdős-Rényi: second moment). Prove Lemma 6.4.12.
Exercise 6.13 (Random binary search tree: property (BST)). Show that the (BST) property is preserved by the algorithm described at the beginning of Section 6.3.1.

Exercise 6.14 (Random binary search tree: limit). Consider the equation (6.3.1).
(i) Show that there exists a unique solution greater than 1.
(ii) Prove that the expression on the left-hand side is strictly decreasing at that solution.

Exercise 6.15 (Random binary search tree: height is well-defined). Let $\mathcal{T}$ be an infinite binary tree. Assign an independent $U[0,1]$ random variable $Z_{v}$ to each vertex $v$ in $\mathcal{T}$, set $S_{\rho}=n$ and then recursively from the root down

$$
S_{v^{\prime}}:=\left\lfloor S_{v} Z_{v}\right\rfloor \quad \text { and } \quad S_{v^{\prime \prime}}:=\left\lfloor S_{v}\left(1-Z_{v}\right)\right\rfloor,
$$

where $v^{\prime}$ and $v^{\prime \prime}$ are the left and right descendants of $v$ in $\mathcal{T}$.
(i) Show that, for any $v$, it holds that $S_{v^{\prime}}+S_{v^{\prime \prime}}=S_{v}-1$ almost surely provided $S_{v} \geq 1$.
(ii) Show that, for any $v$, there is almost surely a descendant $w$ of $v$ (not necessarily immediate) such that $S_{w}=1$.
(iii) Let

$$
H_{n}=\max \left\{h: \exists v \in \mathcal{V}_{h}, S_{v}=1\right\},
$$

where $\mathcal{V}_{h}$ is the set of vertices of $\mathcal{T}$ at topological distance $h$ from the root. Show that $H_{n} \leq n$.

Exercise 6.16 (Ising vs. CFN). Let $\mathcal{T}_{h}$ be a rooted complete binary tree with $h$ levels. Fix $0<p<1 / 2$. Assign to each vertex $v$ a state $\sigma_{v} \in\{+1,-1\}$ at random according to the CFN model described in Section 6.3.2. Show that this distribution is equivalent to a ferromagnetic Ising model on $\mathcal{T}_{h}$ and determine the inverse temperature $\beta$ in terms of $p$. [Hint: Write the distribution of the states under the CFN model as a product over the edges.]

Exercise 6.17 (Monotonicity of $\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}$ ). Let $\mu_{h}^{+}, \mu_{h}^{-}$be as in Section 6.3.2. Show that

$$
\left\|\mu_{h+1}^{+}-\mu_{h+1}^{-}\right\|_{\mathrm{TV}} \leq\left\|\mu_{h}^{+}-\mu_{h}^{-}\right\|_{\mathrm{TV}}
$$

[Hint: Use the Markovian nature of the process.]
Exercise 6.18 (Unsolvability: recursion). Prove Lemma 6.3.15.
Exercise 6.19 (Cayley's formula). Let $\left(Z_{t}\right)$ be a Poisson branching process with offspring mean 1 started at $Z_{0}=1$ and let $T$ be the corresponding Galton-Watson tree. Let $W$ be the total of size of the progeny, that is, the number of vertices in $T$. Recall from Example 6.2.7 that

$$
\mathbb{P}[W=n]=\frac{n^{n-1} e^{-n}}{n!}
$$

(i) Given $W=n$, label the vertices of $T$ uniformly at random with the integers $1, \ldots, n$. Show that every rooted labeled tree on $n$ vertices arises with probability $e^{-n} / n$ !. [Hint: Label the vertices as you grow the tree and observe that a lot of terms cancel out or simplify.]
(ii) Derive Cayley's formula: the number of labeled trees on $n$ vertices is $n^{n-2}$.

Exercise 6.20 (Critical regime: tree components). Let $G_{n} \sim \mathbb{G}_{n, p_{n}}$ where $p_{n}=\frac{1}{n}$.
(i) Let $\gamma_{n, k}$ be the expected number of isolated tree components of size $k$ in $G_{n}$. Justify the formula

$$
\gamma_{n, k}=\binom{n}{k} k^{k-2}\left(\frac{1}{n}\right)^{k-1}\left(1-\frac{1}{n}\right)^{k(n-k)+\binom{k}{2}-(k-1)} .
$$

[Hint: We did a related calculation in Section 2.3.2.]
(ii) Show that, if $k=\omega(1)$ and $k=o\left(n^{3 / 4}\right)$,

$$
\gamma_{n, k} \sim n \frac{k^{-5 / 2}}{\sqrt{2 \pi}} \exp \left(-\frac{k^{3}}{6 n^{2}}\right)
$$

(iii) Conclude that for $0<\delta<1$ the expectation of $U$, the number of isolated tree components of size in $\left[(\delta n)^{2 / 3}, n^{2 / 3}\right]$, is $\Omega\left(\delta^{-1}\right)$ as $\delta \rightarrow 0$.
(iv) For $1 \leq k_{1} \leq k_{2} \leq n-k_{1}$, let $\sigma_{n, k_{1}, k_{2}}$ be the expected number of pairs of isolated tree components where the first one has size $k_{1}$ and the second one has size $k_{2}$. Justify the formula

$$
\begin{aligned}
\sigma_{n, k_{1}, k_{2}}= & \binom{n}{k_{1}} k_{1}^{k_{1}-2}\left(\frac{1}{n}\right)^{k_{1}-1}\left(1-\frac{1}{n}\right)^{k_{1}\left(n-k_{1}\right)+\binom{k_{1}}{2}-\left(k_{1}-1\right)} \\
& \times\binom{ n-k_{1}}{k_{2}} k_{2}^{k_{2}-2}\left(\frac{1}{n}\right)^{k_{2}-1}\left(1-\frac{1}{n}\right)^{k_{2}\left(n-\left(k_{1}+k_{2}\right)\right)+\binom{k_{2}}{2}-\left(k_{2}-1\right)},
\end{aligned}
$$

and show that

$$
\sigma_{n, k_{1}, k_{2}} \leq \gamma_{n, k_{1}} \gamma_{n, k_{2}} .
$$

[Hint: You may need to prove that, for $0<a \leq 1 \leq b$, it holds that $1-a b \leq$ $(1-a)^{b}$.]
(v) Prove that $\operatorname{Var}[U]=O(\mathbb{E}[U])$. [Hint: Use (2.1.6), (iv), and (ii).]

Exercise 6.21 (Critical regime: overshoot bound). The goal of this exercise is to prove Lemma 6.4.17. We use the notation of Section 6.4.4.
(i) Let $W, Z \sim \operatorname{Bin}(n, 1 / n)$ and $0 \leq r \leq n$. Show that $W-r$ conditioned on $W \geq r$ is stochastically dominated by $Z$. [Hint: Use the representation of $W$ as a sum of indicators. Thinking of the partial sums as a Markov chain, consider the first time it reaches $r$.]
(ii) Show that $M_{\tau_{h}^{\prime}}-h$ conditioned on $M_{\tau_{h}^{\prime}} \geq h$ is stochastically dominated by $Z$ from (i). [Hint: By the tower property, it suffices to show that

$$
\mathbb{P}\left[M_{\tau_{h}^{\prime}}-h \geq z \mid \tau_{h}^{\prime}=\ell, M_{\ell-1}=h-r, M_{\ell} \geq h\right] \leq \mathbb{P}[Z \geq z],
$$

for the relevant $\ell, r, z$.
(iii) Use (ii) to prove Lemma 6.4.17.

## Bibliographic remarks

Section 6.1 See [Dur 10, Section 5.3.4] for a quick introduction to branching processes. A more detailed overview relating to its use in discrete probability can be found in [vdH17, Chapter 3]. A classical reference on branching processes is [AN04]. The Kesten-Stigum Theorem is due to Kesten and Stigum [KS66b]. Our proof of a weaker version with the second moment condition follows [Dur 10, Example 5.4.3]. Section 6.1.4 is based loosely on [AN04, Chapter V]. A proof of Theorem 6.1.18 can be found in [Har63]. A good reference for the PerronFrobenius Theorem (Theorem 6.1.17 as well as more general versions) is [HJ13, Chapter 8]. The central limit theorem for $\rho \geq \lambda^{2}$ referred to at the end of Section 6.1.4 is due to Kesten and Stigum [KS66a, KS67]. The critical percolation threshold for percolation on Galton-Watson trees is due to R. Lyons [Lyo90].

Section 6.2 The exploration process in Section 6.2.1 dates back to [ML86] and [Kar90]. The hitting-time theorem (Theorem 6.2.5) in the case $\ell=1$ was first proved in [Ott49]. For alternative proofs, see for example [vdHK08] or [Wen75]. Spitzer's combinatorial lemma (Lemma 6.2.8) is from [Spi56]. See also [Fel71, Section XII.6]. The presentation in Section 6.2.4 follows [vdH10]. See also [Dur85].

Section 6.3 Section 6.3.1 follows [Dev98, Section 2.1] from the excellent volume [HMRAR98]. Section 6.3.2 is partly a simplified version of [BCMR06]. Further applications in phylogenetics, specifically to the sample complexity of phylogeny inference algorithms, can be found in, for example, [Mos04, Mos03, Roc 10, DMR11, RS17]. The reconstruction problem also has applications in community detection [MNS15b]. See [Abb18] for a survey.

Section 6.4 The phase transtion of the Erdős-Rényi graph model was first studied in [ER60]. For much more, see for example [vdH17, Chapter 4], [JLR11, Chapter 5] and [Bol01, Chapter 6]. In particular a central limit theorem for the giant component, proved by several authors including Martin-Löf [ML98], Pittel [Pit90], and Barraez, Boucheron, and de la Vega [BBFdIV00], is established in [vdH17, Section 4.5]. Section 6.4.4 is based on [NP10]. See also [Per09, Sections 2 and 3]. Much more is known about the critical regime; see, e.g., [Ald97, Bol84, Lu90, LuPW94]. Section 6.4.5 is based partly on [Dur06, Section 6.5]. For a lot more on random walk on random graphs (not just Erdős-Rényi), see [Dur06, Chapter 6]. For more on the spectral properties of random graphs, see [CL06].

## Appendix A

## Useful combinatorial formulas

Recall the following facts about factorials and binomial coefficients:

$$
\begin{gathered}
\frac{n^{n}}{e^{n-1}} \leq n!\leq \frac{n^{n+1}}{e^{n-1}} \\
\frac{n^{k}}{k^{k}} \leq\binom{ n}{k} \leq \frac{e^{k} n^{k}}{k^{k}}, \\
(x+y)^{n}=\sum_{k=0}^{n}\binom{n}{k} x^{k} y^{n-k} \\
\sum_{k=0}^{d}\binom{n}{k} \leq\left(\frac{e n}{d}\right)^{d} \\
n!\sim \sqrt{2 \pi n}\left(\frac{n}{e}\right)^{n} \\
\binom{2 n}{n}=(1+o(1)) \frac{4^{n}}{\sqrt{\pi n}}
\end{gathered}
$$

and

$$
\log \binom{n}{k}=(1+o(1)) n H(k / n)
$$

where $H(p):=-p \log p-(1-p) \log (1-p)$. The third one is the binomial theorem. The fifth one is Stirling's formula.

[^11]
## Appendix B

## Measure-theoretic foundations

This appendix contains relevant background on measure-theoretic probability. We follow closely the highly recommended [Wi191]. Missing proofs (and a lot more details and examples) can be found there. Another excellent textbook on this topic is [Dur10].

## B. 1 Probability spaces

Let $S$ be a set. In general it turns out that we cannot assign a probability to every subset of $S$. Here we discuss "well-behaved" collections of subsets. First an algebra on $S$ is a collection of subsets stable under finitely many set operations.

Definition B.1.1 (Algebra on $S$ ). A collection $\Sigma_{0}$ of subsets of $S$ is an algebra on $S$ if the following conditions hold:
(i) $S \in \Sigma_{0}$;
(ii) $F \in \Sigma_{0}$ implies $F^{c} \in \Sigma_{0}$;
(iii) $F, G \in \Sigma_{0}$ implies $F \cup G \in \Sigma_{0}$.

That, of course, implies that the empty set as well as all pairwise intersections are also in $\Sigma_{0}$. The collection $\Sigma_{0}$ is an actual algebra (i.e., a vector space with a bilinear product) with the symmetric difference as its "sum," the intersection as its "product" and the underlying field being the field with two elements.

[^12]Example B.1.2. On $\mathbb{R}$, sets of the form

$$
\bigcup_{i=1}^{k}\left(a_{i}, b_{i}\right]
$$

where the union is disjoint with $k<+\infty$ and $-\infty \leq a_{i} \leq b_{i} \leq+\infty$ form an algebra.

Finite set operations are not enough for our purposes. For instance, we want to be able to take limits. A $\sigma$-algebra is stable under countably many set operations.

Definition B.1.3 ( $\sigma$-algebra on $S$ ). A collection $\Sigma$ of subsets of $S$ is a $\sigma$-algebra on $S$ (or $\sigma$-field on $S$ ) if
(i) $S \in \Sigma$;
(ii) $F \in \Sigma$ implies $F^{c} \in \Sigma$;
(iii) $F_{n} \in \Sigma$, $\forall n$ implies $\cup_{n} F_{n} \in \Sigma$.

Example B.1.4. $2^{S}$ is a trivial example.
To give a nontrivial example, we need the following definition. We begin with a lemma.

Lemma B.1.5 (Intersection of $\sigma$-algebras). Let $\mathcal{F}_{i}, i \in I$, be $\sigma$-algebras on $S$ where $I$ is arbitrary. Then $\cap_{i} \mathcal{F}_{i}$ is a $\sigma$-algebra.

Proof. We prove only one of the conditions. The other ones are similar. Suppose $A \in \mathcal{F}_{i}$ for all $i$. Then $A^{c}$ is in $\mathcal{F}_{i}$ for all $i$ since each $\mathcal{F}_{i}$ is itself a $\sigma$-algebra.

Definition B.1.6 ( $\sigma$-algebra generated by $\mathcal{C}$ ). Let $\mathcal{C}$ be a collection of subsets of $S$. Then we let $\sigma(\mathcal{C})$ be the smallest $\sigma$-algebra containing $\mathcal{C}$, defined as the intersection of all such $\sigma$-algebras (including in particular $2^{S}$ ).

Example B.1.7. The smallest $\sigma$-algebra containing all open sets in $\mathbb{R}$, denoted $\mathcal{B}(\mathbb{R})$, is called the Borel $\sigma$-algebra. This is a non-trivial $\sigma$-algebra in the sense that it can be proved that there exist subsets of $\mathbb{R}$ that are not in $\mathcal{B}$, but that any "reasonable" set is in $\mathcal{B}$. In particular, it contains the algebra in Example B.1.2.

Example B.1.8. The $\sigma$-algebra generated by the algebra in Example B.1.2 is $\mathcal{B}(\mathbb{R})$. This follows from the fact that all open sets of $\mathbb{R}$ can be written as a countable union of open intervals. (Indeed, for $x \in O$ an open set, let $I_{x}$ be the largest open interval contained in $O$ and containing $x$. If $I_{x} \cap I_{y} \neq \emptyset$ then $I_{x}=I_{y}$ by maximality (i.e., take the union). Then $O=\cup_{x} I_{x}$ and there are only countably many disjoint ones because each one contains a rational.)

We now define measures.
Definition B.1.9 (Additivity and $\sigma$-additivity). A non-negative set function on an algebra $\Sigma_{0}$

$$
\mu_{0}: \Sigma_{0} \rightarrow[0,+\infty]
$$

is additive if
(i) $\mu_{0}(\emptyset)=0$;
(ii) $F, G \in \Sigma_{0}$ with $F \cap G=\emptyset$ implies $\mu_{0}(F \cup G)=\mu_{0}(F)+\mu_{0}(G)$.

Moreover $\mu_{0}$ is said to be $\sigma$-additive if condition (ii) is true for any countable collection of disjoint sets whose union is in $\Sigma_{0}$, that is, if $F_{n} \in \Sigma_{0}, n \geq 0$, all pairwise disjoint with $\cup_{n} F_{n} \in \Sigma_{0}$, then $\mu_{0}\left(\cup_{n} F_{n}\right)=\sum_{n} \mu_{0}\left(F_{n}\right)$.

Example B.1.10. For the algebra in the Example B.1.2, the set function

$$
\lambda_{0}\left(\bigcup_{i=1}^{k}\left(a_{i}, b_{i}\right]\right)=\sum_{i=1}^{k}\left(b_{i}-a_{i}\right)
$$

is additive. (In fact, it is also $\sigma$-additive. We will show this later.)
Definition B.1.11 (Measure space). Let $\Sigma$ be a $\sigma$-algebra on $S$. Then $(S, \Sigma)$ is a measurable space. A $\sigma$-additive function $\mu$ on $\Sigma$ is called a measure and $(S, \Sigma, \mu)$ is called a measure space.
Definition B.1.12 (Probability space). If $(\Omega, \mathcal{F}, \mathbb{P})$ is a measure space with $\mathbb{P}(\Omega)=$
probability space 1 then $\mathbb{P}$ is called a probability measure and $(\Omega, \mathcal{F}, \mathbb{P})$ is called a probability space (or probability triple).

The sets in $\mathcal{F}$ are referred to as events.
To define a measure on $\mathcal{B}(\mathbb{R})$ we need the following tools from abstract measure theory.

Theorem B.1.13 (Caratheodory's extension theorem). Let $\Sigma_{0}$ be an algebra on $S$ and let $\Sigma=\sigma\left(\Sigma_{0}\right)$. If $\mu_{0}$ is $\sigma$-additive on $\Sigma_{0}$ then there exists a measure $\mu$ on $\Sigma$ that agrees with $\mu_{0}$ on $\Sigma_{0}$.

If in addition $\mu_{0}$ is finite, the next lemma implies that the extension is unique.
Lemma B.1.14 (Uniqueness of extensions). Let $\mathcal{I}$ be a $\pi$-system on $S$, that is, a family of subsets closed under finite intersections, and let $\Sigma=\sigma(\mathcal{I})$. If $\mu_{1}, \mu_{2}$ are finite measures on $(S, \Sigma)$ that agree on $\mathcal{I}$, then they agree on $\Sigma$.

Example B.1.15. The sets $(-\infty, x]$ for $x \in \mathbb{R}$ form a $\pi$-system generating $\mathcal{B}(\mathbb{R})$. That is, $\mathcal{B}(\mathbb{R})$ is the smallest $\sigma$-algebra containing that $\pi$-system.

Finally we can define Lebesgue measure. We start with $(0,1]$ and extend to $\mathbb{R}$ in the obvious way. We need the following lemma.

Lemma B.1.16 ( $\sigma$-additivity of $\lambda_{0}$ ). Let $\lambda_{0}$ be the set function defined above in Example B.1.10, restricted to $(0,1]$. Then $\lambda_{0}$ is $\sigma$-additive.

Definition B.1.17 (Lebesgue measure on unit interval). The unique extension of $\lambda_{0}$ (see Example B.1.10) to $(0,1]$ is denoted $\lambda$ and is called Lebesgue measure.

## B. 2 Random variables

Let $(S, \Sigma, \mu)$ be a measure space and let $\mathcal{B}=\mathcal{B}(\mathbb{R})$.
Definition B.2.1 (Measurable function). Suppose $h: S \rightarrow \mathbb{R}$ and define

$$
h^{-1}(A)=\{s \in S: h(s) \in A\} .
$$

The function $h$ is $\Sigma$-measurable if $h^{-1}(B) \in \Sigma$ for all $B \in \mathcal{B}$. We denote by $m \Sigma$ (resp., $\left.(m \Sigma)^{+}, b \Sigma\right)$ the $\Sigma$-measurable functions (resp., that are non-negative, bounded).

In the probabilistic case:
Definition B.2.2. A random variable is a measurable function on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

The behavior of a random variable is characterized by its distribution function.
Definition B.2.3 (Distribution function). Let $X$ be a random variable on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The law of $X$ is

$$
\mathcal{L}_{X}=\mathbb{P} \circ X^{-1}
$$

which is a probability measure on $(\mathbb{R}, \mathcal{B})$. By Lemma B.1.14, $\mathcal{L}_{X}$ is determined by the distribution function (DF) of $X$
distribution
function

Example B.2.4. The distribution function of a constant random variable is a jump of size 1 at the value it takes almost surely. The distribution function of a random variable with law equal to Lebesgue measure on $(0,1]$ is

$$
F_{X}(x)= \begin{cases}x & x \in(0,1] \\ 0 & x \leq 0 \\ 1 & x>1\end{cases}
$$

We refer to such as random variable as a uniform random variable over $(0,1]$.
Distribution functions are characterized by a few simple properties.
Proposition B.2.5. Suppose $F=F_{X}$ is the distribution function of a random variable $X$ on $(\Omega, \mathcal{F}, \mathbb{P})$. Then the following hold:
(i) $F$ is non-decreasing;
(ii) $\lim _{x \rightarrow+\infty} F(x)=1, \lim _{x \rightarrow-\infty} F(x)=0$;
(iii) $F$ is right-continuous.

Proof. The first property follows from the monotonicity of probability measure (which itself follows immediately from $\sigma$-additivity).

For the second property, note that the limit exists by the first property. The value of the limit follows from the following important lemma.
Lemma B.2.6 (Monotone convergence properties of measures). Let ( $S, \Sigma, \mu$ ) be a measure space.
(i) If $F_{n} \in \Sigma$, $n \geq 1$, with $F_{n} \uparrow F$, then $\mu\left(F_{n}\right) \uparrow \mu(F)$.
(ii) If $G_{n} \in \Sigma, n \geq 1$, with $G_{n} \downarrow G$ and $\mu\left(G_{k}\right)<+\infty$ for some $k$, then $\mu\left(G_{n}\right) \downarrow \mu(G)$.

Proof. Clearly $F=\cup_{n} F_{n} \in \Sigma$. For $n \geq 1$, write $H_{n}=F_{n} \backslash F_{n-1}$ (with $F_{0}=\emptyset$ ). Then by disjointness

$$
\mu\left(F_{n}\right)=\sum_{k \leq n} \mu\left(H_{k}\right) \uparrow \sum_{k<+\infty} \mu\left(H_{k}\right)=\mu(F) .
$$

The second statement is similar.
Similarly, for the third property, by Lemma B.2.6 again

$$
\mathbb{P}\left[X \leq x_{n}\right] \downarrow \mathbb{P}[X \leq x],
$$

if $x_{n} \downarrow x$.

It turns out that the properties above characterize distribution functions in the following sense.

Theorem B.2.7 (Skorokhod representation). Let $F$ satisfy the three properties above in Proposition B.2.5. Then there is a random variable $X$ on

$$
(\Omega, \mathcal{F}, \mathbb{P})=((0,1], \mathcal{B}(0,1], \lambda)
$$

with distribution function $F$. The law of $X$ is called the Lebesgue-Stieltjes measure associated to $F$.

The result says that all real random variables can be generated from uniform random variables over $(0,1]$.

Proof. Assume first that $F$ is continuous and strictly increasing. Define $X(\omega)=$ $F^{-1}(\omega)$ for all $\omega \in \Omega$. Then, $\forall x \in \mathbb{R}$,

$$
\mathbb{P}[X \leq x]=\mathbb{P}\left[\left\{\omega: F^{-1}(\omega) \leq x\right\}\right]=\mathbb{P}[\{\omega: \omega \leq F(x)\}]=F(x)
$$

In general, let

$$
X(\omega)=\inf \{x: F(x) \geq \omega\}
$$

It suffices to prove that

$$
X(\omega) \leq x \quad \Longleftrightarrow \quad \omega \leq F(x)
$$

The $\Leftarrow$ direction is clear by definition of $X$. On the other hand, by the rightcontinuity of $F$, we have that $\omega \leq F(X(\omega))$. Therefore, by monotonicity of $F$,

$$
X(\omega) \leq x \quad \Rightarrow \quad \omega \leq F(X(\omega)) \leq F(x)
$$

That proves the claim.
Turning measurability on its head, we get the following important definition.
Definition B.2.8. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $Y_{\gamma}, \gamma \in \Gamma$, be a collection of maps from $\Omega$ to $\mathbb{R}$. We let

$$
\sigma\left(Y_{\gamma}, \gamma \in \Gamma\right)
$$

be the smallest $\sigma$-algebra on which the $Y_{\gamma}$ 's are measurable.
In a sense, the above $\sigma$-algebra corresponds to "the partial information available when the $Y_{\gamma}$ s are observed."

Example B.2.9. Suppose we flip two unbiased coins and let $X$ be the number of heads observed. Then, denoting heads by H and tails by T ,

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

which is coarser than the full $\sigma$-algebra $2^{\Omega}$.
Note that $h^{-1}$ preserves all set operations. For example, $h^{-1}(A \cup B)=$ $h^{-1}(A) \cup h^{-1}(B)$. This gives the following important lemma.

Lemma B.2.10 (Sufficient condition for measurability). Suppose $\mathcal{C} \subseteq \mathcal{B}$ with $\sigma(\mathcal{C})=\mathcal{B}$. Then $h^{-1}: \mathcal{C} \rightarrow \Sigma$ implies $h \in m \Sigma$. That is, it suffices to check measurability on a collection generating $\mathcal{B}$.

Proof. Let $\mathcal{E}$ be the sets such that $h^{-1}(B) \in \Sigma$. By the observation before the statement, $\mathcal{E}$ is a $\sigma$-algebra. But $\mathcal{C} \subseteq \mathcal{E}$ which implies $\sigma(\mathcal{C}) \subseteq \mathcal{E}$ by minimality.

As a consequence we get the following properties of measurable functions.
Proposition B.2.11 (Properties of measurable functions). Let $h, h_{n}, n \geq 1$, be in $m \Sigma$ and $f \in m \mathcal{B}$.
(i) $f \circ h \in m \Sigma$.
(ii) If $S$ is a topological space and $h$ is continuous, then $h$ is $\mathcal{B}(S)$-measurable, where $\mathcal{B}(S)$ is generated by the open sets of $S$.
(iii) The function $g: S \rightarrow \mathbb{R}$ is in $m \Sigma$ if for all $c \in \mathbb{R}$,

$$
\{g \leq c\} \in \Sigma
$$

(iv) $\forall \alpha \in \mathbb{R}, h_{1}+h_{2}, h_{1} h_{2}, \alpha h \in m \Sigma$.
(v) $\inf h_{n}, \sup h_{n}, \lim \inf h_{n}, \lim \sup h_{n}$ are in $m \Sigma$.
(vi) The set

$$
\left\{s: \lim h_{n}(s) \text { exists in } \mathbb{R}\right\},
$$

is measurable.
Proof. We sketch the proof of a few of them.
(ii) This follows from Lemma B.2.10 by taking $\mathcal{C}$ as the open sets of $\mathbb{R}$.
(iii) Similarly, take $\mathcal{C}$ to be the sets of the form $(-\infty, c]$.
(iv) This follows from (iii). For example note that, writing the left-hand side as $h_{1}>c-h_{2}$,

$$
\left\{h_{1}+h_{2}>c\right\}=\cup_{q \in \mathbb{Q}}\left[\left\{h_{1}>q\right\} \cap\left\{q>c-h_{2}\right\}\right],
$$

which is a countable union of measurable sets by assumption.
(v) Note that

$$
\left\{\sup h_{n} \leq c\right\}=\cap_{n}\left\{h_{n} \leq c\right\} .
$$

Further, note that liminf is the sup of an inf.

## B. 3 Independence

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.
Definition B.3.1 (Independence). Sub- $\sigma$-algebras $\mathcal{G}_{1}, \mathcal{G}_{2}, \ldots$ of $\mathcal{F}$ are independent if: for all $G_{i} \in \mathcal{G}_{i}, i \geq 1$, and distinct $i_{1}, \ldots, i_{n}$ we have

$$
\mathbb{P}\left[G_{i_{1}} \cap \cdots \cap G_{i_{n}}\right]=\prod_{j=1}^{n} \mathbb{P}\left[G_{i_{j}}\right] .
$$

Specializing to events and random variables:
Definition B.3.2 (Independent random variables). Random variables $X_{1}, X_{2}, \ldots$ are independent if the $\sigma$-algebras $\sigma\left(X_{1}\right), \sigma\left(X_{2}\right), \ldots$ are independent.

Definition B.3.3 (Independent events). Events $E_{1}, E_{2}, \ldots$ are independent if the $\sigma$-algebras

$$
\mathcal{E}_{i}=\left\{\emptyset, E_{i}, E_{i}^{c}, \Omega\right\}, \quad i \geq 1,
$$

are independent.
Recall the more familiar definitions.
Theorem B.3.4 (Independent random variables: familiar definition). Random variables $X, Y$ are independent if and only if for all $x, y \in \mathbb{R}$

$$
\mathbb{P}[X \leq x, Y \leq y]=\mathbb{P}[X \leq x] \mathbb{P}[Y \leq y]
$$

Theorem B.3.5 (Independent events: familiar definition). Events $E_{1}, E_{2}$ are independent if and only if

$$
\mathbb{P}\left[E_{1} \cap E_{2}\right]=\mathbb{P}\left[E_{1}\right] \mathbb{P}\left[E_{2}\right] .
$$

The proofs of these characterizations follow immediately from the following lemma.

Lemma B.3.6 (Independence and $\pi$-systems). Suppose that $\mathcal{G}$ and $\mathcal{H}$ are sub- $\sigma$ algebras and that $\mathcal{I}$ and $\mathcal{J}$ are $\pi$-systems such that

$$
\sigma(\mathcal{I})=\mathcal{G}, \quad \sigma(\mathcal{J})=\mathcal{H} .
$$

Then $\mathcal{G}$ and $\mathcal{H}$ are independent if and only if $\mathcal{I}$ and $\mathcal{J}$ are as well, that is,

$$
\mathbb{P}[I \cap J]=\mathbb{P}[I] \mathbb{P}[J], \quad \forall I \in \mathcal{I}, J \in \mathcal{J} .
$$

Proof. Suppose $\mathcal{I}$ and $\mathcal{J}$ are independent. For fixed $I \in \mathcal{I}$, the measures $\mathbb{P}[I \cap H]$ and $\mathbb{P}[I] \mathbb{P}[H]$ are equal for $H \in \mathcal{J}$ and have total mass $\mathbb{P}[I]<+\infty$. By the Uniqueness of Extensions Lemma (Lemma B.1.14) the above measures agree on $\sigma(\mathcal{J})=\mathcal{H}$.

Repeat the argument. Fix $H \in \mathcal{H}$. Then the measures $\mathbb{P}[G \cap H]$ and $\mathbb{P}[G] \mathbb{P}[H]$ agree on $\mathcal{I}$ and have total mass $\mathbb{P}[H]<+\infty$. Therefore they must agree on $\sigma(\mathcal{I})=$ $\mathcal{G}$.

We give a standard construction of an infinite sequence of independent random variables with prescribed distributions.

Let $(\Omega, \mathcal{F}, \mathbb{P})=((0,1], \mathcal{B}(0,1], \lambda)$ and for $\omega \in \Omega$ consider the binary expansion

$$
\omega=0 . \omega_{1} \omega_{2} \ldots
$$

(For dyadic rationals, use the all-1 ending and note that the dyadic rationals have measure 0 by countability.) This construction produces a sequence of independent so-called Bernoulli trials. That is, under $\lambda$, each bit is Bernoulli( $1 / 2$ ) and any finite collection is independent.

To get two independent uniform random variables, consider the following construction:

$$
\begin{aligned}
& U_{1}=0 . \omega_{1} \omega_{3} \omega_{5} \ldots \\
& U_{2}=0 . \omega_{2} \omega_{4} \omega_{6} \ldots
\end{aligned}
$$

Let $\mathcal{A}_{1}$ (resp. $\mathcal{A}_{2}$ ) be the $\pi$-system consisting of all finite intersections of events of the form $\left\{\omega_{i} \in H_{i}\right\}$ for odd $i$ (resp. even $i$ ). By Lemma B.3.6, the $\sigma$-fields $\sigma\left(\mathcal{A}_{1}\right)$ and $\sigma\left(\mathcal{A}_{2}\right)$ are independent.

More generally, let

$$
\begin{aligned}
V_{1} & =0 . \omega_{1} \omega_{3} \omega_{6} \cdots \\
V_{2} & =0 . \omega_{2} \omega_{5} \omega_{9} \cdots \\
V_{3} & =0 . \omega_{4} \omega_{8} \omega_{13} \cdots \\
\vdots & =\ddots
\end{aligned}
$$

that is, fill up the array diagonally. By the argument above, the $V_{i}$ 's are independent and Bernoulli( $1 / 2$ ).

Finally let $\mu_{n}, n \geq 1$, be a sequence of probability measures with distribution functions $F_{n}, n \geq 1$. For each $n$, define

$$
X_{n}(\omega)=\inf \left\{x: F_{n}(x) \geq V_{n}(\omega)\right\}
$$

By the (proof of the) Skorokhod Representation (Theorem B.2.7), $X_{n}$ has distribution function $F_{n}$.

Definition B.3.7 (I.i.d. random variables). A sequence of independent random variables $\left(X_{n}\right)$ as above is independent and identically distributed (i.i.d.) if $F_{n}=F$ for some $n$.

Alternatively, we have the following more general result.
Theorem B.3.8 (Kolmogorov's extension theorem). Suppose we are given probability measures $\mu_{n}$ on $\left(\mathbb{R}^{n}, \mathcal{B}\left(\mathbb{R}^{n}\right)\right)$ that are consistent, that is,

$$
\mu_{n+1}\left(\left(a_{1}, b_{1}\right] \times \cdots \times\left(a_{n}, b_{n}\right] \times \mathbb{R}\right)=\mu_{n}\left(\left(a_{1}, b_{1}\right] \times \cdots \times\left(a_{n}, b_{n}\right]\right) .
$$

Then there exists a unique probability measure $\mathbb{P}$ on $\left(\mathbb{R}^{\mathbb{N}}, \mathcal{R}^{\mathbb{N}}\right)$ with

$$
\mathbb{P}\left[\omega: \omega_{i} \in\left(a_{i}, b_{i}\right], 1 \leq i \leq n\right]=\mu_{n}\left(\left(a_{1}, b_{1}\right] \times \cdots \times\left(a_{n}, b_{n}\right]\right) .
$$

Here $\mathcal{R}^{\mathbb{N}}$ is the product $\sigma$-algebra, that is, the $\sigma$-algebra generated by finite-dimensional rectangles.

Next, we discuss a first non-trivial result about independent sequences.
Definition B.3.9 (Tail $\sigma$-algebra). Let $X_{1}, X_{2}, \ldots$ be random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Define

$$
\mathcal{T}=\bigcap_{n \geq 1} \mathcal{T}_{n}
$$

where

$$
\mathcal{T}_{n}=\sigma\left(X_{n+1}, X_{n+2}, \ldots\right) .
$$

As an intersection of $\sigma$-algebras, $\mathcal{T}$ is a $\sigma$-algebra. It is called the tail $\sigma$-algebra of the sequence $\left(X_{n}\right)$.

Intuitively, an event is in the tail if changing a finite number of values does not affect its occurence.

Example B.3.10. If $S_{n}=\sum_{k \leq n} X_{k}$, then

$$
\begin{gathered}
\left\{\lim _{n} S_{n} \text { exists }\right\} \in \mathcal{T}, \\
\left\{\limsup _{n} n^{-1} S_{n}>0\right\} \in \mathcal{T},
\end{gathered}
$$

but

$$
\left\{\limsup _{n} S_{n}>0\right\} \notin \mathcal{T} .
$$

Theorem B.3.11 (Kolmogorov's 0-1 law). Let $\left(X_{n}\right)$ be a sequence of independent random variables with tail $\sigma$-algebra $\mathcal{T}$. Then $\mathcal{T}$ is $\mathbb{P}$-trivial, that is, for all $A \in \mathcal{T}$ we have either $\mathbb{P}[A]=0$ or 1 .

Proof. Let $\mathcal{X}_{n}=\sigma\left(X_{1}, \ldots, X_{n}\right)$. Note that $\mathcal{X}_{n}$ and $\mathcal{T}_{n}$ are independent. Moreover, since $\mathcal{T} \subseteq \mathcal{T}_{n}$ we have that $\mathcal{X}_{n}$ is independent of $\mathcal{T}$. Now let

$$
\mathcal{X}_{\infty}=\sigma\left(X_{n}, n \geq 1\right) .
$$

Note that

$$
\mathcal{K}_{\infty}=\bigcup_{n \geq 1} \mathcal{X}_{n}
$$

is a $\pi$-system generating $\mathcal{X}_{\infty}$. Therefore, by Lemma B.3.6, $\mathcal{X}_{\infty}$ is independent of $\mathcal{T}$. But $\mathcal{T} \subseteq \mathcal{X}_{\infty}$ and therefore $\mathcal{T}$ is independent of itself! Hence if $A \in \mathcal{T}$,

$$
\mathbb{P}[A]=\mathbb{P}[A \cap A]=\mathbb{P}[A]^{2},
$$

which can occur only if $\mathbb{P}[A] \in\{0,1\}$.

## B. 4 Expectation

Let $(S, \Sigma, \mu)$ be a measure space. We denote by $\mathbf{1}_{A}$ the indicator of a set $A$, that is,

$$
\mathbf{1}_{A}(s)= \begin{cases}1, & \text { if } s \in A \\ 0, & \text { o.w. }\end{cases}
$$

Definition B.4.1 (Simple functions). A simple function is a function of the form

$$
f=\sum_{k=1}^{m} a_{k} \mathbf{1}_{A_{k}}
$$

where $a_{k} \in[0,+\infty]$ and $A_{k} \in \Sigma$ for all $k$. We denote the set of all such functions by $\mathrm{SF}^{+}$. We define the integral of $f$ by

$$
\mu(f):=\sum_{k=1}^{m} a_{k} \mu\left(A_{k}\right) \leq+\infty
$$

We also write $\mu f=\mu(f)$.
The following is left as a (somewhat tedious but) immediate exercise.
Proposition B.4.2. Let $f, g \in \mathrm{SF}^{+}$.
(i) If $\mu(f \neq g)=0$, then $\mu f=\mu g$. [Hint: Rewrite $f$ and $g$ over the same disjoint sets.]
(ii) For all $c \geq 0, f+g, c f \in \mathrm{SF}^{+}$and

$$
\mu(f+g)=\mu f+\mu g, \quad \mu(c f)=c \mu f
$$

[Hint: This one is obvious by definition.]
(iii) If $f \leq g$ then $\mu f \leq \mu g$. [Hint: Show that $g-f \in \mathrm{SF}^{+}$and use linearity.]

The main definition and theorem of integration theory follows.
Definition B.4.3 (Nonnegative functions). Let $f \in(\mathrm{~m} \Sigma)^{+}$. Then the integral of $f$ is defined by

$$
\mu(f)=\sup \left\{\mu(h): h \in \mathrm{SF}^{+}, h \leq f\right\}
$$

Again we also write $\mu f=\mu(f)$.
Theorem B.4.4 (Monotone convergence theorem). If $f_{n}, f \in(\mathrm{~m} \Sigma)^{+}, n \geq 1$, with $f_{n} \uparrow f$, then

$$
\mu f_{n} \uparrow \mu f
$$

Many theorems in integration follow from the monotone convergence theorem. In that context, the following approximation is useful.

Definition B.4.5 (Staircase function). For $f \in(\mathrm{~m} \Sigma)^{+}$and $r \geq 1$, the $r$-th staircase function $\alpha^{(r)}$ is

$$
\alpha^{(r)}(x)= \begin{cases}0, & \text { if } x=0, \\ (i-1) 2^{-r}, & \text { if }(i-1) 2^{-r}<x \leq i 2^{-r} \leq r, \\ r, & \text { if } x>r,\end{cases}
$$

We let $f^{(r)}=\alpha^{(r)}(f)$. Note that $f^{(r)} \in \mathrm{SF}^{+}$and $f^{(r)} \uparrow f$ as $r \rightarrow+\infty$.

Using the previous definition, we get for example the following properties.
Proposition B.4.6. Let $f, g \in(\mathrm{~m} \Sigma)^{+}$.
(i) If $\mu(f \neq g)=0$, then $\mu(f)=\mu(g)$.
(ii) For all $c \geq 0, f+g, c f \in(\mathrm{~m} \Sigma)^{+}$and

$$
\mu(f+g)=\mu f+\mu g, \quad \mu(c f)=c \mu f .
$$

(iii) If $f \leq g$ then $\mu f \leq \mu g$.

For a function $f$, let $f^{+}$and $f^{-}$be the positive and negative parts of $f$, that is,

$$
f^{+}(s)=f(s) \vee 0, \quad f^{-}(s)=(-f(s)) \vee 0 .
$$

Note that $|f|=f^{+}+f^{-}$. Finally we define

$$
\mu(f):=\mu\left(f^{+}\right)-\mu\left(f^{-}\right),
$$

provided $\mu\left(f^{+}\right)+\mu\left(f^{-}\right)<+\infty$, in which case we write $f \in L^{1}(S, \Sigma, \mu)$. Proposition B.4.6 can be generalized naturally to this definition. Moreover we have the following.

Theorem B.4.7 (Dominated convergence theorem). If $f_{n}, f \in \mathrm{~m} \Sigma, n \geq 1$, with $f_{n}(s) \rightarrow f(s)$ for all $s \in S$, and there is a nonnegative function $g \in L^{1}(S, \Sigma, \mu)$ such that $\left|f_{n}\right| \leq g$, then

$$
\mu\left(\left|f_{n}-f\right|\right) \rightarrow 0
$$

and in particular

$$
\mu f_{n} \rightarrow \mu f,
$$

as $n \rightarrow \infty$.
More generally, for $0<p<+\infty$, the space $L^{p}(S, \Sigma, \mu)$ contains all functions $f: S \rightarrow \mathbb{R}$ such that $\|f\|_{p}<+\infty$, where

$$
\|f\|_{p}:=\mu\left(|f|^{p}\right)^{1 / p},
$$

up to equality almost everywhere. We state the following results without proof.
Theorem B.4.8 (Hölder's inequality). Let $1<p, q<+\infty$ such that $p^{-1}+$ $q^{-1}=1$. Then, for any $f \in L^{p}(S, \Sigma, \mu)$ and $g \in L^{q}(S, \Sigma, \mu)$, it holds that $f g \in L^{1}(S, \Sigma, \mu)$ and further

$$
\|f g\|_{1} \leq\|f\|_{p}\|g\|_{q} .
$$

The case $p=q=2$ is known as the Cauchy-Schwarz inequality (or Schwarz inequality).

Theorem B.4.9 (Minkowski's inequality). Let $1<p<+\infty$. Then, for any $f, g \in$ $L^{p}(S, \Sigma, \mu)$, it holds that $f+g \in L^{p}(S, \Sigma, \mu)$ and further

$$
\|f+g\|_{p} \leq\|f\|_{p}+\|g\|_{p}
$$

Theorem B.4.10 ( $L^{p}$ completeness). Let $1 \leq p<+\infty$. If $\left(f_{n}\right)_{n}$ in $L^{p}(S, \Sigma, \mu)$ is Cauchy, that is,

$$
\sup _{n, m \geq k}\left\|f_{n}-f_{m}\right\|_{p} \rightarrow 0
$$

as $k \rightarrow+\infty$, then there exists $f \in L^{p}(S, \Sigma, \mu)$ such that

$$
\left\|f_{n}-f\right\|_{p} \rightarrow 0
$$

as $n \rightarrow+\infty$.
We can now define the expectation. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.
Definition B.4.11 (Expectation). If $X \geq 0$ is a random variable then we define the expectation of $X$, denoted by $\mathbb{E}[X]$, as the integral of $X$ over $\mathbb{P}$. More generally (i.e., not assuming non-negativity), if

$$
\mathbb{E}|X|=\mathbb{E}\left[X^{+}\right]+\mathbb{E}\left[X^{-}\right]<+\infty,
$$

we let

$$
\mathbb{E}[X]=\mathbb{E}\left[X^{+}\right]-\mathbb{E}\left[X^{-}\right] .
$$

We denote the set of all such integrable random variables (up to equality almost surely) by $L^{1}(\Omega, \mathcal{F}, \mathbb{P})$.

The properties of the integral for nonnegative functions (see Proposition B.4.6) extend to the expectation.

Proposition B.4.12. Let $X, X_{1}, X_{2}$ be random variables in $L^{1}(\Omega, \mathcal{F}, \mathbb{P})$.
(LIN) If $a_{1}, a_{2} \in \mathbb{R}$, then $\mathbb{E}\left[a_{1} X_{1}+a_{2} X_{2}\right]=a_{1} \mathbb{E}\left[X_{1}\right]+a_{2} \mathbb{E}\left[X_{2}\right]$.
(POS) If $X \geq 0$, then $\mathbb{E}[X] \geq 0$.
One useful implication of (POS) is that $|X|-X \geq 0$ so that $\mathbb{E}[X] \leq \mathbb{E}|X|$ and, by applying the same argument to $-X$, we have further $|\mathbb{E}[X]| \leq \mathbb{E}|X|$.

The monotone convergence theorem (Theorem B.4.4) implies the following results. We first need a definition.

Definition B.4.13 (Convergence almost sure). We say that $X_{n} \rightarrow X$ almost surely (a.s.) if

$$
\mathbb{P}\left[X_{n} \rightarrow X\right]=1
$$

Proposition B.4.14. Let $X, Y, X_{n}, n \geq 1$, be random variables in $L^{1}(\Omega, \mathcal{F}, \mathbb{P})$.
(MON) If $0 \leq X_{n} \uparrow X$, then $\mathbb{E}\left[X_{n}\right] \uparrow \mathbb{E}[X] \leq+\infty$.
(FATOU) If $X_{n} \geq 0$, then $\mathbb{E}\left[\liminf _{n} X_{n}\right] \leq \liminf _{n} \mathbb{E}\left[X_{n}\right]$.
(DOM) If $\left|X_{n}\right| \leq Y, n \geq 1$, with $\mathbb{E}[Y]<+\infty$ and $X_{n} \rightarrow X$ a.s., then

$$
\mathbb{E}\left|X_{n}-X\right| \rightarrow 0,
$$

and, hence,

$$
\mathbb{E}\left[X_{n}\right] \rightarrow \mathbb{E}[X] .
$$

(Indeed,

$$
\begin{aligned}
\left|\mathbb{E}\left[X_{n}\right]-\mathbb{E}[X]\right| & =\left|\mathbb{E}\left[X_{n}-X\right]\right| \\
& \left.\leq \mathbb{E}\left|X_{n}-X\right| .\right)
\end{aligned}
$$

(SCHEFFE) If $X_{n} \rightarrow X$ a.s. and $\mathbb{E}\left|X_{n}\right| \rightarrow \mathbb{E}|X|$ then

$$
\mathbb{E}\left|X_{n}-X\right| \rightarrow 0 .
$$

(BDD) If $X_{n} \rightarrow X$ a.s. and $\left|X_{n}\right| \leq K<+\infty$ for all $n$ then

$$
\mathbb{E}\left|X_{n}-X\right| \rightarrow 0 .
$$

Proof. We only prove (FATOU). To use (MON) we write the liminf as an increasing limit. Letting $Z_{k}=\inf _{n \geq k} X_{n}$, we have

$$
\liminf _{n} X_{n}=\uparrow \lim _{k} Z_{k},
$$

so that by (MON)

$$
\mathbb{E}\left[\liminf _{n} X_{n}\right]=\uparrow \lim _{k} \mathbb{E}\left[Z_{k}\right] .
$$

For $n \geq k$ we have $X_{n} \geq Z_{k}$ so that $\mathbb{E}\left[X_{n}\right] \geq \mathbb{E}\left[Z_{k}\right]$ hence

$$
\mathbb{E}\left[Z_{k}\right] \leq \inf _{n \geq k} \mathbb{E}\left[X_{n}\right] .
$$

Finally, we get

$$
\mathbb{E}\left[\liminf _{n} X_{n}\right] \leq \uparrow \lim _{k} \inf _{n \geq k} \mathbb{E}\left[X_{n}\right]
$$

The following inequality is often useful. We give an example below.

Theorem B.4.15 (Jensen's inequality). Let $h: G \rightarrow \mathbb{R}$ be a convex function on an open interval $G$ such that $\mathbb{P}[X \in G]=1$ and $X, h(X) \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ then

Jensen's inequality

$$
\mathbb{E}[h(X)] \geq h(\mathbb{E}[X])
$$

The $L^{p}$ norm defined earlier applies to random variables as well. That is, for $p \geq 1$, we let $\|X\|_{p}=\mathbb{E}\left[|X|^{p}\right]^{1 / p}$ and denote by $L^{p}(\Omega, \mathcal{F}, \mathbb{P})$ the collection of random variables $X$ (up to almost sure equality) such that $\|X\|_{p}<+\infty$. Jensen's inequality (Theorem B.4.15) implies the following relationship.

Lemma B.4.16 (Monotonicity of norms). For $1 \leq p \leq r<+\infty$, we have $\|X\|_{p} \leq$ $\|X\|_{r}$.

Proof. For $n \geq 0$, let

$$
X_{n}=(|X| \wedge n)^{p} .
$$

Take $h(x)=x^{r / p}$ which is convex on $(0,+\infty)$. Then, by Jensen's inequality,

$$
\left(\mathbb{E}\left[X_{n}\right]\right)^{r / p} \leq \mathbb{E}\left[\left(X_{n}\right)^{r / p}\right]=\mathbb{E}\left[(|X| \wedge n)^{r}\right] \leq \mathbb{E}\left[|X|^{r}\right] .
$$

Take $n \rightarrow \infty$ and use (MON).
This latter inequality is useful among other things to argue about the convergence of expectations. We say that $X_{n}$ converges to $X_{\infty}$ in $L^{p}$ if $\left\|X_{n}-X_{\infty}\right\|_{p} \rightarrow 0$. By the previous lemma, convergence on $L^{r}$ implies convergence in $L^{p}$ for $r \geq p \geq 1$. Further we have:

Lemma B.4.17 (Convergence of expectations). Assume $X_{n}, X_{\infty} \in L^{1}$. Then

$$
\left\|X_{n}-X_{\infty}\right\|_{1} \rightarrow 0
$$

implies

$$
\mathbb{E}\left[X_{n}\right] \rightarrow \mathbb{E}\left[X_{\infty}\right]
$$

Proof. Note that

$$
\left|\mathbb{E}\left[X_{n}\right]-\mathbb{E}\left[X_{\infty}\right]\right| \leq \mathbb{E}\left|X_{n}-X_{\infty}\right| \rightarrow 0 .
$$

So, a fortiori, convergence in $L^{p}, p \geq 1$, implies convergence of expectations.
Square integrable random variables have a nice geometry by virtue of forming a Hilbert space.

Definition B.4.18 (Square integrable variables). Recall that $L^{2}(\Omega, \mathcal{F}, \mathbb{P})$ denotes the set of all square integrable random variables (up to equality almost surely), that is, those $X$ with $\mathbb{E}\left[X^{2}\right]<+\infty$ For $X, Y \in L^{2}(\Omega, \mathcal{F}, \mathbb{P})$, define the inner product $\langle X, Y\rangle:=\mathbb{E}[X Y]$. Then the $L^{2}$ norm is $\|X\|_{2}=\sqrt{\langle X, X\rangle}$.
Theorem B.4.19 (Cauchy-Schwarz inequality). If $X, Y \in L^{2}(\Omega, \mathcal{F}, \mathbb{P})$, then $X Y \in$ $L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and

$$
\mathbb{E}|X Y| \leq \sqrt{\mathbb{E}\left[X^{2}\right] \mathbb{E}\left[Y^{2}\right]}
$$

or put differently

$$
|\langle X, Y\rangle| \leq\|X\|_{2}\|Y\|_{2} .
$$

Theorem B.4.20 (Parallelogram law). If $X, Y \in L^{2}(\Omega, \mathcal{F}, \mathbb{P})$, then

$$
\|X+Y\|_{2}^{2}+\|X-Y\|_{2}^{2}=2\|X\|_{2}^{2}+2\|Y\|_{2}^{2}
$$

## B. 5 Fubini's theorem

We now define product measures and state (without proof) Fubini's Theorem.
Definition B.5.1 (Product $\sigma$-algebra). Let $\left(S_{1}, \Sigma_{1}\right)$ and $\left(S_{2}, \Sigma_{2}\right)$ be measure spaces. Let $S=S_{1} \times S_{2}$ be the Cartesian product of $S_{1}$ and $S_{2}$. For $i=1,2$, let $\pi_{i}: S \rightarrow S_{i}$ be the projection on the $i$-th coordinate, that is,

$$
\pi_{i}\left(s_{1}, s_{2}\right)=s_{i}
$$

The product $\sigma$-algebra $\Sigma=\Sigma_{1} \times \Sigma_{2}$ is defined as

$$
\Sigma=\sigma\left(\pi_{1}, \pi_{2}\right)
$$

In words, it is the smallest $\sigma$-algebra that makes coordinate maps measurable. It is generated by sets of the form

$$
\pi_{1}^{-1}\left(B_{1}\right)=B_{1} \times S_{2}, \quad \pi_{2}^{-1}\left(B_{2}\right)=S_{1} \times B_{2}, \quad B_{1} \in \Sigma_{1}, B_{2} \in \Sigma_{2}
$$

Theorem B.5.2 (Fubini's Theorem). For $F \in \Sigma$, let $f=\mathbf{1}_{F}$ and define

Fubini's
Theorem

$$
\mu(F):=\int_{S_{1}} I_{1}^{f}\left(s_{1}\right) \mu_{1}\left(\mathrm{~d} s_{1}\right)=\int_{S_{2}} I_{2}^{f}\left(s_{2}\right) \mu_{2}\left(\mathrm{~d} s_{2}\right),
$$

where

$$
\begin{aligned}
& I_{1}^{f}\left(s_{1}\right):=\int_{S_{2}} f\left(s_{1}, s_{2}\right) \mu_{2}\left(\mathrm{~d} s_{2}\right) \in \mathrm{b} \Sigma_{1}, \\
& I_{2}^{f}\left(s_{2}\right):=\int_{S_{1}} f\left(s_{1}, s_{2}\right) \mu_{1}\left(\mathrm{~d} s_{1}\right) \in \mathrm{b} \Sigma_{2} .
\end{aligned}
$$

(The equality and inclusions above are part of the statement.) The set function $\mu$ is a measure on $(S, \Sigma)$ called the product measure of $\mu_{1}$ and $\mu_{2}$ and we write $\mu=\mu_{1} \times \mu_{2}$ and

$$
(S, \Sigma, \mu)=\left(S_{1}, \Sigma_{1}, \mu_{1}\right) \times\left(S_{2}, \Sigma_{2}, \mu_{2}\right) .
$$

Moreover $\mu$ is the unique measure on $(S, \Sigma)$ for which

$$
\mu\left(A_{1} \times A_{2}\right)=\mu\left(A_{1}\right) \mu\left(A_{2}\right), \quad A_{i} \in \Sigma_{i} .
$$

If $f \in(\mathrm{~m} \Sigma)^{+}$then

$$
\mu(f)=\int_{S_{1}} I_{1}^{f}\left(s_{1}\right) \mu_{1}\left(\mathrm{~d} s_{1}\right)=\int_{S_{2}} I_{2}^{f}\left(s_{2}\right) \mu_{2}\left(\mathrm{~d} s_{2}\right),
$$

where $I_{1}^{f}, I_{2}^{f}$ are defined as before (i.e., as the sup over bounded functions from below). The same is valid if $f \in \mathrm{~m} \Sigma$ and $\mu(|f|)<+\infty$.

Some applications of Fubini's Theorem (Theorem B.5.2) follow. We first recall the following useful formula.

Theorem B.5.3 (Change-of-variables formula). Let $X$ be a random variable with law $\mathcal{L}$. If $f: \mathbb{R} \rightarrow \mathbb{R}$ is such that either $f \geq 0$ or $\mathbb{E}|f(X)|<+\infty$ then

$$
\mathbb{E}[f(X)]=\int_{\mathbb{R}} f(y) \mathcal{L}(\mathrm{d} y)
$$

Proof. We use the standard machinery.

1. For $f=\mathbf{1}_{B}$ with $B \in \mathcal{B}$,

$$
\mathbb{E}\left[\mathbf{1}_{B}(X)\right]=\mathcal{L}(B)=\int_{\mathbb{R}} \mathbf{1}_{B}(y) \mathcal{L}(\mathrm{d} y) .
$$

2. If $f=\sum_{k=1}^{m} a_{k} \mathbf{1}_{A_{k}}$ is a simple function, then by (LIN)

$$
\mathbb{E}[f(X)]=\sum_{k=1}^{m} a_{k} \mathbb{E}\left[\mathbf{1}_{A_{k}}(X)\right]=\sum_{k=1}^{m} a_{k} \int_{\mathbb{R}} \mathbf{1}_{A_{k}}(y) \mathcal{L}(\mathrm{d} y)=\int_{\mathbb{R}} f(y) \mathcal{L}(\mathrm{d} y) .
$$

3. Let $f \geq 0$ and approximate $f$ by a sequence $\left\{f_{n}\right\}$ of increasing simple functions. By (MON)

$$
\mathbb{E}[f(X)]=\lim _{n} \mathbb{E}\left[f_{n}(X)\right]=\lim _{n} \int_{\mathbb{R}} f_{n}(y) \mathcal{L}(\mathrm{d} y)=\int_{\mathbb{R}} f(y) \mathcal{L}(\mathrm{d} y) .
$$

4. Finally, assume that $f$ is such that $\mathbb{E}|f(X)|<+\infty$. Then by (LIN)

$$
\begin{aligned}
\mathbb{E}[f(X)] & =\mathbb{E}\left[f^{+}(X)\right]-\mathbb{E}\left[f^{-}(X)\right] \\
& =\int_{\mathbb{R}} f^{+}(y) \mathcal{L}(\mathrm{d} y)-\int_{\mathbb{R}} f^{-}(y) \mathcal{L}(\mathrm{d} y) \\
& =\int_{\mathbb{R}} f(y) \mathcal{L}(\mathrm{d} y) .
\end{aligned}
$$

Theorem B.5.4. Let $X$ and $Y$ be independent random variables with respective laws $\mu$ and $\nu$. Let $f$ and $g$ be measurable functions such that either $f, g \geq 0$ or $\mathbb{E}|f(X)|, \mathbb{E}|g(Y)|<+\infty$. Then

$$
\mathbb{E}[f(X) g(Y)]=\mathbb{E}[f(X)] \mathbb{E}[g(Y)] .
$$

Proof. From the change-of-variables formula (Theorem B.5.3) and Fubini's Theorem (Theorem B.5.2), we get

$$
\begin{aligned}
\mathbb{E}[f(X) g(Y)] & =\int_{\mathbb{R}^{2}} f(x) g(y)(\mu \times \nu)(\mathrm{d} x \times \mathrm{d} y) \\
& =\int_{\mathbb{R}}\left(\int_{\mathbb{R}} f(x) g(y) \mu(\mathrm{d} x)\right) \nu(\mathrm{d} y) \\
& =\int_{\mathbb{R}}(g(y) \mathbb{E}[f(X)]) \nu(\mathrm{d} y) \\
& =\mathbb{E}[f(X)] \mathbb{E}[g(Y)] .
\end{aligned}
$$

Definition B.5.5 (Density). Let $X$ be a random variable with law $\mu$. We say that $X$ has density $f_{X}$ if for all $B \in \mathcal{B}(\mathbb{R})$

$$
\mu(B)=\mathbb{P}[X \in B]=\int_{B} f_{X}(x) \lambda(\mathrm{d} x) .
$$

Theorem B.5.6 (Convolution). Let $X$ and $Y$ be independent random variables with distribution functions $F$ and $G$ respectively. Then the distribution function, $H$, of $X+Y$ is

$$
H(z)=\int F(z-y) \mathrm{d} G(y) .
$$

This is called the convolution of $F$ and $G$. Moreover, if $X$ and $Y$ have densities $f$ and $g$ respectively, then $X+Y$ has density

$$
h(z)=\int f(z-y) g(y) \mathrm{d} y .
$$

Proof. From Fubini's Theorem (Theorem B.5.3), denoting the laws of $X$ and $Y$ by $\mu$ and $\nu$ respectively,

$$
\begin{aligned}
\mathbb{P}[X+Y \leq z] & =\iint \mathbf{1}_{\{x+y \leq z\}} \mu(\mathrm{d} x) \nu(\mathrm{d} y) \\
& =\int F(z-y) \nu(\mathrm{d} y) \\
& =\int F(z-y) \mathrm{d} G(y) \\
& =\int\left(\int_{-\infty}^{z} f(x-y) \mathrm{d} x\right) \mathrm{d} G(y) \\
& =\int_{-\infty}^{z}\left(\int f(x-y) \mathrm{d} G(y)\right) \mathrm{d} x \\
& =\int_{-\infty}^{z}\left(\int f(x-y) g(y) \mathrm{d} y\right) \mathrm{d} x .
\end{aligned}
$$

See Exercise 2.1 for a proof of the following standard formula.
Theorem B.5.7 (Moments of nonegative random variables). For any nonnegative random variable $X$ and positive integer $k$,

$$
\begin{equation*}
\mathbb{E}\left[X^{k}\right]=\int_{0}^{+\infty} k x^{k-1} \mathbb{P}[X>x] \mathrm{d} x \tag{B.5.1}
\end{equation*}
$$

## B. 6 Conditional expectation

Before defining the conditional expectation, we recall some elementary concepts. For two events $A, B$, the conditional probability of $A$ given $B$ is defined as

$$
\mathbb{P}[A \mid B]=\frac{\mathbb{P}[A \cap B]}{\mathbb{P}[B]},
$$

where we assume $\mathbb{P}[B]>0$.
Now let $X$ and $Z$ be random variables taking values $x_{1}, \ldots, x_{m}$ and $z_{1}, \ldots, z_{n}$ respectively. The conditional expectation of $X$ given $Z=z_{j}$ is defined as

$$
y_{j}=\mathbb{E}\left[X \mid Z=z_{j}\right]=\sum_{i} x_{i} \mathbb{P}\left[X=x_{i} \mid Z=z_{j}\right],
$$

where we assume $\mathbb{P}\left[Z=z_{j}\right]>0$ for all $j$. As motivation for the general definition, we make the following observations.

- We can think of the conditional expectation as a random variable $Y=$ $\mathbb{E}[X \mid Z]$ defined as follows

$$
Y(\omega)=y_{j} \text { on } G_{j}=\left\{\omega: Z(\omega)=z_{j}\right\} .
$$

- Then $Y$ is $\mathcal{G}$-measurable where $\mathcal{G}=\sigma(Z)$.
- On sets in $\mathcal{G}$, the expectation of $Y$ agrees with the expectation of $X$. Indeed, note first that

$$
\begin{aligned}
\mathbb{E}\left[Y ; G_{j}\right] & =y_{j} \mathbb{P}\left[G_{j}\right] \\
& =\sum_{i} x_{i} \mathbb{P}\left[X=x_{i} \mid Z=z_{j}\right] \mathbb{P}\left[Z=z_{j}\right] \\
& =\sum_{i} x_{i} \mathbb{P}\left[X=x_{i}, Z=z_{j}\right] \\
& =\mathbb{E}\left[X ; G_{j}\right] .
\end{aligned}
$$

This is also true for all $G \in \mathcal{G}$ by summation over $j$.
We are ready to state the general definition of the conditional expectation. Its existence and uniqueness follow from the next theorem.

Theorem B.6.1 (Conditional expectation). Let $X \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{G} \subseteq \mathcal{F} a$ sub- $\sigma$-algebra. Then:
(i) (Existence) There exists a random variable $Y \in L^{1}(\Omega, \mathcal{G}, \mathbb{P})$ such that

$$
\begin{equation*}
\mathbb{E}[Y ; G]=\mathbb{E}[X ; G], \forall G \in \mathcal{G} \tag{B.6.1}
\end{equation*}
$$

Such $a Y$ is called a version of the conditional expectation of $X$ given $\mathcal{G}$ and is denoted by $\mathbb{E}[X \mid \mathcal{G}]$.
(ii) (Uniqueness) It is unique in the sense that, if $Y$ and $Y^{\prime}$ are two versions of the conditional expectation, then $Y=Y^{\prime}$ almost surely.

When $\mathcal{G}=\sigma(Z)$, we sometimes use the notation $\mathbb{E}[X \mid Z]:=\mathbb{E}[X \mid \mathcal{G}]$. A similar convention applies to collections of random variables, for example, $\mathbb{E}\left[X \mid Z_{1}, Z_{2}\right]:=$ $\mathbb{E}\left[X \mid \sigma\left(Z_{1}, Z_{2}\right)\right]$ and so on.

We first prove uniqueness. Existence is proved below after some more concepts are introduced.

Proof of Theorem B.6.1 (ii). By way of contradiction, let $Y, Y^{\prime}$ be two versions of $\mathbb{E}[X \mid G]$ such that without loss of generality $\mathbb{P}\left[Y>Y^{\prime}\right]>0$. By monotonicity, there is $n \geq 1$ with $G=\left\{Y>Y^{\prime}+n^{-1}\right\} \in \mathcal{G}$ such that $\mathbb{P}[G]>0$. Then, by definition,

$$
0=\mathbb{E}\left[Y-Y^{\prime} ; G\right]>n^{-1} \mathbb{P}[G]>0
$$

which gives a contradiction.
To prove existence, we use the $L^{2}$ method. In $L^{2}(\Omega, \mathcal{F}, \mathbb{P})$, the conditional expectation reduces to an orthogonal projection.

Theorem B.6.2 (Conditional expectation: $L^{2}$ case). Let $X \in L^{2}(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{G} \subseteq \mathcal{F}$ a sub- $\sigma$-algebra. Then there exists an (almost surely) unique $Y \in L^{2}(\Omega, \mathcal{G}, \mathbb{P})$ such that

$$
\|X-Y\|_{2}=\Delta:=\inf \left\{\|X-W\|_{2}: W \in L^{2}(\Omega, \mathcal{G}, \mathbb{P})\right\}
$$

and, moreover, $\langle Z, X-Y\rangle=0, \forall Z \in L^{2}(\Omega, \mathcal{G}, \mathbb{P})$. In particular, it satisfies (B.6.1). Such a $Y$ is called the orthogonal projection of $X$ on $L^{2}(\Omega, \mathcal{G}, \mathbb{P})$.

Proof. Take $\left(Y_{n}\right)$ such that $\left\|X-Y_{n}\right\|_{2} \rightarrow \Delta$. We use the fact that $L^{2}(\Omega, \mathcal{G}, \mathbb{P})$ is complete (Theorem B.4.10) and first seek to prove that $\left(Y_{n}\right)$ is Cauchy. Using the parallelogram law (Theorem B.4.20), note that

$$
\left\|X-Y_{r}\right\|_{2}^{2}+\left\|X-Y_{s}\right\|_{2}^{2}=2\left\|X-\frac{1}{2}\left(Y_{r}+Y_{s}\right)\right\|_{2}^{2}+2\left\|\frac{1}{2}\left(Y_{r}-Y_{s}\right)\right\|_{2}^{2} .
$$

The first term on the right-hand side is $\geq 2 \Delta^{2}$ by definition of $\Delta$, so taking limits $r, s \rightarrow+\infty$ we have what we need, that is, that $\left(Y_{n}\right)$ is indeed Cauchy.

Let $Y$ be the limit of $\left(Y_{n}\right)$ in $L^{2}(\Omega, \mathcal{G}, \mathbb{P})$. Note that by the triangle inequality

$$
\Delta \leq\|X-Y\|_{2} \leq\left\|X-Y_{n}\right\|_{2}+\left\|Y_{n}-Y\right\|_{2} \rightarrow \Delta
$$

as $n \rightarrow+\infty$. As a result, for any $Z \in L^{2}(\Omega, \mathcal{G}, \mathbb{P})$ and $t \in \mathbb{R}$,

$$
\|X-Y-t Z\|_{2}^{2} \geq \Delta^{2}=\|X-Y\|_{2}^{2}
$$

so that, expanding and rearranging, we have

$$
-2 t\langle Z, X-Y\rangle+t^{2}\|Z\|_{2}^{2} \geq 0,
$$

which is only possible for every $t \in \mathbb{R}$ if the first term is 0 .
Uniqueness follows from the parallelogram law and the definition of $\Delta$.

We return to the proof of existence of the conditional expectation. We use the standard machinery.

Proof of Theorem B.6.1 (i). The previous theorem implies that conditional expectations exist for indicators and simple functions. Now take $X \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and write $X=X^{+}-X^{-}$, so we can assume $X$ is in fact nonnegative without loss of generality. Using the staircase function

$$
X^{(r)}= \begin{cases}0, & \text { if } X=0 \\ (i-1) 2^{-r}, & \text { if }(i-1) 2^{-r}<X \leq i 2^{-r} \leq r \\ r, & \text { if } X>r,\end{cases}
$$

we have $0 \leq X^{(r)} \uparrow X$. Let $Y^{(r)}=\mathbb{E}\left[X^{(r)} \mid \mathcal{G}\right]$. Using an argument similar to the proof of uniqueness, it follows that $U \geq 0$ implies $\mathbb{E}[U \mid \mathcal{G}] \geq 0$ for a simple function $U$. Using linearity (which is immediate from the definition), we then have $Y^{(r)} \uparrow Y:=\lim \sup Y^{(r)}$ which is measurable in $\mathcal{G}$. By (MON),

$$
\mathbb{E}[Y ; G]=\mathbb{E}[X ; G], \forall G \in \mathcal{G}
$$

That concludes the proof.
Before deriving some properties, we give a few examples.
Example B.6.3. If $X \in L^{1}(\Omega, \mathcal{G}, \mathbb{P})$ then $\mathbb{E}[X \mid \mathcal{G}]=X$ almost surely trivially.
Example B.6.4. If $\mathcal{G}=\{\emptyset, \Omega\}$, then $\mathbb{E}[X \mid \mathcal{G}]=\mathbb{E}[X]$.
Example B.6.5. Let $A, B \in \mathcal{F}$ with $0<\mathbb{P}[B]<1$. If $\mathcal{G}=\left\{\emptyset, B, B^{c}, \Omega\right\}$ and $X=\mathbf{1}_{A}$, then

$$
\mathbb{P}[A \mid \mathcal{G}]= \begin{cases}\frac{\mathbb{P}[A \cap B]}{\mathbb{P}[B]}, & \text { on } \omega \in B, \\ \frac{\mathbb{P}\left[A \cap B^{c}\right]}{\mathbb{P}\left[B^{c}\right]}, & \text { on } \omega \in B^{c} .\end{cases}
$$

Intuition about the conditional expectation sometimes breaks down.
Example B.6.6. On $(\Omega, \mathcal{F}, \mathbb{P})=((0,1], \mathcal{B}(0,1], \lambda)$, let $\mathcal{G}$ be the $\sigma$-algebra of all countable and co-countable (i.e., whose complement in ( 0,1 ] is countable) subsets of $(0,1]$. Then $\mathbb{P}[G] \in\{0,1\}$ for all $G \in \mathcal{G}$ and

$$
\mathbb{E}[X ; G]=\mathbb{E}[\mathbb{E}[X] ; G]=\mathbb{E}[X] \mathbb{P}[G],
$$

so that $\mathbb{E}[X \mid \mathcal{G}]=\mathbb{E}[X]$. Yet, $\mathcal{G}$ contains all singletons and we seemingly have "full information," which would lead to the wrong guess $\mathbb{E}[X \mid \mathcal{G}]=X$.

We show that the conditional expectation behaves similarly to the ordinary expectation. Below all $X$ and $X_{i} \mathrm{~s}$ are in $L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{G}$ is a sub $\sigma$-algebra of $\mathcal{F}$.

Lemma B.6.7 (cLIN). If $a_{1}, a_{2} \in \mathbb{R}$, then $\mathbb{E}\left[a_{1} X_{1}+a_{2} X_{2} \mid \mathcal{G}\right]=a_{1} \mathbb{E}\left[X_{1} \mid \mathcal{G}\right]+$ $a_{2} \mathbb{E}\left[X_{2} \mid \mathcal{G}\right]$ a.s.

Proof. Use the linearity of expectation and the fact that a linear combination of random variables in $\mathcal{G}$ is also in $\mathcal{G}$.

Lemma B.6.8 (cPOS). If $X \geq 0$ then $\mathbb{E}[X \mid \mathcal{G}] \geq 0$ a.s.
Proof. Let $Y=\mathbb{E}[X \mid \mathcal{G}]$ and assume for contradiction that $\mathbb{P}[Y<0]>0$. There is $n \geq 1$ such that $\mathbb{P}\left[Y<-n^{-1}\right]>0$. But that implies, for $G=\left\{Y<-n^{-1}\right\}$,

$$
\mathbb{E}[X ; G]=\mathbb{E}[Y ; G]<-n^{-1} \mathbb{P}[G]<0,
$$

a contradiction.
Lemma B.6.9 (cMON). If $0 \leq X_{n} \uparrow X$ then $\mathbb{E}\left[X_{n} \mid \mathcal{G}\right] \uparrow \mathbb{E}[X \mid \mathcal{G}]$ a.s.
Proof. Let $Y_{n}=\mathbb{E}\left[X_{n} \mid \mathcal{G}\right]$. By (cLIN) and (cPOS), $0 \leq Y_{n} \uparrow$. Then letting $Y=\limsup Y_{n}$, by (MON),

$$
\mathbb{E}[X ; G]=\mathbb{E}[Y ; G],
$$

for all $G \in \mathcal{G}$.
Lemma B.6.10 (cFATOU). If $X_{n} \geq 0$ then $\mathbb{E}\left[\liminf X_{n} \mid \mathcal{G}\right] \leq \liminf \mathbb{E}\left[X_{n} \mid \mathcal{G}\right]$ a.s.

Proof. Note that, for $n \geq m$,

$$
X_{n} \geq Z_{m}:=\inf _{k \geq m} X_{k} \uparrow \in \mathcal{G},
$$

so that $\inf _{n \geq m} \mathbb{E}\left[X_{n} \mid \mathcal{G}\right] \geq \mathbb{E}\left[Z_{m} \mid \mathcal{G}\right]$. Applying (cMON)

$$
\mathbb{E}\left[\lim Z_{m} \mid \mathcal{G}\right]=\lim \mathbb{E}\left[Z_{m} \mid \mathcal{G}\right] \leq \lim \inf _{n \geq m} \mathbb{E}\left[X_{n} \mid \mathcal{G}\right] .
$$

Lemma B.6.11 (cDOM). If $X_{n} \leq V \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and $X_{n} \rightarrow X$ a.s., then

$$
\mathbb{E}\left[X_{n} \mid \mathcal{G}\right] \rightarrow \mathbb{E}[X \mid \mathcal{G}] \text { a.s. }
$$

Proof. Applying (cFATOU) to $W_{n}:=2 V-\left|X_{n}-X\right| \geq 0$,

$$
\begin{aligned}
\mathbb{E}[2 V \mid \mathcal{G}] & =\mathbb{E}\left[\liminf _{n} W_{n} \mid \mathcal{G}\right] \\
& \leq \liminf _{n} \mathbb{E}\left[W_{n} \mid \mathcal{G}\right] \\
& =\mathbb{E}[2 V \mid \mathcal{G}]-\underset{n}{\liminf } \mathbb{E}\left[\left|X_{n}-X\right| \mid \mathcal{G}\right],
\end{aligned}
$$

so we must have

$$
\liminf _{n} \mathbb{E}\left[\left|X_{n}-X\right| \mid \mathcal{G}\right]=0
$$

Now use that $\left|\mathbb{E}\left[X_{n}-X \mid \mathcal{G}\right]\right| \leq \mathbb{E}\left[\left|X_{n}-X\right| \mid \mathcal{G}\right]$ (which follows from (cPOS)).
Lemma B.6.12 (cJENSEN). If $f$ is convex and $\mathbb{E}[|f(X)|]<+\infty$ then

$$
f(\mathbb{E}[X \mid \mathcal{G}]) \leq \mathbb{E}[f(X) \mid \mathcal{G}] .
$$

In addition, we highlight (without proof) the following important properties of the conditional expectation.

Lemma B.6.13 (Taking out what is known). If $X \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ and $Z \in \mathrm{~m} \mathrm{\mathcal{G}}$ is bounded or if $X$ is bounded and $Z \in L^{1}(\Omega, \mathcal{G}, \mathbb{P})$, then $\mathbb{E}[Z X \mid \mathcal{G}]=Z \mathbb{E}[X \mid \mathcal{G}]$. This is also true if $X, Z \geq 0, \mathbb{E}[X]<+\infty$ and $\mathbb{E}[Z X]<+\infty$, or $X \in$ $L^{2}(\Omega, \mathcal{F}, \mathbb{P})$ and $Z \in L^{2}(\Omega, \mathcal{G}, \mathbb{P})$.
Lemma B. 6.14 (Role of independence). If $X \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$ is independent of $\mathcal{H}$ then $\mathbb{E}[X \mid \mathcal{H}]=\mathbb{E}[X]$. In fact, if $\mathcal{H}$ is independent of $\sigma(\sigma(X), \mathcal{G})$, then $\mathbb{E}[X \mid \sigma(\mathcal{G}, \mathcal{H})]=\mathbb{E}[X \mid \mathcal{G}]$.
Lemma B.6.15 (Conditioning on an independent random variable). Suppose $X, Y$ are independent. Let $\phi$ be a function with $\mathbb{E}|\phi(X, Y)|<+\infty$ and let $g(x)=$ $\mathbb{E}(\phi(x, Y))$. Then,

$$
\mathbb{E}(\phi(X, Y) \mid X)=g(X)
$$

Lemma B.6.16 (Tower property). If $\mathcal{H} \subseteq \mathcal{G}$ is a $\sigma$-algebra and $X \in L^{1}(\Omega, \mathcal{F}, \mathbb{P})$

$$
\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}] \mid \mathcal{H}]=\mathbb{E}[\mathbb{E}[X \mid \mathcal{H}] \mid \mathcal{G}]=\mathbb{E}[X \mid \mathcal{H}] .
$$

That is, the "smallest $\sigma$-algebra wins."
An important special case of the latter, also known as the law of total probability or the law of total expectation, is $\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}]]=\mathbb{E}[X]$.

One last useful property:
Lemma B.6.17. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. If $Y_{1}=Y_{2}$ a.s. on $B \in \mathcal{F}$ then $\mathbb{E}\left[Y_{1} \mid \mathcal{F}\right]=\mathbb{E}\left[Y_{2} \mid \mathcal{F}\right]$ a.s. on $B$.

## B. 7 Filtered spaces

Finally we define stochastic processes. Let $E$ be a set and let $\mathcal{E}$ be a $\sigma$-algebra defined over $E$.

Definition B.7.1. $A$ stochastic process (or process) is a collection $\left\{X_{t}\right\}_{t \in \mathscr{T}}$ of $(E, \mathcal{E})$-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\mathscr{T}$ is an arbitrary index set.

Here is a typical example.
Example B.7.2. When $\mathscr{T}=\mathbb{Z}_{+}($or $\mathscr{T}=\mathbb{N}$ or $\mathscr{T}=\mathbb{Z})$ we have a discretetime process, in which case we often write the process as a sequence $\left(X_{t}\right)_{t \geq 0}$. For instance:

- $X_{0}, X_{1}, X_{2}, \ldots$ i.i.d. random variables;
- $\left(S_{t}\right)_{t \geq 0}$ where $S_{t}=\sum_{i \leq t} X_{i}$ with $X_{i}$ as above.

We let

$$
\mathcal{F}_{t}=\sigma\left(X_{0}, X_{1}, \ldots, X_{t}\right),
$$

which can be thought of as "the information known up to time $t$." For a fixed $\omega \in \Omega,\left(X_{t}(\omega): t \in \mathscr{T}\right)$ is called a sample path.
Definition B.7.3. $A$ random walk on $\mathbb{R}^{d}$ is a process of the form:

$$
S_{t}=S_{0}+\sum_{i=1}^{t} X_{i}, \quad t \geq 1
$$

where the $X_{i}$ s are i.i.d. in $\mathbb{R}^{d}$, independent of $S_{0}$. The case $X_{i}$ uniform in $\{-1,+1\}$ is called simple random walk on $\mathbb{Z}$.

Filtered spaces provide a formal framework for time-indexed processes. We restrict ourselves to discrete time. (We will not discuss continuous-time processes in this book.)

Definition B.7.4. A filtered space is a tuple $\left(\Omega, \mathcal{F},\left(\mathcal{F}_{t}\right)_{t \in \mathbb{Z}_{+}}, \mathbb{P}\right)$ where:

- $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space;
- $\left(\mathcal{F}_{t}\right)_{t \in \mathbb{Z}_{+}}$is a filtration, that is,

$$
\mathcal{F}_{0} \subseteq \mathcal{F}_{1} \subseteq \cdots \subseteq \mathcal{F}_{\infty}:=\sigma\left(\cup_{t} \mathcal{F}_{t}\right) \subseteq \mathcal{F}
$$

where each $\mathcal{F}_{i}$ is a $\sigma$-algebra.

Definition B.7.5. Fix $\left(\Omega, \mathcal{F},\left(\mathcal{F}_{t}\right)_{t \in \mathbb{Z}_{+}}, \mathbb{P}\right)$. A process $\left(W_{t}\right)_{t \geq 0}$ is adapted if $W_{t} \in$ $\mathcal{F}_{t}$ for all $t$.

Intuitively, in the previous definition, the value of $W_{t}$ is "known at time $t$."
Definition B.7.6. A process $\left(C_{t}\right)_{t \geq 1}$ is predictable if $C_{t} \in \mathcal{F}_{t-1}$ for all $t \geq 1$.
Example B.7.7. Continuing Example B.7.2. The collection $\left(\mathcal{F}_{t}\right)_{t \geq 0}$ forms a filtration. The process $\left(S_{t}\right)_{t \geq 0}$ is adapted. On the other hand, the process $C_{t}=$ $\mathbf{1}\left\{S_{t-1} \leq k\right\}$ is predictable.

## Bibliography

[Abb18] Emmanuel Abbe. Community Detection and Stochastic Block Models. Foundations and Trends® in Communications and Information Theory, 14(1-2):1-162, June 2018. Publisher: Now Publishers, Inc.
[ABH16] Emmanuel Abbe, Afonso S. Bandeira, and Georgina Hall. Exact Recovery in the Stochastic Block Model. IEEE Transactions on Information Theory, 62(1):471-487, January 2016. Conference Name: IEEE Transactions on Information Theory.
[Ach03] Dimitris Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. J. Comput. Syst. Sci., 66(4):671-687, 2003.
[AD78] Rudolf Ahlswede and David E. Daykin. An inequality for the weights of two families of sets, their unions and intersections. Z. Wahrsch. Verw. Gebiete, 43(3):183-185, 1978.
[AF] David Aldous and James Allen Fill. Reversible Markov chains and random walks on graphs. http://www.stat.berkeley.edu/ ~aldous/RWG/book.html.
[AGG89] R. Arratia, L. Goldstein, and L. Gordon. Two Moments Suffice for Poisson Approximations: The Chen-Stein Method. Annals of Probability, 17(1):9-25, January 1989. Publisher: Institute of Mathematical Statistics.
[AJKS22] Alekh Agarwal, Nan Jiang, Sham M. Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. Available at https: //rltheorybook.github.io/rltheorybook_AJKS.pdf, 2022.
[AK97] Noga Alon and Michael Krivelevich. The concentration of the chromatic number of random graphs. Combinatorica, 17(3):303-313, 1997.
[Ald83] David Aldous. Random walks on finite groups and rapidly mixing Markov chains. In Seminar on probability, XVII, volume 986 of Lecture Notes in Math., pages 243-297. Springer, Berlin, 1983.
[Ald90] David J. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. SIAM J. Discrete Math., 3(4):450-465, 1990.
[Ald97] David Aldous. Brownian excursions, critical random graphs and the multiplicative coalescent. Ann. Probab., 25(2):812-854, 1997.
[Alo03] Noga Alon. Problems and results in extremal combinatorics. I. Discrete Math., 273(1-3):31-53, 2003. EuroComb'01 (Barcelona).
[AMS09] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Explo-ration-exploitation tradeoff using variance estimates in multi-armed bandits. Theoretical Computer Science, 410(19):1876-1902, April 2009.
[AN04] K. B. Athreya and P. E. Ney. Branching processes. Dover Publications, Inc., Mineola, NY, 2004. Reprint of the 1972 original [Springer, New York; MR0373040].
[ANP05] Dimitris Achlioptas, Assaf Naor, and Yuval Peres. Rigorous location of phase transitions in hard optimization problems. Nature, 435:759-764, 2005.
[AS11] N. Alon and J.H. Spencer. The Probabilistic Method. Wiley Series in Discrete Mathematics and Optimization. Wiley, 2011.
[AS15] Emmanuel Abbe and Colin Sandon. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In Venkatesan Guruswami, editor, IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015, pages 670-688. IEEE Computer Society, 2015.
[Axl15] Sheldon Axler. Linear algebra done right. Undergraduate Texts in Mathematics. Springer, Cham, third edition, 2015.
[AZ18] Martin Aigner and Günter M. Ziegler. Proofs from The Book. Springer, Berlin, sixth edition, 2018. See corrected reprint of the 1998 original [ MR1723092], Including illustrations by Karl H. Hofmann.
[Azu67] Kazuoki Azuma. Weighted sums of certain dependent random variables. Tôhoku Math. J. (2), 19:357-367, 1967.
[BA99] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. Science, 286(5439):509-512, 1999.
[BBFdlV00] D. Barraez, S. Boucheron, and W. Fernandez de la Vega. On the fluctuations of the giant component. Combin. Probab. Comput., 9(4):287-304, 2000.
[BC03] Bo Brinkman and Moses Charikar. On the Impossibility of Dimension Reduction in L1. In Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science, 2003.
[BCB12] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. Now Publishers, 2012. Google-Books-ID: R12skwEACAAJ.
[BCMR06] Christian Borgs, Jennifer T. Chayes, Elchanan Mossel, and Sébastien Roch. The Kesten-Stigum reconstruction bound is tight for roughly symmetric binary channels. In FOCS, pages 518-530, 2006.
[BD97] Russ Bubley and Martin E. Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In 38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997, pages 223-231. IEEE Computer Society, 1997.
[BDDW08] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. Constr. Approx., 28(3):253-263, 2008.
[BDJ99] Jinho Baik, Percy Deift, and Kurt Johansson. On the distribution of the length of the longest increasing subsequence of random permutations. J. Amer. Math. Soc., 12(4):1119-1178, 1999.
[Ber46] S.N. Bernstein. Probability Theory (in russian). M.-L. Gostechizdat, 1946.
[Ber14] Nathanaël Berestycki. Lectures on mixing times: A crossroad between probability, analysis and geometry. Available at https: //homepage.univie.ac.at/nathanael.berestycki/ wp-content/uploads/2022/05/mixing3.pdf, 2014.
[BH57] S. R. Broadbent and J. M. Hammersley. Percolation processes. I. Crystals and mazes. Proc. Cambridge Philos. Soc., 53:629-641, 1957.
[BH16] Afonso S. Bandeira and Ramon van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. The Annals of Probability, 44(4):2479-2506, July 2016. Publisher: Institute of Mathematical Statistics.
[Bil12] P. Billingsley. Probability and Measure. Wiley Series in Probability and Statistics. Wiley, 2012.
[BKW14] Itai Benjamini, Gady Kozma, and Nicholas Wormald. The mixing time of the giant component of a random graph. Random Structures Algorithms, 45(3):383-407, 2014.
[BLM13] S. Boucheron, G. Lugosi, and P. Massart. Concentration Inequalities: A Nonasymptotic Theory of Independence. OUP Oxford, 2013.
[Bol81] Béla Bollobás. Random graphs. In Combinatorics (Swansea, 1981), volume 52 of London Math. Soc. Lecture Note Ser., pages 80-102. Cambridge Univ. Press, Cambridge-New York, 1981.
[Bol84] Béla Bollobás. The evolution of random graphs. Trans. Amer. Math. Soc., 286(1):257-274, 1984.
[Bol98] Béla Bollobás. Modern graph theory, volume 184 of Graduate Texts in Mathematics. Springer-Verlag, New York, 1998.
[Bol01] Béla Bollobás. Random graphs, volume 73 of Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, second edition, 2001.
[BR06a] Béla Bollobás and Oliver Riordan. Percolation. Cambridge University Press, New York, 2006.
[BR06b] Béla Bollobás and Oliver Riordan. A short proof of the HarrisKesten theorem. Bull. London Math. Soc., 38(3):470-484, 2006.
[Bre17] Pierre Bremaud. Discrete probability models and methods, volume 78 of Probability Theory and Stochastic Modelling. Springer, Cham, 2017. Probability on graphs and trees, Markov chains and random fields, entropy and coding.
[Bre20] Pierre Bremaud. Markov chains-Gibbs fields, Monte Carlo simulation and queues, volume 31 of Texts in Applied Mathematics. Springer, Cham, 2020. Second edition [of 1689633].
[Bro89] Andrei Z. Broder. Generating random spanning trees. In FOCS, pages 442-447. IEEE Computer Society, 1989.
[BRST01] Béla Bollobás, Oliver Riordan, Joel Spencer, and Gábor Tusnády. The degree sequence of a scale-free random graph process. Random Structures Algorithms, 18(3):279-290, 2001.
[BS89] Ravi Boppona and Joel Spencer. A useful elementary correlation inequality. J. Combin. Theory Ser. A, 50(2):305-307, 1989.
[Bub10] Sébastien Bubeck. Bandits Games and Clustering Foundations. phdthesis, Université des Sciences et Technologie de Lille - Lille I, June 2010.
[BV04] S.P. Boyd and L. Vandenberghe. Convex Optimization. Berichte über verteilte messysteme. Cambridge University Press, 2004.
[Car85] Thomas Keith Carne. A transmutation formula for Markov chains. Bull. Sci. Math. (2), 109(4):399-405, 1985.
[CDL ${ }^{+}$12] P. Cuff, J. Ding, O. Louidor, E. Lubetzky, Y. Peres, and A. Sly. Glauber dynamics for the mean-field Potts model. J. Stat. Phys., 149(3):432-477, 2012.
[CF07] Colin Cooper and Alan Frieze. The cover time of sparse random graphs. Random Structures Algorithms, 30(1-2):1-16, 2007.
[Che52] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. Ann. Math. Statistics, 23:493-507, 1952.
[Che70] Jeff Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In Problems in analysis (Sympos. in honor of Salomon Bochner, Princeton Univ., Princeton, N.J., 1969), pages 195-199. Princeton Univ. Press, Princeton, N.J., 1970.
[Che75] Louis H. Y. Chen. Poisson Approximation for Dependent Trials. The Annals of Probability, 3(3):534-545, 1975. Publisher: Institute of Mathematical Statistics.
[Chu97] Fan R. K. Chung. Spectral graph theory, volume 92 of CBMS Regional Conference Series in Mathematics. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 1997.
[CL06] Fan Chung and Linyuan Lu. Complex graphs and networks, volume 107 of CBMS Regional Conference Series in Mathematics. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2006.
[CR92] V. Chvatal and B. Reed. Mick gets some (the odds are on his side) [satisfiability]. In Foundations of Computer Science, 1992. Proceedings., 33rd Annual Symposium on, pages 620-627, Oct 1992.
[Cra38] H. Cramér. Sur un nouveau théorème-limite de la théorie des probabilités. Actualités Scientifiques et Industrielles, 736:5-23, 1938.
[CRR ${ }^{+}$89] Ashok K. Chandra, Prabhakar Raghavan, Walter L. Ruzzo, Roman Smolensky, and Prasoon Tiwari. The electrical resistance of a graph captures its commute and cover times (detailed abstract). In David S. Johnson, editor, STOC, pages 574-586. ACM, 1989.
[CRT06a] Emmanuel J. Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inform. Theory, 52(2):489-509, 2006.
[CRT06b] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. Comm. Pure Appl. Math., 59(8):1207-1223, 2006.
[CT05] Emmanuel J. Candès and Terence Tao. Decoding by linear programming. IEEE Trans. Inform. Theory, 51(12):4203-4215, 2005.
[CW08] E.J. Candès and M.B. Wakin. An introduction to compressive sampling. Signal Processing Magazine, IEEE, 25(2):21-30, March 2008.
[Dev98] Luc Devroye. Branching processes and their applications in the analysis of tree structures and tree algorithms. In Michel Habib,

Colin McDiarmid, Jorge Ramirez-Alfonsin, and Bruce Reed, editors, Probabilistic Methods for Algorithmic Discrete Mathematics, volume 16 of Algorithms and Combinatorics, pages 249-314. Springer Berlin Heidelberg, 1998.
[Dey] Partha Dey. Lecture notes on "Stein-Chen method for Poisson approximation". https://faculty.math.illinois.edu/ ~psdey/414CourseNotes.pdf.
[ $\left.\mathrm{DGG}^{+} 00\right]$ Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, Mark Jerrum, and Michael Mitzenmacher. An extension of path coupling and its application to the Glauber dynamics for graph colourings (extended abstract). In Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 2000), pages 616-624. ACM, New York, 2000.
[dH] Frank den Hollander. Probability theory: The coupling method, 2012. http://websites.math.leidenuniv.nl/probability/ lecturenotes/CouplingLectures.pdf.
[Dia88] Persi Diaconis. Group representations in probability and statistics, volume 11 of Institute of Mathematical Statistics Lecture NotesMonograph Series. Institute of Mathematical Statistics, Hayward, CA, 1988.
[Dia09] Persi Diaconis. The Markov chain Monte Carlo revolution. Bull. Amer. Math. Soc. (N.S.), 46(2):179-205, 2009.
[Die10] Reinhard Diestel. Graph theory, volume 173 of Graduate Texts in Mathematics. Springer, Heidelberg, fourth edition, 2010.
[DKLP11] Jian Ding, Jeong Han Kim, Eyal Lubetzky, and Yuval Peres. Anatomy of a young giant component in the random graph. Random Structures Algorithms, 39(2):139-178, 2011.
[DMR11] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Evolutionary trees and the ising model on the bethe lattice: a proof of steel's conjecture. Probability Theory and Related Fields, 149:149-189, 2011. 10.1007/s00440-009-0246-2.
[DMS00] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin. Structure of growing networks with preferential linking. Phys. Rev. Lett., 85:4633-4636, Nov 2000.
[Doe38] Wolfgang Doeblin. Exposé de la théorie des chaînes simples constantes de markoff à un nombre fini d'états. Rev. Math. Union Interbalkan, 2:77-105, 1938.
[Don06] David L. Donoho. Compressed sensing. IEEE Trans. Inform. Theory, 52(4):1289-1306, 2006.
[Doo01] J.L. Doob. Classical Potential Theory and Its Probabilistic Counterpart. Classics in Mathematics. Springer Berlin Heidelberg, 2001.
[DP11] Jian Ding and Yuval Peres. Mixing time for the Ising model: a uniform lower bound for all graphs. Ann. Inst. Henri Poincaré Probab. Stat., 47(4):1020-1028, 2011.
[DS84] P.G. Doyle and J.L. Snell. Random walks and electric networks. Carus mathematical monographs. Mathematical Association of America, 1984.
[DS91] Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of Markov chains. Ann. Appl. Probab., 1(1):36-61, 1991.
[DSC93a] Persi Diaconis and Laurent Saloff-Coste. Comparison techniques for random walk on finite groups. Ann. Probab., 21(4):2131-2156, 1993.
[DSC93b] Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. Ann. Appl. Probab., 3(3):696-730, 1993.
[DSS22] Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large $k$. Ann. of Math. (2), 196(1):1-388, 2022.
[Dur85] Richard Durrett. Some general results concerning the critical exponents of percolation processes. Z. Wahrsch. Verw. Gebiete, 69(3):421-437, 1985.
[Dur06] R. Durrett. Random Graph Dynamics. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2006.
[Dur10] R. Durrett. Probability: Theory and Examples. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2010.
[Dur12] Richard Durrett. Essentials of stochastic processes. Springer Texts in Statistics. Springer, New York, second edition, 2012.
[DZ10] Amir Dembo and Ofer Zeitouni. Large deviations techniques and applications, volume 38 of Stochastic Modelling and Applied Probability. Springer-Verlag, Berlin, 2010. Corrected reprint of the second (1998) edition.
[Ebe] Andreas Eberle. Markov Processes. 2021.https://uni-bonn. sciebo.de/s/kzTUFff5FrWGAay.
[EKPS00] W. S. Evans, C. Kenyon, Y. Peres, and L. J. Schulman. Broadcasting on trees and the Ising model. Ann. Appl. Probab., 10(2):410-433, 2000.
[ER59] P. Erdős and A. Rényi. On random graphs. I. Publ. Math. Debrecen, 6:290-297, 1959.
[ER60] P. Erdős and A. Rényi. On the evolution of random graphs. Magyar Tud. Akad. Mat. Kutató Int. Közl., 5:17-61, 1960.
[Fel71] William Feller. An introduction to probability theory and its applications. Vol. II. Second edition. John Wiley \& Sons, Inc., New York-London-Sydney, 1971.
[FK16] Alan Frieze and MichałKaroński. Introduction to random graphs. Cambridge University Press, Cambridge, 2016.
[FKG71] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. Comm. Math. Phys., 22:89-103, 1971.
[FM88] P. Frankl and H. Maehara. The Johnson-Lindenstrauss lemma and the sphericity of some graphs. J. Combin. Theory Ser. B, 44(3):355362, 1988.
[Fos53] F. G. Foster. On the stochastic matrices associated with certain queuing processes. Ann. Math. Statistics, 24:355-360, 1953.
[FR98] AlanM. Frieze and Bruce Reed. Probabilistic analysis of algorithms. In Michel Habib, Colin McDiarmid, Jorge Ramirez-Alfonsin, and Bruce Reed, editors, Probabilistic Methods for Algorithmic Discrete Mathematics, volume 16 of Algorithms and Combinatorics, pages 36-92. Springer Berlin Heidelberg, 1998.
[FR08] N. Fountoulakis and B. A. Reed. The evolution of the mixing rate of a simple random walk on the giant component of a random graph. Random Structures Algorithms, 33(1):68-86, 2008.
[FR13] Simon Foucart and Holger Rauhut. A Mathematical Introduction to Compressive Sensing. Applied and Numerical Harmonic Analysis. Birkhäuser Basel, 2013.
[FV18] S. Friedli and Y. Velenik. Statistical mechanics of lattice systems. Cambridge University Press, Cambridge, 2018. A concrete mathematical introduction.
[GC11] Aurélien Garivier and Olivier Cappé. The KL-UCB Algorithm for Bounded Stochastic Bandits and Beyond. In Proceedings of the 24th Annual Conference on Learning Theory, pages 359-376. JMLR Workshop and Conference Proceedings, December 2011. ISSN: 1938-7228.
[GCS ${ }^{+}$14] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. Bayesian data analysis. Texts in Statistical Science Series. CRC Press, Boca Raton, FL, third edition, 2014.
[Gi159] E. N. Gilbert. Random graphs. Ann. Math. Statist., 30:1141-1144, 1959.
[GL06] Dani Gamerman and Hedibert Freitas Lopes. Markov chain Monte Carlo. Texts in Statistical Science Series. Chapman \& Hall/CRC, Boca Raton, FL, second edition, 2006. Stochastic simulation for Bayesian inference.
[Gri97] Geoffrey Grimmett. Percolation and disordered systems. In Lectures on probability theory and statistics (Saint-Flour, 1996), volume 1665 of Lecture Notes in Math., pages 153-300. Springer, Berlin, 1997.
[Gri10a] Geoffrey Grimmett. Probability on graphs, volume 1 of Institute of Mathematical Statistics Textbooks. Cambridge University Press, Cambridge, 2010. Random processes on graphs and lattices.
[Gri10b] G.R. Grimmett. Percolation. Grundlehren der mathematischen Wissenschaften. Springer, 2010.
[Gri75] David Griffeath. A maximal coupling for Markov chains. Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, 31:95-106, 1974/75.
[GS20] Geoffrey R. Grimmett and David R. Stirzaker. Probability and random processes. Oxford University Press, Oxford, 2020. Fourth edition [of 0667520].
[Ham57] J. M. Hammersley. Percolation processes. II. The connective constant. Proc. Cambridge Philos. Soc., 53:642-645, 1957.
[Har] Nicholas Harvey. Lecture notes for CPSC 536N: Randomized Algorithms. http://www.cs.ubc.ca/~nickhar/w12/.
[Har60] T. E. Harris. A lower bound for the critical probability in a certain percolation process. Proc. Cambridge Philos. Soc., 56:13-20, 1960.
[Har63] Theodore E. Harris. The theory of branching processes. Die Grundlehren der mathematischen Wissenschaften, Band 119. Springer-Verlag, Berlin; Prentice Hall, Inc., Englewood Cliffs, N.J., 1963.
[Haz16] Elad Hazan. Introduction to Online Convex Optimization. Foundations and Trends® in Optimization, 2(3-4):157-325, August 2016. Publisher: Now Publishers, Inc.
[HJ13] Roger A. Horn and Charles R. Johnson. Matrix analysis. Cambridge University Press, Cambridge, second edition, 2013.
[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. Bull. Amer. Math. Soc. (N.S.), 43(4):439-561 (electronic), 2006.
[HMRAR98] M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, editors. Probabilistic methods for algorithmic discrete mathematics, volume 16 of Algorithms and Combinatorics. Springer-Verlag, Berlin, 1998.
[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. J. Amer. Statist. Assoc., 58:13-30, 1963.
[HS07] Thomas P. Hayes and Alistair Sinclair. A general lower bound for mixing of single-site dynamics on graphs. Ann. Appl. Probab., 17(3):931-952, 2007.
[IM98] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In Jeffrey Scott Vitter, editor, STOC, pages 604-613. ACM, 1998.
[Jan90] Svante Janson. Poisson approximation for large deviations. Random Structures Algorithms, 1(2):221-229, 1990.
[JH01] Galin L. Jones and James P. Hobert. Honest exploration of intractable probability distributions via Markov chain Monte Carlo. Statist. Sci., 16(4):312-334, 2001.
[JL84] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In Conference in modern analysis and probability (New Haven, Conn., 1982), volume 26 of Contemp. Math., pages 189-206. Amer. Math. Soc., Providence, RI, 1984.
[JLR11] S. Janson, T. Luczak, and A. Rucinski. Random Graphs. Wiley Series in Discrete Mathematics and Optimization. Wiley, 2011.
[JS89] Mark Jerrum and Alistair Sinclair. Approximating the permanent. SIAM J. Comput., 18(6):1149-1178, 1989.
[Kan86] Masahiko Kanai. Rough isometries and the parabolicity of Riemannian manifolds. J. Math. Soc. Japan, 38(2):227-238, 1986.
[Kar90] Richard M. Karp. The transitive closure of a random digraph. Random Structures Algorithms, 1(1):73-93, 1990.
[Kes80] Harry Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. Comm. Math. Phys., 74(1):41-59, 1980.
[Kes82] Harry Kesten. Percolation theory for mathematicians, volume 2 of Progress in Probability and Statistics. Birkhäuser, Boston, Mass., 1982.
[KP] Júlia Komjáthy and Yuval Peres. Lecture notes for Markov chains: mixing times, hitting times, and cover times, 2012. SaintPetersburg Summer School. http://www.win.tue.nl/~jkomjath/ SPBlecturenotes.pdf.
[KRS] Michael J. Kozdron, Larissa M. Richards, and Daniel W. Stroock. Determinants, their applications to markov processes, and a random walk proof of kirchhoff's matrix tree theorem, 2013. Preprint available at http://arxiv.org/abs/1306.2059.
[KS66a] H. Kesten and B. P. Stigum. Additional limit theorems for indecomposable multidimensional Galton-Watson processes. Ann. Math. Statist., 37:1463-1481, 1966.
[KS66b] H. Kesten and B. P. Stigum. A limit theorem for multidimensional Galton-Watson processes. Ann. Math. Statist., 37:1211-1223, 1966.
[KS67] H. Kesten and B. P. Stigum. Limit theorems for decomposable multi-dimensional Galton-Watson processes. J. Math. Anal. Appl., 17:309-338, 1967.
[KS05] Gil Kalai and Shmuel Safra. Threshold Phenomena and Influence: Perspectives from Mathematics, Computer Science, and Economics. In Computational Complexity and Statistical Physics. Oxford University Press, December 2005. _eprint: https://academic.oup.com/book/0/chapter/354512033/chapter-pdf/43716844/isbn-9780195177374-book-part-8.pdf.
[KSK76] John G. Kemeny, J. Laurie Snell, and Anthony W. Knapp. Denumerable Markov chains. Springer-Verlag, New York-Heidelberg-Berlin, second edition, 1976. With a chapter on Markov random fields, by David Griffeath, Graduate Texts in Mathematics, No. 40.
[Law05] Gregory F. Lawler. Conformally invariant processes in the plane, volume 114 of Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2005.
[Law06] Gregory F. Lawler. Introduction to stochastic processes. Chapman \& Hall/CRC, Boca Raton, FL, second edition, 2006.
[Led01] M. Ledoux. The Concentration of Measure Phenomenon. Mathematical surveys and monographs. American Mathematical Society, 2001.
[Lin02] Torgny Lindvall. Lectures on the coupling method. Dover Publications, Inc., Mineola, NY, 2002. Corrected reprint of the 1992 original.
[LL10] G.F. Lawler and V. Limic. Random Walk: A Modern Introduction. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2010.
[Lov83] L. Lovász. Submodular functions and convexity. In Mathematical programming: the state of the art (Bonn, 1982), pages 235-257. Springer, Berlin, 1983.
[Lov12] László Lovász. Large networks and graph limits, volume 60 of American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2012.
[LP16] Russell Lyons and Yuval Peres. Probability on trees and networks, volume 42 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, 2016.
[LP17] David A. Levin and Yuval Peres. Markov chains and mixing times. American Mathematical Society, Providence, RI, 2017. Second edition of [ MR2466937], With contributions by Elizabeth L. Wilmer, With a chapter on "Coupling from the past" by James G. Propp and David B. Wilson.
[LPW06] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. Markov chains and mixing times. American Mathematical Society, 2006.
[LR85] T.L Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. Advances in Applied Mathematics, 6(1):4-22, March 1985.
[LS88] Gregory F. Lawler and Alan D. Sokal. Bounds on the $L^{2}$ spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. Trans. Amer. Math. Soc., 309(2):557-580, 1988.
[LS12] Eyal Lubetzky and Allan Sly. Critical Ising on the square lattice mixes in polynomial time. Comm. Math. Phys., 313(3):815-836, 2012.
[LS20] Tor Lattimore and Csaba Szepesvári. Bandit Algorithms. Cambridge University Press, 2020.
[Lu90] Tomasz Ł uczak. Component behavior near the critical point of the random graph process. Random Structures Algorithms, 1(3):287310, 1990.
[Lug] Gabor Lugosi. Concentration-of-measure inequalities, 2004. Available at http://www.econ.upf.edu/~1ugosi/anu.pdf.
[LuPW94] Tomasz Ł uczak, Boris Pittel, and John C. Wierman. The structure of a random graph at the point of the phase transition. Trans. Amer. Math. Soc., 341(2):721-748, 1994.
[Lyo83] Terry Lyons. A simple criterion for transience of a reversible Markov chain. Ann. Probab., 11(2):393-402, 1983.
[Lyo90] Russell Lyons. Random walks and percolation on trees. Ann. Probab., 18(3):931-958, 1990.
[Mat88] Peter Matthews. Covering problems for Markov chains. Ann. Probab., 16(3):1215-1228, 1988.
[Mau79] Bernard Maurey. Construction de suites symétriques. C. R. Acad. Sci. Paris Sér. A-B, 288(14):A679-A681, 1979.
[McD89] Colin McDiarmid. On the method of bounded differences. In Surveys in combinatorics, 1989 (Norwich, 1989), volume 141 of London Math. Soc. Lecture Note Ser., pages 148-188. Cambridge Univ. Press, Cambridge, 1989.
[ML86] Anders Martin-Löf. Symmetric sampling procedures, general epidemic processes and their threshold limit theorems. J. Appl. Probab., 23(2):265-282, 1986.
[ML98] Anders Martin-Löf. The final size of a nearly critical epidemic, and the first passage time of a Wiener process to a parabolic barrier. J. Appl. Probab., 35(3):671-682, 1998.
[MNS15a] Elchanan Mossel, Joe Neeman, and Allan Sly. Consistency thresholds for the planted bisection model. In Rocco A. Servedio and Ronitt Rubinfeld, editors, Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pages 69-75. ACM, 2015.
[MNS15b] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. Probab. Theory Related Fields, 162(3-4):431-461, 2015.
[Mor05] Frank Morgan. Real analysis. American Mathematical Society, Providence, RI, 2005.
[Mos01] E. Mossel. Reconstruction on trees: beating the second eigenvalue. Ann. Appl. Probab., 11(1):285-300, 2001.
[Mos03] E. Mossel. On the impossibility of reconstructing ancestral data and phylogenies. J. Comput. Biol., 10(5):669-678, 2003.
[Mos04] E. Mossel. Phase transitions in phylogeny. Trans. Amer. Math. Soc., 356(6):2379-2404, 2004.
[MP03] E. Mossel and Y. Peres. Information flow on trees. Ann. Appl. Probab., 13(3):817-844, 2003.
[MR95] Rajeev Motwani and Prabhakar Raghavan. Randomized algorithms. Cambridge University Press, Cambridge, 1995.
[MS86] Vitali D. Milman and Gideon Schechtman. Asymptotic theory of finite-dimensional normed spaces, volume 1200 of Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.
[MT06] Ravi Montenegro and Prasad Tetali. Mathematical aspects of mixing times in Markov chains. Found. Trends Theor. Comput. Sci., 1(3): $x+121,2006$.
[MT09] Sean Meyn and Richard L. Tweedie. Markov chains and stochastic stability. Cambridge University Press, Cambridge, second edition, 2009. With a prologue by Peter W. Glynn.
[MU05] Michael Mitzenmacher and Eli Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, New York, NY, USA, 2005.
[Nic18] Bogdan Nica. A brief introduction to spectral graph theory. EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2018.
[Nor98] J. R. Norris. Markov chains, volume 2 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1998. Reprint of 1997 original.
[NP10] Asaf Nachmias and Yuval Peres. The critical random graph, with martingales. Israel J. Math., 176:29-41, 2010.
[NW59] C. St. J. A. Nash-Williams. Random walk and electric currents in networks. Proc. Cambridge Philos. Soc., 55:181-194, 1959.
[Ott49] Richard Otter. The multiplicative process. Ann. Math. Statistics, 20:206-224, 1949.
[Pem91] Robin Pemantle. Choosing a spanning tree for the integer lattice uniformly. Ann. Probab., 19(4):1559-1574, 1991.
[Pem00] Robin Pemantle. Towards a theory of negative dependence. J. Math. Phys., 41(3):1371-1390, 2000. Probabilistic techniques in equilibrium and nonequilibrium statistical physics.
[Per] Yuval Peres. Course notes on Probability on trees and networks, 2004. http://stat-www.berkeley.edu/~peres/notes1.pdf.
[Per09] Yuval Peres. The unreasonable effectiveness of martingales. In Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09, pages 997-1000, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
[Pet] Gábor Pete. Probability and geometry on groups. Lecture notes for a graduate course. http://www.math.bme.hu/~gabor/PGG.html.
[Pey08] Rémi Peyre. A probabilistic approach to Carne's bound. Potential Anal., 29(1):17-36, 2008.
[Pit76] J. W. Pitman. On coupling of Markov chains. Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, 35(4):315-322, 1976.
[Pit90] Boris Pittel. On tree census and the giant component in sparse random graphs. Random Structures Algorithms, 1(3):311-342, 1990.
[RAS15] Firas Rassoul-Agha and Timo Seppäläinen. A course on large deviations with an introduction to Gibbs measures, volume 162 of Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2015.
[RC04] Christian P. Robert and George Casella. Monte Carlo statistical methods. Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2004.
[Res92] Sidney Resnick. Adventures in stochastic processes. Birkhäuser Boston, Inc., Boston, MA, 1992.
[Roc10] Sebastien Roch. Toward Extracting All Phylogenetic Information from Matrices of Evolutionary Distances. Science, 327(5971):1376-1379, 2010.
[Rom15] Dan Romik. The surprising mathematics of longest increasing subsequences, volume 4 of Institute of Mathematical Statistics Textbooks. Cambridge University Press, New York, 2015.
[RS17] Sebastien Roch and Allan Sly. Phase transition in the sample complexity of likelihood-based phylogeny inference. Probability Theory and Related Fields, 169(1):3-62, Oct 2017.
[RT87] WanSoo T. Rhee and Michel Talagrand. Martingale inequalities and NP-complete problems. Math. Oper. Res., 12(1):177-181, 1987.
[Rud73] Walter Rudin. Functional analysis. McGraw-Hill Series in Higher Mathematics. McGraw-Hill Book Co., New York-DüsseldorfJohannesburg, 1973.
[Rus78] Lucio Russo. A note on percolation. Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, 43(1):39-48, 1978.
[SE64] M. F. Sykes and J. W. Essam. Exact critical percolation probabilities for site and bond problems in two dimensions. J. Mathematical Phys., 5:1117-1127, 1964.
[SJ89] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. Inform. and Comput., 82(1):93-133, 1989.
[Spi56] Frank Spitzer. A combinatorial lemma and its application to probability theory. Trans. Amer. Math. Soc., 82:323-339, 1956.
[Spi12] Daniel A. Spielman. Lecture notes on spectral graph theory. https://www.cs.yale.edu/homes/spielman/ 561/2012/index.html, 2012.
[SS87] E. Shamir and J. Spencer. Sharp concentration of the chromatic number on random graphs $G_{n, p}$. Combinatorica, 7(1):121-129, 1987.
[SS03] Elias M. Stein and Rami Shakarchi. Fourier analysis, volume 1 of Princeton Lectures in Analysis. Princeton University Press, Princeton, NJ, 2003. An introduction.
[SS05] Elias M. Stein and Rami Shakarchi. Real analysis, volume 3 of Princeton Lectures in Analysis. Princeton University Press, Princeton, NJ, 2005. Measure theory, integration, and Hilbert spaces.
[SSBD14] S. Shalev-Shwartz and S. Ben-David. Understanding Machine Learning: From Theory to Algorithms. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014.
[Ste] J. E. Steif. A mini course on percolation theory, 2009. http: //www. math.chalmers.se/~steif/perc.pdf.
[Ste72] Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. II: Probability theory, pages 583-602, 1972.
[Ste97] J. Michael Steele. Probability theory and combinatorial optimization, volume 69 of CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.
[Ste98] G. W. Stewart. Matrix algorithms. Vol. I. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1998. Basic decompositions.
[Str65] V. Strassen. The existence of probability measures with given marginals. Ann. Math. Statist., 36:423-439, 1965.
[Str14] Daniel W. Stroock. Doeblin’s Theory for Markov Chains. In An Introduction to Markov Processes, pages 25-47. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
[SW78] P. D. Seymour and D. J. A. Welsh. Percolation probabilities on the square lattice. Ann. Discrete Math., 3:227-245, 1978. Advances in graph theory (Cambridge Combinatorial Conf., Trinity College, Cambridge, 1977).
[Tao] Terence Tao. Open question: deterministic UUP matrices. https://terrytao.wordpress.com/2007/07/02/ open-question-deterministic-uup-matrices/.
[Var85] Nicholas Th. Varopoulos. Long range estimates for Markov chains. Bull. Sci. Math. (2), 109(3):225-252, 1985.
[vdH10] Remco van der Hofstad. Percolation and random graphs. In New perspectives in stochastic geometry, pages 173-247. Oxford Univ. Press, Oxford, 2010.
[vdH17] Remco van der Hofstad. Random graphs and complex networks. Vol. 1. Cambridge Series in Statistical and Probabilistic Mathematics, [43]. Cambridge University Press, Cambridge, 2017.
[vdHK08] Remco van der Hofstad and Michael Keane. An elementary proof of the hitting time theorem. Amer. Math. Monthly, 115(8):753-756, 2008.
[Vem04] Santosh S. Vempala. The random projection method. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 65. American Mathematical Society, Providence, RI, 2004. With a foreword by Christos H. Papadimitriou.
[Ver18] Roman Vershynin. High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge University Press, September 2018. Google-Books-ID: TahxDwAAQBAJ.
[vH16] Ramon van Handel. Probability in high dimension. http://www. princeton.edu/~rvan/APC550.pdf, 2016.
[Vil09] Cédric Villani. Optimal transport, volume 338 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 2009. Old and new.
[Wen75] J. G. Wendel. Left-continuous random walk and the Lagrange expansion. Amer. Math. Monthly, 82:494-499, 1975.
[Whi32] Hassler Whitney. Non-separable and planar graphs. Trans. Amer. Math. Soc., 34(2):339-362, 1932.
[Wil91] David Williams. Probability with martingales. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge, 1991.
[Wil96] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In Gary L. Miller, editor, STOC, pages 296-303. ACM, 1996.
[YP14] Se-Young Yun and Alexandre Proutière. Community detection via random and adaptive sampling. In Maria-Florina Balcan, Vitaly Feldman, and Csaba Szepesvári, editors, Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014, volume 35 of JMLR Workshop and Conference Proceedings, pages 138-175. JMLR.org, 2014.
[Yur76] V. V. Yurinskin̆. Exponential inequalities for sums of random vectors. J. Multivariate Anal., 6(4):473-499, 1976.

## Index

$\varepsilon$-net, see epsilon-net

Azuma-Hoeffding inequality, see tail bounds
balancing vectors, 34
ballot theorem, 126
balls and bins, 154
bandits
definition, 170
optimism in the face of uncertainty, 171
upper confidence bound, 171
Bernstein's inequality, see tail bounds
Berry-Esséen theorem, 157
binary classification
definitions, 102-103
empirical risk minimizer, 103
no free lunch, 103
binary search tree, 442
binomial coefficients
binomial theorem, 502
bounds, 23, 502
definition, x
bond percolation
definition, 18
FKG, 265
Bonferroni inequalities, 116
Boole's inequality, see union bound
Boolean functions
random $k$-SAT problem, 37
bounded differences inequality, 152
branching number, 58
branching processes
binomial offspring, 438
duality principle, 432
exploration process, 431
extinction, 417, 418
Galton-Watson process, 416
Galton-Watson tree, 416, 422
geometric offspring, 495
history, 432, 496
infinite line of descent, 496
linear functionals, 427
multitype, see multitype branching processes
Poisson offspring, 420, 433, 435
positive regular case, 425
random-walk representation, 429, 431, 433
canonical paths, 406
Cauchy-Schwarz inequality, 515, 519
Cavender-Farris-Neyman (CFN) model, 451
Cayley's formula, see trees
chaining method, 91
Chebyshev polynomials, 366
Chebyshev's association inequality, 320
Chebyshev's inequality, see tail bounds
Cheeger's inequality, 385, 491
Chen-Stein method
dissociated case, 305, 318
main result, 301
negatively related variables, 324
positively related variables, 324
Stein coupling, 302, 306
Stein equation, 301
Chernoff-Cramér bound, see tail bounds
Chernoff-Cramér method, 67, 117, 146
community recovery, 345
compressed sensing, 96
concentration inequalities, see tail bounds, 120
concentration of measure, 156
conditional expectation
definition, 522-525
examples, 525-526
law of total probability, 527
properties, 526-527
tower property, 527
congestion ratio, 406
convergence in probability, 32
convex duality
Lagrangian, 203
weak duality, 203
coupling
coalescence time, 282
coalescing, 282
coupling inequality, 241, 247
definition, 235
Erdős-Rényi graph model, 466, 483
independent, 235
Markovian, 246, 282
maximal coupling, 242
monotone, 235, 236, 251, 254, 257
path coupling, 293, 299
splitting, 283
stochastic domination, 251
Strassen's theorem, 254
coupon collector, 31, 286
Courant-Fischer theorem, 331, 336, 340, 341, 384
covariance, 30
covering number, 87,111
cumulant-generating function, 64, 66
Curie-Weiss model, 401
cutset, 58
Davis-Kahan theorem, 343, 349
dependency graph, 46
dimension reduction, 93
Dirichlet
energy, 383, 406
form, 383
principle, 206, 229
problem, 183
drift condition, 190
Dudley's inequality, 91, 111
edge expansion constant, see networks
Efron-Stein inequality, 151
electrical networks
definitions, 192-194
Dirichlet energy, see Dirichlet
effective conductance, 200, 206
effective resistance, 200, 214, 229
flow, see flows
Kirchhoff's cycle law, 193
Kirchhoff's node law, 193
Kirchhoff's resistance formula, 218
Nash-Williams inequality, 206, 229
Ohm's law, 193, 219
parallel law, 196
Rayleigh's principle, 209, 219
series law, 196
Thomson's principle, 202, 208
voltage, 192
empirical measure, 110,112
epsilon-net, 86, 89, 97-101
Erdős-Rényi graph model
chromatic number, 158
clique number, $48,318,321$
cluster, 468
connectivity, 52
definition, 18
degree sequence, 247
evolution, 466
exploration process, 468
FKG, 265
giant component, 467, 481
isolated vertices, 51
largest connected component, 467
maximum degree, 78
positive associations, 263
random walk, 490
subgraph containment, 269
threshold function, 47-55
exhaustive sequence, 229
expander graphs
( $d, \alpha$ )-expander family, 394
Pinsker's model, 395
factorials
bounds, 22, 502
definition, $x$
Stirling's formula, 502
Fenchel-Legendre dual, 66
first moment method, 36-41, 46, 48, 49, 52-54, 56, 144, 161
first moment principle, 33, 36, 104
first-step analysis, 186
FKG
condition, 264, 321
inequality, 265, 321, 325
measure, 264, 322
flows
current, 193
definition, 6
energy, 202, 209, 212
flow-conservation constraints, 193, 210
max-flow min-cut theorem, 7, 254
strength, 193
to $\infty, 209,211$
gambler's ruin, 138-141, 195, 201
Gibbs random fields, 19
Glauber dynamics
definition, 21
fast mixing, 297, 401
graph Laplacian
connectivity, 334
definition, 333
degree, 335
eigenvalues, 334
Fiedler vector, 334
network, 339
normalized, 340
quadratic form, 333
graphs
3-1 tree, see trees
$b$-ary tree $\widehat{\mathbb{T}}_{b}^{\ell}, 5,290$
$n$-clique, 382
adjacency matrix, see matrices
bridge, 218
Cayley's formula, see trees
chromatic number, 8,158
clique, 3 , 48
clique number, 48,318
coloring, 8
complete graph $K_{n}, 5$
cutset, 58, 206
cycle $C_{n}, 5,285$
definitions, 2-9
degree, 78
diameter, 368
directed, 8,9
expander, see expander graphs
flow, see flows
graph distance, 4
hypercube $\mathbb{Z}_{2}^{n}$, 5, 286
incidence matrix, see matrices
independent set, 8,34
infinite, 5
infinite binary tree, 451

Laplacian, see graph Laplacian
matching, 8
matrix representation, see matrices
multigraph, 2
network, see networks
oriented incidence matrix, see matrices
perfect matching, 8
spanning arborescence, 222
torus $\mathbb{L}_{n}^{d}, 5$
tree, see trees
Turán graphs, 35
Green function, see Markov chains
Hölder's inequality, 515
Hamming distance, 178
harmonic functions, see Markov chains
Harper's vertex isoperimetric theorem, 158
Harris' inequality, 320, 325
Harris' theorem, see percolation
hitting time, see stopping time
Hoeffding's inequality, see tail bounds
Hoeffding's lemma, see tail bounds, 146
Holley's inequality, 266, 321, 325
increasing event, see posets
indicator trick, 36
inherited property, 421
Ising model
boundary conditions, 260
complete graph, 401
definition, 19
FKG, 265
Glauber dynamics, 286, 297, 400
magnetization, 401
random cluster, 451
trees, 451, 499
isoperimetric inequality, 381
Janson's inequality, 271, 325

Jensen's inequality, 518
Johnson-Lindenstrauss
distributional lemma, 93
lemma, 93-96
Kesten's theorem, see percolation
knapsack problem, 79
Kolmogorov's maximal inequality, 141
Kullback-Leibler divergence, 68
Laplacian
graphs, see graph Laplacian
Markov chains, 183, 204, 323
networks, see graph Laplacian
large deviations, 69, 120
law of total probability, see conditional expectation
laws of large numbers, 32-33
Lipschitz
condition, 178-180
process, 88
Lyapounov function, see Markov chains
Markov chain Monte Carlo, 370
Markov chain tree theorem, 231
Markov chains
average occupation time, 186
birth-death, 194, 228, 376
bottleneck ratio, see networks
Chapman-Kolmogorov, 10
commute time, 214
commute time identity, 214, 230, 290
construction, 10
cover time, 124
decomposition theorem, 127
definitions, 9-17
Doeblin's condition, 283, 322
escape probability, 197, 205
examples, 9-16
exit law, 186
exit probability, 186
first return time, 124
first visit time, 124, 181
Green function, 129, 186, 197
harmonic functions, 181
hitting times, 313
irreducible set, 127
lower bound, 285
Lyapounov function, 190, 294
Markov property, 10
martingales, 134
Matthews' cover time bounds, 132, 230
mean exit time, 186
Metropolis algorithm, 15
mixing time, see mixing times
positive recurrence, 127
potential theory, 185
recurrence, 199, 209-214
recurrent state, 127
relaxation time, 358
reversibility, 181, 365
splitting, see coupling
stationary measure, 128
stochastic domination, 260, 267
stochastic monotonicity, 258
strong Markov property, 124
uniform geometric ergodicity, 284
Varopoulos-Carne bound, 365, 414
Markov's inequality, see tail bounds, 64
martingales, 456
Azuma-Hoeffding inequality, see tail bounds
convergence theorem, 142, 417
definition, 133
Doob martingale, 135, 147, 158
Doob's submartingale inequality, 141, 146
edge exposure martingale, 159
exposure martingale, 158
hitting time, 230

Markov chain, see Markov chains
martingale difference, 147
optional stopping theorem, 137
orthogonality of increments, 143, 148
stopped process, 136
submartingale, 133
supermartingale, 133
vertex exposure martingale, 158
matrices
2-norm, 88
adjacency, 2, 332
block, 328
diagonal, 328
graph Laplacian, see graph Laplacian
incidence, 3
oriented incidence, 9
orthogonal, 328
spectral norm, $88,179,341$
spectral radius, 410, 425
spectrum, 410
stochastic matrix, 9
symmetric, 328
maximal Azuma-Hoeffding inequality, see tail bounds
maximum principle, $14,24,183,229$
method of bounded differences, see tail bounds
method of moments, 119,120
method of random paths, 216, 229
Metropolis algorithm, see Markov chains
minimum bisection problem, 347
Minkowski's inequality, 516
mixing times
$b$-ary tree, 290, 391
cutoff, 289, 364, 411
cycle, $322,359,392$
definition, 17
diameter, 369
diameter bound, 369
distinguishing statistic, 289, 322
hypercube, 362, 393, 413
lower bound, 284, 287, 297, 322, 368, 369
random walk on cycle, 285
random walk on hypercube, 286
separation distance, 357
upper bound, 297
moment-generating functions
$\chi^{2}, 73$
definition, 28
Gaussian, 65
Poisson, 67
Rademacher, 65
moments, 28
exponential moment, see momentgenerating functions
multi-armed bandits, see bandits
multitype branching processes
definitions, 423-425
Kesten-Stigum bound, 455
mean matrix, 424
nonsingular case, 424
negative associations, 219
networks
cut, 382
definition, 8
edge boundary, 381
edge expansion, 382
vertex boundary, 394
no free lunch, see binary classification
notation, ix-xi
operator
compact, 376
norm, 377
spectral radius, 377
optimal transport, 323
optional stopping theorem, see martingales

Pólya's theorem, see random walk
Pólya's urn, 142
packing number, 87
Pajor's lemma, 114
parity functions, 363
Parseval's identity, 352
pattern matching, 155
peeling method, see slicing method
percolation
contour lemma, 42
critical exponents, 441
critical value, $40,55,272,423,438$
dual lattice, 41, 273
Galton-Watson tree, 422
Harris' theorem, 273, 325
Kesten's theorem, 273, 325
on $\mathbb{L}^{2}, 40,272$
on $\mathbb{L}^{d}, 254$
on a graph, 236
on infinite trees, 55, 144
percolation function, 40, 55, 254, 272, 438
RSW lemma, 273, 322, 325
permutations
Erdős-Szekeres Theorem, 39
longest increasing subsequence, 38
random, 38
Perron-Frobenius theory
Perron vector, 426
theorem, 12, 426, 456
Poincaré inequality, 152, 385, 386, 406
Poisson approximation, 245, 269
Poisson equation, 187
Poisson trials, 69
posets
decreasing event, 264
definition, 253
increasing event, 254, 320
positive associations
definition, 263
strong, 321
positively correlated events, 264
probabilistic method, 33-36, 93, 103, 112
probability generating function, 418
probability spaces
definitions, 503-506
distribution function, 506
expectation, 513-519
filtered spaces, 122, 528
Fubini's theorem, 519
independence, 510-513
process, 528
random variables, 506-510
pseudo-regret, 171
random graphs
Erdős-Rényi, see Erdős-Rényi graph model
preferential attachment, 19, 163
stochastic blockmodel, 345
random projection, 93
random target lemma, 184
random variables
$\chi^{2}, 73,95$
Bernoulli, 68, 235, 244, 250, 307
binomial, 68, 287, 320
Gaussian, 30, 70, 73, 93
geometric, 32, 284
Poisson, 67, 241, 244, 250, 252, 420, 496
Rademacher, 65, 70
uncorrelated, 32, 117, 143
uniform, 31
random walk
b-ary tree, 290
asymmetric random walk on $\mathbb{Z}, 380$
biased random walk on $\mathbb{Z}, 139,236$
cycle, 285,359
hypercube, 286, 362
lazy, 16, 238, 284
loop erasure, 220
on a graph, $10-15$
on a network, 20
Pólya's theorem, 215
reflection principle, 125
simple random walk on $\mathbb{Z}, 11,127$, $138,183,211,365,528$
simple random walk on $\mathbb{Z}^{d}, 238$
simple random walk on a graph, 20
tree, 239
Wald's identities, 137-141
Rayleigh quotient, 330, 336, 384
reconstruction problem
definition, 451
MAP estimator, 453
solvability, 452
reflection principle, see random walk
relaxation times
cycle, 361
hypercube, 364
restricted isometry property, $97,101-102$
rough embedding, 210, 211
rough equivalence, 211,230
rough isometry, 230
RSW lemma, see percolation
Sauer's lemma, 108, 110, 112
second moment method, $45,46,48,50$, $52,54,57,58,60$
set balancing, 66
shattering, 108
simple random walk on a graph, see random walk
slicing method, 176
span, 331
sparse signal recovery, 97
sparsity, 96
spectral clustering, 349
spectral gap, 358
spectral theorem, 328,340
Spitzer's combinatorial lemma, 436
Stirling's formula, see factorials
stochastic bandits, see bandits
stochastic domination
binomial, 470
coupling, see coupling
Markov chains, 258
monotonicity, 320
posets, 254
real random variables, 250
stochastic processes
adapted, 529
definition, 528
filtration, 528
predictable, 529
sample path, 528
supremum, 85-93
stopping time
cover time, see Markov chains
definition, 123
first return time, see Markov chains
first visit time, see Markov chains
hitting time, 123
strong Markov property, see Markov chains
Strassen's theorem, 325
sub-exponential variable, see tail bounds sub-Gaussian increments, 91
sub-Gaussian variable, see tail bounds, 146
submodularity, 325
symmetrization, 71, 106
tail bounds
Azuma-Hoeffding inequality, 145, $158,159,227,380$
Bernstein's inequality, 78, 493

Bernstein's inequality for bounded variables, 77
Chebyshev's inequality, 29, 45, 247, 481
Chernoff bound, 69, 155, 365
Chernoff-Cramér bound, 64, 75, 85
definitions, 28-29
general Bernstein inequality, 76
general Hoeffding inequality, 70, 89
Hoeffding's inequality, 82, 112
Hoeffding's inequality for bounded variables, 71
Hoeffding's lemma, 72
Markov's inequality, 29, 36, 104, 141, 146
McDiarmid's inequality, 153, 159, 164, 178, 179
method of bounded differences, 153 , 156, 163
Paley-Zygmund inequality, 46
sub-exponential, 74
sub-Gaussian, 70, 85, 107
Talagrand's inequality, 179
Talagrand's inequality, see tail bounds
threshold phenomena, 37, 41, 47
tilting, 73
total variation distance, 16
tower property, see conditional expectation
trees
3-1 tree, 61
Cayley's formula, 5, 26, 54, 499
characterization, 4-5
definition, 4
infinite, 55, 438
uniform spanning tree, 218-226
type, see recurrence
uniform spanning trees, see trees
union bound, 36, 116

Varopoulos-Carne bound, see Markov chains
VC dimension, 108, 111
Wald's identities, see random walk
Wasserstein distance, 323
Weyl's inequality, 342


[^0]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

[^1]:    *More commonly, "sub-exponential" refers to the case $\alpha=\sqrt{\nu}$.

[^2]:    ${ }^{\dagger}$ Technically, it suffices to assume that there is a countable $\mathscr{T}_{0} \subseteq \mathscr{T}$ such that $\sup _{t \in \mathscr{T}} X_{t}=$ $\sup _{t \in \mathscr{T}_{0}} X_{t}$ almost surely.

[^3]:    ${ }^{\dagger}$ If the $\varepsilon$-net $N$ is not included in $\mathscr{T}$, the Lipschitz condition has to hold on a larger subset that includes both.

[^4]:    ${ }^{\S}$ Technically, $\mathscr{T}$ could be part of a larger countable space by the discussion above.

[^5]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

[^6]:    *The above argument is more natural in the setting of continuous-time Markov chains, but we will not introduced them here.

[^7]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

[^8]:    *It is also called "conductance," but that terminology clashes with our use of the term.

[^9]:    ${ }^{\dagger}$ In terms of random walk, this corresponds to choosing a neighbor uniformly at random and taking the stationary measure equal to the degree. Note that scaling the stationary measure does not affect the edge expansion constant.

[^10]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

[^11]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

[^12]:    Version: December 20, 2023
    Modern Discrete Probability: An Essential Toolkit
    Copyright © 2023 Sébastien Roch

