

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Perspectives on Multi-Agent Systems: From Characterization to Intervention

Permalink

<https://escholarship.org/uc/item/1jh2638f>

ISBN

9798293893010

Author

Zhao, Geng

Publication Date

2025-08-01

Peer reviewed|Thesis/dissertation

Perspectives on Multi-Agent Systems: From Characterization to Intervention

By

Geng Zhao

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Christian Borgs, Co-Chair
Assistant Professor Jiantao Jiao, Co-Chair
Professor Jennifer Chayes
Professor Itai Ashlagi

Summer 2025

Perspectives on Multi-Agent Systems: From Characterization to Intervention

Copyright 2025
by
Geng Zhao

Abstract

Perspectives on Multi-Agent Systems: From Characterization to Intervention

by

Geng Zhao

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Christian Borgs, Co-Chair

Assistant Professor Jiantao Jiao, Co-Chair

Many modern applications of computing, ranging from school choices to epidemic modeling, are fundamentally multi-agent in nature, involving interactions among strategic and often adaptive entities. This thesis presents theoretical methodologies for understanding macroscopic behaviors in such multi-agent systems and establishing insights that inform both inference and learning within such systems as well as the design and improvement of mechanisms and interventions.

Across a spectrum of settings, we showcase how mathematical analysis reveals novel understanding on the structure and dynamics of various multi-agent systems shaped by strategic interactions and network structures. In Stackelberg games with information asymmetry, we demonstrate that expert follower behavior can paradoxically hinder the leader’s learning process, highlighting subtle challenges in multi-agent learning dynamics. In large two-sided matching markets, we introduce a notion of agent fitness under heterogeneous preferences and prove that stable outcomes exhibit universal statistical structure. Such structural characterization lends further insights on school choice systems, as we show that the widely used deferred acceptance mechanism is, with high probability, far from Pareto efficient. We then turn to the problem of estimating global parameters in networks based on local samples and propose a new robustness condition under which certain local estimations become viable. In another setting where network structure is crucial, we study the control of epidemics in networks with community structure, assessing the effectiveness of common intervention strategies, such as social distancing and travel restrictions. Finally, we highlight a simple token-based local policy for service systems such as kidney exchanges, and show that it effectively maintains stability of the system.

Together, the approaches we develop form a methodological toolkit for analyzing multi-agent systems with limited information, drawing on concepts from information theory, probability,

random graph theory, and game theory. They demonstrate how mathematical analysis of local information and agent behavior can yield predictive and design-relevant insights into the global dynamics and performance of large-scale systems.

To my parents and my wife, Zoey

Contents

Contents	ii
List of Figures	v
1 Introduction	1
1.1 Organization	2
1.2 Notations	3
2 Online Learning in Stackelberg Games with an Omniscient Follower	4
2.1 Introduction	4
2.2 Formulation	8
2.3 Linear Stackelberg Games: Curse of Expertise	9
2.4 UCB With Side Observations	11
2.5 Beyond UCB	16
2.6 Conclusion	19
3 Fitness and Welfare in Random Matching Markets	20
3.1 Introduction	20
3.2 Model and Preliminaries	23
3.3 Warm-Up: Uniformly Random Preferences	26
3.4 Main Results: Large Market Behavior	27
3.5 Proof Ideas	29
3.6 Approximating the Probability of Stability	32
3.7 Uniform Influence Profiles	36
3.8 Empirical Distribution of Values and Ranks	37
3.9 Conclusion	38
4 Efficiency and Envy in School Choices	39
4.1 Introduction	39
4.2 Model	42
4.3 Identifying (Un)Improvable Students	44
4.4 Quantifying (Un)Improvable Students	45

4.5	Conclusion	52
5	Local Estimation of the Size of k-Core in Networks	53
5.1	Introduction	53
5.2	Preliminaries	61
5.3	Emergence of a Non-Empty 2-Core in the Giant Above Criticality	63
5.4	Continuity of Percolation Function for the 2-Core	68
5.5	Convergence of the Relative Size of 2-Core in Percolation	70
5.6	Robustness of Algorithm 2	77
5.7	Discussion: Beyond the 2-Core	79
6	Epidemics on Networks with Community Structures: Dynamics and Interventions	81
6.1	Introduction	81
6.2	An SIR Epidemic Model on Two Communities With Travel	84
6.3	Results	86
6.4	Proof Overview	90
6.5	Discussion	93
7	Power of Two Choices in Token Systems	96
7.1	Introduction	96
7.2	Model	100
7.3	Two Agents	101
7.4	The General Case	102
7.5	Strategic Considerations	109
7.6	Conclusion	110
	Bibliography	112
A	Appendix to Chapter 2	128
A.1	Proofs in Section 2.3	128
A.2	Proofs in Section 2.4	129
A.3	Proofs in Section 2.5	132
B	Appendix to Chapter 3	133
B.1	Some Concentration Inequalities	133
B.2	Typical Behaviors of Values in Stable Matching	140
B.3	Number of Stable Matchings: Proof of Proposition 3.6.3	152
B.4	Proof of Proposition 3.7.1	154
B.5	Proof of Proposition B.3.1	158
B.6	Empirical Distribution of Values: Proof of Theorem 3.8.1	162
B.7	Ranks Scale (Approximately) With Values: Proof of Theorem 3.8.3	165

C	Appendix to Chapter 4	171
C.1	Many-to-One Matching	171
C.2	Correlation in Preferences	174
D	Appendix to Chapter 5	176
D.1	Local Algorithm for the Size of 2-Core	176
D.2	Size of the Giant Component	176
E	Appendix to Chapter 6	180
E.1	Preliminaries	180
E.2	Branching Process Upper Bounds	188
E.3	Travelers and Non-Travelers	193
E.4	Proof of Main Theorems	201
F	Appendix to Chapter 7	203
F.1	Proofs from Section 7.2.1	203
F.2	Proofs from Section 7.3	205
F.3	Proof of Theorem 7.4.1	208
F.4	Proof of Theorem 7.4.2	211
F.5	Proofs from Section 7.4	216

List of Figures

3.1	Histograms (counts) of ranks in four arbitrarily chosen stable matchings for a market instance with uniformly random preferences with $n = 10\,000$ men and women. Red vertical dashed lines indicate average ranks on each side; orange dashed lines show the theoretical histogram according to an exponential distribution.	26
4.1	Average fraction of nodes in the largest (blue) and second-largest (red) SCCs (average over 2 000 random envy digraphs for each n).	51
4.2	Average fraction of unimprovable students (average over 2 000 random problems for each n).	51
5.1	The forest $G \setminus H$, oriented in the downstream direction, with vertices annotated with their depths towards H . For the choice of $\ell = 4$, the red vertices cannot possibly be in $C^{(2,\ell)}(G)$	68
5.2	The coloring scheme applied to $G \setminus H$ with the choice of $\ell = 4$. Vertices are annotated with their depths. Gray segments are presented with dashed edges, indicating that they are ignored during the subsequent sprinkling argument. Note that each red or purple segment has size at least ℓ , and at any non-gray vertex, the colored segments that are (completely) contained in its (inclusive) upstream constitute at least $1/3$ of its regular (inclusive) upstream (see Lemma 5.5.2). Evidently, assuming $H \supseteq C^{(2)}(G) \neq \emptyset$, the inclusion of any vertex into a cycle or a path between cycles during sprinkling will result in its entire downstream being included in the 2-core.	72
6.1	Schematic illustration of epidemic trajectory in the two communities (not to scale). τ_1 and τ_2 mark the major “waves” in the two communities: most infections occur within the narrow bands around them. $\tau_{1 \rightarrow 2}$ marks the initial infection of any type-2 individual. τ_{end} indicates the end of the epidemic, i.e., when the number of infectious individual first drops to zero.	88
6.2	Simulation on two communities of 1 million individuals each, with an underlying network generated from a configuration model with geometrically distributed degrees (parameter chosen such the basic reproduction number $R_0 = 2$. Other parameters are: $\beta = 1.5$, $\gamma = 3$, $\rho_{\text{T}} = 10^{-3}$, and $\rho_{\text{H}} = 1$	95

C.1	Average fraction of nodes in the largest (blue) and second-largest (red) SCCs (average over 1,000 random problems for each n), preferences and priorities uniform i.i.d.	174
C.2	Average fraction of nodes in the largest (blue) and second-largest (red) SCCs in a two-tiered model (average over 1 000 random envy digraphs for each n).	175
F.1	(F.20) as a function of β for several values of M when the agents are symmetric and $d = 2$	209

Acknowledgments

Completing this PhD has been a journey made only possible by the support, generosity, and guidance of many people to whom I owe deep gratitude.

First and foremost, I want to express my deepest gratitude towards my co-advisors, Christian Borgs and Jiantao Jiao. I feel fortunate and deeply honored to have learned from and been inspired by both of them. Throughout the years, I have had numerous memorable discussions with Christian, often hours long, on topics ranging from mathematics and research philosophy to career paths and life in general. His engagement, patience, wide-ranging curiosity, and attention to detail has shaped the way I approached research. Jiantao has consistently guided me toward what truly matters—impact and real-world significance. From our very first conversation, I have been continually inspired by his depth, vision, and passion. What I especially appreciate about both of them is the flexibility and freedom they gave me to pursue my own interest in research and career directions. They have both offered me invaluable advice, resources, and opportunities that I feel forever grateful for.

I am also deeply grateful to Jennifer Chayes and Itai Ashlagi for serving as committee members and offering invaluable advice and support over the years. I feel incredibly fortunate and privileged to have Jennifer on my committee, especially considering her many responsibilities and tight schedule. Besides her sharp and helpful feedback, I feel truly inspired by her unbounded enthusiasm and visionary thinking, in both research and leadership. My connection with Itai spans nearly a decade, beginning with an undergraduate research project and continuing steadily through my time at Stanford, experiences in industry, and throughout graduate school. I have learned so much and enjoyed every opportunity to collaborate with him, so much so that we once (half-jokingly) discussed the possibility of completing my Berkeley PhD remotely from across the Bay. I am forever thankful for all the advice, encouragement, critiques, ideas, and opportunities he has offered me.

I also want to thank Ramesh Johari and Amin Saberi for their advice and support throughout the years, dating back to my time at Stanford. Their mentorship was instrumental in leading me to Berkeley and in shaping the course of my academic journey. Many other faculty members generously provided advice, feedback, and encouragement. I am indebted to Mark Braverman, Michael I. Jordan, Nika Haghtalab, Shanghua Teng, Adam Tauman Kalai, Lexin Li, and David Wagner. Their influence on my work and thinking has been profound.

Over the years, I've been incredibly fortunate to work alongside many brilliant peers and collaborators. I would like to thank Hannah Li, Banghua Zhu, Karissa Huang, Julien Piet, Zhanhao Hu, Xiao (Sam) Huang, Chawin Sitawarin, and Süleyman Kerimov, Omer Tamuz, Josué Ortega, among many others. Collaborating with them has been an enriching experience filled with learning, inspiration, and moments of genuine enjoyment. These collaborations have not only shaped my research but also led to lasting friendships that I deeply value.

I am also thankful to my friends, whose steady support, encouragement, and good humor helped me stay grounded throughout this journey. Wen Zhang has been a steadfast friend

since our college years. He introduced me to many interesting ideas and fun problems and has been my go-to person whenever I got stuck debugging my program or research. Kunhe Yang and Hanlin Zhu, my wonderful labmates and close neighbors, brought endless support and joy to my everyday life over the four years. I also feel lucky to have had the friendship of so many others, both at Berkeley and around the world, including Christian Ikeokwu, Ricardo Jose Sandoval, Eric Zhao, Nivasini Ananthakrishnan, Tianhao Wu, Baihe Huang, Brian Lee, Mark Bedaywi, Haozhe Jiang, Ishaq Adan-Ali, Xinyan Hu, Zhanghao Fan, Jiaxin Guan, Shunhua Jiang, Shuangping Li, Shuangning Li, and Ruiming Cao. For all the support, wisdom, and friendship I have received—I am deeply thankful.

Above all, none of this would have been achievable without the unwavering support of my family. My parents encouraged my decision to study abroad, which first led me to Stanford for undergraduate studies and later inspired me to pursue research seriously and embark on a PhD. My wife Zoey's constant love, patience, and belief in me carried me through moments of doubt and helped make this accomplishment possible. It is difficult to imagine having completed this journey without them.

Chapter 1

Introduction

Modern systems of computation and algorithms are becoming increasingly large-scale and decentralized, often deployed in complex, dynamic environments with multiple agents. Algorithms now play central roles in tasks such as school assignment and resource allocation, where individual preferences must be reconciled with global constraints; massive networks that capture social interactions, infrastructure, or communication are analyzed to track contagion, detect vulnerabilities, and shape information flows; looking ahead, fleets of autonomous vehicles interacting with human drivers may soon operate in real time with limited coordination and partial observability. Understanding, designing, and improving such systems presents fundamental challenges.

Complex as they are, such systems can often be abstracted into concrete, albeit difficult, mathematical objects. Algorithms, mechanisms, control policies, estimators, etc. can be thought of as functions, and under appropriate models of the noisy environment, they are just random variables or stochastic processes. For these, we have plenty of tools to study them.

This thesis aims to showcase a few examples of how mathematical analysis can lead to new insights on the structure and dynamics of different kinds of multi-agent systems, ranging from matching markets to epidemics over contact networks. We will address questions such as characterizing welfare and efficiency in matching markets and evaluating intervention measures to control the spread of epidemics. Even though our models are often stylized for tractability reasons (as is common in theoretical studies), we believe that the insights revealed by the theorems are novel and fundamental and should go beyond theory.

Our analytical approach integrates tools from information theory, game theory, random graphs, and stochastic processes. A unifying theme across the chapters is that local properties—of structure or behavior—can collectively give rise to sharp global phenomena. For instance, Chapter 5 establishes conditions under which global parameters of a graph can be estimated from local neighborhood samples; Chapter 7 demonstrates the remarkable effectiveness of a simple local-improvement policy for maintaining stability in token economies.

1.1 Organization

Each of the following chapters examines a different facet of the theme of large multi-agent systems, where we analyze stylized models to characterize outcomes, evaluate policies, and establish fundamental limits of learning and inference.

Chapter 2 begins with the foundational problem of online learning under information asymmetry in Stackelberg games. Here, we study a setting when a leader must learn from responses of an omniscient follower. This setting reveals challenges such as the “curse of expertise,” where the follower’s optimal responses can obscure useful feedback, highlighting the tension between strategic interactions and efficient learning. The chapter is based on [214], a joint work with Banghua Zhu, Jiantao Jiao, and Michael I. Jordan.

Chapter 3 shifts focus to characterizing welfare and fitness in large matching markets. We study random two-sided matching markets and show that, despite the complexity and randomness, the empirical distribution of welfare exhibits surprisingly regular behaviors. This analysis highlights techniques from concentration and large deviation theory, as well as ideas about fixed points and maximum entropy. The chapter is based on a joint work with Itai Ashlagi and Mark Braverman [20].

Chapter 4 builds on the topic of matching markets by addressing the efficiency limitations of stable mechanisms. Using tools from random graph theory, we demonstrate how stable matchings, often obtained through deferred acceptance in practice, suffer substantial inefficiencies. The analysis is carried out with graph theoretic tools applied to the concept of an envy digraph, which helps identify which students can be helped by alternative mechanisms. The chapter is based on a joint work with Josué Ortega, Gabriel Ziegler, and R. Pablo Arribillaga [165].

Chapter 5 moves to large-scale networked systems and considers local algorithms for structural inference. Focusing on the k -core structure in sparse random graphs, this chapter analyzes when and how local neighborhood information can be used to estimate global properties such as the size of the 2-core. We introduce the notion of structural robustness and use ideas from random graph and percolation theory to establish theoretical guarantees on natural local estimators. This chapter is based on a joint work with Christian Borgs [48].

Chapter 6 focuses on the problem of epidemic control in networks with community structures. We analyze SIR epidemic spread across weakly connected communities, accounting for dynamic travel and internal structure. Our results quantify how interventions such as travel bans or social distancing alter the epidemic trajectory, providing new perspectives for understanding the efficacy of epidemic intervention measures. This chapter is based on a joint work with Christian Borgs and Karissa Huang.

Finally, Chapter 7 demonstrates how a simple local rule can sustain system-level stability in token systems. Motivated by applications such as national kidney exchange systems, we explore how participants can use a policy reflecting the idea of the “power of two choices” [156] to coordinate resource sharing. Leveraging Lyapunov theory and fluid approximations for stochastic processes, we establish several desirable properties of the resulting dynamics, even in markets with thin supply. This chapter is based on a joint work with Itai Ashlagi,

Süleyman Kerimov, and Omer Tamuz [24].

1.2 Notations

We summarize some key notations used repeatedly in this thesis.

Sets, vectors, and matrices. $\mathbb{N}$, $\mathbb{Z}$, and $\mathbb{R}$ denote the sets of natural numbers, integers, and real numbers, respectively. $\mathbb{R}_+$ denotes non-negative reals. For a set S , $|S|$ denotes its cardinality. For $n \in \mathbb{N}$, we write $[n] = \{1, \dots, n\}$.

For a vector $\mathbf{v} \in \mathbb{R}^n$, we use $\|\mathbf{v}\|_p$ to denote its ℓ^p norm for $p \in (0, \infty]$. The notation $\text{diag}(\mathbf{v})$ denotes the $n \times n$ diagonal matrix with entries of $\mathbf{v}$ on the diagonal.

Asymptotics. We use standard notations to describe asymptotic behavior. For functions $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, we say $f = O(g)$ or $f \lesssim g$ (resp. $\Omega(g)$ or $f \gtrsim g$) if $\limsup_{x \rightarrow \infty} |f(x)|/g(x) < \infty$ (resp. $\liminf_{x \rightarrow \infty} |f(x)|/g(x) > 0$); $f = o(g)$ or $f \ll g$ if $f(x)/g(x) \rightarrow 0$ as $x \rightarrow \infty$; and $f = \Theta(g)$ or $f \asymp g$ if $f = O(g)$ and $f = \Omega(g)$. We write $\tilde{O}(g)$ to hide polylogarithmic factors, e.g., $\tilde{O}(g) := O(g \log^k g)$ for some constant k .

Probability. $\text{Bern}(p)$ denotes the Bernoulli distribution with success probability $p \in [0, 1]$, and $\text{Bin}(n, p)$ denotes the binomial distribution from n i.i.d. trials each with success probability $p \in [0, 1]$. $\text{Exp}(\lambda)$ and $\text{Poi}(\lambda)$ denote, respectively, the exponential distribution and the Poisson distribution with rate λ . For real-valued random variables X and Y , $X \preceq Y$ denotes stochastic domination of X by Y , i.e., $\mathbb{P}(X > t) \leq \mathbb{P}(Y > t)$ for all $t \in \mathbb{R}$. We say an event $\mathcal{E}_n$ happens with high probability (whp) if $\mathbb{P}(\mathcal{E}_n) \rightarrow 1$ as $n \rightarrow \infty$.

Chapter 2

Online Learning in Stackelberg Games with an Omniscient Follower

Understanding and controlling multi-agent systems necessitate a theoretical foundation for how agents interact strategically. The literature on this topic is abound. In this chapter, we examine a benchmark scenario: online learning in Stackelberg games. This setting highlights the essential challenges that arise from strategic interactions, information asymmetry, and decentralized decision-making. Through analyzing regret and sample complexity in this simple game-theoretic framework, we provide insight into the interplay between information availability and strategic learning.

2.1 Introduction

The multi-agent learning problem [82, 209, 83, 213] has received significant attention reflecting its wide variety of real-world applications, including autonomous driving [192, 187] and human-robot interaction [126, 140, 93, 210]. In a multi-agent system, it is natural to assume that each agent possesses a different set of information due to its different viewpoint and history of actions. This phenomenon is commonly referred to as the property of *information asymmetry* [211]. Such information asymmetry poses challenges to the coordination and cooperation between learning agents. In this chapter, we study how the information asymmetry affects the sample complexity of learning a two-player decentralized cooperative repeated Stackelberg game, with a focus on the setting where the follower is omniscient and myopic, and always best-responds to the leader's actions.

Consider an illustrative example in human-robot interaction where a robot is required to collaborate with a human to achieve some shared objective. This can be formulated as a repeated Stackelberg game where the interactions between human and robot happen in multiple rounds, and the human is an omniscient expert who knows the exact target and how to achieve it. In each round, the robot, as the leader who hopes to learn the world model and human behavior from scratch, first takes some action. After seeing the robot's action,

the human, as an expert follower who possesses perfect information about the world, always best-responds to the robot’s action to maximize their reward. The robot hopes to use as few as possible interactions to learn the world model and human behavior, and eventually find the optimal action that maximizes a shared reward.

Concretely, during each round t of the interaction, the leader first plays an action $a_t \in \mathcal{A}$, and the follower plays another action $b_t \in \mathcal{B}$ upon (perfectly) observing a_t . We assume that the two players share a reward, $r_t = h^*(a_t, b_t) + z_t$, where $z_t \in \mathbb{R}$ is some zero-mean sub-Gaussian noise, h^* belongs to a family $\mathcal{H}$. We also assume that the follower has full knowledge of the reward and always best responds with $b_t \in \arg\max_{b \in \mathcal{B}} h^*(a_t, b)$. However, the leader does not know h^* and can only explore via taking actions a_t and making inferences from past observations $(a_1, b_1, r_1), \dots, (a_{t-1}, b_{t-1}, r_{t-1})$.¹ We are interested in providing tight bound for the Stackelberg regret, defined as

$$\mathcal{R}(T) = \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \left(\max_{b \in \mathcal{B}} h^*(a, b) - \max_{b_t \in \mathcal{B}} h^*(a_t, b_t) \right) \right].$$

The Stackelberg regret characterizes the gap between the reward achieved from the optimal leader action and the reward from the actual leader action a_t .

Compared with the traditional bandit problem, the extra observation of b_t can be viewed as side information accompanying the usual action-reward pair. Depending on how the function family $\mathcal{H}$ and side information b are designed, the complexity of learning for the leader may vary. Here we briefly summarize several illustrative examples where the follower may help or harm the leader’s learning process. We will present a general formalization that encompasses these examples in the next section.

1. **Curse of expertise.** Imagine that in a driving system, the self-driving vehicle (leader) and the human driver (follower) work together to avoid collisions. For most of the aggressive actions the leader takes, the final reward for non-collision is high since the human driver will consistently exert efforts to evade the self-driving vehicle in order to prevent collisions. From the leader’s point of view, aggressive actions lead to similar outcomes as safe actions. The expertise of the human prevents the leader from learning from failure cases.
2. **Imitation Learning.** Consider an assembly robot (leader) that learns to move goods to a destination with a human expert (follower). This can be modeled by the robot choosing a drop-off location, from which the human expert continues to the correct destination. In this simple example, the robot and the human expert cooperate in a “linear” fashion—the expert can complete whatever the robot leaves undone, and upon observation of the expert’s move the robot should simply imitate the behavior of the human expert in the future. This corresponds to an “imitation-based” interaction that can greatly accelerate the learning process.

¹For simplicity, we assume in the introduction that the leader can see $b_1, \dots, b_{t-1}$ without noise. Later we generalize to the case when the observed b_t is also noisy.

3. **Expert-guided learning.** In most cases, the self-driving vehicle may have some target that is similar but not exactly the same as the human driver. For example, they both aim to avoid collision while heading to a different target. In this case, a pure imitation-based learning will fail. But the self-driving vehicle can still glean good driving standards from the human driver. With the extra observation of the behavior of human driver, the self-driving vehicle can learn much faster.

Extending beyond robotics applications, our framework is potentially applicable in various repeated cooperative game settings where direct communication is hard, unreliable, or forbidden. For instance, it captures the learning aspect of language models adjusting to human preferences, personalized digital healthcare, or other settings of AI-human interaction where explicit revelation of utility function is difficult: in such settings, the AI system (e.g., the language model, or the digital “doctor”) works with a human user to achieve a common goal without direct communication of the true preferences or needs; instead, the system must learn them through repeated interactions with the users.

In this chapter, we abstract and formalize these three scenarios into a simple linear Stackelberg game and analyze the sample complexity of this game. We briefly overview our main results in the next section.

2.1.1 Main Results

Contrary to the traditional literature on linear bandits, we show that the worst-case sample complexity for achieving ϵ -Stackelberg regret is at least exponential even when h^* belongs to the linear family $\mathcal{H}_\phi = \{\theta \cdot \phi(a, b)\}$. The hard instance corresponds to the ‘curse of expertise’ example discussed above, where the follower’s best response hurts the observation, and thus harms the whole learning process.

Theorem 2.1.1 (Curse of expertise, informal). *There exists some ϕ such that for any algorithm, we can find some $h^* \in \mathcal{H}_\phi$ with the regret being $\Omega(T^{(d-3)/(d-2)})$.*

This shows that the leader needs an exponential number of samples to learn a good policy even when the reward is linear. We also present an upper bound $O(T^{(d+1)/(d+2)})$ for linear rewards in Theorem 2.3.3.

On the other hand, the side information b_t can also greatly improve the sample complexity when the linear family is structured. We provide an Upper Confidence Bound (UCB) based algorithm [26] that leads to an improved bound in this setting. In particular, we recover the rate for imitation learning when the leader can simply mimic the behavior of the follower.

Theorem 2.1.2 (Imitation learning, informal). *There exists some ϕ such that for any $h^* \in \mathcal{H}_\phi$, when b_t is observed, the leader can achieve regret $O(\log^2(T))$ by imitating the follower behavior. However, when b_t is not observed, the regret is $\Theta(\sqrt{T})$.*

Similarly, we can also design cases where observing b_t helps reduce the problem to a traditional linear bandit, while not observing b_t suffers from exponential sample complexity.

Theorem 2.1.3 (Expert-guided, informal). *There exists some ϕ such that for any $h^* \in \mathcal{H}_\phi$, when b_t is observed, the leader can achieve regret $O(\sqrt{T})$. However, when b_t is not observed, the regret is $\Omega(T^{(d-4)/(d-2)})$.*

In addition to these three examples, we discuss more complicated scenarios where UCB fails and we show that a careful analysis is necessary to achieve a near-optimal rate. In particular, we establish such a rate for polynomial bandits, where the best-response corresponds to a lower degree polynomial, which helps improve the rate when the noise level for reward and the observed follower behavior is similar.

Theorem 2.1.4 (Polynomial bandit, informal). *There exists a family of $2k$ -degree polynomial, such that the regret is $\Theta(\sqrt{d^{2k-1}T})$ when b_t is observed, and $\Theta(\sqrt{d^{2k}T})$ when b_t is not observed.*

2.1.2 Related Work

Decentralized Stackelberg Games. The problem of repeated Stackelberg games has been studied extensively [194, 147, 135, 116], in a standard setting where the leader leads and the myopic follower follows with its best response for the current round.

Kao, Wei, and Subramanian [116] and Lauffer et al. [135] study a similar setting to ours, in which a leader and a follower interact through a cooperative Stackelberg game that comprises a two-stage bandit problem. However, Kao, Wei, and Subramanian [116] restrict their focus to the tabular case where both $\mathcal{A}$ and $\mathcal{B}$ are finite and the reward h^* is uncorrelated for different actions (a, b) . They also assume that both the leader and the agent are running regret-minimization algorithms independently. They show that the classic upper confidence bound (UCB) algorithm for the multi-arm bandit problem can be used for both the leader and the agent, respectively, to achieve asymptotically optimal performance (i.e., no-regret). However, it is unclear that such results can generalize to bandits with function approximation and the case of omniscient agents. Indeed, our results show that the general case (or even just the linear case) is not always statistically tractable. Note also that Lauffer et al. [135] show that the regret can depend exponentially on the dimension of the agent's utility.

Other examples of Stackelberg games include Stackelberg security games [63, 196], strategic learning [96, 72, 143], dynamic task pricing [121] and online contract design [100, 216]. The problem of online learning in contract theory considers a decentralized general-sum Stackelberg game with omniscient agents. It focuses on a special case where the rewards for the leader and the agent are both linear. It is shown in Zhu et al. [216] that one has to pay exponential sample complexity in this setting to achieve small regret in the worst case.

Centralized Stackelberg Game. Centralized Stackelberg games are also well studied in the literature [215, 29, 89, 212], where the machine learning algorithm has control over both the leader and the follower. Bai et al. [29] consider the repeated Stackelberg game where both the leader and the agent learn their optimal actions (a Stackelberg equilibrium) from samples. However, they assume a central controller that can determine the actions of both the leader and the agent. Moreover, they rely on an assumption of a bounded gap between the optimal response and an ϵ -approximate best response. In contrast, in our framework, we assume that the agent's utility is unknown, and that the agent always takes the best response.

Bandit with side information. There has been significant effort in studying bandits with side information [207, 133, 85]. Such side information is generally assumed to be available before a decision. Foster et al. [85] also consider the case when an extra observation is available after taking the actions. However, they mainly focus on the setting of reinforcement learning where the extra observation is the trajectory. Although our observation of follower behavior can also be viewed as side information, it also alters the reward in the Stackelberg game, which changes the structure of the multi-agent problem.

2.2 Formulation

We consider a two-player cooperative Stackelberg bandit game with an omniscient follower.

Let $\mathcal{A} \subseteq \mathbb{R}^{d_1}$ and $\mathcal{B} \subseteq \mathbb{R}^{d_2}$ be compact sets. Up to a scaling factor, we will assume that $\mathcal{A}$ and $\mathcal{B}$ reside inside the unit ball centered at the origin. During each round $t \in [T]$ of interaction, the leader plays an action $a_t \in \mathcal{A}$, and the follower plays $b_t \in \mathcal{B}$ upon (perfectly) observing a_t . The two players both receive a reward $r_t = h^*(a_t, b_t) + z_t$, where $z_t \in \mathbb{R}$ is zero-mean σ_r -sub-Gaussian and is independent of all past events. We will make the realizability assumption that h^* belongs to a (known) family $\mathcal{H}$ of real-valued functions on $\mathbb{B}^{d_1} \times \mathbb{B}^{d_2}$. As is common in the study of bandits, we assume that reward function is bounded, i.e., there exists $C \in (0, \infty)$ such that $0 \leq h \leq C$ for all $h \in \mathcal{H}$. We assume $C = 1$ throughout the chapter unless stated otherwise.

We will assume that the follower, modeled after an expert human player, has full knowledge of the game and can always best respond with an optimal action $b_t \in \operatorname{argmax}_{b \in \mathcal{B}} h^*(a_t, b)$. The leader then makes a noisy observation of b_t , given by $\hat{b}_t = b_t + w_t$, where $w_t \in \mathbb{R}^{d_2}$ is zero-mean σ_b -sub-Gaussian (e.g., component-wise σ_b -sub-Gaussian with independent zero-mean coordinates) and independent of all past events.

For convenience, we denote the set of best responses to leader's action a when the ground truth reward function is h by $b_h^*(a)$. Denote $\bar{h}(a) := \max_{b \in \mathcal{B}} h(a, b)$. The optimal action, unbeknownst to the leader, is denoted $a^* := \operatorname{argmax}_{a \in \mathcal{A}} \bar{h}^*(a)$.

The leader's objective is to minimize the *regret* during T rounds of interactions, defined

as

$$\mathcal{R}(T) = \max_{a \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \bar{h}^*(a) - \bar{h}^*(a_t) \right]. \quad (2.1)$$

We will also focus on the sample complexity of achieving low (average) regret; that is, for some $\epsilon, \delta \in [0, 1]$, the minimal $T \in \mathbb{N}$ such that $\mathcal{R}(T) \leq \epsilon T$.

Notations. We use calligraphic letters for sets and operators, e.g., $\mathcal{A}$. Given a set $\mathcal{A}$, we write $|\mathcal{A}|$ for the cardinality of $\mathcal{A}$. $\mathbb{B}^d$ and $\mathbb{S}^{d-1}$ denote the unit ball and the unit sphere, both centered at the origin, in d -dimensional Euclidean space. Vectors are assumed to be column vectors except for the probability and measure vectors. For a vector $v \in \mathbb{R}^d$ and an integer $i \in \mathbb{N}$, we use v_i to denote the i -th element of v , and v_{-i} to denote the vector of all elements in v except for v_i . For two n -dimensional vectors x and y , we use $x \cdot y = x^\top y$ to denote their inner product.

Parameterized family. In subsequent discussions, we will consider the parameterized case when $\mathcal{H}$ admits a parameterization over a compact parameter space Θ . The class is denoted by $\mathcal{H}_\Theta = \{h_\theta | \theta \in \Theta\}$. When the parameterization is linear, that is,

$$h_\theta(a, b) = \theta \cdot \phi(a, b) \quad (2.2)$$

for some feature function $\phi : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{B}^d$, we will denote the class by $\mathcal{H}_{\Theta, \phi}$. We denote the true parameter by θ^* . For instance, when $\mathcal{A}$ and $\mathcal{B}$ are the sets of standard basis vectors in $\mathbb{R}^{|\mathcal{A}|}$ and $\mathbb{R}^{|\mathcal{B}|}$ with $\phi(a, b) = ab^\top$ and θ is bounded in $\mathbb{R}^{|\mathcal{A}| \times |\mathcal{B}|}$, we recover the tabular case model in Kao, Wei, and Subramanian [116] with finite action sets. In general, however, we will focus on cases with infinite action sets.

2.3 Linear Stackelberg Games: Curse of Expertise

In this section, we study the sample complexity of learning in linear Stackelberg game, where the family of reward is restricted to $\mathcal{H}_{\Theta, \phi}$ for some given Θ and ϕ .

2.3.1 An Exponential Lower Bound

It is well known that the regret for traditional linear bandits grows as $\Theta(d\sqrt{T})$ [1]. In the case of a linear Stackelberg game, we present a worst-case lower bound on the regret that is exponential in dimensionality for the linear family. This suggests that the leader cannot learn the task well unless in possession of an exponential number of samples even when we restrict to linear Stackelberg games.

Assume the leader makes perfect observations of the follower's responses (i.e., $\sigma_b = 0$). We have the following lower bound.

Theorem 2.3.1. *For any $d \geq 4$, there exists some ϕ such that, for any algorithm that the leader runs, one can find some instance with $h_\theta \in \mathcal{H}_{\Theta, \phi}$ such that*

$$\mathcal{R}(T) \gtrsim T^{(d-4)/(d-2)}. \quad (2.3)$$

In other words, the sample complexity for achieving ϵ (average) regret is at least $\Omega((1/\epsilon)^{\frac{d-2}{2}})$.

The proof is detailed in Appendix A.1.1. The worst-case instance presented below can be reduced to the ReLU bandit problem shown below, which is known to suffer from the exponential sample complexity [73].

Example 2.3.2. *Let $\mathcal{A} = \mathbb{B}^{d-1}$, $\mathcal{B} = [0, 1]$ and $\Theta = \{\theta \mid \theta_{-d} \in \mathbb{S}^{d-2}, \theta_d = 1 - \Delta\}$ for some $\Delta \in (0, 1)$. Let the feature function be $\phi(a, b) = ((1 - b)a, b)$.*

One can verify that in this case, one has

$$\bar{h}_\theta(a) = \max\{1 - \Delta, \theta_{-d} \cdot a\}. \quad (2.4)$$

Thus when a is chosen far from θ_{-d} , the reward will remain constant.

Theorem 2.3.1 is no mystery mathematically: the best response may destroy linearity for the leader’s observations, imposing a toll. Conceptually, however, the message from the theorem is striking: it highlights a “curse of expertise”; i.e., the potential difficulty to learn with an expert on a decentralized bandit learning task with a large action space. From the classic single-agent bandit learning perspective, the task the two agents aim to solve is straightforward: a linear bandit on an action space $\phi(\mathcal{A}, \mathcal{B})$. In other words, if the expert follower lets the novice leader control the choice of b , the average regret would steadily decrease at a rate of $\tilde{O}(d\sqrt{T})$. On the other hand, with a myopic focus, the follower’s expertise in best responding ironically results in a significantly higher regret, as it deprives the learner of the ability to explore.

In the context of autonomous driving, for example, this can manifest in scenarios where the autonomous vehicle takes a poor action (e.g., an aggressive lane change) yet other vehicles or pedestrian immediately respond by slowing down or steering away to avoid a possible collision, thereby hiding the potential negative consequences of the action. The lack of coordination and the constant best response from the follower, both common in practice, makes it hard for the leader to efficiently learn the reward landscape or improve their current policy.

2.3.2 An Exponential Upper Bound

For any class $\mathcal{H}$ of reward functions on a pair of actions (a, b) , an upper bound on the sample complexity (and regret) can be obtained using a covering argument.

Theorem 2.3.3. *Let $N(\epsilon) = N(\mathcal{H}, \epsilon, \|\cdot\|_\infty)$ denote the ℓ^∞ covering number of $\mathcal{H}$ with radius $\epsilon > 0$. Then we can achieve*

$$\mathcal{R}(T) \lesssim \inf_{\epsilon > 0} \epsilon T + \sqrt{N(\epsilon)T}. \quad (2.5)$$

To achieve this, simply compute an ϵ -covering of $\mathcal{H}$ and let the leader play no-regret algorithms on the ϵ -covering set. Note that although the covering is constructed for pair of actions $(a, b) \in \mathcal{A}_\epsilon \times \mathcal{B}_\epsilon$, it suffices for the leader to run no-regret algorithms on actions $\mathcal{A}_\epsilon$. The detailed algorithm and proof are given in Appendix A.1.2.

This upper bound is achieved when the leader does not even utilize the observations of the follower's responses. Indeed, in the worst case (e.g., in Example 2.3.2), the responses will not provide information.

As a corollary, in the linear regime with $\mathcal{H}_{\Theta, \phi}$, the covering number is $N(\epsilon) = N(\Theta, \epsilon, \|\cdot\|) \leq \exp\left(O\left(d \log \frac{1}{\epsilon}\right)\right)$ [206]. Choosing $\epsilon \asymp T^{-1/(d+2)}$, Theorem 2.3.3 reduces to the following upper bound in the linearly parameterized case.

Corollary 2.3.4. *In the linear case, we can achieve $\mathcal{R}(T) \lesssim T^{(d+1)/(d+2)}$.*

In other words, the sample complexity for achieving average regret equal to ϵ is upper bounded by $O\left((1/\epsilon)^{d+2}\right)$. This upper bound is agnostic to any structural property of the feature function ϕ , such as smoothness or even continuity.

2.4 UCB With Side Observations

Although the worst-case sample complexity for linear Stackelberg games is exponential, it is possible to obtain a fine-grained analysis and improved rate for the family $\mathcal{H}_{\Theta, \phi}$ when ϕ is better structured. A natural choice of algorithm for the leader is some variant of UCB that incorporates observations of the follower's actions. In this section, we will describe a general recipe for a family of UCB algorithms to incorporate the side information as well as the challenge in their design.

2.4.1 Algorithm Description

We consider the following variant of UCB that uses the follower's responses as side information to improve the confidence set.

Algorithm 1: UCB with side information from expert

Input: Regression oracles $\text{Reg}^{(b)}$, $\text{Reg}^{(r)}$ on reward and response; confidence radii $\{\alpha_t\}_{t=1}^T$, $\{\beta_t\}_{t=1}^T$; action set $\mathcal{A}$; horizon T

Output: Sequence of actions $a_1, \dots, a_T$

```

1 for  $t \leftarrow 1$  to  $T$  do
2    $h_t^{(b)} \leftarrow \text{Reg}^{(b)}(\hat{b}_1, \dots, \hat{b}_{t-1})$ 
3    $h_t^{(r)} \leftarrow \text{Reg}^{(r)}(r_1, \dots, r_{t-1})$ 
4    $\mathcal{H}_t^{(b)} := \left\{ h : \sum_{i=1}^{t-1} \|b_h^*(a_i) - b_{h_t^{(b)}}^*(a_i)\|^2 \leq \alpha_t^2 \right\}$ 
5    $\mathcal{H}_t^{(r)} := \left\{ h : \sum_{i=1}^{t-1} (\hat{h}^{(r)}(a_i) - \hat{h}_t^{(r)}(a_i))^2 \leq \beta_t^2 \right\}$ 
6    $\mathcal{H}_t \leftarrow \mathcal{H}_t^{(b)} \cap \mathcal{H}_t^{(r)}$ 
7    $a_t \leftarrow \arg \max_{a \in \mathcal{A}} \sup_{h \in \mathcal{H}_t} \hat{h}^{(r)}(a)$  // optimistic response estimate
8   Observe reward  $r_t$  and response  $\hat{b}_t$ 
    
```

Remark 2.4.1. The regression oracles and the sequences $\{\alpha_t\}_{t \in [T]}$, $\{\beta_t\}_{t \in [T]}$ must be chosen appropriately so that the following condition holds: Given an error tolerance $\delta \in (0, 1)$, we require $h^* \in \bigcap_{t=1}^T \mathcal{H}_t$ with probability at least $1 - \delta$.

Remark 2.4.2. A common choice for $\text{Reg}^{(b)}$ and $\text{Reg}^{(r)}$ is the least-squares regression oracle that computes

$$h_t^{(b)} \in \underset{h \in \mathcal{H}}{\operatorname{argmin}} \sum_{i=1}^{t-1} \|b_h^*(a_i) - \hat{b}_i\|^2 \quad (2.6)$$

and

$$h_t^{(r)} \in \underset{h \in \mathcal{H}}{\operatorname{argmin}} \sum_{i=1}^{t-1} (\bar{h}(a_i) - r_i)^2. \quad (2.7)$$

When the least-squares computation becomes infeasible under complex response-reward structures (this is common for (2.6)), custom oracles need to be designed. A more intricate approach may be to jointly construct the estimate using both $\{\hat{b}_\tau\}_{\tau \in [t-1]}$ and $\{r_\tau\}_{\tau \in [t-1]}$. We leave it for future research to study systematic designs of the oracles and the confidence sets.

Remark 2.4.3. When the responses are unobserved or ignored (e.g., by choosing $\alpha_t = \infty$), Algorithm 1 reduces to the classic Eluder UCB using the least-squares (reward) oracle with $\mathcal{H}_t = \mathcal{H}_t^{(r)}$ [185].

The choices of $\{\alpha_t\}_{t \in \mathbb{N}}$ and $\{\beta_t\}_{t \in \mathbb{N}}$ can pose another challenge. A naive attempt to get a generic upper bound on α_t is to use a covering argument as in Russo and Van Roy [185] using the following measurement between two functions $h, h' \in \mathcal{H}$: $d^{(b)}(h, h') = \sup_a \|b_h^*(a) -$

$b_h^*(a)\|$. But note that this does not necessarily define a norm, and further the covering number of $\mathcal{H}$ in this sense can be infinite when the best response is discontinuous in the leader's action a . Thus, such an approach is often not useful and one may have to determine α_t on a per instance basis.

2.4.2 Examples

While Theorem 2.3.1 shows that the involvement of the omniscient follower can lead to “curse of expertise,” a stark deterioration in the sample complexity, there are many scenarios where the leader's observation of the follower's responses can expedite learning significantly. In this section, we will explore a few such examples.

2.4.2.1 An Imitation-Based Example

Let us consider a setting where the leader achieves efficient learning through imitation. Heuristically, imitation arises when the optimal action for the leader is equal to the best response for the omniscient follower or a function of it. This may capture, for instance, real-world robotics applications where the actions of the robot and the human expert are exchangeable and the true goal can be easily inferred from the expert's action. A simple scenario is when the robot and the human expert are supposed to carry out the same task perfectly, in which case the robot should simply treat the expert as a role model and imitate. The following is a concrete example.

Example 2.4.4. Let $\mathcal{A} = \mathcal{B} = \Theta = \mathbb{S}^{d-1}$ (or $\mathbb{B}^d$ equivalently)². Consider the linearly parameterized function class $\mathcal{H}_{\Theta, \phi}$ with feature function

$$\phi(a, b) = a + b. \quad (2.8)$$

Here, the optimal response $b_\theta^* \equiv \theta$ is independent of a , and $\bar{h}_\theta(a) = \theta \cdot a + 1$.

Construction of confidence sets. The (noisy) observations of the follower's best responses simplify the problem into an imitation learning task. A simple oracle for the best-response observations is to take the $\mathcal{A}$ -projected empirical average of responses, i.e., $\theta_t^{(b)} = \Pi_{\mathcal{A}}(\frac{1}{t-1} \sum_{i=1}^{t-1} \hat{b}_i)$.³ The response-based confidence set reduces to

$$\Theta_t^{(b)} = \left\{ \theta \in \Theta \mid \|\theta - \theta_t^{(b)}\| \leq \frac{\alpha_t}{\sqrt{t-1}} \right\}.$$

²While it is customary to consider $\Theta = \mathbb{B}^d$, we will observe below that the imitation-based algorithm does not crucially rely on $\|\theta^*\|$ and only incurs smaller regret if $\|\theta^*\| < 1$. This is because the algorithm asymptotically relies solely on the response observations, which are invariant under scaling of θ^* . It is also without loss of generality to restrict all actions to the sphere.

³Define the projection of $y \in \mathbb{R}^d$ onto a closed set $\mathcal{X} \subseteq \mathbb{R}^d$ as $\Pi_{\mathcal{X}}(y) := \operatorname{argmin}_{x \in \mathcal{X}} \|y - x\|$, breaking ties arbitrarily when the minimizer is not unique.

Standard sub-Gaussian concentration results suggest that the (Euclidean) radius of this confidence set shrinks at a rate of $t^{-1/2}$.

Lemma 2.4.5. *To ensure $\theta^* \in \bigcap_{t \in [T]} \Theta_t$ with probability at least $1 - \delta$, it suffices to choose $\alpha_t = \Theta(\sigma_b \sqrt{d + \log \frac{T}{\delta}})$.*

UCB chooses actions on $\mathbb{S}^{d-1}$ increasingly close to the empirical estimate $\theta_t^{(b)}$.⁴ The regret bound follows from these choices of confidence sets.

Proposition 2.4.6. *In Example 2.4.4, UCB achieves a regret bound*

$$\mathcal{R}_{UCB}(T) \lesssim \sigma_b^2 \log T \cdot (d + \log T). \quad (2.9)$$

In other words, the average regret decays at a rate of $\tilde{O}(\sigma_b^2 d/T)$. This has also been analyzed in the setting of imitation learning [179], and the results are consistent.

Remark 2.4.7. *When the follower's responses are unobserved (still assumed to be best responses), this is simply a linear bandit, and the minimax regret is $\Omega(\sigma_b d \sqrt{T}) \gg O(\sigma_b^2 d \log^2 T)$. This indicates the value of the b_t observations. When the follower's response is noiseless, one can see that a single sample suffices to find the optimal response since one always observes $b_\theta^* = \theta$.*

Remark 2.4.8. *Note the gap in the $\Theta(\log T)$ regret when the response observations are used and the $\Theta(\sqrt{T})$ regret when they are ignored or unavailable, showing the value of those response observations. In fact, it is easy to modify this example slightly (e.g., taking $\phi(a, b) = \max\{|\theta^\top a|, \Delta\}b$ for some $\Delta \in (0, 1)$) to create an even larger gap: When the leader uses the response observations, the regret is $\tilde{O}(d \log T)$ with sample complexity $\tilde{O}(d \log \frac{1}{\epsilon})$; When the response observations are unavailable, the sample complexity increases to $\Omega(\epsilon^{-d})$.*

2.4.2.2 Expert-Guided Exploration

In many scenarios, the omniscient follower's actions may not directly reveal the exact state of the world but still provide crucial information. The next example illustrates a simple setting where the follower's response can significantly reduce the sample complexity.

Example 2.4.9. *Let $\mathcal{A} = \mathcal{B} = \mathbb{S}^{d-1}$ and*

$$\Theta = \{(\theta_a, \theta_b) \in \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \mid \theta_a \cdot \theta_b \geq \zeta\}$$

for some $\zeta \in (0, 1)$. Consider the parameterized family of functions $\mathcal{H}_\Theta = \{h_\theta \mid \theta \in \Theta\}$ where

$$h_\theta(a, b) = \text{ReLU}(\theta_a \cdot a - \Delta) + \theta_b \cdot b,$$

for some $\Delta \in (0, 1)$. For simplicity, we will assume that the response observations are noiseless (i.e., $\sigma_b = 0$), although the noisy case can be analyzed analogously.

⁴Even simpler, the leader can play the $\mathcal{A}$ -projected empirical average of responses. Under our choice of constant α , the analysis will be the same, with the result differ by at most a constant factor.

Confidences sets. The best response is $b_\theta^* \equiv \theta_b$, again independent of the leader's action. Upon observing $b_1 = \theta_b$, the leader should construct confidence sets $\Theta_t^{(b)} = \{\theta_a \in \mathbb{S}^{d-1} | \theta_a \cdot b_1 \geq \zeta\} \times \{b_1\}$, while $\Theta_t^{(r)}$ is chosen as in linear UCB. As a result, all subsequent actions the leader takes must fall into

$$\mathcal{A}_1 := \{a \in \mathcal{A} | a \cdot b_1 \geq \zeta\}. \quad (2.10)$$

This refinement of the action set will reduce the sample complexity, and depending on the size of ζ relative to Δ , the reduction can be significant.

Strong reduction. When $1 - \zeta \leq (1 - \Delta)/4$, the leader learns that $\theta_a \cdot b_1 \geq \zeta$. In particular, any action $a \in \mathcal{A}_1$ must satisfy

$$\begin{aligned} \theta_a \cdot a &= \frac{2 - \|\theta_a - a\|^2}{2} \geq \frac{2 - (\|\theta_a - b_1\| + \|a - b_1\|)^2}{2} \\ &\geq \frac{2 - (2\sqrt{2 - 2\zeta})^2}{2} = 1 - 4(1 - \zeta) \geq \Delta, \end{aligned} \quad (2.11)$$

and thus $\bar{h}(a) = \theta_a \cdot a - \Delta + 1$ behaves as a linear function within $\mathcal{A}_1$. By playing UCB within $\mathcal{A}_1$, the leader reduces the problem to a linear bandit instance and thus achieves the following regret bound.

Proposition 2.4.10. *Assume $1 - \zeta \leq (1 - \Delta)/4$ in Example 2.4.9. UCB achieves*

$$\mathcal{R}_{UCB}(T) \leq \tilde{O}(d\sqrt{T}). \quad (2.12)$$

This leads to a sample complexity of $\tilde{O}(d^2/\epsilon^2)$, in contrast to the exponential sample complexity $\exp(O(d \log \frac{1}{\epsilon}))$ if the responses were unobserved. Information from the follower's response guides the leader's exploration to the well conditioned part of the action space. Given the $\Omega(d\sqrt{T})$ sample complexity of linear bandits, the upper bound (2.12) is tight (up to logarithmic terms).

Weak reduction. When ζ is small relative to Δ , the problem does not immediately reduce to a linear bandit, but we have the following improved upper bound.

Proposition 2.4.11. *There exists an algorithm Alg that achieves*

$$\mathcal{R}_{\text{Alg}}(T) \leq O((C_\zeta^d T^{d+1})^{\frac{1}{d+2}}), \quad (2.13)$$

where $C_\zeta := \sqrt{1 - \zeta^2} \in (0, 1)$.

This bound improves as ζ decreases. The sample complexity is therefore $\tilde{O}(C_\zeta^d \epsilon^{-d-2})$, a C_ζ^d reduction compared with the original complexity without observing the responses in Corollary 2.3.4.

Since the reduced problem is still a ReLU bandit, UCB will not be suitable. Instead, (2.13) can be achieved through discretization of $\mathcal{A}_1$ as the upper bound in Theorem 2.3.3.

2.5 Beyond UCB

Although the UCB algorithm gives a near-optimal rate in most of the above examples. We also provide two cases where UCB fails to achieve the optimal rate. This necessitates a tailored algorithm design in specific settings.

2.5.1 Nonlinear (Polynomial) Family

UCB is known to fail to achieve the optimal rate in the case of the polynomial bandit family [104], where the reward is a polynomial activation on top of a linear family. We construct an example which utilizes the structure of the polynomial bandit, formally defined below.

Example 2.5.1 (Polynomial bandit). *Consider the convex function $f(x) = x^{2k}$ for some $k \in \mathbb{Z}_+$. Let*

$$\mathcal{A} = \mathbb{B}^{d-1}, \mathcal{B} = [-1, 1], \Theta = \mathbb{B}^{d-1} \times \{1\}, \quad (2.14)$$

and

$$\phi(a, b) = (2kba, -f^*(2kb)), \quad (2.15)$$

where f^* is the convex conjugate of f . Consider the nonlinearly parameterized family

$$\mathcal{H}_\Theta := \{h_\theta(a, b) = f(\theta \cdot \phi(a, b)) \mid \theta \in \Theta\}. \quad (2.16)$$

By properties of the convex conjugate,

$$\bar{h}_\theta(a) = f(\theta_{-d} \cdot a) = (\theta_{-d} \cdot a)^{2k} \quad (2.17)$$

with the best response

$$\begin{aligned} b_\theta^*(a) &= \operatorname{argmax}_{-1 \leq b \leq 1} 2kb\theta_{-d} \cdot a - f^*(2kb) \\ &= \frac{f'(\theta_{-d} \cdot a)}{2k} = (\theta_{-d} \cdot a)^{2k-1} \in [-1, 1]. \end{aligned}$$

This observation allows us to apply results on polynomial bandits [104].

Response-regret structure. Observe the following properties of the best response function in Example 2.5.1.

1. The expected reward is a function of the best response, independent of the true parameter. Namely,

$$\bar{h}_\theta(a) = b_\theta^*(a)^{\frac{2k}{2k-1}}. \quad (2.18)$$

This mapping is Lipschitz:

$$|\bar{h}_\theta(a) - \bar{h}_\theta(a')| \leq \frac{2k}{2k-1} |b_\theta^*(a) - b_\theta^*(a')|, \quad (2.19)$$

and further

$$\operatorname{argmax}_{a \in \mathcal{A}} b_\theta^*(a) = \theta \in \operatorname{argmax}_{a \in \mathcal{A}} \bar{h}_\theta(a), \quad (2.20)$$

with both maxima being 1.

2. The response observation, as a degree $2k - 1$ polynomial, is more informative than the reward observation, a degree $2k$ polynomial, when the noise levels are the same and $\theta_{-d} \cdot a$ is small.

Based on these two observations, the leader may view the response b_t as a *proxy reward* and aim to minimize the *proxy regret*

$$\hat{\mathcal{R}}(T) := \sum_{t=1}^T 1 - b_\theta^*(a_t). \quad (2.21)$$

This is consistent with minimizing the true regret $\mathcal{R}(T)$, which differs from the proxy regret $\hat{\mathcal{R}}(T)$ by at most a constant factor by (2.19).

Regret bound. Using the response observations exclusively to minimize the proxy regret $\hat{\mathcal{R}}(T) = \sum_{t=1}^T 1 - b_\theta^*(a_t)$, the leader reduces her task to a polynomial bandit problem with a degree $2k - 1$ polynomial activation function. By (2.19), we may focus on bounding the proxy regret. Corollary 3.16 from Huang et al. [104] suggests that

$$\hat{\mathcal{R}}(T) \leq \tilde{O}(\sqrt{d^{2k-1}T}), \quad (2.22)$$

or equivalently the sample complexity is $\tilde{O}(d^{2k-1}/\epsilon^2)$ for achieving ϵ average proxy regret. The following bound on the true regret follows from (2.19) and (2.22).

Proposition 2.5.2. *In example 2.5.1, there exists an algorithm Alg, using the response observations exclusively, that achieves*

$$\mathcal{R}_{\text{Alg}}(T) \leq O(\sqrt{d^{2k-1}T}). \quad (2.23)$$

Proposition 2.5.2 suggests an $\tilde{O}(d^{2k-1}/\epsilon^2)$ sample complexity. For instance, the leader can achieve this regret with the zeroth-order algorithm proposed in Huang et al. [104, Algorithm 6].

Remark 2.5.3 (Lower bound). *Since the reward observations have a higher signal-to-noise-ratio, we should expect that the sample complexity of Example 2.5.1 to be the same order as the sample complexity of achieving ϵ average regret in a degree $2k - 1$ polynomial bandit. Huang et al. [104] shows that this is lower bounded by $\Omega(d^{2k-1}/\epsilon^2)$. Thus, (2.23) is essentially optimal.*

Remark 2.5.4 (Benefit of observing responses). *If the leader does not observe the responses, the problem is equivalent to a degree $2k$ polynomial bandit. The optimal regret without observing the experts actions will lead to an $\tilde{O}(d^{2k}/\epsilon^2)$ sample complexity. Thus, the response observations contribute to shaving of a factor of d , which can be significant when the dimensionality is high.*

Remark 2.5.5 (Suboptimality of UCB). *Using the traditional Eluder UCB algorithm leads to a suboptimal sample complexity of $\tilde{O}(d^{2k}/\epsilon^2)$ when the leader solely uses the response observations. Still, this is a factor d improvement compared to what she can achieve with UCB without the response observations.*

2.5.2 Failure of the Optimism Principle

The next example is adapted from the ReLU bandit in Example 2.3.2, and shows that optimism-based method can have dramatic suboptimality in certain problems.

Example 2.5.6. *Let $\mathcal{A} = \mathbb{B}^{d-1}$, $\mathcal{B} = \mathbb{B}^{d-1} \times [0, 1]$, and*

$$\Theta = \{(\theta_{-d}, \theta_d) \mid \theta_{-d} \in \mathbb{B}^d, \theta_d = 1 - \Delta\} \quad (2.24)$$

for some $\Delta \in (0, 1)$. Consider the linear family $\mathcal{H}_{\Theta, \phi}$ with

$$\phi(a, b) = \|a\|((1 - b_d)a, b_d - \|b_{-d}\|) + \frac{1 - \|a\|}{2}(b_{-d}, 0). \quad (2.25)$$

For any $\theta \in \Theta$ with $\theta_{-d} \in \mathbb{S}^{d-1}$, the optimal action for the leader is θ_{-d} , with the follower best responding $(0, 0)$ and achieving unit expected reward.

When $\|a\| = 1$, this function behaves exactly as in Example 2.3.2, where $b_\theta^*(a) = (0, 1)$ whenever $\theta_{-d} \cdot a < 1 - \Delta$; When $a = 0$, the best response is $b_\theta^*(0) = (\theta_{-d}, b_d)$. Thus, if the response observations are noiseless, the leader learns the true parameter and hence the optimal action in one round by playing $a_1 = 0$.

However, any optimism-based method such as UCB will not achieve such efficient learning, even when the response are noiselessly observed. It is straightforward to verify that, for any action a with $\|a\| < 1$, the optimistic reward satisfies

$$\sup_{\theta \in \Theta} \bar{h}_\theta(a) < 1. \quad (2.26)$$

Thus, as long as the confidence set contains some θ with $\theta_{-d} \in \mathbb{S}^{d-1}$, which holds under our initial condition, optimism causes the leader to only take actions $a \in \mathbb{S}^{d-1}$, reducing the problem to the worst-case Example 2.3.2.

2.6 Conclusion

We have studied a model of online learning in decentralized cooperative Stackelberg games. We showed that, even with an omniscient follower who always best responds (myopically), the worst-case sample complexity for a linear family can be as large as $\exp(\Theta(d \log \frac{1}{\epsilon}))$. This “curse of expertise” highlights the challenge caused by miscoordinated exploration. This also raises the question of how a non-myopic expert follower should respond to the leader’s actions (without knowing the leader’s exact algorithm) to expedite their learning and maximize their long-term reward.

We considered the UCB-type algorithm that incorporates response observations. A few examples of various hardness were considered, ranging from efficient learning through imitation and guided exploration to the worst-case linear family example with an exponential sample complexity.

Besides the examples considered in the chapter, there are numerous scenarios where the roles of the leader and the follower are more complex to reason about. This poses unique challenges for both the learning process of the leader and the subsequent analysis of regret, indicating a fertile ground for future research. Specifically, our current template of Algorithm 1 requires designing the confidence sets based on the specific response-reward structure of each problem. It remains open to find a general design (or prove the lack thereof) that systematically synthesizes the response and reward observations. A general framework of analysis that can provide a unified yet sharp upper bound on the examples is also valuable.

Chapter 3

Fitness and Welfare in Random Matching Markets

Understanding large multi-agent systems is challenging in general, especially with stochasticity and heterogeneity in agent profiles. But in many settings, structure emerges despite these difficulties, allowing us to understand key aspects of the system. This chapter illustrates such a case in the context of two-sided matching markets, where we use tools and insights from probability and information theory to establish a sharp characterization of welfare distribution and individual fitness under the classic stable matching framework.

3.1 Introduction

In this chapter, we study distributional outcomes in random two-sided matching markets. The market we will consider follows the marriage model of Gale and Shapley [88], where agents on each side of the market have preferences over potential partners on the other side. The market is assumed to be stable, meaning that there are no blocking pairs of agents who would rather match to each other over their assigned partners. In the markets that we consider, agents have latent preferences generated according to observed characteristics. We are interested in the empirical distributional outcomes of each side of the market and the relation between outcomes of each side of the market.

We study these questions in large randomly generated markets, with both vertical and horizontal differentiation. Every agent has an observed personal score for every other agent in the market, and her preferences follow a logit model based on these scores. We assume that no agent is a-priori overwhelmingly more desirable than any other agent. It is shown that the observed characteristics alone determine the empirical welfare distribution on each side of the market. Moreover, the joint surplus in the market is fixed, and the average welfare of one side of the market is a sufficient statistic to determine the empirical welfare distribution on both sides of the market.

The model we consider has an equal number of men and women. For every man m_i

and every woman $\mathbf{w}_j$, we are given non-negative scores a_{ij} and b_{ji} , which can be viewed as generated from observed characteristics. Each man and each woman have strict preference rankings generated independently and proportionally to these latent scores, as in the Logit model.¹ Equivalently, each man $\mathbf{m}_i$ has a latent value from matching with woman $\mathbf{w}_j$ that is distributed exponentially with rate a_{ij} (smaller values are considered better).² Women's latent values for men are generated similarly.³

We identify an intrinsic fitness for each agent that represents her relative competitiveness in the market, independent of the realized stable outcome. For every pair of agents on opposing sides of the market, we can obtain a mutual score of the pair's match. If we write these scores in a matrix, the intrinsic fitness values correspond to scaling coefficients that make the mutual matrix bi-stochastic.⁴ Intuitively, this bi-stochastic mutual matrix can be thought of as consisting of *a-priori* probabilities of each pair matching. This representation captures the interactions between the sides of the market.

We find that the welfare, or agents' ranks, when scaled by their intrinsic fitness, have an approximately exponential empirical distribution on each side of the market. Moreover, the average welfare of agents on one side of the market is sufficient to determine the average on the other side. Overall, each agent's welfare can be seen as determined by a global parameter, her intrinsic fitness, and an extrinsic factor with exponential distribution across the population. This characterization holds with high probability in every stable matching. In fact, this structure extends to matchings that are only approximately stable, which allow for a vanishing fraction of blocking pairs.

As the result applies to all stable matchings, the proof approach follows a union bound argument. The probability of each matching being stable can be derived from the integral formula by Knuth [123]. The heterogeneity of preferences, however, brings great difficulty, which we overcome with a truncation technique to accommodate heavy tails of agents' outcomes and a fixed-point argument on the eigenspace of the characterizing matrix of the market. As we show, the number of stable matchings is small, and for each stable matching the exponential empirical distribution arises intuitively as it has the largest entropy among all non-negative distributions with a given mean.

Closely related to this work are Menzel [153] and Peski [169], which analyze matching large markets and find that the joint surplus in the market is unique.⁵ Both papers characterize the matching rates between agents of different types while our focus is on evaluating the welfare and rank distributions for each side of the market.

Menzel [153] uses a fixed point argument in a supply-demand model generalizing and

¹Numerous empirical papers that study two-sided matching market assume agents' preferences follow a logit model (see e.g., [8, 99]).

²One can view the utility of an agent for her match to be the negative of the corresponding latent value.

³Special cases of this general model are markets with uniformly random preferences [124, 172, 125, 170, 21] or when agents have common public scores [151, 25].

⁴This representation is valid since preferences are invariant under such transformations.

⁵Peski [169] considers a large roommate model and analyzes proper stable matchings, and the results apply also to two-sided matching markets.

building on insights from Dagsvik [65]. A technical assumption on outside options allows Menzel [153] to establish uniqueness by showing that rejecting chains must be short.

Our proof approach is closer to Peski [169], which directly bounds the number of stable matchings that (do not) satisfy certain properties, through handling large deviations. In Peski [169], the desired property regards the distribution over match rates, while we require that the empirical welfare on each side follows an exponential distribution.

Peski [169] shows that the distribution of matched types is the one that maximizes the sum of the average welfare and the entropy of the distribution. Our result implies that stability also maximizes entropy but in the finer rank or welfare distribution on side of the market. A technical contribution of our work is the notion of (anti-)fitness, obtained via Sinkhorn matrix balancing, which constructs an explicit non-trivial mapping from types to values.

Finally, there are some differences in the assumptions over agents' preferences. Menzel [153] establishes that, at the limit, agents choose partners according to a logit model from opportunity sets, while we analyze large markets and assume that agents have logit based preferences. Peski [169] also has more general preferences that can allow for correlations in idiosyncratic shocks.

3.1.1 Related Works

The analysis of random two-sided matching markets goes back to Knuth, Motwani, and Pittel [124] and Pittel [172, 170], who consider markets with uniformly random complete preference lists. These papers establish the number of stable matchings as well as the average ranks on each side. A key finding is that the product of the average rank of agents on each side of the market is approximately the size of the market [170], implying that stable matchings essentially lie on a hyperbola. Our findings generalize these findings to markets to random logit markets. We also expand these findings to describe the distributional outcomes in the market.

Several papers consider markets with uniformly drawn preferences with an unequal number of agents on each side of market [21, 171, 51]. The key finding is that stability is essentially unique agents on the short side have a substantial advantage. Our work contributes to this literature by considering more general preferences and offer structural insights about approximately stable outcomes.

There have also been studies of random matching markets with non-uniform preferences [107, 127, 136, 19, 92, 25]. These papers study either the size of the core and manipulation in matching markets and the typical finding is that in large markets, the outcomes are essentially unique and frictions are limited.

3.1.2 Notations

In this chapter, we follow the convention where boldface letters denote vectors (lower case) and matrices (upper case), e.g., $\mathbf{x} = (x_i)_{i \in [n]}$ and $\mathbf{A} = (a_{ij})_{i \in [n], j \in [m]}$, and capital letters

denote random variables.

For two identically shaped matrices (or vectors) $\mathbf{M}$ and $\mathbf{N}$, $\mathbf{M} \circ \mathbf{N}$ denotes their Hadamard (entry-wise) product. For a vector $\mathbf{x} \in \mathbb{R}^n$ with non-zero entries, denote its coordinate-wise inverse by $\mathbf{x}^{-1}$.

We denote the probability density function (pdf) and cumulative distribution function (CDF) of $\text{Exp}(\lambda)$ by f_λ and F_λ , respectively. For distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ over spaces $\mathcal{X}$ and $\mathcal{Y}$, $\mathcal{D}_1 \times \mathcal{D}_2$ denotes their product distribution over $\mathcal{X} \times \mathcal{Y}$. $\hat{\mathcal{F}}(\mathbf{x})$ denotes the empirical distribution function for the components of a vector $\mathbf{x}$, treated as a function from $\mathbb{R}$ to $[0, 1]$. $\mathcal{F}(\mathcal{D})$ denotes the CDF of a distribution $\mathcal{D}$ on $\mathbb{R}$. For real-valued random variables X and Y , $X \preceq Y$ denotes stochastic domination of X by Y .

We say $f = o_\alpha(g)$ if $f/g \rightarrow 0$ as $\alpha \rightarrow 0$ (uniformly over all other parameters, such as n). For example, $\sqrt{\epsilon} = o_\epsilon(1)$.

3.2 Model and Preliminaries

3.2.1 A Matching Market

We begin with describing a general two-sided matching market.

Matching markets with random preferences. We consider a two-sided matching market consisting of two equally sized sets of agents: men $\mathcal{M}$ and women $\mathcal{W}$, where $|\mathcal{M}| = |\mathcal{W}| = n$. The men are labeled $\mathbf{m}_1, \dots, \mathbf{m}_n$ and the women $\mathbf{w}_1, \dots, \mathbf{w}_n$. Each man (resp. woman) has a complete strict preference ordering $\succ_{\mathbf{m}_i}$ (resp. $\succ_{\mathbf{w}_j}$) over the agents on the opposite side drawn from some distribution.

A *matching* is a bijection $\mu : \mathcal{M} \rightarrow \mathcal{W}$. For notational simplicity, we identify men and women with the index set $[n]$ and write $\mu : [n] \rightarrow [n]$, where $\mu(i) = j$ (equivalently, $\mu^{-1}(j) = i$) indicates that $\mathbf{m}_i$ is matched with $\mathbf{w}_j$.

Stability. A matching μ is *unstable* if there exists a pair $(\mathbf{m}_i, \mathbf{w}_j)$ such that $\mathbf{w}_j \succ_{\mathbf{m}_i} \mathbf{w}_{\mu(i)}$ and $\mathbf{m}_i \succ_{\mathbf{w}_j} \mathbf{m}_{\mu^{-1}(j)}$. That is, $\mathbf{m}_i$ and $\mathbf{w}_j$ both prefer each other over their assigned partners, forming a *blocking pair*. A matching is *stable* if no such blocking pair exists. Denote the set of stable matching by $\mathcal{S}$. It is well known that $\mathcal{S}$ is nonempty.

Ranks and their distribution. Given a stable matching μ , the *rank* of a man $\mathbf{m}_i$, denoted by $R_i(\mu)$, is the position of his assigned partner $\mu(i)$ in his preference list. The vector $\mathbf{R}(\mu) = (R_i(\mu))_{i \in [n]}$ captures the ranks of all men in the matching.

3.2.2 Logit-Based Preference Model

We consider a general model of matching markets with logit-based preferences. Preferences in the market are randomly generated as follows. For each man i , we are given a stochastic

vector $\hat{\mathbf{a}}_i = (\hat{a}_{ij})_{j \in [n]} \in \mathbb{R}_+^n$. Then, $\mathbf{m}_i$'s preference list is generated from a logit model based on $\hat{\mathbf{a}}_i$. In particular, let $\mathcal{D}_i$ be the distribution on $\mathcal{W}$ that places probability $\hat{a}_{ij}$ on $\mathbf{w}_j$; then $\mathbf{m}_i$ samples from $\mathcal{D}_i$ his favorite partner, and repeatedly sample from it without replacement his next favorite partner until completing his list. Similarly, each $\mathbf{w}_j$'s preference list is generated from a logit model based on a given stochastic vector $\hat{\mathbf{b}}_j = (\hat{b}_{ji})_{i \in [n]}$. Denote by $\hat{\mathbf{A}} = (\hat{a}_{ij})_{i,j \in [n]}$ and $\hat{\mathbf{B}} = (\hat{b}_{ji})_{j,i \in [n]}$ the row-stochastic matrices.

Example 3.2.1 (Markets with public fitness). *A matching market is said to have public fitness if all agents on the same side of the market share an identical preference distribution. Formally, this means that for all $i \in [n]$, the stochastic vector $\hat{\mathbf{a}}_i$ is identical across agents, i.e., $\hat{\mathbf{a}}_i = (\hat{a}_j)_{j \in [n]} \in \mathbb{R}_+^n$, and similarly, for all $j \in [n]$, $\hat{\mathbf{b}}_j = (\hat{b}_i)_{i \in [n]} \in \mathbb{R}_+^n$. This can be interpreted as having for each agent j a common and observed fitness a_j .*

Beyond the public fitness model, defining fitness for agents is less clear. For instance, it is unclear how the prior scores women assign to men could be meaningfully aggregated into scalars that captures mens' fitness independent of the realized outcome.

Matrix balancing and intrinsic fitness. To analyze the market, it will be helpful to describe it using an alternative scaling scheme. We look for scaling vectors $\phi, \psi \in \mathbb{R}_+^n$ for the rows of $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ such that $\mathbf{M} = n^{-1} \mathbf{A} \circ \mathbf{B}^\top$ is bistochastic⁶, where $\mathbf{A} = \text{diag}(\phi) \hat{\mathbf{A}}$ and $\mathbf{B} = \text{diag}(\psi) \hat{\mathbf{B}}$. As shown in Sinkhorn [193, Theorem 1], under mild conditions that we will discuss later, such a bistochastic matrix $\mathbf{M}$ always exists, and the scaling vectors ϕ and ψ are unique up to constant rescaling. That is, ϕ and ψ jointly solve

$$\frac{1}{n} \text{diag}(\phi) (\hat{\mathbf{A}} \circ \hat{\mathbf{B}}^\top) \text{diag}(\psi) \mathbf{1} = \mathbf{1} \quad \text{and} \quad \frac{1}{n} \text{diag}(\psi) (\hat{\mathbf{B}} \circ \hat{\mathbf{A}}^\top) \text{diag}(\phi) \mathbf{1} = \mathbf{1}, \quad (3.1)$$

where $\mathbf{1}$ is the vector consisting of all 1's. The matrix $\mathbf{M}$ will be referred to as the *mutual matrix*.

The bistochasticity constraint implies that a uniform increase of $\mathbf{m}_i$'s popularity among women results in a proportional decrease in ϕ_i . Thus, ϕ_i can be interpreted as the “average anti-popularity” of man $\mathbf{m}_i$ among women; Note that since $\hat{\mathbf{A}} = \text{diag}(\phi)^{-1} \mathbf{A}$ is row-stochastic, we conveniently have $\phi_i = \sum_{j=1}^n a_{ij}$, and similarly $\psi_j = \sum_{i=1}^n b_{ji}$ in the balanced form.⁷

We refer to the vector ϕ and ψ as the men's and women's *anti-fitness* vector, respectively (and note that a smaller intrinsic fitness value means the agent is more competitive). In the matching market with public scores (Example 3.2.1), the fitness vectors are $\phi = \hat{\mathbf{b}}^{-1}$ and $\psi = \hat{\mathbf{a}}^{-1}$, where the inverse is taken component-wise and the mutual matrix $\mathbf{M} = \mathbf{J} := (n^{-1})_{i,j \in [n]}$.

⁶A nonnegative matrix $\mathbf{M}$ is bistochastic if the sum of entries in each row and each column is one.

⁷One can construct examples where every woman prefers $\mathbf{m}_1$ over $\mathbf{m}_2$ by some margin (say, $\hat{b}_{j1} \geq 1.1\hat{b}_{j2}$ across all $j \in [n]$), yet $\mathbf{m}_2$ has a lower anti-fitness value. While seemingly counterintuitive, this arises because an individual's fitness depends not only on others' preferences but also on her own. In short, a man who prefers women who also favor him may appear more fit than one who prefers those less likely to reciprocate.

Latent values. Preferences based on the logit preference model can be generated equivalently as follows. Let $\mathbf{X}, \mathbf{Y} \in \mathbb{R}_+^{n \times n}$ be two random matrices with independent entries X_{ij} (resp. Y_{ji}) sampled from $\text{Exp}(a_{ij})$ (resp. $\text{Exp}(b_{ji})$). The preference profile is then derived from $\mathbf{X}$ and $\mathbf{Y}$ as follows:

$$\begin{aligned} \mathbf{w}_{j_1} \succeq_{\mathbf{m}_i} \mathbf{w}_{j_2} &\iff X_{ij_1} < X_{ij_2}, \\ \mathbf{m}_{i_1} \succeq_{\mathbf{w}_j} \mathbf{m}_{i_2} &\iff Y_{ji_1} < Y_{ji_2}. \end{aligned}$$

We refer to each X_{ij} (resp. Y_{ji}) for $i, j \in [n]$ as the *latent value* (or simply *value*) of $\mathbf{m}_i$ (resp. $\mathbf{w}_j$) if matched with $\mathbf{w}_j$ (resp. $\mathbf{m}_i$). Note that a lower rank corresponds to a lower latent value.

Market connectivity. As we study the asymptotic behavior of matchings in large markets, we require that the market is “sufficiently connected”. A counter example is given by

$$\hat{\mathbf{A}} = \hat{\mathbf{B}} = \frac{2}{n} \begin{pmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{pmatrix}, \quad (3.2)$$

which represents two disjoint markets, each with uniform independent preferences among its participants. In this case, even the anti-fitness vectors are ill-defined because their scaling can vary between the two sub-markets. Intuitively we require that no agent finds any other agent disproportionately more favorable than any other agent. We formalize this:

Assumption 3.2.2 (Strong connectivity). *There exist an absolute constant $C \in [1, \infty)$ such that*

$$\frac{1}{C} \leq \min_{i,j \in [n]} \min\{a_{ij}, b_{ji}, nm_{ij}\} \leq \max_{i,j \in [n]} \max\{a_{ij}, b_{ji}, nm_{ij}\} \leq C \quad \text{for all } n. \quad (3.3)$$

Remark 3.2.3. Assumption 3.2.2 imposes strong connectivity on the market by requiring the attractiveness among all pairs of men and women to vary by at most a constant factor as the market grows. We expect the results to hold under a weaker assumption, which can be described through the spectral gap of the matrix $\mathbf{M}$.⁸ The spectral gap is a common notion in the studies of the structure of networks and communities.⁹ Assumption 3.2.2 substantially simplifies the analysis and exposition.

⁸Recall that, as a bistochastic matrix, $\mathbf{M}$ has a largest eigenvalue of 1 and all other eigenvalues of magnitude at most 1. One may think of the market as contiguous in this weaker sense if the spectral gap of $\mathbf{M}$, given by $1 - |\lambda_{\max}(\mathbf{M} - \mathbf{J})|$, is bounded away from zero as the market grows.

⁹In the model, the spectral gap of $\mathbf{M}$ describes the extent to which the market interconnects globally or decomposes into multiple sub-markets. A larger spectral gap means that the market is more cohesive, with more uniform or homogeneous preferences. For instance, the uniform market with $\mathbf{M} = \mathbf{J}$ has a unit spectral gap, the maximum possible value. A smaller spectral gap means that the market is more clustered, with a clearer boundary between communities. For instance, any block-diagonal bistochastic matrix (with at least two blocks) has a zero spectral gap, and corresponds to a market that decomposes into two or more independent sub-markets; One cannot hope to have a uniform structure result in such markets.

3.3 Warm-Up: Uniformly Random Preferences

Before analyzing the general model, we consider a special case where agents have preferences generated independently and uniformly at random and use numerical simulations to generate intuition.

Figure 3.1 plots histograms of the ranks for both sides in four arbitrarily chosen stable matchings in such a market with $n = 10\,000$ agents on each side. Two patterns emerge from the simulations: (i) The empirical distributions of ranks exhibit an approximately exponential shape, and (ii) the product of the average ranks on the two sides is close to n .




Figure 3.1: Histograms (counts) of ranks in four arbitrarily chosen stable matchings for a market instance with uniformly random preferences with $n = 10\,000$ men and women. Red vertical dashed lines indicate average ranks on each side; orange dashed lines show the theoretical histogram according to an exponential distribution.

The second observation (ii) is a manifestation of the “law of hyperbola,” a result first shown by Pittel [170]. It highlights a fundamental trade-off: if one side enjoys lower average ranks, the other side must compensate with proportionately higher average ranks. Observation (i) appears to be novel. It suggests a natural hypothesis: in large markets, the empirical distribution of ranks in a typical stable matching follows an exponential law. As the market size grows, one may expect the empirical rank distribution to “converge” to a limiting exponential distribution in an appropriate sense. Formally, for any fixed error tolerance $\epsilon > 0$, we expect that with high probability, for any stable matching $\mu \in \mathcal{S}$, there exists a positive parameter λ such that

$$\|\hat{\mathcal{F}}(\mathbf{R}(\mu)) - F_{\lambda/n}\|_{\infty} \leq \epsilon,$$

where $\hat{\mathcal{F}}(\mathbf{v})$ denotes the empirical cumulative distribution function (CDF) of the components of a vector $\mathbf{v}$, and $F_\lambda(t) = 1 - e^{-\lambda t}$ is the CDF of an $\text{Exp}(\lambda)$ distribution. Furthermore, a natural choice for λ should be the average rank of women in μ . In the remainder of this chapter, we establish that a similar behavior arises in a general logit-based preference model, which allows intrinsic qualities.

3.4 Main Results: Large Market Behavior

In this section we characterize the empirical distributional outcomes under any stable matching.

For a matching μ , we denote by $\mathbf{X}(\mu)$ the vector of realized latent values of the men under the matching; that is, $\mathbf{X}(\mu) = (X_{i,\mu(i)})_{i \in [n]}$. Similarly, we denote by $\mathbf{Y}(\mu)$ the vector of realized values of women. Recall that $\mathbf{R}(\mu)$ denotes men's rank vector under μ . Denote by $\hat{\mathcal{F}}(\mathbf{v})$ the empirical distribution of the components of a vector $\mathbf{v}$ and by F_λ the CDF of $\text{Exp}(\lambda)$.

Theorem 3.4.1 (Empirical distribution of values). *Fix any $\epsilon > 0$. With high probability, for any stable matching $\mu \in \mathcal{S}$ there exists $\lambda_\mu > 0$ such that*

$$\|\hat{\mathcal{F}}(\mathbf{X}(\mu)) - F_{\lambda_\mu}\|_\infty \leq \epsilon. \quad (3.4)$$

The choice of λ_μ in (3.4) can be chosen as the truncated mean of the women's values.

Theorem 3.4.2 (Empirical distribution of fitness-adjusted ranks). *Fixed $\epsilon > 0$. With high probability, for any stable matching $\mu \in \mathcal{S}$ there exists $\lambda_\mu > 0$ such that*

$$\|\hat{\mathcal{F}}(\phi^{-1} \circ \mathbf{R}(\mu)) - F_{\lambda_\mu/n}\|_\infty \leq \epsilon, \quad (3.5)$$

where ϕ is the anti-fitness vector.

Again, the choice of λ_μ in (3.5) can be chosen as the truncated mean of the women's fitness-adjusted ranks. Corollary 3.4.2 has the advantage that the ranks, unlike latent values, are directly observable from the matching outcome.

3.4.1 Discussion

Theorem 3.4.1 implies that the average value of one side of the market is essentially sufficient to determine the average value of the other side. Roughly, for a given stable matching μ , the value of λ in (3.4) and (3.5) is approximately the sum of the women's values in μ .¹⁰

¹⁰Technically, the choice of λ can be taken as the sum of the values of women after excluding a small fraction δ of the women who are the least satisfied (those with the highest latent values) under the matching μ . This truncation, which is also done for technical reasons, avoids outliers, and in fact shows that the predictions still hold under even weaker notions of stability. Such trimming need not be necessary with a more careful analysis.

This suggests that the average value of men is approximately $1/\lambda \approx 1/\|\mathbf{Y}(\mu)\|_1$. Therefore, the product of the average values of the two sides of the market gives approximately $1/n$ simultaneously in all stable matchings with high probability. We believe that, with a refined analysis, one can show that $\sup_{\mu \in \mathcal{S}} |n^{-1} \|\mathbf{X}(\mu)\|_1 \|\mathbf{Y}(\mu)\|_1 - 1| \xrightarrow{p} 0$.

Moreover, the average value of men is also sufficient to predict the empirical value distribution on each side of the market. For example, if we find that 30% of the men have value h or higher, then we should expect 9% to have value $2h$ or higher.

Values are not always observable, however, Theorem 3.4.2 offers a similar characterization with respect to ranks. It implies that the product of the average scaled ranks of men and women should be asymptotically n , and the average rank on each side determines the empirical rank distributions.

Observe that the scaling in (3.5) is consistent with the intuition that $\phi_i = \sum_{j=1}^n a_{ij}$ is the average fitness of $\mathbf{m}_i$. Within a stable matching, a more popular man should, on average, achieve a better (smaller) rank than a less popular one. For instance, in a market with bounded public scores (Example 3.2.1), the number of proposals each man receives during the woman-proposing deferred acceptance algorithm is approximately inversely proportional to his anti-fitness. This implies that his rank is proportional to ϕ_i in the woman optimal stable matching.

Remark 3.4.3. *To emphasize the role of stability in the results, it is helpful to contrast these with the outcome generated by the random serial dictatorship (RSD). For simplicity, assume preferences are generated uniformly at random and women choose their favorite partner sequentially. By a coupon collector argument, the average rank of the women is of order $\log n$. On the other hand, the fraction of women who get their top choice is asymptotically $1 - e^{-1}$. Therefore, the women's rank distribution must have a heavy tail, thus departing from an exponential or geometric one. The men's rank distribution again fails to be exponential: the matching is completely independent of men's preferences and thus the rank distribution must be uniform, although the value distribution will be simply $\text{Exp}(1)$. Note also the failure of the law of hyperbola in this case: the average rank is of order $\log n$ for the women and linear for the men.*

Robustness. The results apply also to approximately stable matchings, defined as follows.

Definition 3.4.4 (Sub-market approximate stability). *We say that a matching μ between $\mathcal{M}$ and $\mathcal{W}$ is α -stable for some $0 < \alpha < 1$ if there exists a sub-market of size at least $(1 - \alpha)n$ on which μ is stable; that is, there exist subsets $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$ both with cardinality $|\mathcal{M}'| = |\mathcal{W}'| \geq (1 - \alpha)n$ such that $\mu(\mathcal{M}') = \mathcal{W}'$ and the partial matching induced by μ between $\mathcal{M}'$ and $\mathcal{W}'$ is stable (within this sub-market). We refer to the stable sub-matching between $\mathcal{M}'$ and $\mathcal{W}'$ as the stable part of μ . Denote the set of α -stable matchings by $\mathcal{S}_\alpha$.*

The following can be derived from the quantitative versions of Theorem 3.4.1 and Corollary 3.4.2, which we will present in Section 3.8.

Corollary 3.4.5. *Assume $\alpha < n^{-\eta}$ for some constant $\eta > 1/2$. Then, as $n \rightarrow \infty$,*

$$\max_{\mu \in \mathcal{S}_\alpha} \inf_{\lambda \in \mathbb{R}_+} \|\hat{\mathcal{F}}(\mathbf{X}(\mu)) - F_\lambda\|_\infty \xrightarrow{p} 0 \quad \text{and} \quad \max_{\mu \in \mathcal{S}_\alpha} \inf_{\lambda \in \mathbb{R}_+} \|\hat{\mathcal{F}}(\phi^{-1} \circ \mathbf{R}(\mu)) - F_\lambda\|_\infty \xrightarrow{p} 0. \quad (3.6)$$

The approximately exponential empirical distribution applies to any matching that is stable except for $o(\sqrt{n})$ agents. The key observation is that, if the market satisfies the contiguity assumption, then each sub-market is also contiguous, and we can apply a union bound over all sub-markets of size $(1 - \alpha)n$.

Similarly, stable matchings should have a similar behavior in markets that are slightly imbalanced (if there is a linear imbalance agents on the short side will get on average a constant rank, contradicting the exponential empirical distribution).

3.5 Proof Ideas

3.5.1 Intuition

We begin with offering a high-level intuition for the result. The proof does not follow exactly the intuition due to technical difficulties. We would like to establish, using a union bound argument (over the $n!$ matchings), that

$$\sum_{\mu} \mathbb{P}(\mu \in \mathcal{S} \text{ and the conditions of (3.4) fail}) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (3.7)$$

Establishing (3.7) directly appears to be difficult. A more approachable statement is to first bound the expected number of stable matchings:

$$\sum_{\mu} \mathbb{P}(\mu \in \mathcal{S}) = \exp(o(n)), \quad (3.8)$$

and then show that for an arbitrary matching μ , the probability (3.4) failing (i.e., the empirical distribution of $\mathbf{X}(\mu)$ is not exponential) conditional on stability of μ is exponentially small.

To better understand this, consider a matching μ selected randomly and its realized value vectors $\mathbf{X} = \mathbf{X}(\mu)$ and $\mathbf{Y} = \mathbf{Y}(\mu)$. Note that the probability of an unmatched pair $\mathbf{m}_i$ and $\mathbf{w}_j$ to block in μ is approximately $a_{ij}b_{ji}X_iY_j = nm_{ij}X_iY_j$, and therefore $n\mathbf{X}^\top \mathbf{M} \mathbf{Y}$ is approximately the expected number of blocking pairs. Without imposing stability, this would be of order n^2 . However, stability would require this to be of much smaller order, and therefore a large deviation event.

Consider the case of uniform preferences, where $n\mathbf{X}^\top \mathbf{M} \mathbf{Y} = \|\mathbf{X}\|_1 \|\mathbf{Y}\|_1$ since $\mathbf{M} = \mathbf{J} := (n^{-1})_{i,j=1}^n$. In this case, the sums $\|\mathbf{X}\|_1$ and $\|\mathbf{Y}\|_1$ are essentially sufficient for determining stability. Condition on these two quantities, the exponential empirical distribution is a consequence of maximum entropy positive distribution conditional on a fixed mean. In the

non-uniform preferences, the large deviation event combined with the strong connectivity assumption (Assumption 3.2.2) assures that these quantities are sufficient for similar reasoning about stability.

The proof requires careful analysis to handle the tail events and make the argument go through.

The intuition above reveals another reason why assumptions like Assumption 3.2.2 is necessary apart from our discussion in Section 3.2.2. Observe that even (3.8) can fail when we relax Assumption 3.2.2. Consider a market consisting of $n/2$ small markets with just 2 men and 2 women in each. Under the uniform preferences within each market, the expected number of stable matchings is $9/8$, thus

$$\sum_{\mu} \mathbb{P}(\mu \in \mathcal{S}) = (9/8)^{n/2} \neq \exp(o(n)).$$

3.5.2 Proof Sketch

Case with uniformly random preferences. Let us first look at simple case, in which agents' preferences are generated independently and uniformly at random, i.e., all canonical scores equal $1/n$. The proof in this case is more straightforward due to the symmetry among agents and the established high probability bound on the number of stable matchings [172]. We keep however the discussion informal in this section.

For a given matching μ , we study the conditional distribution of the value vectors $\mathbf{X}(\mu), \mathbf{Y}(\mu) \in \mathbb{R}^n$ conditional on stability of μ . Conditional also on women's value vector $\mathbf{Y}(\mu) = \mathbf{y}$, man $\mathbf{m}_i$'s value $X_i(\mu)$ must satisfy $X_i(\mu) < X_{ij}$ for each $j \neq \mu(i)$ with $Y_{ji} < y_j$ in order not to form a blocking pair. Since X_{ij} and Y_{ji} are i.i.d. samples from $\text{Exp}(1)$ for each $j \neq \mu(i)$, one should expect $X_i(\mu)$ to be effectively less than the minimum of about $\sum_{j \neq \mu(i)} 1 - e^{-y_j} \approx \|\mathbf{y}\|_1$ number of $\text{Exp}(1)$ random variables. Such a constraint acts independently on each X_i (conditional on $\mathbf{Y}(\mu) = \mathbf{y}$), and therefore in the posterior distribution one should expect $\mathbf{X}(\mu)$ to behave like i.i.d. samples from $\text{Exp}(\|\mathbf{y}\|_1)$.

Concretely, conditional on $\mathbf{U}(\mu) = \mathbf{u} = 1 - F(\mathbf{x})$ and $\mathbf{V}(\mu) = \mathbf{v} = 1 - F(\mathbf{y})$, where we recall that $F(z) = 1 - e^{-z}$ is the CDF of $\text{Exp}(1)$ and is applied component-wise on the value vectors. The likelihood of $(\mathbf{x}, \mathbf{y})$ when μ is stable is $\prod_{j \neq \mu(i)} (1 - u_i v_j)$ [172]. With some crude first order approximation (namely, $1 - z \approx e^{-z}$), this expression can be approximated by

$$\prod_{j \neq \mu(i)} (1 - u_i v_j) \approx \exp \left(- \sum_{i,j \in [n]} x_i y_j \right) = \exp(-\|\mathbf{x}\|_1 \|\mathbf{y}\|_1),$$

where we also put in the terms $x_i y_{\mu(i)}$ for $i \in [n]$ despite their absence in the original product. By Bayes' rule, conditional on $\mathbf{Y}(\mu) = \mathbf{y}$ and that μ is stable, the distribution of $\mathbf{X}(\mu)$ is approximately $p(\mathbf{x} | \mu \in \mathcal{S}, \mathbf{Y} = \mathbf{y}) \propto \exp(-\|\mathbf{x}\|_1 \|\mathbf{y}\|_1) \cdot \prod_{i=1}^n e^{-x_i} = \prod_{i=1}^n \exp(-(1 + \|\mathbf{y}\|_1)x_i)$. Note that this is the joint density of the n -fold product of $\text{Exp}(1 + \|\mathbf{y}\|_1)$. Our main theorems in the case with uniformly random preferences follow directly from the convergence of empirical measures.

The general case. The entire result can be viewed abstractly from the following lens: For any matching μ , we expect the value vectors $(\mathbf{X}(\mu), \mathbf{Y}(\mu))$ to behave “nicely” with very high probability conditional on stability of μ , so that even if we apply union bound on all stable matchings (which we will show to be “rare” separately) it is still unlikely to see any “bad” value vectors. To do so requires a careful analysis on the conditional distribution of value (given stability) $\mathcal{D}_\mu \in \Delta(\mathbb{R}_+^n \times \mathbb{R}_+^n)$, which depends on both the (unconditional) preference distribution (the “prior”) and the conditional probability that μ is stable given a pair of value vector $\mathbf{X}(\mu) = \mathbf{x}$ and $\mathbf{Y}(\mu) = \mathbf{y}$, which we will denote by $p_\mu(\mathbf{x}, \mathbf{y})$. We will define the ultimate “nice” event to be the ϵ -proximity of empirical distribution of value (or fitness-adjusted rank) to some exponential distribution, but unsurprisingly it is hard to analyze this event directly from $\mathcal{D}_\mu$, which is complicated itself, in one single step. Instead, we will follow a “bootstrapping” approach. Namely, we will construct a nested sequence of subsets $\mathbb{R}_+^n \times \mathbb{R}_+^n = \Omega_0 \supseteq \Omega_1 \supseteq \cdots \supseteq \Omega_K$ representing events on the joint value vectors of a stable matching, with Ω_K the desired event that the empirical distribution of value for men is ϵ -close to some exponential distribution. Step by step, we will show that a stable matching, conditional on its value vectors in Ω_i , must have value vectors in Ω_{i+1} with very high probability. Here is the roadmap to establishing these increasingly “nice” events:

- (a) As a first step, we approximate $p_\mu(\mathbf{x}, \mathbf{y})$, the likelihood of value vectors $\mathbf{x}$ and $\mathbf{y}$ in a stable matching μ , by the function $q(\mathbf{x}, \mathbf{y}) = \exp(-n\mathbf{x}^\top \mathbf{M}\mathbf{y})$. That is, the log likelihood of value vectors in a stable matching is approximately bilinear in the two vectors. To establish this, we identify a weak regularity condition on the value vectors of all stable matchings in terms of the first and second moments and extremal quantiles of the value vectors, under which the approximation holds (see Section 3.6.3). Such a condition is met by all stable matchings with high probability (see Appendix B.2 for details). The proof primarily consists of standard analysis of the deferred acceptance algorithm and careful use of first- and second-order approximation of $p_\mu(\mathbf{x}, \mathbf{y})$. Here we use the fact that the men-proposing and women-proposing deferred acceptance algorithms output the extremal outcomes with respect to the two sides’ values among all possible stable matchings.
- (b) In the expression for $q(\mathbf{x}, \mathbf{y})$, the value vectors relate through the matching matrix $\mathbf{M}$. However, we show next that, in stable matchings, we can further simplify things by approximately factoring $n\mathbf{X}(\mu)^\top \mathbf{M}\mathbf{Y}(\mu)$ into a product $\|\mathbf{X}(\mu)\|_1 \|\mathbf{Y}(\mu)\|_1$ of sums of values on the two sides; since $n\mathbf{M}$ is bistochastic, this essentially says that only the first “principal component” of $\mathbf{M}$ matters. More specifically, both $\mathbf{M}\mathbf{Y}(\mu)$ and $\mathbf{M}^\top \mathbf{X}(\mu)$ lie near the principal eigenspace of $\mathbf{M}$. To prove this, we use a fixed point argument to deduce that $\mathbf{M}\mathbf{Y}(\mu)$ depends almost deterministically on $\mathbf{M}^\top \mathbf{X}(\mu)$ and, symmetrically, $\mathbf{M}^\top \mathbf{X}(\mu)$ on $\mathbf{M}\mathbf{Y}(\mu)$; For such relationship to hold, both quantities need to lie near the principal eigenspace of $\mathbf{M}$, which is the span of $\mathbf{1}$ under Assumption 3.2.2 (see Section 3.7). Along the way, we also deduce an upper bound for the (unconditional) probability for μ to be stable (see Section 3.6.4), suggesting an sub-exponential upper bound on the typical number of stable matchings (Corollary 3.6.3).

- (c) Under the previous event, the men's values behaves approximately like i.i.d. exponential samples with rate $\|\mathbf{Y}(\mu)\|_1$ conditional on stability of μ and $\mathbf{Y}(\mu)$ – in fact, they are conditionally independent and nearly identically distributed. The result on the empirical distribution of men's values follows immediately from a concentration inequality of Dvoretzky–Kiefer–Wolfowitz (DKW) type, generalized for nearly identically distributed independent random variables (Lemma B.1.3).
- (d) Finally, we translate values into ranks. Using the classic first- and second-moment method, we show that for majority of the agents in the market, the rescaled rank (based on one's own scores) lies close to the value. This implies Theorem 3.4.2.

There is one caveat, however: In (a), a second order expansion of $p_\mu(\mathbf{x}, \mathbf{y})$ is required in order to justify the approximation with $q(\mathbf{x}, \mathbf{y})$. As a result, we need to control second order behavior of value, i.e., $\|\mathbf{X}(\mu)\|_2^2 = \sum_{i=1}^n X_i(\mu)^2$, in any stable matching μ . However, the second moment cannot be easily controlled due to the heavy tail of $\text{Exp}(1)$ —the moment generating function for X^2 does not even exist for $X \sim \text{Exp}(1)$. To resolve this issue, we use a truncation technique in the upper $\delta/2$ -quantile of the values on each side. By choosing δ sufficiently small, we can ensure that the truncation only affects the empirical distribution by an arbitrarily small amount in ℓ^∞ norm; this small sacrifice can be easily absorbed into the ϵ error term in (3.4). The real challenge, however, is that in (b) and (c) we will have to deal with not just all stable matchings, but all stable *partial* matchings: that is, matchings that are stable when restricted to some $(1 - \delta)$ -fraction of the market. See Section 3.6.2 for the technical definition of truncated and partial matchings.

3.6 Approximating the Probability of Stability

The goal is to understand the likely ranks (or values) of agents that can be realized when a matching μ is stable. Since the prior distribution of the values is known (as given by the matrices $\mathbf{A}$ and $\mathbf{B}$), it is enough to understand under what configurations of values μ is likely to be stable. Furthermore, knowing the probability of stability for a matching allows to compute the overall probability that μ is stable. This section provides a closed form expression for the probability of stability and its approximation. The heuristic analysis will suggest that the log probability of stability behaves as a bilinear function in the value vectors of the two sides through the mutual matrix. The validity of such an approximation will be formally established in Section 3.6.3. For the precise argument, we also need to extend the definition of stability to partial matchings, i.e., matchings between a large fraction of the market, which emerge naturally from stable matchings once attention is restricted to a well-behaving part of the market.

3.6.1 Probability of Stability and Its Bilinear Approximation

As mentioned above, a central quantity to our analysis is the probability of stability, which is defined for each matching μ as a function $p_\mu : \mathbb{R}_+^n \times \mathbb{R}_+^n \rightarrow [0, 1]$ to be the probability that μ is stable given values of men and women in μ . That is

$$p_\mu(\mathbf{x}, \mathbf{y}) = \mathbb{P}(\mu \in \mathcal{S} | \mathbf{X}(\mu) = \mathbf{x}, \mathbf{Y}(\mu) = \mathbf{y}).$$

Just like the integral formula used by Knuth [123] and Pittel [172, 170] to study matching markets with uniformly random preferences, the probability of a matching μ being stable can be similarly characterized by an integral

$$\begin{aligned} \mathbb{P}(\mu \in \mathcal{S}) &= \mathbb{E}_{\mathbf{X} \sim \prod_{i=1}^n \text{Exp}(a_{i,\mu(i)}), \mathbf{Y} \sim \prod_{i=1}^n \text{Exp}(b_{i,\mu^{-1}(i)})} [p_\mu(\mathbf{X}, \mathbf{Y})] \\ &= \int_{\mathbb{R}_+^n \times \mathbb{R}_+^n} p_\mu(\mathbf{x}, \mathbf{y}) \prod_{i=1}^n f_{a_{i,\mu(i)}}(x_i) f_{b_{i,\mu^{-1}(i)}}(y_i) d\mathbf{x} d\mathbf{y}. \end{aligned}$$

The function p_μ can be further expressed in closed form. Condition on the value vector $\mathbf{X}(\mu) = \mathbf{x}$ and $\mathbf{Y}(\mu) = \mathbf{y}$ and sample the rest of the values X_{ij} and Y_{ji} for all $j \neq \mu(i)$. Each pair of $i, j \in [n]$ with $j \neq \mu(i)$ may form a blocking pair when $X_{ij} < x_i$ and $Y_{ji} < y_j$, which event happens with probability $(1 - \exp(-a_{ij}x_i))(1 - \exp(-b_{ji}y_j))$. For μ to be stable, there must be no blocking pairs and thus

$$p_\mu(\mathbf{x}, \mathbf{y}) = \prod_{\substack{i,j \in [n] \\ j \neq \mu(i)}} (1 - (1 - e^{-a_{ij}x_i})(1 - e^{-b_{ji}y_j})). \quad (3.9)$$

While exact, the closed form of p_μ is hard to use; approximations (including upper and lower bounds) of it will often simplify our analysis. Under Assumption 3.2.2, i.e., $\max_{i,j_1,j_2} a_{ij_1}/a_{ij_2} \leq C^2$ and $\max_{j,i_1,i_2} b_{ji_1}/b_{ji_2} \leq C^2$, we observe a simple upper bound

$$p_\mu(\mathbf{x}, \mathbf{y}) \leq \prod_{\mu(i) \neq j} \left(1 - (1 - e^{-\hat{x}_i/C^2})(1 - e^{-\hat{y}_j/C^2}) \right), \quad (3.10)$$

where $\hat{x}_i = x_i a_{i,\mu(i)}$ and $\hat{y}_j = y_j b_{j,\mu^{-1}(j)}$ for $i, j \in [n]$ are the renormalized values (thus named because they have unit mean). This bound is fairly conservative and crude for our final purpose, but will prove useful for establishing preliminary results.

To further simplify the analysis, we recognize that, through first order approximation,

$$p_\mu(\mathbf{x}, \mathbf{y}) \approx \prod_{i \neq j} (1 - a_{ij} b_{ji} x_i y_j) \leq \exp \left(- \sum_{i \neq j} a_{ij} b_{ji} x_i y_j \right) \approx \exp(-n \mathbf{x}^\top \mathbf{M} \mathbf{y}). \quad (3.11)$$

As mentioned in our proof sketch, we write

$$q(\mathbf{x}, \mathbf{y}) := \exp(-n \mathbf{x}^\top \mathbf{M} \mathbf{y}) = \exp \left(-n \sum_{i,j=1}^n m_{ij} x_i y_j \right)$$

for this approximated expression for the probability of stability. In other words, the log probability of achieving stability with realized value vectors $\mathbf{x}$ and $\mathbf{y}$ is approximately $\mathbf{x}^\top \mathbf{M} \mathbf{y}$, a bilinear function in both vectors. The bilinearity will allow to decouple the two sides of the market, and analyze each side while conditioning on the other.

Our analysis will show next that $q(\mathbf{x}, \mathbf{y})$ is indeed “good” approximation of $p_\mu(\mathbf{x}, \mathbf{y})$. Before formalizing this statement, we do need to discuss some technicalities.

3.6.2 Partial Matchings and Truncation

A challenge in arguing that $q(\mathbf{x}, \mathbf{y})$ serves as a good approximation of $p_\mu(\mathbf{x}, \mathbf{y})$ is that the second-order behaviors of the men’s and women’s values, which are relaxed in (3.11), are hard to control due to the heavy tail of exponential distribution. However, outliers in the upper δ -quantile of the values can affect the empirical CDF by at most δ in ℓ^∞ -distance: this can be absorbed into our error tolerance. We address the concern of outliers using a truncation technique, where we consider matchings that are stable within a significant subset of the market. It will later turn out that this approach also allows us to study approximately stable matchings.

We now formalize a general partial matching and then describe a particular way to form stable partial matchings. Let $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$ be subsets of the men and women with cardinality $|\mathcal{M}'| = |\mathcal{W}'| = n'$. A *partial matching* $\mu' : \mathcal{M}' \rightarrow \mathcal{W}'$ is a bijection between $\mathcal{M}'$ and $\mathcal{W}'$. Denote the values of men among $\mathcal{M}'$ and women among $\mathcal{W}'$ in the partial matching μ' by $\mathbf{X}_{\mathcal{M}'}(\mu')$ and $\mathbf{Y}_{\mathcal{W}'}(\mu')$, respectively. Although it may be natural to view $\mathbf{X}_{\mathcal{M}'}(\mu')$ and $\mathbf{Y}_{\mathcal{W}'}(\mu')$ as n' -dimensional vector, we choose to view them as n -dimensional vectors where the components corresponding to men in $\mathcal{M} \setminus \mathcal{M}'$ and women in $\mathcal{W} \setminus \mathcal{W}'$ are zero (recall that since smaller values are better outcomes, zero is the best possible latent value). Therefore, conditional on $\mathbf{X}_{\mathcal{M}'}(\mu') = \mathbf{x}'$ and $\mathbf{Y}_{\mathcal{W}'}(\mu') = \mathbf{y}'$ for $\mathbf{x}', \mathbf{y}' \in \mathbb{R}^n$ supported on $\mathcal{M}'$ and $\mathcal{W}'$, respectively, the probability that μ' is stable (as a matching between $\mathcal{M}'$ and $\mathcal{W}'$) is simply $p_{\mu'}(\mathbf{x}', \mathbf{y}')$.

Given a full stable matching μ and any $\delta > 0$, we define the following routine to construct a stable partial matching of size $n - \lfloor \delta n \rfloor$: Let $\bar{\mathcal{M}}_{\mu, \delta/2} \subseteq \mathcal{M}$ be the subset of $\lfloor \delta n / 2 \rfloor$ men with the largest value (i.e., the least happy men) in μ , and similarly let $\bar{\mathcal{W}}_{\mu, \delta/2} \subseteq \mathcal{W}$ be the set of $\lfloor \delta n / 2 \rfloor$ least happy women. Construct $\mathcal{M}'_{\mu, \delta} \subseteq \mathcal{M} \setminus (\bar{\mathcal{M}}_{\mu, \delta/2} \cup \mu(\bar{\mathcal{W}}_{\mu, \delta/2}))$ of cardinality $n - \lfloor \delta n \rfloor$. This is always possible because $|\bar{\mathcal{M}}_{\mu, \delta/2} \cup \mu(\bar{\mathcal{W}}_{\mu, \delta/2})| \leq 2\lfloor \delta n / 2 \rfloor \leq \lfloor \delta n \rfloor$ and in fact here can be multiple ways to choose $\mathcal{M}'_{\mu, \delta}$. The specific way $\mathcal{M}'_{\mu, \delta}$ is chosen (when $|\mathcal{M} \setminus (\bar{\mathcal{M}}_{\mu, \delta/2} \cup \mu(\bar{\mathcal{W}}_{\mu, \delta/2}))| > n - \lfloor \delta n / 2 \rfloor$) is irrelevant to our discussion, but it may be helpful to assume that the choice is made based on some canonical ordering of the men so there is no extra randomness. Let $\mu_\delta : \mathcal{M}'_{\mu, \delta} \rightarrow \mu(\mathcal{M}'_{\mu, \delta})$ be the partial matching induced by μ on $\mathcal{M}'_{\mu, \delta}$ and their partners. Define the δ -truncate value vector for μ to be $\mathbf{X}_\delta(\mu) := \mathbf{X}_{\mathcal{M}'_{\mu, \delta}}(\mu_\delta)$ and $\mathbf{Y}_\delta(\mu) := \mathbf{Y}_{\mu(\mathcal{M}'_{\mu, \delta})}(\mu_\delta)$.

3.6.3 Validity of the Approximation and Reduction to Partial Matchings

We now show that for any stable matching μ , our $q(\mathbf{x}, \mathbf{y})$ serves as a good approximation to $p_\mu(\mathbf{x}, \mathbf{y})$ for its truncated value vectors $\mathbf{x} = \mathbf{X}_\delta(\mu)$ and $\mathbf{y} = \mathbf{Y}_\delta(\mu)$.

Proposition 3.6.1. *For any $\delta > 0$ and $c \in (0, 1/2)$, with probability at least $1 - \exp(-n^c)$, for all stable matchings μ , the truncated value vectors $\mathbf{X}_\delta(\mu)$ and $\mathbf{Y}_\delta(\mu)$ satisfies*

$$\frac{p_\mu(\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu))}{q(\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu))} \leq \exp \left(O \left(\frac{n}{(\log n)^{1/2}} \right) \right). \quad (\star)$$

We will establish the bound $(\star)$ in two steps. First, we identify a family of conditions on the value vectors, and show that together they guarantee $(\star)$. We then work with each individual event and show with a union bound argument that, with high probability, no stable matchings will violate it. The construction of such conditions is technical. See Corollary B.2.13 in Appendix B.2.

Some of these conditions are particularly useful in our proofs and provide analogous bounds to classic results in the setting with uniform preferences. The next lemma describes some of them.

Proposition 3.6.2. *Let $\delta > 0$ and $c \in (0, 1/2)$ be fixed. With probability at least $1 - \exp(-n^c)$, for all stable matchings μ , the truncated value vectors $\mathbf{X}_\delta(\mu)$ and $\mathbf{Y}_\delta(\mu)$ satisfy*

$$\Omega(\log n) \leq \|\mathbf{X}_\delta(\mu)\|_1, \|\mathbf{Y}_\delta(\mu)\|_1 \leq O(n(\log n)^{-7/8}) \quad (3.12)$$

and

$$\mathbf{X}_\delta(\mu)^\top \mathbf{M} \mathbf{Y}_\delta(\mu) \leq O(\log n)^{1/8}, \quad (3.13)$$

where $O(\cdot)$ and $\Omega(\cdot)$ hides appropriately chosen constants independent of n .

It is helpful to compare the conditions (3.12) and (3.13) to properties of stable matchings in the classic uniform preference case: the optimal and pessimal average ranks on each side of order $\Omega(\log n)$ and $O(n/\log n)$ [172], respectively, and the “law of hyperbola” that characterizes the product of the average ranks on the two side as asymptotically n in all stable matchings [170]. Our conditions are close parallel. However, due to the heterogeneity of preferences in our model, we have to pay a small price of an extra multiplicative factor of up to $(\log n)^c$ with an arbitrarily small constant $c > 0$ in some of these conditions.

3.6.4 Estimating the (Marginal) Probability of Stability

With the approximate expression $q(\mathbf{x}, \mathbf{y})$ for the probability function of stability, we can derive the following sub-exponential (in probability) upper bound on the number of stable matchings by integrating $\mathbf{x}$ and $\mathbf{y}$ over the likely region.

Proposition 3.6.3 (Number of stable matchings). *Fix any $\delta > 0$. Let N_δ denote the number of stable partial matchings of size $n - \lfloor \delta n \rfloor$ in a random instance of the matching market. Then, $N_\delta \leq \exp(o_\delta(n))$ whp. Further, the number N of stable matchings satisfies $N \leq \exp(o_\delta(n))$ whp.*

To derive the bound on the number of full stable matchings from that of stable *partial* matchings, we use the following observation: any such partial matching can be extended to a stable *full* matching in at most $\exp(\Theta(\delta n))$ ways. This is because any such extension requires the remaining $\lfloor \delta n \rfloor$ men and women to be at least stably matched among themselves, but this sub-market of size $\lfloor \delta n \rfloor$ admits at most $\exp(\Theta(\delta n))$ stable matchings [166]. It therefore suffices to bound the number of stable partial matchings. We do so by classifying these partial matchings according to whether their realized value vectors $(\mathbf{x}, \mathbf{y})$ fall into a likely region (namely $\mathcal{R}^*$, formally defined in Corollary B.2.13). We show that (1) stable partial matchings with likely value vectors total in expectation of order $\exp(o_\delta(n))$ and (2) those with unlikely value vectors, not surprisingly, typically do not exist. The proof is deferred to Appendix B.3.

3.7 Uniform Influence Profiles

While the approximation $q(\mathbf{x}, \mathbf{y})$ is simpler to analyze than $p_\mu(\mathbf{x}, \mathbf{y})$, the heterogeneity in preferences suggest that men's typical outcomes may depend on the women's values in a complicated way. Recall the intuition that $n\mathbf{x}^\top \mathbf{M}\mathbf{y} = -\log q(\mathbf{x}, \mathbf{y})$ gives an estimate of the expected number of block pairs in a matching if the realized value vectors are $\mathbf{x}$ and $\mathbf{y}$, and we expect this quantity to be small for stable matchings. We can think of $\mathbf{M}\mathbf{y}$ as the “influence” of the women on the men: a higher value in the i -th component of $\mathbf{M}\mathbf{y}$ (compared to other components) would require the i -th component of $\mathbf{x}$ to be further suppressed. It is a priori unclear that such influence profile should appear mostly uniform across men, i.e., $\mathbf{M}\mathbf{y}$ is typically near a scalar multiple of $\mathbf{1}$. This is formalized in the next Proposition.

Proposition 3.7.1 (Approximately uniform influence profiles). *For $\delta > 0$ sufficiently small, there exists a choice of $\zeta = \zeta(\delta) > 0$ such that $\zeta \rightarrow 0$ as $\delta \rightarrow 0$ and, with probability at least $1 - O(e^{-n^c})$, for any stable matching $\mu \in \mathcal{S}$, there exists $t_\mu > 0$ such that*

$$\sum_{i=1}^n \mathbb{1} \{ |(\mathbf{M}\mathbf{y})_i - t_\mu| \geq \zeta t_\mu \} \leq \zeta n. \quad (3.14)$$

Note that this result holds trivially in a market with uniform preferences, where $\mathbf{M} = \mathbf{J}$ has a 1-dimensional column space in the direction of $\mathbf{1}$. For general logit-based preferences, however, the claimed alignment is far less obvious. The intuition here is related to the bounded spectral gap of $\mathbf{M}$: namely, Assumption 3.2.2 ensures that $\mathbf{M}$ has a strict smoothing effect on the components when applied to a vector. Through the lens of large deviations, we show that, conditioned on $\mathbf{y}$, a typical value of $\mathbf{M}^\top \mathbf{x}$ exhibits strictly smaller

(relative) fluctuation across its components than $\mathbf{M}\mathbf{y}$, and vice versa. This leads to a pair of approximate fixed point equations, which jointly push both vectors $\mathbf{M}^\top \mathbf{x}$ and $\mathbf{M}\mathbf{y}$ towards scalar multiples of $\mathbf{1}$. This fixed-point perspective is conceptually similar to the techniques developed by Menzel [153] and Peski [169].

3.8 Empirical Distribution of Values and Ranks

We now establish our main results on the structure of values and ranks from the approximate uniformity of influence profiles in stable matchings. The next theorem quantifies the proximity of the empirical distribution of values in stable matchings to some exponential distribution, and essentially follows from the heuristics that the men's values $\mathbf{X}$, conditioned on stability and the women's values, behave like i.i.d. samples from $\text{Exp}(\|\mathbf{Y}\|_1)$.

Theorem 3.8.1 (Empirical distribution of value). *Fix any $\epsilon > 0$. Then*

$$\mathbb{P}\left(\max_{\mu \in \mathcal{S}} \inf_{\lambda \in \mathbb{R}_+} \|\hat{\mathcal{F}}(\mathbf{X}(\mu)) - F_\lambda\|_\infty > \epsilon\right) \lesssim e^{-n^c} \quad (3.15)$$

asymptotically as $n \rightarrow \infty$. In particular, the infimum over λ in (3.15) can be replaced with the choice of $\lambda = \|\mathbf{Y}_\delta(\mu)\|_1$ for δ sufficiently small.

The next theorem is a generalization of the “law of hyperbola” in Pittel [170]. It can be seen as a corollary of the approximately $\text{Exp}(\|\mathbf{Y}\|_1)$ empirical distribution of the vector $\mathbf{X}$, except that a δ -truncation is again needed to avoid sensitivity to outliers in the upper tail. The proof is essentially the same as Theorem 3.8.1 and shall be omitted; the only difference is to replace the DKW-type concentration inequality for the empirical distribution with a Bernstein-type concentration for the empirical average.

Proposition 3.8.2 (Law of hyperbola in truncated matching). *For any fixed $\epsilon, \delta > 0$ and $0 < c < 1/2$,*

$$\mathbb{P}\left(\max_{\mu \in \mathcal{S}} |n^{-1} \|\mathbf{X}_\delta(\mu)\|_1 \|\mathbf{Y}_\delta(\mu)\|_1 - 1| > \epsilon\right) \lesssim e^{-n^c}. \quad (3.16)$$

Our final result characterizes the empirical distribution of fitness-adjusted ranks in stable matchings, and follows naturally from the value distribution in Theorem 3.8.1. Heuristically, we would expect the rank $R_i(\mu)$ for man i to be proportional to his value $X_i(\mu)$ conditioned on stability of μ and the values $X_i(\mu) = x_i$ and $\mathbf{Y}(\mu) = \mathbf{y}$. Indeed, a woman w_j with $j \neq \mu(i)$ stands ahead of $w_{\mu(i)}$ in the preference of man m_i exactly when $X_{ij} < x_i$ and $Y_{ji} > y_{\mu(i)}$. As X_{ij} and Y_{ji} jointly follow the product distribution $\text{Exp}(a_{ij}) \times \text{Exp}(b_{ji})$ conditional on the event that $X_{ij} < x_i$ and $Y_{ji} < y_{\mu(i)}$ do not simultaneously happen, the conditional probability that $\mathbf{w}_j \succeq_{\mathbf{m}_i} \mathbf{w}_{\mu(i)}$ is $\frac{(1-e^{-a_{ij}x_i})e^{-b_{ji}y_j}}{1-(1-e^{-a_{ij}x_i})(1-e^{-b_{ji}y_j})} \approx a_{ij}x_i$. By summing over $j \neq \mu(i)$, we should expect that the rank $R_i(\mu)$ to be in expectation close to $x_i \sum_{j \neq i} a_{ij} \approx x_i \phi_i$; further, as the sum of independent Bernoulli random variables, $R_i(\mu)$ should concentrate at

its expectation, therefore leading to $R_i(\mu) \approx x_i \phi_i$ simultaneously for most $i \in [n]$. The proof in Appendix B.7 formalizes this intuition.

Theorem 3.8.3 (Empirical distribution of fitness-adjusted ranks). *For any fixed $\epsilon > 0$,*

$$\mathbb{P}\left(\max_{\mu \in \mathcal{S}} \inf_{\lambda \in \mathbb{R}_+} \|\hat{\mathcal{F}}(\phi^{-1} \circ \mathbf{R}(\mu)) - F_\lambda\|_\infty > \epsilon\right) \lesssim e^{-n^\epsilon} \quad (3.17)$$

asymptotically as $n \rightarrow \infty$, where ϕ is men's fitness vector. As in Theorem 3.8.1, the infimum over λ in (3.15) can be replaced with the choice of $\lambda = \|\mathbf{Y}_\delta(\mu)\|_1$ for δ sufficiently small.

3.9 Conclusion

We studied the welfare structure in stable outcomes in large two-sided matching markets with logit-based preferences. Under a contiguity condition that prevents agents from disproportionately favoring or unfavoring other agents, we characterize stable and almost stable matchings outcomes in terms of the empirical distributions of latent values and ranks of each side of the market.

Our results suggest that the welfare in a stable matching of a random market can be decomposed into three parts: a global parameter that determines the trade-off between the two sides, a personal intrinsic fitness computed from the systematic scores, and an exogenous factor behaving as a standard exponential random variable. In other words, the average fitness-adjusted rank of one side of the market is a sufficient statistic for the outcome distribution.

Chapter 4

Efficiency and Envy in School Choices

The concept of stable matching, as studied in the previous chapter, has seen wide applications across market design and allocation problems. One key limitation is that stable outcomes are not necessarily efficient. Drawing on random graph theory, this chapter shows that typical stable matchings can exhibit substantial inefficiency. The resulting structural insights help inform the design of improved mechanisms or policies.

4.1 Introduction

The deferred acceptance (DA) algorithm, first proposed by Gale and Shapley in their study of the stable matching problem, has been widely adopted for assignment problems, as it ensures stability and encourages truthful preference reporting. Many school choice systems, such as those in Boston [6] and New York City [3], use a student-proposing version of DA to assign students to schools; this chapter focuses on such school choices settings. However, DA is known to result in outcomes that are Pareto-inefficient for students; in some cases, it may assign most students to one of their least-preferred schools [120]. Its inefficiency has also been documented empirically. Abdulkadiroğlu, Pathak, and Roth [4] show that approximately 2% of students in the New York City high school match (using a version of DA) could have been assigned to a more preferred school without harming others.

To address this inefficiency while preserving weaker forms of stability and strategy-proofness, several mechanisms that Pareto dominate DA have been proposed. One such mechanism, Kesten’s Efficiency-Adjusted Deferred Acceptance (EADA), has shown exceptional promise in both theory and the lab. Yet empirical work suggests that EADA’s efficiency gains may be unevenly distributed (see Section 4.1.1). Consequently, two natural theoretical questions emerge: *which and how many students can improve their DA placement under EADA, and more generally, under efficient mechanisms that Pareto dominate DA?* We tackle these two questions in this chapter.

We begin by identifying *which* students are improvable. Our analysis introduces a novel analytical framework: DA’s *envy digraph*, a directed graph in which nodes represent students

who point to each other if they are envious of their DA assignment, irrespective of whether the envy is justified. Proposition 4.3.2 shows that a student is improvable if and only if they lie on a trading cycle in this graph. This characterization allows us to identify improvable students using a graphical approach and directly implies that sources and sink nodes in the envy digraph are unimprovable. While sink nodes correspond to those assigned to their top choice under DA (who trivially envy nobody), source nodes (those whom nobody envies) always include every unassigned student or those matched to their least preferred choice (Proposition 4.3.3). This observation has an important policy implication: students who are poorly assigned or unassigned cannot be helped by any mechanism that preserves DA’s outcome as a baseline.

Next, we focus on *how many* students are unimprovable. Using the probabilistic framework of random matching markets with uniformly and independently distributed preferences, we analyze the proportion of unimprovable students in a typical large market, and show that this fraction becomes arbitrarily small as market size grows large. The proof leverages the density of DA’s envy digraph, as most students (resp. schools) make (resp. receive) a logarithmic number of applications. This density implies that the probability of two sizeable groups of students not pointing to each other becomes vanishingly small in large markets. The envy relations in school choice markets under DA exhibit a characteristic property of random homogeneous directed graphs—the emergence of a unique giant strongly connected component where most nodes participate in at least one cycle.

Although DA’s inefficiency was previously known, Theorem 4.4.1 reveals that its pervasiveness is not confined to pathological examples or specific datasets but emerges as a systematic, expected property in random markets. The consequences are so severe that the vast majority of students could find ways to improve their placement without harming others. Our result has important implications: while the most disadvantaged students cannot benefit from enhanced efficiency, the vast majority of students are improvable, revealing both the substantial scope for welfare improvements and the fundamental inefficiency of DA in large markets.

Our final result, Theorem 4.4.6, addresses a related policy question: which efficient mechanism dominating DA should be chosen to maximize the number of improved students? By connecting student improvability to the classical cycle-packing problem in graph theory, we prove that in large markets, any mechanism that Pareto dominates DA will improve approximately the same number of students—almost all of them. This surprising asymptotic equivalence occurs because the high density of DA’s envy digraph ensures that if a student is not included in one particular trading cycle, they will likely participate in an alternative trade. This result has important implications for market design: while different mechanisms that Pareto dominate DA may perform differently in small markets [122], they are virtually indistinguishable in terms of size of improvement in large markets. Thus, policymakers can focus on other desirable properties—such as minimal manipulation incentives or weak stability, both areas in which EADA excels—when selecting among mechanisms that Pareto dominate DA.

In summary, our contributions are threefold. First, we introduce the novel envy digraph

framework that provides a precise characterization of improvable students in school choice. Second, we prove that almost all students are improvable in large markets, revealing the pervasive inefficiency of DA and the substantial scope for welfare improvements. Third, we establish the asymptotic equivalence of all efficient mechanisms that Pareto dominate DA, demonstrating that the choice among such mechanisms becomes increasingly irrelevant as markets grow.

4.1.1 Related Literature

Our work contributes to several strands of literature on matching mechanisms and random market models. We organize our discussion around four interconnected themes.

On DA’s Pareto-inefficiency. The Pareto-inefficiency of DA (for students) is well-known [5]. Theoretical bounds on DA’s inefficiency were established by [120]. Numerous empirical studies have documented this inefficiency in practice [3, 4, 54, 164].

Efficient mechanisms that dominate DA. The most prominent mechanism in this class is Efficiency-Adjusted Deferred Acceptance (EADA) [120]. Over the past decade, EADA’s properties and implementation have been extensively studied [35, 197, 74, 201, 202, 77, 198, 71, 180, 56], demonstrating that it is possible to achieve an efficient improvement over DA while maintaining relatively low instability and manipulability. Another well-known efficient mechanism that Pareto-dominates DA is DA+TTC, which applies the top trading cycles (TTC) procedure to the allocation obtained by DA [10].

Two seminal papers have studied the family of mechanisms that Pareto dominate DA: [16] provide a comprehensive study of stable-dominating rules, which include any efficient mechanism that improves on student-proposing DA, and [197], who introduce the concept of improvable students and under-demanded schools. While these foundational works identify the existence of improvable students, they do not provide a complete characterization of who is unimprovable or quantify their prevalence—gaps we address by introducing the envy digraph framework and deriving asymptotic bounds.¹

Empirical evaluations of mechanisms that dominate DA. [68] conduct an extensive comparison across course allocation datasets, finding that the difference between DA and EADA rank distributions is not statistically significant in 7 out of 19 datasets. They also show that EADA and DA always match the same number of students. [164] compute counterfactual EADA allocations for Budapest’s school choice system and find that, while EADA improves the average student placement, it does not help unassigned students or the worst-placed pupil (assigned to his 13th ranked school under both mechanisms) or those

¹Tang and Yu [197, p. 548] conjectured that the set of unimprovable students is large. Our work demonstrates that this is not the case in large markets.

unassigned. In laboratory experiments, [52] find that EADA generates Pareto-efficient allocations more frequently than DA in the lab but note that the distribution of efficiency improvements is uneven among participants.

Random matching markets. Our analysis of the proportion of unimprovable students builds on a rich literature on random matching markets at the intersection of economics and computer science [208, 123, 172, 170, 108, 127, 53, 137, 136, 142, 21, 55, 54, 177, 25, 20, 163, 164, 182]. Our work advances this literature by providing a characterization and quantification of unimprovable students, introducing new analytical tools that connect matching theory to random graph theory. The emergence of giant strongly connected components in random digraphs is a well-studied phenomenon in graph theory, but its application to understanding efficiency in matching markets represents a novel theoretical contribution.

Finally, our Theorem 4.4.6 is related to a number of asymptotic results in matching markets. While previous equivalence theorems focus on anonymous statistics such as normalized rank distributions or average utilitarian efficiency [53, 137, 142, 177, 55], we prove equivalence in terms of the number of students who improve over DA. Since this statistic is non-anonymous and mechanisms that Pareto dominate DA are not strategy-proof, existing equivalence results, in particular Pycia’s comprehensive framework, do not apply to our setting, making Theorem 4.4.6 a distinct theoretical contribution.

4.2 Model

Following [5], a school choice problem P consists of a finite set of students I and a finite set of schools S . Each student i has a strict preference $\succ_i$ over the schools. Each school s has a quota of available seats q_s and a strict priority over the students $\triangleright_s$, determined by local educational regulations. To allow for unassigned students, we allow the existence of a null school (denoted $s_\emptyset$), which has unlimited capacity. We use $n = |S|$ to denote the number of schools.

For a given school choice problem P , a *matching* μ is a mapping from I to S such that no school is matched to more students than its quota. We denote by μ_i the school to which student i is assigned and by μ_s^{-1} the set of students assigned to school s . With this notation, a matching needs to satisfy $|\mu_s^{-1}| \leq q_s$ for every $s \in S$. We call every student i with $\mu_i = s_\emptyset$ *unassigned*.

The function $\text{rk}_i : S \rightarrow [n]$ specifies the rank of school s according to the preference profile $\succ_i$ of student i :

$$\text{rk}_i(s) = |\{s' \in S : s' \succ_i s\}| + 1,$$

so that the most desirable option gets a rank of 1, whereas the least desirable gets a rank of n . With some abuse of notation, we use the same rank function to specify the students’ rank per the priority profile of schools, i.e. $\text{rk}_s(i) = |\{j \in I : i \triangleright_s j\}| + 1$.

A matching μ *weakly Pareto-dominates* matching ν if, for every student $i \in I$, $\text{rk}_i(\mu_i) \leq \text{rk}_i(\nu_i)$. A matching μ *Pareto-dominates* matching ν if μ weakly dominates ν and there exists a student $j \in I$ with $\text{rk}_j(\mu_j) < \text{rk}_j(\nu_j)$. A matching is *Pareto-dominated* if there exists a matching that Pareto-dominates it and is *Pareto-efficient* if it is not Pareto-dominated.

Student i *desires* school s in matching μ if $\text{rk}_i(s) < \text{rk}_i(\mu_i)$ and he *envies* student j at matching μ if $\text{rk}_i(\mu_j) < \text{rk}_i(\mu_i)$. We say that student j *violates* student i 's priority at school s in matching μ if i desires s , $\mu_j = s$, and $\text{rk}_s(i) < \text{rk}_s(j)$. A matching μ is *non-wasteful* if every school s that is desired by some student in μ satisfies $|\mu_s^{-1}| = q_s$. A matching μ is *stable* if it is non-wasteful and no student's priority at any school is violated in μ .

A *mechanism* associates a matching to every school choice problem. We are mainly interested in the *student-proposing deferred acceptance (DA) mechanism* [88], which works as follows:

Round 1: Every student applies to her most preferred school. Every school tentatively accepts the best students according to its priority, up to its capacity, and rejects the rest.

Round k : Every student rejected in the previous round applies to her next best school. Among both new applicants and previously accepted students, every school accepts the best students according to its priority, up to its capacity, and rejects the rest.

The procedure stops when there is a round without any new rejection. We use $\text{DA}(P)$ to denote the unique resulting matching generated by DA in school choice problem P . $\text{DA}_i(P)$ denotes the school to which student i is assigned. $\text{DA}_s^{-1}(P)$ denotes the set of students assigned to school s in $\text{DA}(P)$.

Similarly, we will use $\mathbf{M}$ to denote an arbitrary mechanism that returns a Pareto-efficient allocation that dominates DA. We use $\mathcal{M}(P)$ to denote the class of such mechanisms for school choice problem P , which include Efficiency-Adjusted Deferred Acceptance (EADA) and Top Trading Cycles using DA's allocation as endowments (DA+TTC).

We now introduce a key concept in this chapter: unimprovable students.

Definition 4.2.1. A student $i \in I$ is *unimprovable* if, for every matching $\mathbf{M}(P) \in \mathcal{M}(P)$, we have $\mathbf{M}_i(P) = \text{DA}_i(P)$.

Unimprovable students do not necessarily need to be in an adverse situation. For example, every student who is assigned to their most preferred school in DA is unimprovable. We denote by $U(P)$ the set of unimprovable students in school choice problem P . It is well-known that, for any school choice problem, there is at least one unimprovable student [197].

4.3 Identifying (Un)Improvable Students

Recall that a digraph G is given by a set of nodes V and a set of directed edges E . We define DA's envy graph as follows.

Definition 4.3.1. *Given a school choice problem P , we define DA's envy digraph $G^{\text{DA}(P)} = (I, E(\text{DA}(P)))$ where each node corresponds to a student and edge (i, j) exists if i envies j in DA.*

DA's envy digraph gives us a full characterization of unimprovable students by identifying cycles in $G^{\text{DA}(P)}$. A *cycle* of a digraph is a sequence of nodes $i_0, i_1, \dots, i_j, i_0$ such that there is an edge between any consecutive nodes and no edge is repeated. A *trading cycle* is a cycle in which every node appears only once, except for i_0 .

Proposition 4.3.2. *A student is unimprovable if and only if he does not belong to any cycle in $G^{\text{DA}(P)}$.*

Proof. The if direction, stating that a student is unimprovable if he does not belong to any cycle, is well-known in the literature [120, 197, 78].

For the only if part: if the cycle in which he participates is a trading cycle, then simply execute such exchange so that each student node is now assigned to the school assigned to the student node he points to. By construction of the graph, every student node involved in the swap is better off, while all other students remain unaffected.

If the cycle is not a trading cycle, then at least one node besides i_0 appears multiple times in the sequence. Let i_k be any such node. Then, rewrite the sequence so that every node appearing before i_k first appears remains, and every node appearing after i_k last appears remains (in other words, remove everything between i_k first and last appearance). Note that the modified cycle that obtains is still a cycle in the original graph. If any repeated node remains, repeat the previous step until there are no other repeated nodes left. Because the length of the cycle is finite, we end up with a trading cycle in which students may execute the exchange prescribed by the edge's direction to improve their placement without affecting others, concluding the proof. $\square$

This characterization will be key for us to quantify unimprovable students in a probabilistic framework in the next section using a graph-theoretical approach. We can use a similar approach as before to find students who are always unimprovable, as follows.

Proposition 4.3.3. *For any school choice problem P , if there is a student i such that*

i) $\text{rk}_i[\text{DA}_i(P)] = n$ and there is at least one school that is under-demanded, or

ii) $\text{DA}_i(P) = s_\emptyset$,

then i is unimprovable.

Proof. For part i): if a student i is assigned to his least preferred school and some school is under-demanded, then $\text{DA}_i(P)$ must be the only under-demanded school because every other institution rejected student i , and since any students assigned to an under-demanded school is unimprovable [197], the conclusion follows. Part ii) is proven by [16], but we provide a short proof nonetheless. If a student is not assigned, then the fictitious school $s_\emptyset$ is under-demanded, as it did not reject any student due to its infinite capacity. Consequently, per the previous argument, such student is unimprovable. $\square$

The characterization of unimprovable students through cycles of the envy digraph provides us with a clear identification of specific categories of students who cannot benefit from any mechanism that Pareto dominates DA. However, it remains unclear how many students fall into these categories in a typical school choice problem. In the next section, we turn to quantifying the proportion of unimprovable students using a probabilistic framework, which will allow us to understand the scope and distribution of potential welfare improvements through mechanisms that Pareto dominate DA.

4.4 Quantifying (Un)Improvable Students

Now we turn to the question of identifying how many students are expected to be unimprovable on an average school choice problem. We consider a *random school choice problem* where strict preferences and priorities are drawn independently and uniformly at random, with an equal number of students and schools (and each school's quota is one). We will use this framework assuming (as the majority of the aforementioned literature) that the number of students and schools is equal and given by n . The benefit of imposing this strong assumption is that it allows us to obtain tractable results. We denote by P_n a random instance of a school choice problem of size n and by $G^{\text{DA}}(P_n)$ the corresponding random envy digraph generated by DA in such problem.²

In such markets, it can be shown using a coupon collector argument that in expectation about n/H_n (resp. H_n) students envy (resp. are envied by) nobody asymptotically as $n \rightarrow \infty$, where $H_n = \Theta(\log n)$ is the n -th harmonic number; see Ortega et al. [165] for a formal analysis. Importantly, the fraction of students who are unenvied or who envy nobody vanishes (albeit slowly) as n grows large. This observation naturally leads us to a broader question: What is the asymptotic behavior of the total fraction of *unimprovable* students? To answer this, we introduce a novel approach using strongly connected components in graph theory, which enable us to characterize the global structure of DA's envy digraph as market size grows.

²We discuss extensions to many-to-one matching and correlated preferences in Appendix C.1.

4.4.1 The Giant Strongly Connected Component

A digraph is *strongly connected* if there is a path in each direction between each pair of nodes. That is, a path exists from the first node in the pair to the second, and another path exists from the second node to the first. A *strongly connected component (SCC)* of a directed graph $G = (I, E)$ is a subgraph $G' = (I' \subseteq I, E' \subseteq E)$ that is strongly connected and is maximal with this property: no additional edges or vertices from G can be included in the subgraph without breaking its property of being strongly connected. The *size* of an SCC is the number of nodes in it. An SCC is *trivial* if it only includes one node. An SCC is *giant* when the fraction of nodes it contains is bounded away from zero asymptotically.

In our envy digraph, improvable students and non-trivial SCCs are related by Proposition 4.3.2. First, note that if a student belongs to a non-trivial SCC, it must be part of a cycle with other nodes in the SCC and therefore is improvable. Conversely, if a student is improvable, they must belong to a cycle and every node in this cycle can reach the others and be reached by them, thereby forming (part of) a non-trivial SCC, which the student we started out with belongs to.

While the geometric out-degree distribution of $G^{\text{DA}(P_n)}$ differs from the Poisson distribution characteristic of the standard Erdős–Rényi random digraph model, we demonstrate below that a unique giant SCC also emerges in DA’s random envy digraph with high probability.

Theorem 4.4.1. *Fix an arbitrary $\epsilon > 0$. With high probability, $G^{\text{DA}(P_n)}$ has a unique giant SCC containing at least $(1 - \epsilon)n$ students.*

4.4.2 Proof of Theorem 4.4.1

The proof strategy hinges on establishing that DA’s envy random digraph is sufficiently dense to preclude fragmentation into two large disjoint components lacking directed edges between them. Such a partition would imply that no student in the first component had applied to any school matched to students in the second component during the execution of DA, an event whose probability decreases super-polynomially as n approaches infinity. We formalize this intuition through a series of lemmas. First, we demonstrate that, with high probability, any substantial subset of students generate a significant number of applications during DA. Subsequently, the probability of the absence of cross-component applications between large partitions is vanishingly small. The proof concludes by applying the union bound across all possible partitions, establishing the desired result.

Bounding the number of applications. We use a simple fact about DA in large markets: with high probability, the total number of applications made by a sizable subset of students is at least on the order of $n \log n$. This is because any single student, with high probability, needs $\Theta(\log n)$ applications to become matched. We formalize the claim that *every* set of students of size at least ϵn collectively makes at least $\delta n \log n$ applications for some constant $\delta > 0$ depending on ϵ .

Lemma 4.4.2. *Fix $\epsilon > 0$ and $L \leq \infty$. With high probability, all but at most ϵn of the students each make at least L applications in the DA process.*

Proof. It suffices to show that a *single* student i makes at least L applications with probability at least $1 - \delta$, for some δ that can be arbitrarily small as $n \rightarrow \infty$. Once this has been established, we can let X_i be the indicator that student i makes fewer than L applications. Then $\mathbb{E}[\sum_i X_i] \leq \delta n$. By Markov's inequality, the probability that more than ϵn students make fewer than L applications is bounded above by δ/ϵ for large enough n . Taking δ to be sufficiently small relative to ϵ implies that with high probability, at most ϵn students make fewer than L applications.

That any student makes $\Omega(\log n)$ applications in DA in probability is a well-known result in random matching markets [123, 172].³ $\square$

Partition argument and probability of no cross-applications. Suppose that $G^{\text{DA}(P_n)}$ has two disjoint subsets $I_0, I_1 \subseteq I$, each of size at least ϵn , such that there is no directed edge from I_0 to I_1 in $G^{\text{DA}(P_n)}$. By definition, $(i \rightarrow j)$ means student i prefers $\text{DA}_j(P_n)$ to $\text{DA}_i(P_n)$. So if there is no edge from I_0 to I_1 , it means no student in I_0 prefers the school matched to any student in I_1 over its own match. Equivalently, no student in I_0 applied to any school $\text{DA}_j(P_n)$ with $j \in I_1$ in the DA algorithm; otherwise, that application would produce an envy edge $i \rightarrow j$ if that school ultimately ended up matched to j .

Hence, an equivalent statement is: If $G^{\text{DA}(P_n)}$ has an SCC of size *at most* $(1 - 2\epsilon)n$, we can split I into two blocks I_0, I_1 , each at least ϵn , so that no student in I_0 ever applied to any school in $S_1 = \{\text{DA}_i(P_n) : i \in I_1\}$. We first show that for a *fixed* partition (I_0, S_1) , the probability of no-application across blocks is very small.

Lemma 4.4.3. *Let $I_0 \subseteq I$ and $S_1 \subseteq S$ be subsets of students and schools with $|I_0|, |S_1| \geq \epsilon n$. Then the probability that the students in I_0 in combination make at least Ln applications yet no one applies to any school in S_1 is at most $(1 - \epsilon)^{Ln}$.*

Proof. If applications were all mutually independent, the claim would be a simple consequence of standard binomial concentration. In DA, however, the applications depend on the history of rejections and hence see slight correlation. Instead, we consider a version of deferred acceptance with redundant applications (what [123] calls the *amnesiac algorithm*) where we generate students' preferences in a deferred manner and direct each application-to-be-made to a uniformly random school independently. That is, we allow a student to re-apply to a school that has already rejected her (and hence will do so again). To get her true preferences, we simply remove the duplicates from the sequence of applications she will ever make. It is easy to see that this version of amnesiac DA can be coupled exactly to the conventional DA to yield the same outcome.

³See also Ashlagi, Braverman, and Zhao [20, Theorem 3.2] for a characterization of rank distribution in any stable outcome.

In the amnesiac algorithm, each student makes weakly more applications. Thus, it suffices to consider the event that the students in I_0 in combination make at least Ln applications in the amnesiac algorithm yet no one applies to any school in S_1 .

To understand this probability, we use the following coupling trick to handle the preferences (and applications) of students in I_0 : First, generate an (infinite) sequence of i.i.d. uniform samples $\sigma = (\sigma_1, \sigma_2, \dots)$ from the set of schools S independent of the preferences of all schools and all students outside I ; Then, run DA with redundant applications, and whenever an application is to be sampled from a student $i \in I_0$, supply with the next element from σ . We can verify inductively that, at any stage, this coupling yields the same distribution (as in the amnesiac algorithm) for the next applications to make—simply uniform and independent of all previous events. Our target no-cross-application event in the lemma statement implies that (1) at least Ln elements from σ are read before DA terminates and (2) none of these elements belong to the set S_1 ; these two combined imply that $\sigma_j \notin S_1$ for all $j = 1, \dots, Ln$. By construction, the sequence $\sigma_1, \dots, \sigma_{Ln}$ are i.i.d. samples independent of all other sources of randomness, and thus $\mathbb{P}(\sigma_j \notin S_1 \ \forall j \in [Ln]) = (\mathbb{P}(\sigma_1 \notin S_1))^{Ln}$. Since $|S_1| \geq \epsilon n$, we have $\mathbb{P}(\sigma_1 \notin S_1) \leq 1 - \epsilon$, and our desired upper bound of $(1 - \epsilon)^{Ln}$ follows immediately. $\square$

Union bound over all subsets We now combine Lemma 4.4.3 with a union bound over all possible realizations of I_0 and S_1 to obtain the following corollary, from which the proof of Theorem 4.4.1 follows naturally.

Corollary 4.4.4. *Fix $\epsilon > 0$. With high probability, there does not exist a subset of students $I_0 \subseteq I$ and a subset of schools $S_1 \subseteq S$ such that (1) both subsets are of size at least ϵn , and (2) no one in I_0 ever applied to any school in S_1 in DA.*

Proof. For any $I_0 \subseteq I$ and $S_1 \subseteq S$ satisfying the cardinality condition, let $\mathcal{E}_{I_0, S_1}$ denote the event that no one in I_0 ever applied to any school in S_1 in DA. Proving the corollary reduces to bounding the probability of the union event

$$p := \mathbb{P}\left(\bigcup_{|I_0|, |S_1| \geq \epsilon n} \mathcal{E}_{I_0, S_1}\right).$$

Let $\mathcal{A}_{I_0}$ denote the event that students in I_0 makes a combined number of at least Ln applications (including redundant ones) for some constant L to be specified later. By a union bound, we have

$$p \leq \mathbb{P}\left(\bigcup_{|I_0|, |S_1| \geq \epsilon n} (\mathcal{E}_{I_0, S_1} \cap \mathcal{A}_{I_0}) \cup \mathcal{A}_{I_0}^c\right) \leq \sum_{|I_0|, |S_1| \geq \epsilon n} \mathbb{P}(\mathcal{E}_{I_0, S_1} \cap \mathcal{A}_{I_0}) + \mathbb{P}\left(\bigcup_{|I_0| \geq \epsilon n} \mathcal{A}_{I_0}^c\right). \quad (4.1)$$

By Lemma 4.4.3, $\mathbb{P}(\mathcal{E}_{I_0, S_1} \cap \mathcal{A}_{I_0}) \leq (1 - \epsilon)^{Ln}$. The number of ways to choose any $I_0 \subseteq I$ is at most 2^n and same for $S_1 \subseteq S$. We only need to consider partitions where $|I_0| \geq \epsilon n$ and

$|S_1| \geq \epsilon n$, so the total number of relevant partitions is strictly smaller. Thus, the summation term in the right-hand side of expression (4.1) is at most

$$(2^n)^2 \cdot (1 - \epsilon)^{Ln} = \exp((2 \log 2 + L \log(1 - \epsilon))n) \rightarrow 0$$

as $n \rightarrow \infty$, granted that L is sufficiently large compared to ϵ . By Lemma 4.4.2, with high probability, all but $\epsilon n/2$ students each make at least $2L/\epsilon$ applications, and as a result any ϵn students make in total at least $2L/\epsilon \cdot \epsilon n/2 = Ln$ applications, implying the event $\bigcap_{|I_0| \geq \epsilon n} \mathcal{A}_{I_0}$. Thus, $\mathbb{P}(\bigcup_{|I_0| \geq \epsilon n} \mathcal{A}_{I_0}^c) \rightarrow 0$ as $n \rightarrow \infty$. Combining the bounds on the two terms gives an upper bound on p that tends to zero as $n \rightarrow \infty$, finishing the proof. $\square$

We now prove Theorem 4.4.1 using Corollary 4.4.4.

Proof of Theorem 4.4.1. Note that DA with redundant applications produces the exact same matching and envy graph (modulo duplicate edges due to redundancy). Thus, we may use DA with redundant applications as our underlying data generation algorithm.

Given that ϵ is arbitrary, it suffices to show that the giant SCCs contains at least $(1 - 2\epsilon)n$ students with high probability. When this is not the case, there must exist a set of SCCs of combined size at least ϵn that either cannot reach or cannot be reached from the largest SCC. In either case, there exist disjoint subsets $I_0, I_1 \subseteq I$ both of size at least ϵn where there are no edges from I_0 to I_1 , and equivalently no students in I_0 have applied to any school in $S_1 := \{\text{DA}_i(P) : i \in I_1\}$. This is, however, unlikely: By Corollary 4.4.4, the probability of being able to find such sizable subsets I_0 and S_1 with no applications between them tends to zero. Hence, the size of the giant SCCs is at least $(1 - 2\epsilon)n$ with high probability. $\square$

Theorem 4.4.1 yields a surprising corollary, namely:

Corollary 4.4.5. *The fraction of students who are unimprovable converges to zero in large markets.*

Theorem 4.4.1 implies that the set of improvable students forms a substantial majority of the entire student population, but is silent regarding how many of them can be improved by the same mechanism, nor does it specify which mechanism chooses the maximum improvement over DA. Our next Theorem tackles both questions.

4.4.3 Equivalence Among Mechanisms That Dominate DA

We define a *cycle packing* H as a union of pairwise-disjoint cycles in G , with the property that no additional cycles exist in the subgraph $G \setminus H$. $V(H)$ denotes the vertex set of H , or equivalently, the set of nodes that are in some cycle in H . Note that, by Proposition 4.3.2, every efficient mechanism corresponds to a cycle packing of DA's envy digraph. Our second Theorem shows that, with high probability, every efficient mechanism improves almost all nodes.

Theorem 4.4.6. *With high probability, every cycle packing of the envy digraph $G^{\text{DA}(P_n)}$ covers at least $(1 - \epsilon)n$ nodes, for any arbitrary constant $\epsilon > 0$.*

Proof. Let $H \subseteq G^{\text{DA}(P_n)}$ be a cycle packing. Consider the induced subgraph $G^{\text{DA}(P_n)} \setminus H = G^{\text{DA}(P_n)}[I \setminus V(H)]$, which by definition, contains no cycles.

Suppose, for contradiction, that $|V(H)| < (1 - \epsilon)n$. Then, the directed acyclic subgraph $G^{\text{DA}(P_n)} \setminus H$ contains $N = |I \setminus V(H)| > \epsilon n$ vertices. Hence, there exists a topological ordering $i_1, \dots, i_N$ of vertices in $I \setminus V(H)$ such that no directed edge from i_k to i_ℓ for $1 \leq k < \ell \leq N$ is present in $G^{\text{DA}(P_n)} \setminus H$ (and hence $G^{\text{DA}(P_n)}$).

Partitioning these nodes into $I_0 = \{i_1, \dots, i_{\lfloor N/2 \rfloor}\}$ and $I_1 = \{i_{\lfloor N/2 \rfloor + 1}, \dots, i_N\}$ yields two subsets, each with size at least $\lfloor \epsilon n / 2 \rfloor$, and crucially, no edges in $G^{\text{DA}(P_n)}$ go from I_0 into I_1 . This scenario is identical to the one in Theorem 4.4.1, which is improbable due to Corollary 4.4.4. Consequently, we must have $|V(H)| \geq (1 - \epsilon)n$ with high probability. $\square$

Theorem 4.4.6 establishes not merely the prevalence of improvable students but also demonstrates that nearly all such students can concurrently benefit from improvement via non-intersecting trading cycles. Furthermore, Theorem 4.4.6 establishes the asymptotic equivalence across all such mechanisms regarding the proportion of students they improve. This result has significant implications for policy implementation. When policymakers deliberate on which efficiency-enhancing mechanism to adopt as an alternative to DA, concerns about differential distributions of efficiency gains become largely irrelevant, as all such mechanisms ultimately benefit approximately the same number of students in large markets. This equivalence result is particularly striking given that in carefully constructed small instances, different efficient mechanisms that dominate DA may significantly differ in their corresponding number of improved students [122].

4.4.4 Simulations and Extensions

The simulations depicted in Figure 4.1 show the emergence of a unique giant component in $G^{\text{DA}(P_n)}$ for even relatively small values of n .

The simulations in Figure 4.2 illustrate the convergence of unimprovable students to zero as market size increases. The convergence rate is approximately $O(1/\log n)$, directly corresponding to the rate at which students receive their top choice under DA (as shown in Corollary 1). This theoretical relationship explains the gradual nature of the decline.

In the Supplementary Appendix, we discuss two extensions of our work.

1. Many-to-one matching: Our probabilistic analysis has assumed that the number of schools and students is the same, yet our main results (Theorems 4.4.1 and 4.4.6)) can be extended to a model with n schools, each with a quota q , and qn students. In this model, each student makes fewer applications, but each rejection makes him envy q instead of only 1 student. This trade-off between application frequency and envy multiplicity results in an envy digraph that becomes even more densely connected. Consequently, the emergence of a unique giant strongly connected component remains




Figure 4.1: Average fraction of nodes in the largest (blue) and second-largest (red) SCCs (average over 2000 random envy digraphs for each n).




Figure 4.2: Average fraction of unimprovable students (average over 2000 random problems for each n).

a robust phenomenon, as well as the asymptotic size equivalence among efficient mechanisms that dominate DA. However, convergence occurs more slowly as q increases. We provide proofs and simulations in Appendix C.1.

2. Correlated preferences: In Appendix C.2, we introduce correlation on students' preferences by dividing the set of schools into hard tiers, similar to the setting in [25]. We find that several giant SCCs form, in fact as many as tiers, including a sizeable fraction of students. Simulations suggest that the fraction of nodes in a giant SCC converges to one asymptotically.

4.5 Conclusion

DA's inefficiency affects almost all students in large markets, with nearly all students becoming improvable through trading cycles in the envy digraph as market size grows. While our findings suggest that an efficiency adjustment should be implemented, our equivalence theorem says such debate should be guided by other properties and not by the number of students who benefit from a better school placement.

Chapter 5

Local Estimation of the Size of k -Core in Networks

Building on the graph-theoretic reasoning in the previous chapter, we now turn to a different domain—networks with limited observability—but retain a shared analytical lens. Our goal in this chapter is to understand when global properties such as the size of the giant SCC can be inferred from local information (e.g., via sampling), a question relevant for inference and intervention on large, complex networks. We introduce the notion of structural robustness for graphs and examine the possibility to locally estimate the size of k -cores.

5.1 Introduction

Real world networks are often enormous. For example, the Facebook network involves over 3 billion users as of 2024 [154], and the internet easily has more than 15 billion nodes [109]. The sheer sizes of these real world networks makes it hard for an analyst to load the entire graph into memory, or even to fully determine the entire network (e.g., contact tracing over the entire population for the study of epidemics is impossible). As a result, testing global properties or estimating global quantities can often be hard. Examples of such global quantities include the size of the largest component, the size of the k -core, or the probability and size of an epidemic outbreak under a given epidemic model for the spread.

One possible solution to this challenge is to develop local algorithms, that is, algorithms depending only on bounded neighborhoods. The general framework for local algorithms is to sample a vertex v in the graph G , explore its radius r neighborhood (which we denote by $\mathcal{N}_r(v; G)$), and then compute a local estimate for the desired quantity. This subroutine is repeated for T times and the average is taken as the estimate. Advantages of local algorithms include reduced memory use, parallelizability for distributed computing settings (each machine doing computation on an independent neighborhood), and scalability (run time not scaling with the network size).

For instance, a natural local algorithm to estimate the (relative) size of the largest con-

nected component (often called the “giant”) in a graph G operates as follows: Repeatedly draw random vertices v , and output the fraction of them that are in connected components whose size exceeds a certain threshold (checked, e.g., using breadth first search). Clearly, such an algorithm does not always give approximately correct answers: it has no way to distinguish between a large connected graph G_0 and a graph consisting of two copies of G_0 . However, when G is obtained from some random graph model (e.g., Erdős–Rényi random graphs, inhomogeneous random graphs, configuration models, etc.), it is known that, with appropriately chosen parameters, the local algorithm gives approximately correct answer with high probability. These models are locally tree-like, and can be approximated with a branching process.

In this chapter, we analyze convergence of the (fractional) size of the k -core in converging sequences of graphs, with a special interest in the case of $k = 2$. The k -core $C^{(k)}(G)$ of a graph G is the maximum subgraph of G where all vertices have degree at least k . For many classic random graph models such as the Erdős–Rényi model [173], the configuration model [113], and inhomogeneous random graphs [181], the k -core can be analyzed using a branching process approximation, suggesting a natural local algorithm exploring the neighborhoods of a random vertex. More precisely, let $\mathcal{A}^{(k)}(\mathcal{N}_r(v; G))$ be the event that $\mathcal{N}_r(v; G)$ contains a k -regular tree of depth r rooted in v . A local algorithm would sample T vertices from the input graphs, and then output the fraction of times the event $\mathcal{A}^{(k)}(\mathcal{N}_r(v; G))$ occurs as an estimate for the relative size of the k -core. With a little bit of extra work, the results of [181] imply that this local algorithm gives indeed a good approximation for the relative size of the k -core, provided the input graphs is drawn from an inhomogeneous random graph model (aka Graphon model).

Note that in addition to using that these graphs are locally tree-like, the analysis of [181] relies heavily on the conditional independence in an inhomogeneous random graph model – once the features are chosen, edges are independent. But real world networks are neither locally tree-like, nor do they have the independence properties of inhomogeneous random graphs. This raises the question:

How robust under model misspecification are these local algorithms?

To address this question, we want to find suitable, fairly general conditions on a growing sequence of graphs G_n such that these algorithms give approximately correct answers with high probability. In particular, the condition should apply to the random graph models mentioned above. Further, a desirable condition should also generalize to non-tree-like models (e.g., the household models [200, 34, 102]), possibly after modifying the algorithm to allow for a large number of cycles.

One natural condition to impose is that the graph instance should “look like” these models *locally*. Without specifying a particular model, this idea is naturally captured by the by now standard notion of local convergence [40, 12], which requires convergence of the distributions of finite neighborhoods in the sequence.

But such a condition is clearly not enough. Indeed, running any local algorithm on a d -regular graph with n vertices, versus running it on $\sqrt{n}$ disjoint copies of d regular size $\sqrt{n}$

will clearly return essentially identical answers, even though the first graph has a linear-sized connected k -core (at least once d is large enough), while the second one does not. To address this second issue, we will need to add a condition which says that the graph is “sufficiently well connected” to not fool local algorithms. A natural such condition would be expansion, but this seems way too strong, since none of the aforementioned models are expanders, even if one restricts oneself to the giant.

This brings us to the contributions of this chapter, in which we provide a positive result for the special case of the 2-core and illustrates the challenges with k -cores for $k \geq 3$. First, we *formulate a natural generalization of the above local algorithm for the 2-core* that makes sense for graphs containing many loops. We then prove that this algorithm *is robust against misspecifications*; more precisely, it is an asymptotically correct approximation algorithm provided that the sequence of graphs $\{G_n\}$ obeys two conditions:

1. The sequence is locally convergent;
2. It enjoys a notion of structural robustness and has bounded average degree.

These conditions can be further generalized to account for the potential presence of smaller connected components; see Assumption 5.1.4 and Theorem 5.1.5.

Our notion of local convergence is the standard one (and will be reviewed when we state our main theorems), and our notion of structural robustness is a weakening of the notion of large set expansion introduced in [14]. Before precisely defining it, we formulate our algorithm.

To motivate it, we note that if v is in the 2-core of the r -neighborhood $\mathcal{N}_r(v; G)$, then it is clearly in the 2-core of G . But even if locally it is not, it could still be in the 2-core of G . In fact, for locally tree-like graphs, we can almost never certify a vertex to be in the 2-core within the local neighborhood. We capture this in the following notation: given $\mathcal{N}_r(v; G)$, we say that v is *potentially in the 2-core* if there exists a graph G' such that $\mathcal{N}_r(v; G) = \mathcal{N}_r(v; G')$ and v is in the 2-core of G' (note that this is easy to check in a breath first search exploration starting from v , see appendix for an algorithm that is linear in the number of edges in $\mathcal{N}_r(v; G)$). Using this subroutine, we describe the following local algorithm for estimating the size of the 2-core in a graph, as well as that of its largest component, $C_{\max}^{(2)}(G)$, where in this work, we define the size $|H|$ of a graph H to be its number of vertices.

Algorithm 2: Local approximation for the size of 2-core

Input: Graph $G = (V, E)$, locality parameter K , iteration number T

Output: Estimated size of $C^{(2)}(G)$ and $C_{\max}^{(2)}(G)$

```

1 for  $t = 1, \dots, T$  do
2   Sample  $v \in V$  uniformly at random
3   Breadth first explore  $\mathcal{N}_R(v; G)$  up to the maximal radius  $R$  such that
       $|\mathcal{N}_R(v; G)| \leq K$ 
4   Set  $I_t^{(2)} = 1$  if  $v$  is potentially in the 2-core given the neighborhood  $\mathcal{N}_R(v; G)$  and
      0 otherwise
5   Set  $I_t^{(2,\infty)} = 1$  if  $I_t^{(2)} = 1$  and  $\mathcal{N}_R(v; G)$  contains a vertex  $u$  at distance  $R$  from  $v$ ,
      and 0 otherwise
6 return  $I^{(2)} := \frac{1}{T} \sum_{t=1}^T I_t^{(2)}$  as the estimated fractional size of  $C^{(2)}(G)$ , and
       $I^{(2,\infty)} := \frac{1}{T} \sum_{t=1}^T I_t^{(2,\infty)}$  as the estimated fractional size of  $C_{\max}^{(2)}(G)$ 
    
```

Remark 5.1.1. (i) The 2-core candidacy (Line 4-5) can be checked while exploring the neighborhood. See Algorithm 3 in Appendix D.1 for an explicit (pseudocode) implementation of the local exploration. (ii) The algorithm can be parallelized by distributing the T independently sampled vertices across different computing nodes, each exploring the neighborhood of one vertex. (iii) In the distributed setting, the time and memory complexity for each computing node is $O(K^2)$ in the worst case, with K typically chosen to be a large constant. Importantly, the complexity does not scale with $|V|$.

Having formulated our algorithm, we now define our weak notion of expansion. We use the notion of an (ϵ, δ) -cut of a graph $G = (V, E)$ from [67], defined as partition of V into two sets A and A^c such that both have size at least $\epsilon|V|$ while the number of edges between A and A^c is at most $\delta|V|$.

Definition 5.1.2 (Structurally robust sequences with bounded average degree). *We say a possibly random graph sequence $\{G_n\}_{n \in \mathbb{N}}$ is a structurally robust (or robust in short) if for any $\epsilon \in (0, 1/2)$ there exists $\delta > 0$ (depending on ϵ) such that with high probability (whp) G_n has no (ϵ, δ) -cut. We say that such a sequence has bounded average degree if there exists a $\bar{d} < \infty$ such that whp, the average degree of G_n is bounded by $\bar{d}$.*

Robustness requires that for any tolerance parameter $\epsilon \in (0, 1/2)$, there exists a “safety margin” $\delta > 0$ such that whp, no adversarial attacker can reduce the size of the giant by more than ϵn vertices unless they remove at least δn edges (robustness against adversarial attacks). Note that this robustness against adversarial attack implies robustness against random edge deletions: Given any $\epsilon \in (0, 1/2)$ there exists $\delta > 0$ such the size of the giant in the percolated graphs $G_n(p)$ is whp larger than $(1 - \epsilon)n$ provided $p > 1 - \frac{\delta}{2\bar{d}}$.

Remark 5.1.3 (Expansion and robustness). *Note that for any robust sequence, $\frac{1}{n}|C_{\max}(G_n)|$ necessarily converges to 1 in probability. The structural robustness can be seen as a weakening*

of the usual notion of (edge) expansion. With the notation

$$\phi(G, \epsilon) := \min_{A \subseteq V: \epsilon|V| \leq |A| \leq |V|/2} \frac{E(A, V \setminus A)}{|A|},$$

we can rewrite the definition equivalently as requiring the existence of $\delta(\epsilon) > 0$ for any $\epsilon \in (0, 1/2)$ such that $\phi(G_n, \epsilon) \geq \delta(\epsilon)$ whp. Interchanging the quantifiers on δ and ϵ weakens the requirement, and gives us the definition of large-set expanders in [14]: a graph sequence $\{G_n\}_{n \in \mathbb{N}}$ is called large-set expanders if there exists $\delta > 0$ such that

$$\phi(G_n, \epsilon) \geq \delta$$

whp for all $\epsilon \in (0, 1/2)$. The only difference is that here δ may not depend on ϵ .

As it turns out, this slight generalization of expansion makes a big difference for the applicability of our results, since in contrast to the large-set expansion condition, our condition holds for the giant of typical random graph models, in particular for the giant components of both the inhomogeneous random graph [45, Lemma 11.3] (thus in particular, Erdős–Rényi graphs, stochastic block models, etc.) and the configuration model [67, Proposition 4.4].

5.1.1 Main Results

Our main result states that for a sequence of graphs G_n , Algorithm 2 gives an asymptotically correct estimate for the relative size of the 2-core when the sequence is a locally convergent robust sequence with bounded average degree. As it turns out, our result holds for more general sequences of graphs, namely those obtained from the former by edge percolation.

To state these results, we need some notation. First, given a locally finite graph G (i.e., a possibly infinite graph with finite degrees) and a vertex v in G , we use $C(v; G)$ to denote the connected component of v in G . Next, given a graph G and a parameter $p \in [0, 1]$, we define a *percolated* graph, $G(p)$, by independently deleting each edge with probability $1 - p$.

Finally, we recall that a sequence of graphs G_n is called *locally convergent*, if the distribution of the r -neighborhoods of a random vertex in G_n converges for all r . Formally, let $\mathcal{G}^*$ be the space of all rooted, locally finite graphs. We will identify two rooted graphs in $\mathcal{G}^*$ if they are isomorphic, with isomorphism defined by the existence of a bijection between the vertex sets that map the root into the root, edges into edges, and non-edges into non-edges. One then defines a distance d_{loc} on $\mathcal{G}^*$ by setting $d_{loc}(G, G') = \frac{1}{1+R}$ where R is the largest integer such that the balls of radius R around the roots in G and G' are isomorphic. This metric turns $\mathcal{G}^*$ into a metric space. One then defines a sequence of G_n to be locally convergent if there exists a measure μ on $\mathcal{G}^*$ such that for all continuous, bounded functions f on $\mathcal{G}^*$, the empirical averages $\frac{1}{n} \sum_{v_n \in G_n} f(G_n, v_n)$ converge¹ to $\mathbb{E}_\mu[f]$, where for notational convenience, we have assumed that G_n has precisely n vertices, a convention we use throughout this chapter.

¹If the sequences G_n is random, this quantity is random, and we require convergence in probability (we will always assume that μ is deterministic, making the right hand side non-random).

Given a sample (G, o) drawn from the limit measure μ , we then consider the percolated graph, $G(p)$, and determine the probability, $\zeta(p; \mu)$, that the component of the root is infinite,

$$\zeta(p; \mu) := \mathbb{P}_{(G, o) \sim \mu}(|C(o; G(p))| = \infty). \quad (5.1)$$

We call ζ the percolation function. Noting that ζ is non-decreasing, we define the percolation threshold $p_c(\mu)$ as $p_c(\mu) := \inf\{p \in [0, 1] : \zeta(p; \mu) > 0\}$. We often drop the μ dependency in ζ and p_c when it is clear from the context. As we will see in Theorem 5.2.3 below, for all $p \neq p_c$, the size of the largest component in the percolated sequence $G_n(p)$ divided by n converges to $\zeta(p)$ when the sequence G_n is locally convergent with bounded average degree and obeys our robustness assumption. This in particular implies that p_c is a threshold for the appearance of a giant component in $G_n(p)$.

With these preparations, we are now ready to state our main result. To the fullest generality, we will make the following assumptions on a sequence $\{G_n\}_{n \in \mathbb{N}}$ of graphs.

Assumption 5.1.4. (a) **Local convergence:** G_n converges locally in probability to $(G, o) \in \mathcal{G}^*$ with law μ ;

(b) **Tight sizes of small components:** The small components have tight sizes in the sense that for any $\epsilon > 0$, there exists $K < \infty$ such that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{v \notin C_{\max}(G_n)} \mathbb{1}\{|C(v; G_n)| \geq K\} < \epsilon$$

in probability.²

(c) **Robust giant with bounded average degree:** In the case when $\zeta(1) > 0$, the largest component $C_{\max}(G_n)$ is robust with bounded average degree in probability.

Theorem 5.1.5 (Local algorithm for the 2-core). *Let $\{G_n\}_{n \in \mathbb{N}}$ be a (possibly random) sequence of graphs satisfying Assumption 5.1.4, and let $G_n(p)$ be the percolated graphs as above. For any fixed error tolerance $\epsilon > 0$ and $p \neq p_c$, there exist constants K, T such that Algorithm 2 estimates the relative size of the 2-core (of the entire graph and/or of its largest connected component) up to an ϵ additive error for the percolated graph $G_n(p)$ with high probability. That is, with input $G_n(p)$, and parameters K, T sufficiently large, the output of Algorithm 2 satisfies*

$$\left| \frac{1}{n} |C^{(2)}(G_n(p))| - I^{(2)} \right| \leq \epsilon \quad \text{and} \quad \left| \frac{1}{n} |C_{\max}^{(2)}(G_n(p))| - I^{(2, \infty)} \right| \leq \epsilon, \quad (5.2)$$

each with probability at least $1 - \epsilon$ for n sufficiently large. Furthermore, the required number of iterations T can be bounded by $C \frac{1}{\epsilon^2} \log \frac{1}{\epsilon}$ for some universal constant C .

²Under the assumption of local convergence, this tightness condition is equivalent to the statement that $\frac{1}{n} |C_{\max}(G_n)| \xrightarrow{P} \zeta(1)$ [101, Theorem 2.2 and Remark 2.3].

To prove the theorem, we will analyze the local neighborhoods of a random vertex v in G_n , and examine under which condition a vertex v that is *potentially* in the 2-core (given the neighborhood $\mathcal{N}_R(v; G)$) is actually in the 2-core. As we will see, asymptotically, as $R \rightarrow \infty$, this is the case for most vertices v if G_n is robust with bounded average degree. For the more general case where G_n consists of a robust giant and small components with tight sizes (Assumption 5.1.4), the heuristic above holds for $C_{\max}(G_n)$ and the small components can be treated using standard techniques on locally convergence. As a byproduct of implementing this proof strategy, we will first prove the following theorem about the asymptotic size of $C^{(2)}(G(p))$ and $C_{\max}^{(2)}(G(p))$.

To state it, we define the analogue of the percolation function ζ for the 2-core. To this end, we consider an infinite, rooted graph (G, o) drawn from the limiting distribution μ of our locally convergent sequence G_n , and consider the 2-core $C^{(2)}(G(p))$ of the percolated graph $G(p)$. It will be convenient to divide $C^{(2)}(G(p))$ into the union of all its finite components, $C^{(2, < \infty)}(G(p))$, and the remainder, $C^{(2, \infty)}(G(p)) = C^{(2)}(G(p)) \setminus C^{(2, < \infty)}(G(p))$. We then define three quantities: the probability that the root is in an infinite component of the 2-core after percolation,

$$\zeta^{(2, \infty)}(p) := \mathbb{P}_{(G, o) \sim \mu}(o \in C^{(2, \infty)}(G(p))), \quad (5.3)$$

the analogue for finite components, $\zeta^{(2, < \infty)}(p; \mu) := \mathbb{P}_{(G, o) \sim \mu}(o \in C^{(2, < \infty)}(G(p)))$, and the sum of the two, which we can write as

$$\zeta^{(2)}(p) := \zeta^{(2, < \infty)}(p) + \zeta^{(2, \infty)}(p) = \mathbb{P}_{(G, o) \sim \mu}(o \in C^{(2)}(G(p))). \quad (5.4)$$

To get an intuition for this quantity, we note that if G_n is locally tree-like, then G is almost surely a tree, and $o \in C^{(2)}(G(p))$ if and only if there are two disjoint path from o to ∞ in $G(p)$. Note that this will in particular imply that the root lies in an infinite component of the 2-core of $G(p)$, i.e., $\zeta^{(2)}(p) = \zeta^{(2, \infty)}(p)$. But in general, we get contributions from both $\zeta^{(2, < \infty)}(p)$ and $\zeta^{(2, \infty)}(p)$. We call $\zeta^{(2, \infty)}(p)$ the percolation function for the *giant* 2-core³.

The key to establishing Theorem 5.1.5 is the following theorem on the convergence of relative sizes of 2-cores to the corresponding percolation functions in robust graphs.

Theorem 5.1.6 (Size of the 2-core in robust graphs). *Let $\{G_n\}_{n \in \mathbb{N}}$ be a sequence of graphs converging locally to μ as $n \rightarrow \infty$. Assume that G_n is robust with bounded average degrees (in probability). Then, for all $p \neq p_c$,*

$$\frac{1}{n} |C^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2)}(p) \quad \text{and} \quad \frac{1}{n} |C_{\max}^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2, \infty)}(p). \quad (5.5)$$

³It is often more instructive to study the emergence of a *giant* (i.e., linear-sized) 2-core in percolation, as the small components of the 2-core depend heavily on local structure of the graphs. For instance, consider the following household model, where we replace each vertex in an Erdős–Rényi random graph with a triangle, attaching existing edges to the three vertices at random independently. For any $p \in (0, 1)$, there will always be a constant fraction of the households that are preserved in percolation with probability p , yet a giant 2-core can only exist above the standard percolation threshold.

Proof ideas. To prove the Theorem 5.1.6 we use a two step sprinkling argument. In the first step, we use sprinkling to establish the existence of a non-empty 2-core inside the giant. This non-empty 2-core then serves as the seed for a linear sized 2-core in a second second sprinkling argument. The second step is based on a novel coloring scheme for the vertices in the tree-part.

k -cores for $k \geq 3$. On the contrary, the natural extension of the algorithm and analysis to k -cores ($k \geq 3$) fails. Unlike 2-cores, k -cores in expanders cannot be estimated through local algorithms like Algorithm 2.

Theorem 5.1.7. *For $k \geq 3$, under Assumption 5.1.4, the k -core equivalence of Algorithm 2 may output an estimate for $n^{-1}|C^{(k)}(G_n)|$ that differs from the true value by a constant with constant probability. In fact, there exist a sequence of bounded-degree expanders $\{G_n\}_{n \in \mathbb{N}}$ on n vertices converging locally to a measure μ as $n \rightarrow \infty$ with diverging k -core proportions, i.e.,*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} |C^{(k)}(G_n)| \neq \limsup_{n \rightarrow \infty} \frac{1}{n} |C^{(k)}(G_n)|. \quad (5.6)$$

Theorem 5.1.7 presents a stark contrast in behaviors between 2-cores and higher order cores in expanders. Indeed, the second statement in the Theorem suggests that, for $k \geq 3$, k -cores in expanders fail to be local, which rules out all hope to locally estimate their sizes. This point is best illustrated with the following simple example: the local operation of removing a single edge from a connected k -regular graph does not affect the local limit or the expansion, yet it completely eliminates the k -core of the graph.

Open directions. From an algorithmic point of view, it would be interesting to consider what other parameters can be well approximated by a local algorithm under our structural robustness condition, such as the size of k -connected components. Second, from a more structural perspective, an interesting question is whether the structural robustness property is preserved (for the giant connected component) whp under percolation.

5.1.2 Related Works

The notion of a k -core (or k -degeneracy) was first introduced in [139, 191] as a way to capture cohesion and robustness of a network. The size of the k -core was studied for several random graph models, including the Erdős–Rényi model [173], inhomogeneous random graphs [181], and the configuration model [113]. These models, being locally tree-like, see a k -core within a giant connected component when the parameters are above a certain critical threshold. For these models, it is also well known that the giant component is unique in edge percolation, with its size given in terms of the survival probability of a suitable birth process (e.g., [80, 45, 86]).

The question whether these results for the giant hold in some model free settings has only been analyzed recently [15, 39, 128, 188]. The main assumptions used in these works

are local convergence (originally introduced in a different context in [12, 40]), and some notion of expansion. Sprinkling arguments are used in several of these papers, even though the technique is of course much older and goes back to Erdős and Rényi [80]. Note that it is a-priori not clear what should replace the survival probability of a birth process when the graphs are not locally tree-like, and indeed, the results for the *size* of the giant in all these works assume that the local limit is supported on trees. The above question was independently answered in [14] and [101], where it was realized that the percolation function (5.1) should take the role of the survival probability when the graph is not locally tree-like. Continuing in the line of research of [15, 39, 128, 188], the authors of [14] then replaced the bounded degree expander assumptions of the previous work by the weaker assumption of large set expanders of bounded average degree, a condition they then established for many more realistic models, including preferential attachment [176, 36]⁴.

However, before this work, very little was known about the size of k -cores in the absence of specific modeling assumption. Our results generalize the condition in [14], and determine the size of both the giant and the 2-core in percolation under this weakened assumption.

Finally, local algorithms on networks have been extensively studied as a solution to the challenge arising from enormous networks. The earliest works date back to [18, 141] from a distributed computing perspective. Lower bounds were established to characterize the tradeoff between locality and approximation quality (e.g., [141, 158, 130, 129, 91]). In recent years, there has also been a significant growth of the local computation algorithm community following the work by Rubinfeld et al. [184], which shares a similar emphasis on locality within massive graphs but employs a different complexity metric. A variety of problems have been considered, often involving combinatorial optimization on massive networks. Examples include maximal independent sets [189, 90], maximal matchings [162, 146], vertex cover [167, 162], distance estimation [148], graph spanners [168, 138], page rank [17], and cluster detection [178, 145].

Most closely related to this work is the work of [13], who consider the problem of local estimation of the probability and size of epidemic outbreaks on networks, giving an algorithm with theoretical guarantees for an epidemic model where the duration of an infection is deterministic, provided the underlying contact network is a large set expander of bounded average degree that converges locally. In spirit our work is closely related to this work.

5.2 Preliminaries

Consequences of local convergence. The following fact offers about percolation on locally converging graph sequences will be used repeated.

Lemma 5.2.1 (Lemma 3.1, [14]). *Let $\{G_n\}_{n \in \mathbb{N}}$ be a sequence of (possibly random) graphs converging locally to $(G, o) \in \mathcal{G}^*$ with law μ . Let $f : \mathcal{G}^* \rightarrow \mathbb{R}$ be a bounded function that*

⁴Unfortunately, the giant component of an Erdős–Rényi random graph generally does not satisfy this condition.

depends only on a finite neighborhood around the root. Then, for all $p \in [0, 1]$,

$$\frac{1}{n} \sum_{v \in V(G_n)} f(G_n(p), v) \xrightarrow{p} \mathbb{E}[f(G(p), o)]. \quad (5.7)$$

By choosing f to be appropriate indicator functions on neighborhoods, we obtain the next lemma as a corollary, which states that almost all vertices have neighborhoods of a bounded size in locally converging graphs.

Lemma 5.2.2 (Bounded neighborhood size). *Let $\{G_n\}_{n \in \mathbb{N}}$ be a locally converging sequence with limit $(G, o) \in \mathcal{G}^*$ with law μ . For any $r \in \mathbb{N}$ and $\epsilon \in (0, 1)$, there exists $N \in \mathbb{N}$ such that*

$$\frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{|\mathcal{N}_r(v)| \geq N\} \leq \epsilon \quad \text{whp.} \quad (5.8)$$

Proof. Let $f_N(G, o) = \mathbb{1}\{|\mathcal{N}_r(o)| \geq N\}$, which is clearly local and bounded. By Lemma 5.2.1,

$$\frac{1}{n} \sum_{v \in V(G_n)} f_N(G_n, v) \xrightarrow{p} \mathbb{E}[f_N(G, o)] = \mathbb{P}(|\mathcal{N}_r(o)| \geq N).$$

Since G is locally finite a.s., for N sufficiently large, we have $\mathbb{P}(|\mathcal{N}_r(o)| \geq N) \leq \epsilon/2$, which completes the proof. $\square$

Size of the giant component in percolation. For a slightly stronger notion of expansion, the following theorem was proved in [14]. It turns out that the proofs in [14] easily generalize to our notion of expansion, yielding the following theorem. We sketch its proof in Appendix D.2.

Theorem 5.2.3. *Let $\{G_n\}_{n \in \mathbb{N}}$ be a locally converging robust sequence with bounded average degree, with limit law μ over $\mathcal{G}^*$. Then $\frac{1}{n}|C_{\max}(G_n(p))| \xrightarrow{p} \zeta(p)$ for all $p \neq p_c$. Furthermore, ζ is continuous except possibly at p_c .*

In view of our remarks on robustness against random edge deletions, one easily sees that under the assumptions of the theorem, $\zeta(p) \rightarrow 1$ as $p \rightarrow 1$. So in particular $p_c < 1$ for such sequences.

Components of moderate size in percolation. The following corollary bounds the combined size of moderately large components outside $C_{\max}(G_n(p))$, and follows directly from the convergence of $|n^{-1}C_{\max}(G_n(p))|$ to $\zeta(p)$.

Corollary 5.2.4. *For any $p \neq p_c$ and any $\epsilon > 0$, there exists K sufficiently large such that*

$$\frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{|C(v; G_n(p))| \geq K, v \notin C_{\max}(G_n(p))\} \leq \epsilon \quad \text{whp.} \quad (5.9)$$

Proof. The inequality is only interesting when $C_{\max}(G_n(p))$ itself is of size at least K —otherwise the left hand side is zero.

When $|C_{\max}(G_n(p))| \geq K$, we have

$$\mathbb{1}\{|C(v; G_n(p))| \geq K, v \notin C_{\max}(G_n(p))\} = \mathbb{1}\{|C(v; G_n(p))| \geq K\} - \mathbb{1}\{v \in C_{\max}(G_n(p))\}.$$

Choose K sufficiently large such that $\zeta_{\geq K}(p) := \mathbb{P}_{(G,o) \sim \mu}(|C(o; G(p))| \geq K) \leq \zeta(p) + \epsilon/2$, which is possible due to the convergence $\zeta(p) = \lim_{K \rightarrow \infty} \zeta_{\geq K}(p)$ (e.g., monotone convergence). By Theorem 5.2.3,

$$\frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{v \in C_{\max}(G_n(p))\} = \frac{1}{n} |C_{\max}(G_n(p))| \xrightarrow{p} \zeta(p),$$

and by Lemma 5.2.1,

$$\frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{|C(v; G_n(p))| \geq K\} \xrightarrow{p} \zeta_{\geq K}(p).$$

Taking the difference between these gives the corollary. □

5.3 Emergence of a Non-Empty 2-Core in the Giant Above Criticality

As a first step towards our main result, we establish the following proposition, which states that for $p > p_c$, with positive probability, the root of a graph drawn from the limiting distribution μ remains in an infinite 2-core even after percolation.

Proposition 5.3.1 (Emergence of infinite 2-core above criticality). *For μ a local limit of robust sequences, we have $\zeta^{(2,\infty)}(p) > 0$ whenever $p > p_c$ and $\zeta^{(2,\infty)}(p) = 0$ whenever $p < p_c$.*

Before proceeding to the proof, we first establish a general lemma that lower bounds the probability of a local event on the root using that on any vertex in a neighborhood of the root. The lemma is a direct consequence of the local convergence (in probability) and the submodularity of the limit.

Lemma 5.3.2. *Let $\mathcal{E}(G, o)$ be a local event. Then for any $\delta > 0$, $r \in \mathbb{N}$, and $p \in [0, 1]$, there exists a constant $D \in \mathbb{N}$ (depending on δ and r) such that*

$$\mathbb{P}(\exists u \in \mathcal{N}_r(o; G) \text{ s.t. } \mathcal{E}(G(p), u)) \leq D\mathbb{P}(\mathcal{E}(G(p), o)) + \delta. \quad (5.10)$$

where all probabilities are understood as over the randomness of $(G, o) \sim \mu$ and the percolation.

Proof. Observe that it suffices to show (5.10) for $r = 1$; for the general case, we simply use induction based on the fact that $\mathbb{P}(\exists u \in \mathcal{N}_r(o; G) \text{ s.t. } \mathcal{E}(G(p), u)) = \mathbb{P}(\exists v \in \mathcal{N}_1(o; G), u \in \mathcal{N}_{r-1}(v; G) \text{ s.t. } \mathcal{E}(G(p), u))$.

Since μ is a distribution on the space of locally finite graphs, we have

$$\deg(o; G) \vee \max_{u \sim o} \deg(u; G) < \infty \quad \text{a.s.,}$$

where the notation $u \sim v$ indicates the adjacency of two vertices $u, v \in V$ (in G). Thus, for any $\delta > 0$, there exists $D < \infty$ such that $\deg(o) < D$ and $\max_{u \sim o} \deg(u; G) < D$ with probability at least $1 - \delta$.

Further, μ as a local limit is unimodular and hence involution invariant [11, Proposition 2.2]. In particular, this implies that

$$\mathbb{E} \left[\sum_{u \sim o} \mathbb{P}(\mathcal{E}(G(p), o), \deg(o) < D | (G, o)) \right] = \mathbb{E} \left[\sum_{u \sim o} \mathbb{P}(\mathcal{E}(G(p), u), \deg(u) < D | (G, o)) \right]. \quad (5.11)$$

(Note that this depends on the fact that $\mathbb{P}(\mathcal{E}(G(p), o), \deg(o) < D | (G, o))$ is a local function.) The right hand side is an upper bound for $\mathbb{P}(\exists u \sim o \text{ s.t. } \mathcal{E}(G, u), \deg(u) < D)$. Thus,

$$\begin{aligned} \mathbb{P}(\exists u \sim o \text{ s.t. } \mathcal{E}(G(p), u)) &\leq \mathbb{P}(\exists u \sim o \text{ s.t. } \deg(u) \geq D) \\ &\quad + \mathbb{E}[\deg(o) \cdot \mathbb{P}(\mathcal{E}(G(p), o), \deg(o) < D | (G, o))] \\ &\leq \delta + D\mathbb{P}(\mathcal{E}(G(p), o)) \end{aligned}$$

as we claimed. $\square$

For the next step, we argue that Lemma 5.3.2 allows us to relate the probability that o lies in an infinite component of $C^{(2)}(G(p))$ to the probability that it lies within a bounded distance to an infinite component of $C^{(2)}(G(p))$. To state this formally, we define the r -neighborhood $\mathcal{N}_r(A; G)$ of a set of vertices A as the induced subgraph on the set of vertices with distance at most r from A .

Corollary 5.3.3. *For $p > p_c$ and any $r \in \mathbb{N}$, the probability $\mathbb{P}(o \in C^{(2, \infty)}(G(p)))$ is non-zero if and only if $\mathbb{P}(o \in \mathcal{N}_r(C^{(2, \infty)}(G(p))))$ is non-zero.*

Proof. The “only if” direction is trivial. For the “if” direction, fix $p > p_c$ and assume $a := \mathbb{P}(o \in \mathcal{N}_r(C^{(2, \infty)}(G(p)))) > 0$. Since $C^{(2, \infty)}(G(p)) \subseteq C^{(2, \ell)}(G(p))$ for any $\ell \in \mathbb{N}$ a.s., we have

$$\mathbb{P}(o \in \mathcal{N}_r(C^{(2, \ell)}(G(p)))) \geq \mathbb{P}(o \in \mathcal{N}_r(C^{(2, \infty)}(G(p)))) = a. \quad (5.12)$$

By Lemma 5.3.2, there exists $D \in \mathbb{N}$ such that

$$\mathbb{P}(o \in C^{(2, \ell)}(G(p))) \geq \frac{1}{D} \left(\mathbb{P}(o \in \mathcal{N}_r(C^{(2, \ell)}(G(p)))) - \frac{a}{2} \right) \geq \frac{a}{2D}. \quad (5.13)$$

Sending $\ell \rightarrow \infty$ gives $\zeta^{(2, \infty)}(p) = \mathbb{P}(o \in C^{(2, \infty)}(G(p))) = \lim_{\ell \rightarrow \infty} \mathbb{P}(o \in C^{(2, \ell)}(G(p))) \geq \frac{a}{2D} > 0$. $\square$

We now show that, for $(G, o) \sim \mu$, the probability of o lying in some finite neighborhood of $C^{(2,\infty)}(G(p))$ is strictly positive.

Lemma 5.3.4. *For any $p > p_c$, there exist some $r \in \mathbb{N}$ such that $\mathbb{P}(o \in \mathcal{N}_r(C^{(2,\infty)}(G(p)))) > 0$.*

Proof. It suffices to show that

$$\mathbb{P}\left(o \in \bigcup_{r \in \mathbb{N}} \mathcal{N}_r(C^{(2,\infty)}(G(p)))\right) = \mathbb{P}(|C(o; G(p))| = \infty), \quad (5.14)$$

which is positive due to the super-criticality assumption. The lemma follows from the fact that the quantity above is the monotone limit of $\mathbb{P}(o \in \mathcal{N}_r(C^{(2,\infty)}(G(p))))$ as $r \rightarrow \infty$.

If $|C(o; G(p))| = \infty$ yet $o \notin \mathcal{N}_r(C^{(2,\infty)}(G(p)))$ for all $r \in \mathbb{N}$, then the 2-core of this infinite component must be empty, and thus there exists a unique ray extending from o to infinity. Let $\mathcal{E}_p$ denote such an event. (5.14) reduces to the claim that $\mathbb{P}(\mathcal{E}_p) = 0$.

Assume that $\mathbb{P}(\mathcal{E}_p) = \epsilon > 0$. Choose $p' \in (p_c, p)$ such that $\zeta(p') > \zeta(p) - \epsilon$, which is possible due to continuity of ζ above p_c (see Proposition D.2.4). Obtain $G(p')$ by keeping the common edges in $G(p)$ and $G(p'/p)$; equivalently, first obtain $G(p)$ and then do percolation with probability p'/p on it. Conditional on $\mathcal{E}_p$, there is a unique ray in $G(p)$ from o to infinity, along which the removal of any edge leaves o in a finite component, and thus, with probability 1, o is in a finite component in $G(p')$. Thus, $\zeta(p') \leq \zeta(p) - \epsilon$, a contradiction. $\square$

The proof of Proposition 5.3.1 is now a 3-line argument.

Proof of Proposition 5.3.1. The second part is trivial, since $\zeta(p) = 0$ for $p < p_c$. Assume thus that $p > p_c$. By Lemma 5.3.4, there exist some $r \in \mathbb{N}$ such that $\mathbb{P}(o \in \mathcal{N}_r(C^{(2,\infty)}(G(p)))) > 0$, which by Corollary 5.3.3 implies that $\zeta^{(2,\infty)}(p) = \mathbb{P}(o \in C^{(2,\infty)}(G(p))) > 0$. $\square$

Given that we already know that above p_c there is a unique giant of linear size (Theorem 5.2.3), the second statement of the theorem suggests that percolation on G_n above criticality should also leave a linear sized giant 2-core, of asymptotic size $n\zeta^{(2,\infty)}(p)$. We will prove this in the next subsection, but here we will only establish the following proposition, with the weaker statement that the giant component contains a non-empty 2-core above p_c . This non-empty 2-core will later seed another sprinkling argument to prove the existence of a linear sized 2-core.

Proposition 5.3.5 (Nontrivial 2-core of the giant). *For $p > p_c$, $C^{(2)}(C_{\max}(G_n(p))) \neq \emptyset$ whp.*

To sketch the proof of this proposition, we need some notation. Recall that $\zeta^{(2,\infty)}$ is defined in terms of $C^{(2,\infty)}(G)$, which is the union of the infinite components of the 2-core of a graph G . Let $\mathcal{N}_\ell^+(v)$ be the graph obtained from $\mathcal{N}_\ell(v)$ by planting an infinite ray on each vertex at distance ℓ from v . We then define $C^{(2,\ell)}(G)$ as the set of vertices v in $V(G)$ such that $v \in C^{(2,\infty)}(\mathcal{N}_\ell^+(v))$. In other words, $v \in C^{(2,\ell)}(G)$ if local information about the ℓ -neighborhood cannot rule it out as part of the infinite 2-core.

Since every vertex in $C^{(2,\ell)}(G_n(p))$ is in a connected component of size at least ℓ , and the combined size of all such components except the largest one is arbitrarily small for ℓ sufficiently large (this is a consequence of the uniqueness of the giant, or more precisely, of the fact that $|n^{-1}C_{\max}(G_n(p))| \rightarrow_p \zeta(p)$), we can guarantee $|C^{(2,\ell)}(G_n(p)) \setminus C_{\max}(G_n(p))| \leq \epsilon n$ for any $\epsilon > 0$ granted that we choose ℓ sufficiently large. The following lower bound on the size of $C^{(2,\ell)}(C_{\max}(G_n(p)))$ follows immediately.

Lemma 5.3.6. *For any $\epsilon > 0$ and $p \in (p_c, 1]$, we have for ℓ sufficiently large*

$$\frac{1}{n}|C^{(2,\ell)}(C_{\max}(G_n(p)))| \geq \zeta^{(2,\infty)}(p) - \epsilon \quad \text{whp.} \quad (5.15)$$

Proof. Observe that $C^{(2,\ell)}(G(p)) \supseteq C^{(2,\infty)}(G(p))$. Thus, $\zeta^{(2,\ell)}(p) \geq \zeta^{(2,\infty)}(p)$, which is positive for $p > p_c$ due to Proposition 5.3.1. By the local convergence of $G_n(p)$ to percolation in the limit (see Lemma 5.2.1) and the local determination of $C^{(2,\ell)}$, we have $\frac{1}{n}|C^{(2,\ell)}(G_n(p))| \xrightarrow{p} \zeta^{(2,\ell)}(p)$. It suffices to show that, for ℓ sufficiently large,

$$\frac{1}{n}|C^{(2,\ell)}(C_{\max}(G_n(p)))| \geq \frac{1}{n}|C^{(2,\ell)}(G_n(p))| - \epsilon \quad \text{whp,} \quad (5.16)$$

or equivalently,

$$\frac{1}{n}|C^{(2,\ell)}(G_n(p) \setminus C_{\max}(G_n(p)))| \leq \epsilon \quad \text{whp,} \quad (5.17)$$

Note that any vertex v in $C^{(2,\ell)}(G_n(p) \setminus C_{\max}(G_n(p)))$ must be in a connected component of size at least ℓ but not the largest component. By Corollary 5.2.4, (5.17) holds for ℓ sufficiently large. $\square$

Lemma 5.3.6 suggests that about a correct fraction of vertices in the giant locally look like part of the 2-core. One might hope to directly argue that most of such vertices are intricately connected and therefore truly in the 2-core. We do not know how to prove this directly, and therefore first establish Proposition 5.3.5 using a first sprinkling argument, and then use the non-empty 2-core in the giant as a seed for a second, more complicated sprinkling argument.

Proof of Proposition 5.3.5. Fix $p' \in (p_c, p)$. We may assume that $C_{\max}(G_n(p'))$ is a tree; otherwise its 2-core is already nonempty. By Lemma 5.3.6, we know that, for sufficiently large ℓ , we have

$$\frac{1}{n}|C^{(2,\ell)}(C_{\max}(G_n(p')))| \geq \epsilon := \zeta^{(2,\infty)}(p')/2 > 0 \quad (5.18)$$

whp, and we may condition on this event. Consider a connected subgraph $H \subseteq C_{\max}(G_n(p'))$ with size $\lfloor \epsilon n/2 \rfloor$, constructed, e.g., through breadth/depth first search from a leaf of tree. Removing H (i.e., all vertices along with edges incident to them) from $C_{\max}(G_n(p'))$ leaves a forest, i.e., a collection of disjoint trees on vertices in $C_{\max}(G_n(p')) \setminus H$.

At this point, let us define the following terminology, which we will find useful later. Consider a finite *connected* graph G and a nonempty connected subgraph $H \subseteq G$ that

contains the 2-core of G (possibly empty). Then $G \setminus H$ is a forest, where each tree has a unique vertex that connects to H in G . We think of this vertex as its root, and orient the tree towards the root (we call this direction downstream). We define the downstream of $u \notin H$, denoted $\mathbf{ds}(u)$, as the set of vertices strictly between u and H , and the upstream of u , denoted $\mathbf{us}(u)$, is the set of vertices $v \notin H$ with $u \in \mathbf{ds}(v)$. The *depth* of $u \notin H$ towards H in G , denoted $\text{depth}_H(u; G)$ is recursively defined as 1 if u is a leaf node (i.e., has degree 1) in G and $1 + \max_{v \in \mathbf{us}(u)} \text{depth}_H(v; G)$ otherwise. Observe that removing H from G leaves a forest. Each tree has a unique vertex with maximal depth h , which happens to be the root. See Figure 5.1 for an illustration.

Now we return to the setting where $G = C_{\max}(G_n(p'))$, with $H \subseteq C_{\max}(G_n(p'))$ of size $\lfloor \epsilon n/2 \rfloor$. Let F be the sub-forest of $C_{\max}(G_n(p')) \setminus H$ consisting of trees of size ℓ or larger. Then any vertex $v \in C_{\max}(G_n(p'))$ not in H or F must have its entire upstream contained in the interior of $\mathcal{N}_\ell(v; G_n(p'))$; by the definition of $C^{(2, \ell)}$, this implies it cannot be in $C^{(2, \ell)}(C_{\max}(G_n(p')))$, see again Figure 5.1 for an illustration. Put differently, all vertices of $C^{(2, \ell)}(C_{\max}(G_n(p')))$ must be either in H or F . Thus,

$$|V(F)| \geq |V(F) \cap C^{(2, \ell)}(C_{\max}(G_n(p')))| \geq |C^{(2, \ell)}(C_{\max}(G_n(p')))| - |V(H)| \geq \epsilon n - \frac{\epsilon n}{2} = \frac{\epsilon n}{2}. \quad (5.19)$$

This leaves us with H and F both containing at least $\epsilon n/3$ vertices. By robustness of G_n , the minimum cut between H and F in G_n is at least δn for some δ depending only on ϵ , which translate through Menger's theorem to δn disjoint paths joining H and F in G_n . Since $|E(G_n)| \leq \bar{d}n/2$ whp, at least $\delta n/2$ of these paths have a length at most $L = \bar{d}/\delta$. On the other hand, F consists of at most n/ℓ trees, and thus the number of disjoint paths from F to H in $G_n(p')$ is at most n/ℓ (recall that $C_{\max}(G_n(p'))$ is assumed to be a tree). Choose $\ell > 4/\delta$ so that $n/\ell < \delta n/4$, leaving at least $\delta n/4$ disjoint paths in G_n that join H with F and are not fully contained in the tree $C_{\max}(G_n(p'))$.

We now sprinkle edges with probability $\beta = 1 - \frac{1-p}{1-p'}$ to obtain $G_n(p)$ from $G_n(p')$; that is, obtain an independent copy of $G_n(\beta)$ and set $G_n(p)$ to contain the union of (edges in) $G_n(p')$ and $G_n(\beta)$. The probability that none of the $\delta n/4$ disjoint paths of length at most L show up in $G_n(\beta)$ is upper bounded by $(1 - \beta^L)^{\delta n/4} \rightarrow 0$ as $n \rightarrow \infty$. Thus there exist, with high probability, two non-identical paths in $C_{\max}(G_n(p))$ between some pair of vertices $u \in H$ and $v \in F$, implying that $C_{\max}(G_n(p))$ cannot be a tree and must have a nonempty 2-core. $\square$




Figure 5.1: The forest $G \setminus H$, oriented in the downstream direction, with vertices annotated with their depths towards H . For the choice of $\ell = 4$, the red vertices cannot possibly be in $C^{(2,\ell)}(G)$.

5.4 Continuity of Percolation Function for the 2-Core

Before proceeding to a refined analysis, we take a quick but necessary digression to study the continuity of the functions $\zeta^{(2)}$ and $\zeta^{(2,\infty)}$ for local limits of robust sequences. The main idea is again similar to the continuity analysis for the percolation function ζ in [14] (see Section 2.3 and Appendix B). For $\zeta^{(2,\infty)}$, we will analyze the infinite connected components in the percolated graph as is done in [41]. The continuity of $\zeta^{(2)}$ will be easy to deduce once we have continuity of $\zeta^{(2,\infty)}$. Throughout this section, we will be concerned with the random rooted graph $(G, o) \in \mathcal{G}^*$ with law μ , which is assumed to be the local limit of a robust sequence.

Lemma 5.4.1. *The function $\zeta^{(2,\infty)}$, like ζ , is continuous everywhere on $[0, 1]$ except possibly at p_c .*

Proof. Continuity below p_c is trivial since $\zeta^{(2,\infty)}(p) \equiv 0$ for $p < p_c$. Consider any $p > p_c$. By Proposition 5.3.1, $\zeta^{(2,\infty)}(p) > 0$. By definition,

$$\zeta^{(2,\infty)}(p) = \mathbb{P}(o \in C^{(2,\infty)}(G(p))). \quad (5.20)$$

It is straightforward to verify that $o \in C^{(2,\infty)}(G(p))$ if and only if $\tilde{H} = C(o; G(p)) \setminus \{o\}$ satisfies one of the following two conditions:

- (I) There exist at least two neighbors u and v of o in $G(p)$ such that $C(u; \tilde{H})$ and $C(v; \tilde{H})$ (possibly the same component) are both infinite;

- (II) There exists exactly one neighbor u of o in $G(p)$ with $|C(u; \tilde{H})| = \infty$ and at least one other neighbor v of o such that $C(o; \tilde{H})$ contains a cycle or contains another neighbor $w \neq u, v$ of o .

By Theorem 6.7 in [11], every infinite connected component of $G(p) \setminus \{o\}$ (and in particular, of $\tilde{H}$) contains an infinite connected component in $G(p') \setminus \{o\}$ for $p' \in (p_c, p)$ a.s.: Otherwise, with a positive probability, we have an infinite component of $G(p) \setminus \{o\}$, connecting to o in $G(p)$ through $k \in \mathbb{N}$ edges, that contains no infinite component of $G(p')$; by resampling the edge percolation on these k edges, we will end up with a positive probability for $G(p)$ to contain an infinite component without an infinite sub-component in $G(p')$, a contradiction with Theorem 6.7 in [11].

Now it suffices to show that, conditional on $o \in C^{(2,\infty)}(G(p))$ (this is valid since the probability of such an event is strictly positive), we have $o \in \bigcup_{p'' < p} C^{(2,\infty)}(G(p''))$ a.s. This translates to

$$\mathbb{P}\left(o \in \bigcup_{p'' < p} C^{(2,\infty)}(G(p'')) \middle| o \in C^{(2,\infty)}(G(p))\right) = \sup_{p'' < p} \frac{\zeta^{(2,\infty)}(p'')}{\zeta^{(2,\infty)}(p)} = 1, \quad (5.21)$$

which immediately suggests left continuity at p , and when combined with right continuity of $\zeta^{(2,\infty)}$ gives our desired result.

To prove the above claim, let $H = C(o; G(p))$, and conditional on $o \in C^{(2,\infty)}(G(p))$ and the realization of H . Note that the cases (I) and (II) are mutually exclusive. First consider case (I). Choose $p' \in (p_c, p)$. We may assume that $C(u; \tilde{H})$ and $C(v; \tilde{H})$ each contains an infinite connected component of $G(p')$, and hence there exist finite paths from u, v to them in $G(p)$. By sending $p'' < p$ sufficiently close to p , these paths along with the edges (o, u) and (o, v) will ultimately be included in $G(p'')$, which suggests $o \in C^{(2,\infty)}(G(p''))$. For case (II), the reasoning for u remains unchanged, and there is a finite path from v to a (finite) cycle or to w in $G(p)$. This path will again be preserved in $G(p'')$ for $p'' < p$ sufficiently close to p , suggesting $o \in C^{(2,\infty)}(G(p''))$. This proves the claim and completes our proof. $\square$

Corollary 5.4.2. *The function $\zeta^{(2)}$, like ζ , is continuous everywhere on $[0, 1]$ except possibly at p_c .*

Proof. To study continuity of $\zeta^{(2)}$, we consider its point-wise difference with $\zeta^{(2,\infty)}$, denoted as $\zeta^{(2,<\infty)}(p) := \zeta^{(2)}(p) - \zeta^{(2,\infty)}(p) = \mathbb{P}(o \in C^{(2,\infty)}(G(p)), |C(o; G(p))| < \infty)$. Since we are not in a finite connected component, it should be clear that, conditional on the realization of $C(o; G(p))$, for $p'' < p$ sufficiently close to p , no edges in $C(o; G(p))$ are lost and hence $o \in C^{(2,<\infty)}(G(p''))$ with $|C(o; G(p''))| < \infty$. Thus, $\zeta^{(2,<\infty)}$ is left continuous on the entire unit interval. Right continuity should also be obvious. Hence, $\zeta^{(2,<\infty)}$ is continuous, and $\zeta^{(2)} = \zeta^{(2,<\infty)} + \zeta^{(2,\infty)}$ is continuous on $[0, 1]$ except possibly at p_c . $\square$

5.5 Convergence of the Relative Size of 2-Core in Percolation

In this section, we sketch the proof of the convergence in probability of the relative size of the giant 2-core of $G_n(p)$. We first focus on the lower bound in Section 5.5.1; the corresponding upper bound follows from standard arguments using the local convergence, requiring no expansion properties, and is presented in Section 5.5.2. The size of the entire 2-core is later treated in Section 5.5.3.

5.5.1 Lower Bound on the Giant 2-Core Size

We now develop our desired lower bound on the size of $C_{\max}^{(2)}(G_n(p))$, given by the following proposition.

Proposition 5.5.1. *For any $\epsilon > 0$ and $p \neq p_c$, we have*

$$\frac{1}{n}|C_{\max}^{(2)}(G_n(p))| \geq \frac{1}{n}|C^{(2)}(C_{\max}(G_n(p)))| \geq \zeta^{(2,\infty)}(p) - \epsilon \quad \text{whp.} \quad (5.22)$$

A coloring scheme. To prove the proposition, we introduce the following color scheme (see Figure 5.2 for an illustration). Let $G = (V, E)$ be a finite connected graph, and let $H \subseteq G$ be a subgraph that contains $C^{(2)}(G)$. Due to Corollary 5.3.5, it suffices to consider $C^{(2)}(G) \neq \emptyset$. Color vertices and edges of G as follows with a localization parameter $\ell \in \mathbb{N}$:

1. Color a vertex $v \notin H$ **red** if $\text{depth}_H(v; G) \leq \ell$; and color it **purple** if $\text{depth}_H(v; G) > 2\ell$ and $(\text{depth}_H(v; G) \bmod 2\ell) \in \{1, 2, \dots, \ell\}$; otherwise we leave it **black**.
2. Color each edge between same-colored vertices with *consecutive* depths using the same color as its endpoints. This creates **red** and **purple** connected components in $G \setminus H$, which we call colored *segments*. Call a segment *complete* if it contains a vertex with depth divisible by ℓ and *incomplete* otherwise.
3. In each incomplete **purple** segment, revert the color of the vertices and edges (to **black**). In each incomplete **red** segment, color the vertices and edges (along with the edge immediately downstream the segment) **gray**, which will be subsequently ignored.

We will call a vertex outside H *colored* if at the end of this process, it is colored **red** or **purple**. Note that by construction, all colored segments in $G \setminus H$ are trees of height exactly ℓ , where the maximally upstream vertices have distance exactly $\ell - 1$ from the unique maximally downstream vertex; in particular, all segments have size at least ℓ . Further, we will call a segment *regular* if it is not colored **gray**. Recall from the proof of Proposition 5.3.5 the definition of upstream and downstream for a vertex; for a vertex u , we call $\overline{\text{us}}(u) := \{u\} \cup \text{us}(u)$ the *inclusive* upstream. The *regular inclusive upstream* at a vertex $v \notin H$, denoted $\overline{\text{us}}^{\text{reg}}(v)$, refers to all regular vertices in $\overline{\text{us}}(v)$; the *colored inclusive upstream* size at v , denoted $\overline{\text{us}}^{\text{col}}(v)$, is the union of colored segments of $G \setminus H$ that are (completely) contained in $\overline{\text{us}}(v)$. The

next lemma associates the size of the colored upstream with that of the regular upstream. Before stating it, we note that any vertex v in $G \setminus H$ with depth ℓ or smaller cannot lie in $C^{(2,\infty)}(\mathcal{N}_\ell^+(v))$, showing that any vertex in $C^{(2,\ell)}(G) \setminus H$ must have depth at least $\ell + 1$, and hence is either black or purple. See Figure 5.2 for an illustration.

Lemma 5.5.2. *For any regular vertices v , $|\overline{\mathbf{us}}^{\text{col}}(v)| \geq \frac{1}{3}|\overline{\mathbf{us}}^{\text{reg}}(v)|$.*

Remark 5.5.3. *Note that $\overline{\mathbf{us}}^{\text{col}}(v)$ is obtained by removing from $\overline{\mathbf{us}}^{\text{reg}}(v)$ black vertices and colored vertices in segments that are not completely contained in $\overline{\mathbf{us}}^{\text{reg}}(v)$. The lemma bounds both effects.*

Proof. It is easy to see the claim when the regular upstream of v is a single path of length $b = \overline{\mathbf{us}}^{\text{reg}}(v) - 1 \geq \ell$: the worst case occurs when $b = 3\ell - 1$, in which case $\overline{\mathbf{us}}^{\text{col}}(v) = \ell$ (due to the **red** segment at the leaf, as the next **purple** segment is incomplete) and the ratio is $\frac{\ell}{3\ell-1} \geq \frac{1}{3}$.



When the regular upstream of v is not a single path, we may decompose it into stem paths, i.e., paths along vertices of consecutive depth, using the following process.

- Initialize $U = \overline{\mathbf{us}}^{\text{reg}}(v)$, inheriting the depth numbers and colors from the original graph.
- While U is non-empty, iteratively find the vertex u with maximum depth (as inherited from G) in U , breaking tie arbitrarily if necessary.
 - Terminate if u is **red**: all remaining vertices in H must be **red**.
 - Find a stem path to u , using any one if multiple exist.⁵ Extract and remove it from U . The length of this stem path is equal to $\text{depth}(u) - 1$, which is at least ℓ since u is not red.


The process above will decompose the upstream of v into multiple stem upstreams of length at least ℓ and some extra red vertices. The $1/3$ lower bound for colored ratio holds for each stem upstream (where we only count relatively complete colored segments in each path), and assembling the paths back into the upstream of v can only increase the number of relatively complete colored segments. $\square$

Proof of the lower bound. Proposition 5.5.1 will follow from two intermediate lemmas. To state the first, we assume $p_c < p' < p$. Choose H to be $C^{(2)}(C_{\max}(G_n(p')))$ if it reaches size $\epsilon n/8$; otherwise, expand $C_{\max}^{(2)}(G_n(p'))$ into a connected subgraph $H \subseteq C_{\max}(G_n(p'))$ of size $\lfloor \frac{\epsilon n}{8} \rfloor$ (in light of Corollary 5.3.5, we may assume that $C_{\max}(G_n(p'))$ has a nontrivial 2-core). We compute the depths of vertices in $C_{\max}(G_n(p'))$ with respect to H , and apply the coloring scheme to $C_{\max}(G_n(p'))$ with parameter ℓ , which we will choose later.

⁵This is always possible: If u has not been removed, neither must be any of its direct upstream w , since w has a strictly smaller depth than u and cannot have been removed as part of a previous path (which must also go through u , the unique direct downstream of w).


 (a) Color vertices purple or red according to their depth modulo 2ℓ .


(b) Color edges between same-color vertices of consecutive depths to create segments.



(c) Revert colors for vertices and edges in incomplete segments.

Figure 5.2: The coloring scheme applied to $G \setminus H$ with the choice of $\ell = 4$. Vertices are annotated with their depths. Gray segments are presented with dashed edges, indicating that they are ignored during the subsequent sprinkling argument. Note that each red or purple segment has size at least ℓ , and at any non-gray vertex, the colored segments that are (completely) contained in its (inclusive) upstream constitute at least $1/3$ of its regular (inclusive) upstream (see Lemma 5.5.2). Evidently, assuming $H \supseteq C^{(2)}(G) \neq \emptyset$, the inclusion of any vertex into a cycle or a path between cycles during sprinkling will result in its entire downstream being included in the 2-core.

Lemma 5.5.4. *Assume $p > p_c$ and $\epsilon < \zeta^{(2,\infty)}(p)$. For (5.22) to hold, it suffices to prove that for $p' < p$ sufficiently close to p and ℓ sufficiently large*

$$\frac{1}{n} |C^{(2,\ell)}(C_{\max}(G_n(p')) \setminus C^{(2)}(C_{\max}(G_n(p))))| \leq \frac{\epsilon}{2} \quad \text{whp.} \quad (5.23)$$

Furthermore, if the event in (5.23) fails to happen, then there exist a set of colored segments S in $C_{\max}(G_n(p'))$ of combined size at least $\frac{\epsilon n}{8}$ with $S \cap C^{(2)}(G_n(p)) = \emptyset$.

Proof. The first part of the lemma is straightforward: By continuity of $\zeta^{(2,\infty)}$ for $p \neq p_c$ from Lemma 5.4.1, we may choose $p' \in (p_c, p)$ such that $\zeta^{(2,\infty)}(p') \geq \zeta^{(2,\infty)}(p) - \frac{\epsilon}{4}$. For ℓ sufficiently large, Corollary 5.3.6 ensures

$$\frac{1}{n} |C^{(2,\ell)}(C_{\max}(G_n(p')))| \geq \zeta^{(2,\infty)}(p') - \frac{\epsilon}{4} \geq \zeta^{(2,\infty)}(p) - \frac{\epsilon}{2} \quad \text{whp,} \quad (5.24)$$

from which the reduction from (5.22) to (5.23) follows.

For the second part of Lemma 5.5.4, assume that the event in (5.23) fails; that is, there are more than $\frac{\epsilon n}{2}$ vertices in $C^{(2,\ell)}(C_{\max}(G_n(p')))$ that fail to join $C^{(2)}(C_{\max}(G_n(p)))$. We claim that at least $\frac{3\epsilon n}{8}$ of these vertices are in $C^{(2,\ell)}(C_{\max}(G_n(p')) \setminus H$ and hence are again colored black or purple. To see this, we consider two cases: (1) if $H = C^{(2)}(C_{\max}(G_n(p')))$, then this follows from $C^{(2)}(C_{\max}(G_n(p))) \supseteq C^{(2)}(C_{\max}(G_n(p'))) = H$; (2) if $H \supsetneq C^{(2)}(C_{\max}(G_n(p')))$, then $|H| = \lfloor \frac{\epsilon n}{8} \rfloor$ per our construction, and at least $\frac{\epsilon n}{2} - |H| \geq \frac{3\epsilon n}{8}$ of these vertices are in $C^{(2,\ell)}(C_{\max}(G_n(p')) \setminus H$, as claimed.

Next, we claim that if a vertex $v \in C^{(2,\ell)}(C_{\max}(G_n(p')) \setminus H$ stays outside $C^{(2)}(G_n(p))$, then so must its entire upstream. To see this, we show the contrapositive by considering $u \in \text{us}(v) \cap C^{(2)}(G_n(p))$ with a case-by-case discussion:

- If u is in a cycle $\Gamma \subseteq G_n(p)$, then there is a (simple) path from $u \in \Gamma$ to some cycle in $C^{(2)}(C_{\max}(G_n(p')))$ that traverses the entire downstream of u . Then v is on a path between cycles (if not directly in the cycle Γ);
- Otherwise, u sits on a path between two cycles $\Gamma_1, \Gamma_2 \subseteq G_n(p)$. Hence, there are two disjoint paths from u to Γ_1 and Γ_2 , and at least one of them does not overlap with the downstream of u to v (otherwise, u is in a cycle). Concatenating this path with the entire downstream of u gives a simple path from $C^{(2)}(C_{\max}(G_n(p')))$, which contains a cycle, to Γ_1 or Γ_2 , again certifying v as on a path between cycles.

By this observation, we can identify all black and purple vertices $v_1, \dots, v_J$ outside $C^{(2)}(G_n(p))$ that are maximally downstream (i.e., no other black or purple vertices outside $C^{(2)}(G_n(p))$ are in the downstream of them), with their downstream vertices all merged into $C^{(2)}(G_n(p))$ while their inclusive upstreams remaining outside. In particular, any black or purple vertex not in $C^{(2)}(G_n(p))$ must be in the inclusive upstream of v_j for some $j \in [J]$; there will be at least $\frac{3\epsilon n}{8}$ of them when the event in (5.23) fails. Thus, we must have $\sum_{j=1}^J |\text{us}^{\text{reg}}(v_j)| \geq \frac{3\epsilon n}{8}$. By Lemma 5.5.2, $\sum_{j=1}^J |\text{us}^{\text{col}}(v_j)| \geq \frac{\epsilon n}{8}$. In other words, there exist a

set of colored segments in $C_{\max}(G_n(p'))$ of combined size $\frac{\epsilon n}{8}$ or more that remain fully outside $C^{(2)}(G_n(p))$, as we claimed. $\square$

The second lemma that we need for the proof of Proposition 5.5.1 will be a consequence of the robustness of G_n and the bounded average degree. The proof again uses a sprinkling argument.

Lemma 5.5.5. *There exists $\delta > 0$ depending on ϵ such that, for $p_c < p' < p$ and for sufficiently large ℓ , the probability that there exist a set of colored segments S in $C_{\max}(G_n(p'))$ of combined size at least $\frac{\epsilon n}{8}$ with $S \cap C^{(2)}(G_n(p)) = \emptyset$ is at most $2^{n/\ell} \cdot (1 - \beta^L)^{\frac{\delta n}{2} - \frac{n}{\ell}}$, where $\beta = 1 - \frac{1-p}{1-p'}$ and $L = \bar{d}/\delta$.*

Proof. With n vertices in total and the size of each colored segment at least ℓ , there are at most n/ℓ colored segments and hence at most $2^{n/\ell}$ possible subsets S of them. For each subset S of colored segments with combined size $\frac{\epsilon n}{8}$ or larger, by robustness of G_n , there exist at least δn disjoint paths between these segments and H in the original graph G_n , with some $\delta > 0$ depending on ϵ ; further, at least half of these paths have a length at most $L = \bar{d}/\delta$. Let P denote the set of such paths. We obtain $G_n(p)$ from $G_n(p')$ by sprinkling, i.e., recovering each edge in $E(G_n)$ with probability β independently, during which each of the paths in P is recovered with probability at least β^L .

Focus on a specific colored segment A in S , which is a (sub-)tree in $C_{\max}(G_n(p')) \setminus H$. Let u be the maximally downstream vertex in A , which is unique, and let v be the direct downstream of u (possibly in H). If sprinkling recovers a path between A and H that does not go through the direct downstream edge (u, v) , then u must be in a cycle and become part of $C^{(2)}(G_n(p))$: To see this, note that the removal of (u, v) from $G_n(p)$ leaves u and v still connected – via the newly recovered path to H and the existing downstream of v .

Now return to the set P of disjoint short paths between S and H in G_n . Since S contains at most n/ℓ colored segments, at most n/ℓ of the paths in P go through a direct downstream edge of a colored segment in S . For $\ell > \frac{2}{\delta}$, at least $\frac{\delta n}{2} - \frac{n}{\ell}$ disjoint paths in P do not go through the direct downstream edges of the segments in S . As is discussed above, if any of these paths is recovered via sprinkling, S cannot be fully outside $C^{(2)}(G_n(p))$. Thus, the probability that S is fully outside $C^{(2)}(G_n(p))$ is at most $(1 - \beta^L)^{\frac{\delta n}{2} - \frac{n}{\ell}}$. Taking a union bound over all possible choices of S , we ensure that the probability of having a set S of colored segments with combined size at least $\frac{\epsilon n}{8}$ that remain fully outside $C^{(2)}(G_n(p))$ is at most $2^{n/\ell} \cdot (1 - \beta^L)^{\frac{\delta n}{2} - \frac{n}{\ell}}$ as we claimed. $\square$

The proof of Proposition 5.5.1 immediately follows from the two lemmas.

Proof of Proposition 5.5.1. We assume without loss of generality that $p > p_c$ and $\zeta^{(2,\infty)}(p) > \epsilon$. As in Lemma 5.5.4, we choose $p' < p$ sufficiently close to p such that $\zeta^{(2,\infty)}(p') \geq \zeta^{(2,\infty)}(p) - \frac{\epsilon}{4}$, and choose ℓ sufficiently large such that $\frac{1}{n}|C^{(2,\ell)}(C_{\max}(G_n(p')))| \geq \zeta^{(2,\infty)}(p') - \frac{\epsilon}{4}$ and $\ell > \max\{\frac{4}{\delta}, \frac{8}{\beta^L \delta}\}$. By Lemma 5.5.5, the probability of having a set S of colored segments

with combined size at least $\frac{\epsilon n}{8}$ that have empty intersection with $C^{(2)}(G_n(p))$ is at most

$$2^{n/\ell} \cdot (1 - \beta^L)^{\frac{\delta n}{2} - \frac{n}{\ell}} \leq e^{\frac{n}{\ell} - \beta^L \frac{\delta n}{4}} \leq e^{-\beta^L \delta n/8} \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (5.25)$$

Thus, by Lemma 5.5.4, (5.23) holds whp and our proof is finished. $\square$

5.5.2 Upper Bound on the Giant 2-Core Size

In this section, we use standard techniques to prove a high-probability upper bound on the relative size of $C^{(2)}(C_{\max}(G_n(p)))$, as stated in the following lemma. The proof mostly does not rely on any expansion assumptions, except for the knowledge of the giant component.

Lemma 5.5.6. *For any $\epsilon > 0$ and $p \in (0, 1]$, we have*

$$\frac{1}{n} |C^{(2)}(C_{\max}(G_n(p)))| \leq \zeta^{(2,\infty)}(p) + \epsilon \quad \text{whp.} \quad (5.26)$$

Proof. Note that this lemma becomes trivial when $\frac{1}{n} |C_{\max}(G_n(p))| < \epsilon$. Thus, we may assume $|C_{\max}(G_n(p))| \geq \epsilon n$.

Recall from Section 5.3 the definition of the localized event that $o \in C^{(2,\ell)}(G)$ for a rooted graph (G, o) ; in particular, we have $C^{(2,\infty)}(G) = \bigcap_{\ell \in \mathbb{N}} C^{(2,\ell)}(G)$. Let $\zeta^{(2,\ell)}(p) = \mathbb{P}_{(G,o)}(o \in C^{(2,\ell)}(G(p)))$, which approaches $\zeta^{(2,\infty)}(p) = \mathbb{P}_{(G,o)}(o \in C^{(2,\infty)}(G(p)))$ from above. By choosing ℓ sufficiently large, we can guarantee that $\zeta^{(2,\ell)}(p) \leq \zeta^{(2,\infty)}(p) + \epsilon/3$.

Since $C^{(2,\ell)}$ describes a local event, by Lemma 5.2.1, the local convergence of G_n implies

$$\frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{v \in C^{(2,\ell)}(G_n(p))\} = \frac{1}{n} |C^{(2,\ell)}(G_n(p))| \xrightarrow{p} \zeta^{(2,\ell)}(p). \quad (5.27)$$

In particular,

$$\frac{1}{n} |C^{(2,\ell)}(G_n(p))| \leq \zeta^{(2,\ell)}(p) + \frac{\epsilon}{3} \leq \zeta^{(2,\infty)}(p) + \frac{2\epsilon}{3} \quad \text{whp.} \quad (5.28)$$

It now suffices to bound $C^{(2)}(C_{\max}(G_n(p))) \setminus C^{(2,\ell)}(G_n(p))$. Observe that a vertex in $C^{(2)}(G_n(p))$ but not in $C^{(2,\ell)}(G_n(p))$ must have an ℓ -neighborhood that contains a full connected component of $G_n(p)$. In particular, $v \in C^{(2)}(C_{\max}(G_n(p))) \setminus C^{(2,\ell)}(G_n(p))$ implies $|\mathcal{N}_\ell(v; G_n)| \geq |C_{\max}(G_n(p))| \geq \epsilon n$. By Lemma 5.2.2, we conclude that

$$\frac{1}{n} |C^{(2)}(C_{\max}(G_n(p))) \setminus C^{(2,\ell)}(G_n(p))| \leq \frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{|\mathcal{N}_\ell(v)| \geq \epsilon n\} \leq \frac{\epsilon}{3} \quad \text{whp.} \quad (5.29)$$

Combining this with (5.28) concludes the proof. $\square$

Combining Lemma 5.5.6 with the matching lower bound from Proposition 5.5.1 yields the in-probability convergence of the relative size of the giant 2-core.

Theorem 5.5.7. *For any $\epsilon > 0$ and $p \neq p_c$, $\frac{1}{n} |C_{\max}^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2,\infty)}(p)$. At $p = p_c$, the right hand side serves as a high-probability upper bound.*

5.5.3 Size of the Entire 2-Core

We now examine the size of the entire 2-core of $G_n(p)$. The main idea is that the components outside the giant $C_{\max}(G_n(p))$ are mostly small and thus the 2-core of them can almost be locally determined.

Concretely, since we already know from Theorem 5.5.7 the in-probability convergence of the relative size $C^{(2)}(C_{\max}(G_n(p)))$ to $\zeta^{(2,\infty)}(p)$ when $p \neq p_c$, we may simply focus on the quantity $C^{(2)}(G_n(p) \setminus C_{\max}(G_n(p)))$, i.e., the 2-core outside the largest connected component of $G_n(p)$.

It suffices to show the following result, from which Theorem 5.5.9 follows.

Proposition 5.5.8. *For any $p \in [0, 1]$,*

$$\frac{1}{n} |C^{(2)}(G_n(p) \setminus C_{\max}(G_n(p)))| \xrightarrow{p} \zeta^{(2)}(p) - \zeta^{(2,\infty)}(p). \quad (5.30)$$

Proof. Given a graph G , let $\zeta^{(2,<\infty)}(p) = \zeta^{(2)}(p) - \zeta^{(2,\infty)}(p)$, and let $Z_{<\ell}(G)$ (resp. $Z_{\geq\ell}(G)$) denote the set of vertices in connected components of G with size less than ℓ (resp. no less than ℓ).

Fix some $\epsilon > 0$. By Corollary 5.2.4, there exists a sufficiently large constant ℓ such that

$$\frac{1}{n} |Z_{\geq\ell}(G_n(p)) \setminus C_{\max}(G_n(p))| \leq \frac{\epsilon}{3} \quad \text{whp.} \quad (5.31)$$

Observe that, when $|C_{\max}(G_n(p))| \geq \ell$, we have

$$G_n(p) \setminus C_{\max}(G_n(p)) = Z_{<\ell}(G_n(p)) \sqcup (Z_{\geq\ell}(G_n(p)) \setminus C_{\max}(G_n(p)))$$

with $\sqcup$ denoting disjoint union, and when $|C_{\max}(G_n(p))| < \ell$, $C^{(2)}(G_n(p) \setminus C_{\max}(G_n(p)))$ differs from $C^{(2)}(Z_{<\ell}(G_n(p)))$ by at most ℓ vertices. In either case, it suffices to show that, by choosing ℓ sufficiently large, we can guarantee

$$\left| \frac{1}{n} |C^{(2)}(Z_{<\ell}(G_n(p)))| - \zeta^{(2,<\infty)}(p) \right| \leq \frac{2\epsilon}{3} \quad \text{whp.} \quad (5.32)$$

Define $\zeta^{(2,<\ell)}(p) := \mathbb{P}_{(G,o) \sim \mu}(o \in C^{(2)}(G(p)), |C(o; G(p))| < \ell)$. Note that the corresponding events form a monotone increasing sequence as $\ell \rightarrow \infty$, and hence, with $(G, o) \sim \mu$,

$$\begin{aligned} \lim_{\ell \rightarrow \infty} \zeta^{(2,<\ell)}(p) &= \lim_{\ell \rightarrow \infty} \mathbb{P}(o \in C^{(2)}(G(p)), |C(o; G(p))| < \ell) \\ &= \mathbb{P}(o \in C^{(2)}(G(p)), |C(o; G(p))| < \infty) = \zeta^{(2,<\infty)}(p). \end{aligned}$$

By choosing ℓ sufficiently large, we can guarantee $|\zeta^{(2,<\ell)}(p) - \zeta^{(2,<\infty)}(p)| \leq \epsilon/3$. Further, the event that $o \in C^{(2)}(G)$ and $|C(o; G)| < \ell$ for a rooted graph (G, o) is local, and by Lemma 5.2.1,

$$\frac{1}{n} |C^{(2)}(Z_{<\ell}(G_n(p)))| = \frac{1}{n} \sum_{v \in V(G_n)} \mathbb{1}\{v \in C^{(2)}(G_n(p)), |C(v; G_n(p))| < \ell\} \xrightarrow{p} \zeta^{(2,<\ell)}(p). \quad (5.33)$$

Thus, (5.32) holds and our proof is complete. $\square$

Combining Proposition 5.5.8 with Theorem 5.5.7 gives the convergence of the relative size of the entire 2-core in $G_n(p)$, as stated in the next Theorem.

Theorem 5.5.9. *For any $\epsilon > 0$ and $p \neq p_c$, $\frac{1}{n}|C^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2)}(p)$. At $p = p_c$, the right hand side serves as a high-probability upper bound.*

5.6 Robustness of Algorithm 2

Finally, based on the convergence results, we analyze the robustness of Algorithm 2 against model misspecification.

First, we prove the following generalization of Theorem 5.1.6 for graph sequences $\{G_n\}_{n \in \mathbb{N}}$ satisfying Assumption 5.1.4 (not necessarily robust themselves).

Proposition 5.6.1. *Let $\{G_n\}_{n \in \mathbb{N}}$ be a sequence of graphs satisfying Assumption 5.1.4. Then, for all $p \neq p_c$,*

$$\frac{1}{n}|C^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2)}(p) \quad \text{and} \quad \frac{1}{n}|C_{\max}^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2,\infty)}(p). \quad (5.34)$$

Proof. To reduce the general case under Assumption 5.1.4 to locally converging robust graphs in Theorem 5.1.6, we consider two separate cases based on whether a giant component is present in G_n .

Case 1: giant component absent. Assume $\zeta(1) = \mathbb{P}(|C(o, G)| = \infty) = 0$, i.e., the limit μ is supported on finite graphs. As a result, G_n has no giant components and $n^{-1}|C_{\max}(G_n)| \xrightarrow{p} 0$. The convergence of $\frac{1}{n}|C_{\max}^{(2)}(G_n(p))|$ to $\zeta^{(2,\infty)}(p) = 0$ follows directly. With no giant components in G_n , the condition 5.9 in Corollary 5.2.4 holds due to (b) of Assumption 5.1.4. The convergence of $\frac{1}{n}|C^{(2)}(G_n(p))|$ to $\zeta^{(2)}(p)$ follows from the same analysis as in the proof of Proposition 5.5.8.

Case 2: giant component present. Assume $\zeta(1) > 0$ for the local limit μ . Under Assumption 5.1.4, the giant component $\tilde{G}_n = C_{\max}(G_n)$ also have a local limit (in probability) $\tilde{\mu}$ obtained from μ by conditioning the root lying in the unique infinite component [101, Theorem 2.5]; Further, $\tilde{\mu}$ share the same p_c as μ . Theorem 5.1.6 applied to $\tilde{G}_n$ implies

$$\frac{1}{|\tilde{G}_n|}|C^{(2)}(\tilde{G}_n(p))| \xrightarrow{p} \zeta^{(2)}(p; \tilde{\mu}) \quad \text{and} \quad \frac{1}{|\tilde{G}_n|}|C_{\max}^{(2)}(\tilde{G}_n(p))| \xrightarrow{p} \zeta^{(2,\infty)}(p; \tilde{\mu}).$$

Since $\zeta^{(2,\infty)}(p; \tilde{\mu}) = \zeta^{(2,\infty)}(p; \mu)/\zeta(1; \mu)$ and $n^{-1}|\tilde{G}_n| \xrightarrow{p} \zeta(1; \mu)$, we have

$$\frac{1}{n}|C_{\max}^{(2)}(\tilde{G}_n(p))| \xrightarrow{p} \zeta^{(2,\infty)}(p; \mu).$$

The components outside $C_{\max}(G_n)$ are all of size $o(n)$ whp; thus, $C_{\max}^{(2)}(\tilde{G}_n(p))$ is the same as $C_{\max}^{(2)}(G_n(p))$ whp and the claimed convergence of $\frac{1}{n}|C_{\max}^{(2)}(G_n(p))|$ to $\zeta^{(2,\infty)}(p; \mu)$ follows.

Similarly, $\zeta^{(2)}(p; \tilde{\mu}) = \mathbb{P}(o \in C^{(2)}(G(p)), |C(o, G)| = \infty) / \zeta(1; \mu)$, and hence

$$\frac{1}{n} |C^{(2)}(\tilde{G}_n(p))| \xrightarrow{p} \mathbb{P}(o \in C^{(2)}(G(p)), |C(o, G)| = \infty).$$

For the convergence of $\frac{1}{n} |C^{(2)}(G_n(p))|$, it suffices to show

$$\frac{1}{n} |C^{(2)}(H_n(p))| \xrightarrow{p} \mathbb{P}(o \in C^{(2)}(G(p)), |C(o, G)| < \infty). \quad (5.35)$$

The proof of (5.35) is similar to the previous case of no giant components and is a standard consequence of local convergence and tightness of the component sizes; the details are omitted. $\square$

Note that, by Proposition 5.6.1, the actual fractional sizes of $C^{(2)}(G_n(p))$ and $C_{\max}^{(2)}(G_n(p))$ are well approximated by the percolation functions $\zeta^{(2)}(p)$ and $\zeta^{(2, \infty)}(p)$, respectively, for $p \neq p_c$. It then suffices to establish that the estimates $I^{(2)}$ and $I^{(2, \infty)}$ also converge to the ϵ -neighborhood of the corresponding percolation functions in the limit. To ensure this, we need to choose the parameter K in the algorithm to be sufficiently large to achieve good quality of local approximation of the 2-core, and choose T to also be sufficiently large to rely on Hoeffding concentration. With this heuristic, we now present the full proof.

Proof of Theorem 5.1.5. Let μ denote the local limit of the sequence $\{G_n\}_{n \in \mathbb{N}}$ as usual. First, by Theorems 5.5.7 and 5.5.9, we know that for all $p \neq p_c(\mu)$,

$$\frac{1}{n} |C^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2)}(p; \mu) \quad \text{and} \quad \frac{1}{n} |C_{\max}^{(2)}(G_n(p))| \xrightarrow{p} \zeta^{(2, \infty)}(p; \mu). \quad (5.36)$$

Thus, it suffices to have $I^{(2)}$ (resp. $I^{(2, \infty)}$) in the $\frac{\epsilon}{2}$ -neighborhood of $\zeta^{(2)}(p)$ (resp. $\zeta^{(2, \infty)}(p)$) with probability $1 - \frac{\epsilon}{2}$. For now, we focus on $I^{(2, \infty)}$ as an estimate of $\zeta^{(2, \infty)}(p)$; the analysis for $I^{(2)}$ is analogous.

Algorithm 2 realizes T independently chosen neighborhood in the percolated graph, and in particular, $\{I_t^{(2, \infty)}\}_{t \in [T]}$ are i.i.d. Bernoulli random variables (conditional on the input graph G_n if it is random). For any $\ell \in \mathbb{N}$ (to be specified later), we can decompose $\mathbb{E}[I_t^{(2, \infty)} | G_n(p)] = \mathbb{P}(v_t \in C^{(2, R_t)}(G_n(p)) | G_n(p))$ into

$$\mathbb{P}(v_t \in C^{(2, R_t)}(G_n(p)), R_t \geq \ell | G_n(p)) + \mathbb{P}(v_t \in C^{(2, R_t)}(G_n(p)), R_t < \ell | G_n(p)). \quad (5.37)$$

For the first term on the right hand side in (5.37), we have

$$\mathbb{P}(v_t \in C^{(2, R_t)}(G_n(p)), R_t \geq \ell | G_n(p)) \leq \mathbb{P}(v_t \in C^{(2, \ell)}(G_n(p)) | G_n(p)). \quad (5.38)$$

By Lemma 5.2.1, $\mathbb{P}(v_t \in C^{(2, \ell)}(G_n(p)) | G_n(p)) \xrightarrow{p} \zeta^{(2, \ell)}(p)$, and since $\lim_{\ell \rightarrow \infty} \zeta^{(2, \ell)}(p) = \zeta^{(2, \infty)}(p)$, by choosing ℓ sufficiently large, we can guarantee that

$$\mathbb{P}(v_t \in C^{(2, R_t)}(G_n(p)), R_t \geq \ell | G_n(p)) \leq \zeta^{(2, \infty)}(p) + \frac{\epsilon}{8} \quad \text{whp.} \quad (5.39)$$

For the second term, note that

$$\mathbb{P}(v_t \in C^{(2,R_t)}(G_n(p)), R_t < \ell | G_n(p)) \leq \mathbb{P}(R_t < \ell | G_n(p)) \leq \mathbb{P}(|\mathcal{N}_\ell(v_t; G_n(p))| > K | G_n(p)). \quad (5.40)$$

By Lemma 5.2.2, for K sufficiently large, this quantity is at most $\epsilon/8$ whp. Thus,

$$\mathbb{E}[I_t^{(2,\infty)} | G_n(p)] \leq \zeta^{(2,\infty)}(p) + \frac{\epsilon}{4} \quad \text{whp.} \quad (5.41)$$

On the other hand, since $R_t \leq K$ and $C^{(2,R_t)}(G_n(p)) \supseteq C^{(2,K)}(G_n(p))$, we have

$$\mathbb{P}(v_t \in C^{(2,R_t)}(G_n(p)) | G_n(p)) \geq \mathbb{P}(v_t \in C^{(2,K)}(G_n(p)) | G_n(p)). \quad (5.42)$$

With analogous reasoning as before, we can guarantee

$$\mathbb{E}[I_t^{(2,\infty)} | G_n(p)] = \mathbb{P}(v_t \in C^{(2,R_t)}(G_n(p)) | G_n(p)) \geq \zeta^{(2,\infty)}(p) - \frac{\epsilon}{4} \quad \text{whp} \quad (5.43)$$

by choosing K sufficiently large. Combining (5.41) with (5.43) suggests $|\mathbb{E}[I_t^{(2,\infty)} | G_n(p)] - \zeta^{(2,\infty)}(p)| \leq \frac{\epsilon}{4}$ whp.

The rest is straightforward. Since $I^{(2,\infty)}$ is the average of T i.i.d. Bernoulli random variables $\{I_t^{(2,\infty)}\}_{t \in [T]}$ (conditional on $G_n(p)$), by the standard Hoeffding's inequality,

$$\mathbb{P}\left(|I^{(2,\infty)} - \mathbb{E}[I_t^{(2,\infty)} | G_n(p)]| \geq \frac{\epsilon}{4} \middle| G_n(p)\right) \leq \frac{\epsilon}{4} \quad (5.44)$$

for $T = O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})$, and marginalizing over G_n preserves the bound. Thus, we finally have $|I^{(2,\infty)} - \zeta^{(2,\infty)}(p)| \leq \epsilon/2$ with probability at least $1 - \epsilon/2$. Combining this with the in-probability convergence of $\frac{1}{n}|C_{\max}^{(2)}(G_n(p))|$ to $\zeta^{(2,\infty)}(p)$ concludes the proof. $\square$

5.7 Discussion: Beyond the 2-Core

Now that we have established the locality of the 2-core in structurally robust graphs using the two-step sprinkling technique, a natural question arises: Can this approach be generalized to study stronger connectivity notions, such as the k -core for $k \geq 3$? Unfortunately, even for expanders, the answer appears negative.

Theorem 5.7.1 (Diverging proportion of the k -core). *For $k \geq 3$, there exist a sequence of bounded-degree expanders $\{G_n\}_{n \in \mathbb{N}}$ on n vertices converging locally to a measure μ as $n \rightarrow \infty$ with diverging k -core proportions, i.e.,*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} |C^{(k)}(G_n)| \neq \limsup_{n \rightarrow \infty} \frac{1}{n} |C^{(k)}(G_n)|. \quad (5.45)$$

For $k \geq 3$, the k -core behaves very differently than the 2-core and does not exhibit locality. Theorem 5.7.1 follows directly from the observation that the removal of a single edge within an expander can trigger a cascading collapse of the k -core. To see this, consider a random k -regular graph with $k \geq 3$, and let G be its giant connected component. It is well known that G is of linear size and is an expander with high probability. By definition, $C^{(k)}(G) = G$. Now we remove an edge from G to obtain G^- , which has barely any impact on the expansion property. However, note that $C^{(k)}(G^-) = \emptyset$, yet G^- and G are so similar that asymptotically they are not even statistically distinguishable.

This simple example illustrates a critical point: the expansion property alone is insufficient to guarantee the locality of higher-order cores. The phenomenon underscores that, unlike the 2-core, the k -core for $k \geq 3$ is highly sensitive to small perturbations, challenging the application of locality arguments in this broader context. This observation also raises further questions of whether we can define higher-order analogues of expansion, or alternatively, what other connectivity notions can be estimated locally.

Chapter 6

Epidemics on Networks with Community Structures: Dynamics and Interventions

In this chapter, we turn to the dynamics of epidemics in populations with community structure. Using tools from stochastic processes and dynamical systems, we study how an epidemic evolves over time as it spreads both within and between communities. The goal is not only to characterize the progression of the outbreak, but also to understand how structural information, such as density of travel and contact, can guide effective intervention under uncertainty.

6.1 Introduction

The rapid spread of infectious diseases has become a global concern, drawing heightened public attention over the last few years in light of the COVID-19 pandemic, and the monkeypox and measles outbreaks. In response to such outbreaks, various non-pharmaceutical interventions, including social distancing and travel bans, have been proposed and implemented to curb transmission. While there have been several empirical studies on the efficacy of social distancing and travel bans in response to various outbreaks, there has been little rigorous quantitative analysis to date. Our work aims to put understanding of these interventions on a rigorous footing.

Social distancing has long been thought of as a tool for preventing the transmission of disease. In the 19th century, public health laws were formalized in response to the 1918 influenza pandemic. Cities that implemented social distancing earlier in the epidemic had lower transmission and mortality rates [46]. [132] found that during the COVID-19 pandemic, individuals living in retirement communities with greatest social distancing had a much lower risk of COVID-19 compared to individuals living in communities with poorer social distancing.

In contrast to social distancing, the efficacy of travel bans as a standalone measure is much more dubious. Mateus et al. [150] found no evidence that travel restrictions would contain influenza. Similarly, Chinazzi et al. [57] found that travel restrictions during COVID-19 only delayed the epidemic, on the order of days, in mainland China, that travel restrictions to and from mainland China only “modestly” affected the epidemic trajectory unless combined with another intervention.

In this chapter, we study a simple model of an SIR epidemic on a dynamic network that models two communities with travel between them, with travel modeled by two Poisson processes, with the rate of return much higher than the rate of travel, so that at any given time, visitors only make up a small proportion in each community. The epidemic spreads separately in each community (between all individual currently there, whether they are residents or guests), and only spreads from one community to another when an infected individual travels before they recover (whether the individual is a guest infected while not at home, or an individual infected at home before traveling). For simplicity, the contact network in each community is modeled by an Erdős–Rényi random graph, even though our results should hold for other network models for the communities, like stochastic block models or configuration models, at least as long the models are such that $R_0 < \infty$.

We start the infection in the first community and are interested in the trajectory of the epidemic in both communities as well as the efficacy of social distancing and travel bans in our setting. In the real world, one may imagine that the two communities are separate geographical regions with flights going between them. We study the sparse regime, where the average degree of the vertices is bounded and where the number of travelers is polynomially small in the number of vertices in the graph.

Our main results are two-fold. First, we characterize the trajectory of the epidemic, pinning down the times at which the infection reaches an ϵ fraction of vertices in the first community, the time at which the infection first passes from the first to the second community, and the time it reaches an ϵ fraction in the second community, proving in particular that it reaches the second community long before reaching an ϵ fraction in the first. Second, we show that social distancing can reduce the final size of the infection, whereas travel bans have negligible impact on the final size of the infection, providing a mathematical understanding of the observational results discussed above.

Related work. The SIR model is a canonical model of disease spread in closed, finite populations and was introduced in the seminal paper in 1927 by Kermack and McKendrick [119]. The initial model and its stochastic counterpart do not take into account underlying network structure induced by contacts in a population. However, the network structure of a population is known to greatly influence the trajectory of an epidemic. As such, there has been a long line of work studying the SIR epidemic on random graphs with bounded average degree, mirroring real-world networks, including the Erdős–Rényi random graph [159], the configuration model [44, 37, 66, 111], graphs with local household structure and tunable clustering [32, 103, 49], and the stochastic block model [47].

Another perspective is the meta-population view, in which a heterogeneous population is divided into large, homogeneous sub-populations, between which the infection can spread at varying rates. Such models go back to at least Hethcote [98], where the inter-group connections are assumed to be static. A little later, Rvachev and Longini Jr [186] proposed a model that explicitly incorporated the effect of air travel and used it to analyze the 1968/69 “Hong Kong” pandemic, using a deterministic, continuous variable model. This work has been extremely influential, leading to a large body of work modeling the influence of travel on various specific outbreaks, from influenza [144, 94, 95] to AIDS [84] and SARS [106].

There are several work studying the model introduced by Rvachev and Longini from a more theoretical point of view, see in particular the series of papers by Colizza et al. [62, 61, 59, 60]. More closely related to our work is the paper of Barthélemy et al. [38] in which the travel between communities is modeled by a stochastic, integer random variable described by a so-called Cox process. Based on a mean-field approximation for the dynamic within the communities and further approximations to study the Cox process describing travel, they then determine the asymptotic behavior of the first time an infected individual arrives in a second community, and in particular show that under the conditions considered in their paper, there is a non-zero probability that the epidemic never spreads to the second community, leading to a percolation type description for the probability of the epidemic jumping into other communities.

More similar to our work in spirit and in its aim for mathematical rigor is the recent work of Ball et al. [33] who study the SIR epidemics in populations composed of large, weakly connected subcommunities. Their setting resembles ours in that the description of the epidemic is fully stochastic, deriving laws of large numbers where needed instead of postulating them, but differs in that the mixing is fully homogeneous within each subpopulation and the connections between communities are much sparser, leading again to an percolation type description for the global spread, now rigorously derived from the underlying fully stochastic model. They characterize whether a global outbreak occurs, but do not examine temporal dynamics or control interventions. Furthermore, their model of inter community connections is static, rather than based on an explicit model of travel.

By contrast, the recent papers do provide mathematically rigorous treatment of dynamic networks [134, 50, 114, 30, 75, 58, 105]. Notably, Lashari, Trapman, and Sirl [134] model partnership dynamics; Britton, Deijfen, and Liljeros [50], Jiang et al. [114], Ball and Britton [30], and Durrett and Yao [75] study SI and SIR dynamics on Erdős–Rényi random graphs and configuration models where individuals drop or rewire connections to infectious neighbors; Cocomello and Ramanan [58] study networks where the infection and recovery rates vary dynamically; and Huang and Röllin [105] study the SIR epidemic on a dynamic inhomogeneous Erdős–Rényi random graph, where edges appear and disappear independently of each other. From a technical point of view, our work differs from this line of work in that we consider both group structure and a dynamic network driven by the movement of individuals, thus making the approaches in the papers above no longer directly applicable. Furthermore, we are not only interested in the time evolution of the epidemic, but also the policy implications and efficacy of interventions.

Our work also differs from previous empirical and theoretical work in that we incorporate “memory” into our travel model, allowing the concept of a home-community and different rates for travelers leaving their community and for those returning, thus giving a more realistic model for travel. Note that while we only present proofs for the simple setting of two communities each having an Erdős–Rényi internal structure, it is clear the model is well defined for more than two communities, as well as more realistic models for the internal community structure, like stochastic block models, degree corrected stochastic block models, or configuration models. We thus believe that in addition to our contribution of a rigorous understanding of the effectiveness of travel restrictions vs social distancing, our model also should advance more realistic modeling for empirical studies of the effect of travel on the spread of an epidemic.

6.2 An SIR Epidemic Model on Two Communities With Travel

Our aim is to understand how travel dynamics impact the spread of an SIR epidemic between two communities connected through travel. We model these dynamics using a dynamic graph $G(t)$, consisting of two densely connected communities of individuals (nodes in the graph), which are sparsely connected through travel. Edges in the graph appear and disappear as individuals travel between the two communities. Individuals are assigned a label representing their home community in order to track whether they are currently traveling or home. Individuals travel infrequently to the other community, staying there for a duration comparable to the infectious period (e.g., days or weeks). Within each community, infection spreads between its residents and the small fraction of visitors, who may bring the infection back to their home community if they do not recover before they return. Thus, both travel to a guest community and home can facilitate spread between the two communities. For simplicity, we model the contact network within each community by an Erdős–Rényi random graph although we believe our analyses would extend if other random graph models with more internal structure (e.g. stochastic block model or configuration model) were used within each community.

Contact Network. Let G be a graph on $2n$ vertices with vertices partitioned into two intrinsic *types*, V_1 and V_2 , each containing n vertices. A vertex’s type represents its home location and does not change over time. In addition, to model a dynamic population with travel, there is a notion of *community* representing the physical location of vertices. At any given time, each vertex resides in one of two physical communities, Q_1 or Q_2 , and vertices can travel between communities. We denote the type of node v by $\text{type}(v) \in \{1, 2\}$, and its physical community at time t by $q_v(t)$. We say a node v is at home at time t if $\text{type}(v) = q_v(t)$, and otherwise we say it is traveling. Let V_1^H denote the set of type 1 vertices in community

1 (i.e., those currently at home), and let V_1^T denote the set of type 1 vertices in community 2 (i.e., those currently traveling). Similarly, define V_2^H and V_2^T for type 2 vertices.

Contacts are represented as edges between individuals. Within each community, edges are formed independently with equal probability $p := c/n$. When any individual travels for the first time, new edges are formed with individuals in the other community with probability c'/n and these edges will remain. For simplicity of notation, we assume that $c = c'$. Preflipping the random coins governing these edges, we equivalently may connect all vertices with probability c/n , whether they have the same label or not, but only activate the edges between vertices of different types when one of them travels.

Travel Dynamics. We model travel states $q_v(t)$ as independent continuous-time Markov processes governing movement between communities. For type 1 and type 2 nodes, the transition matrix for travel is given by $\begin{pmatrix} -\rho_T & \rho_T \\ \rho_H & -\rho_H \end{pmatrix}$ and $\begin{pmatrix} -\rho_H & \rho_H \\ \rho_T & -\rho_T \end{pmatrix}$, respectively. Standard results for Markov processes implies that in the steady state, an individual is at travel with probability $p_T = \frac{\rho_T}{\rho_T + \rho_H}$ and at home with probability $p_H = 1 - p_T$. Throughout the chapter, we will assume that $n^{-1}\rho_H \ll \rho_T \ll \rho_H$ so that travel is rare, i.e., $p_T \ll 1$, but frequent enough to give a substantial risk of transmission between the two communities, i.e., $np_T \gg 1$. An edge between two individuals currently in the same (different) community is said to be active (inactive). We use $G(t)$ to denote the network of *currently active* edges.

SIR Dynamics. The SIR infection on $G(t)$ is a Markov Chain whose state space is given by the sets of susceptible, infected, and recovered vertices in each community at time t . The process starts with a single infected vertex $v_0 \in V_1^H$ chosen uniformly at random, with all other vertices susceptible. At time t , an infected vertex u can infect its neighbor v only if the edge (u, v) is active at time t . Infection events happen at rate β and infected vertices recover at rate γ .

Let $S_k(t), I_k(t), R_k(t)$ denote the number of susceptible, infected, and recovered vertices in V_k at time $t \geq 0$ for $k \in \{1, 2\}$. Initially there is one infected vertex, e.g. $I_1(0) = 1$. We assume that the initially infected vertex is not traveling. Since $S_k(t)$ is non-increasing in t and thus must admit a limit, we define $S_k(\infty) = \lim_{t \rightarrow \infty} S_k(t)$; $R_k(\infty)$ can be defined similarly. We will also denote the number of susceptible type- k vertices who at home or traveling by $S_k^H(t)$ and $S_k^T(t)$, respectively.

Key parameters of SIR Dynamics. The basic reproduction number, denoted as R_0 , is a fundamental parameter that quantifies the average number of secondary infections produced by a typical infectious individual in a fully susceptible population [69]. In our model, R_0 is given by

$$R_0 = \frac{c\beta}{\beta + \gamma} \quad (6.1)$$

in the limit as the rate of travel tends to zero. As we will see later, consistent with the literature, the threshold $R_0 > 1$ coincides with the possibility of a major outbreak in the population. Furthermore, the rate of exponential growth (per unit time) of the initial outbreak is closely related to R_0 up to a constant factor determining the time scale. In particular, this exponential growth is governed by the rate parameter λ , defined as

$$\lambda := c\beta - \beta - \gamma = (\beta + \gamma)(R_0 - 1). \quad (6.2)$$

Note that $\lambda > 0$ precisely when $R_0 > 1$, in which case the initial infection may grow rapidly and lead to a widespread outbreak; when $R_0 < 1$, the infection will die out over time without causing a major outbreak.

In the rest of the chapter, we will assume that R_0 is strictly above, since otherwise the setting becomes trivial.

Assumption 6.2.1. *The basic reproduction number $R_0 = \frac{c\beta}{\beta+\gamma} > 1$.*

Two other parameters of relevance are the asymptotic probability of a large outbreak, π , and its relative size, r_∞ in the case without travel, i.e., in a single community of n individuals with connection probability c/n . They are known to be given as the largest solutions in $[0, 1]$ of the implicit equations

$$1 - \pi = G(\pi) := \gamma \int_0^\infty \exp(-c(1 - e^{-\beta\ell})\pi - \gamma\ell) d\ell \quad (6.3)$$

and

$$1 - r_\infty = e^{-R_0 r_\infty}, \quad (6.4)$$

respectively. It is straightforward to check that in our model r_∞ is strictly positive under the assumption of $R_0 > 1$.

6.3 Results

To assess the impact of interventions such as travel bans and social distancing, we first characterize the spread of an epidemic across two communities in the absence of any control measures (Theorem 6.3.2). We then use this as a baseline to evaluate how different interventions alter this baseline trajectory, focusing on their impact in reducing the final outbreak size, decreasing the probability of large outbreaks, and delaying the onset of the epidemic. We show that travel bans are largely ineffective in limiting outbreak size, yet social distancing, even when implemented late, can substantially reduce the impact in the second (Theorems 6.3.6 and 6.3.7).

In the discussion that follows, it is helpful to formalize what we mean by a substantial outbreak. In particular, we use the following definition to distinguish between cases where the infection dies out quickly and those where it spreads to a significant fraction of the population:

Definition 6.3.1. *An outbreak is called **noticeable** (in a community or in the entire population) if it eventually infects at least ϵn individuals for some small $\epsilon \in (0, r_\infty)$.¹*

Heuristics. Intuitively, one may expect that a travel would be effective if implemented when the outbreak first becomes noticeable in community 1. However, we show that by that time, travel has already seeded a moderate number of infections in community 2—more than enough to sustain its own growth, even if all subsequent travel is halted. Further imported infections have only a secondary effect compared to the internal transmissions.

In contrast, social distancing, modeled as reducing the transmission rate or equivalently a reduction in R_0 , turns out to be an effective intervention. If implemented in community 2 when the outbreak is noticeable in community 1, a noticeable outbreak in community 2 will very likely be prevented. The intuition for this is as follows: at the time when the outbreak is noticeable in community 1, there are only a small number of infections in community 2. Social distancing reduces the rate of transmission among individuals in community 2 such that each newly infected (or imported) case will, on average, infect a bounded number of individuals before they all recover. Meanwhile, the epidemic in community 1 will continue to grow until it reaches herd immunity and later terminates. During this period, the number of imported cases in community 2 is of low order. Thus, the total size of the infection will remain sublinear and the outbreak will never become noticeable in community 2.

Figure 6.1 provides a schematic illustration of the epidemic trajectory in the absence of intervention. The spread across the two communities can be partitioned into four phases:

- **From time 0 to $\tau_{1 \rightarrow 2}$:** The infection starts to spread within community 1. It may die out early, but if it survives, it begins to grow exponentially.
- **From $\tau_{1 \rightarrow 2}$ to τ_1 :** Infections begin seeding community 2 via travel. Both communities experience exponential growth, though community 2 still has only a small number of infections.
- **From τ_1 to τ_2 :** The outbreak in community 1 becomes noticeable and soon passes the herd immunity threshold. Community 2 accumulates a polynomial number of infections. At this stage, a travel ban is no longer meaningful, as enough infections have already been seeded to sustain exponential growth in community 2.
- **From τ_2 to the end:** Community 2 experiences its own noticeable outbreak and eventually reaches herd immunity, regardless of travel policy. However, if social distancing is implemented in community 2 at time τ_1 , the outbreak will stay below noticeable size throughout.

¹The parameter ϵ should be thought of as a numerical constant (independent of n) that is a small fraction of r_∞ , even though mathematically all we need is the strict inequality $\epsilon < r_\infty$.




Figure 6.1: Schematic illustration of epidemic trajectory in the two communities (not to scale). τ_1 and τ_2 mark the major “waves” in the two communities: most infections occur within the narrow bands around them. $\tau_{1 \rightarrow 2}$ marks the initial infection of any type-2 individual. τ_{end} indicates the end of the epidemic, i.e., when the number of infectious individual first drops to zero.

Formal statements. We now formalize our heuristics developed above. To formally state our results, we define the following stopping times for the times at which the epidemic reaches an ϵ -fraction of individuals in community 1 and 2:

$$\tau_1(\epsilon) := \inf\{t > 0 : S_1(t) \leq (1 - \epsilon)n\}, \quad (6.5)$$

$$\tau_2(\epsilon) := \inf\{t > 0 : S_2(t) \leq (1 - \epsilon)n\}. \quad (6.6)$$

Also, define the stopping time for the first time that the infection arrives in community 2 by

$$\tau_{1 \rightarrow 2} := \inf\{t > 0 : I_2(t) > 0\}. \quad (6.7)$$

The following theorem characterizes the epidemic trajectory—typical timing and size of the outbreaks—in both communities in the absence of interventions². Throughout, we will

²Note that we did not define stopping times in terms of the size of the outbreak in a particular physical community, but instead considered the size of the outbreak counting people of a given type. Asymptotically, this makes no difference, since the epidemic dies out in time $O(\ln n)$, implying that during this time only a tiny $\tilde{O}(n^{-\alpha})$ fraction of individuals travel. In fact, our proofs show that all of our theorems remain valid if one defines the stopping times $\tau_1(\epsilon)$ and $\tau_2(\epsilon)$ in terms of $S_1^H + S_2^T$ and $S_2^H + S_1^T$ —or in terms of S_1^H and S_2^T .

assume that $0 < c, \gamma, \beta, \rho_H < \infty$ are fixed constants independent of n , while we scale down ρ_T with n .

Theorem 6.3.2 (SIR epidemic with no interventions). *Consider the SIR epidemic spreading on $G(t)$ with rate of travel given by $\rho_T = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. Then the following hold:*

- (a) *The probability of a large outbreak converges to the largest solution $\pi \in [0, 1]$ of the fixed point equation (6.3), i.e.,*

$$\mathbb{P}(R_1(\infty) + R_2(\infty) \geq \ln n) \rightarrow \pi \text{ as } n \rightarrow \infty.$$

- (b) *Conditioned on not having a large outbreak, the epidemic dies out in finite time, never reaches community 2 and only ever infects a finite number of individuals in community 1; With high probability, the stopping times $\tau_1(\epsilon)$ and $\tau_2(\epsilon)$ are infinite.*
- (c) *Conditioned on a large outbreak, $n^{-1}R_1(\infty)$ and $n^{-1}R_2(\infty)$ converge to the largest solution $r_\infty \in [0, 1]$ of the fixed point equation (6.4). In addition, conditioned on a large outbreak,*

$$\frac{\lambda}{\ln n}(\tau_{1 \rightarrow 2}, \tau_1(\epsilon), \tau_2(\epsilon)) \xrightarrow{P} (\alpha, 1, 1 + \alpha)$$

for $\epsilon \in (0, r_\infty)$ and $\tau_1(\epsilon), \tau_2(\epsilon) \xrightarrow{P} \infty$ for $\epsilon > r_\infty$; the epidemic dies out in time $O(\log n)$.

Remark 6.3.3. *Theorem 6.3.2 suggests that large outbreaks in the two communities are tightly coupled: a large outbreak in the first community will very likely lead to a large outbreak in the second, only with a delay that is logarithmic in the inverse travel rate. Within each community, once the infection begins to spread, the time it takes to reach macroscopic scale concentrates at $\lambda^{-1} \ln n$, and most infections occur in a narrow window around this time. This behavior closely mirrors what one would expect from SIR dynamics on a single (static) Erdős–Rényi network. In other words, aside from seeding the outbreak across communities, the spread within each community behaves almost as if the other did not exist.*

We now turn to evaluating the effects of intervention. Specifically, we consider two types of intervention—travel bans and social distancing—applied at the time when community 1 reaches a noticeable outbreak, and analyze their impact on the trajectory in community 2. In our model of SIR dynamics, these policies are formulated as follows.

Definition 6.3.4 (Travel ban). *A travel ban is implemented at time t_{tb} if for all $t \geq t_{\text{tb}}$ and all $v \in V_1 \cup V_2$, $q_v(t) = q_v(t_{\text{tb}})$.*

Definition 6.3.5 (Social distancing). *Social distancing is implemented in a community at time t_{sd} if for all $t \geq t_{\text{sd}}$ we assume that in that community β is replaced by β' such that $\lambda' = c\beta' - \beta' - \gamma < 0$.*

Theorem 6.3.6 (Effect of travel ban). *Consider the setting of Theorem 6.3.2. Suppose that at time $\tau_1(\epsilon)$ a travel ban is implemented. Then the same conclusions as in Theorem 6.3.2 hold. In particular, the asymptotic probability of a noticeable outbreak remains unchanged and, conditional on such an outbreak, its final size in each community still converges to r_∞ and $\tau_2(\epsilon)/\ln n \xrightarrow{p} 1 + \alpha$ for any $\epsilon < r_\infty$.*

Note that not only does the theorem state that the probability and size of a large outbreak in the second community is asymptotically the same whether a travel ban is implemented or not, but it also says that a delay of onset (in Community 2) due to travel distancing is of lower order, consistent with empirical observations.

Theorem 6.3.7 (Effect of social distancing). *Consider the setting of Theorem 6.3.2. Suppose that at time $\tau_1(\epsilon)$ social distancing is implemented. Then the final size in the second community is sublinear in the number of vertices:*

$$n^{-1}R_2(\infty) \xrightarrow{p} 0.$$

In particular, with probability 1 the infection size in the second community is $\tilde{O}(n^{1-\alpha})$ and the number of additional infections in the second community after time $\tau_1(\epsilon)$ is $o(n)$.

To understand the ideas behind the proof of our theorems, it is useful to first understand the dynamics of an epidemic in a single community.

6.4 Proof Overview

We will prove Theorems 6.3.2, 6.3.6 and 6.3.7 in Appendix E. Here we give a brief overview, focusing mainly on the heuristic why one should expect these theorems to hold. We start with Theorem 6.3.2.

To understand the early phases of an epidemic seeded in community 1, we note that as long as the epidemic is small, say of order at most $O(\ln n)$, the chance that one of the infected individuals travels after being infected is bounded by $\tilde{O}(n^{-\alpha})$. Thus, in the early stages, the epidemic evolves essentially as if there was just one community. In particular, the probability that it grows to size $\ln n$ is essential given by the solution π to (6.3), the probability of a large outbreak in a single community. Furthermore, in these early stages, the infections grows at the single community rate λ , showing that as long as $\lambda t \ll \ln n^\alpha$, the size in community 1 is much smaller than n^α , making it still unlikely that any of the infected individuals will travel.

But once the infection passes size n^α , it becomes likely that one of the infected individuals travels after being infected (or that the infection hits a traveler who then returns to their home community). In either case, these infected individuals serve as seeds in community 2, each having constant probability of contributing to $I_2(t)$. When the number of these attempts diverges, say $(\ln n)^2$, we will see whp infections inside the second community from these travelers, explaining why $\tau_{1 \rightarrow 2}$ behaves like $\frac{1}{\lambda} \ln n^\alpha$.

After this point, the infection continues to grow exponentially in community 1, but meanwhile it also starts to grow exponentially in community 2, in both cases at rate λ . One might expect travel in the opposite direction to create new seeds in community 1. When this happens, the epidemic in community 1 is already large, and the contribution from these new seeds is negligible compared to the overall growth in community 1, showing that it will take roughly time $\frac{1}{\lambda} \ln(\epsilon n) = \frac{1}{\lambda} \ln n(1 + o(1))$ to grow to size ϵn . Thus, we should expect $\tau_1(\epsilon)$ to be roughly equal to $\frac{1}{\lambda} \ln n$.

Between time $\tau_{1 \rightarrow 2}$ and time $\tau_1(\epsilon)$, the epidemic grows exponentially in both communities at rate λ , with the size in community 2 behind by a factor of $n^{-\alpha}$. Thus both seeds by travels (which happens at rate $\theta(n^{-\alpha})$) and the exponential growth of the original seeds contribute the same order of magnitude. Since this phase takes time of order $\frac{1-\alpha}{\lambda} \ln n$, one might expect slightly faster growth in community 2, possibly multiplying the size of the epidemic by a factor of order $\ln n$. However, once the epidemic in community 1 has reached size ϵn , it will reach herd immunity in constant time, and die out, at which point all which happens is exponential growth in community 2 at rate λ . Thus it should take between time roughly $\frac{1}{\lambda} \ln\left(\frac{\epsilon n}{\ln n}\right)$ and $\frac{1}{\lambda} \ln(\epsilon n)$ to grow the epidemic from size of order 1 to size ϵn in community 2. Since the difference between the two is of lower order, we expect $\tau_2(\epsilon)$ to be roughly $\tau_{1 \rightarrow 2} + \frac{1}{\lambda} \ln n$, explaining the scaling of $\tau_2(\epsilon)$.

Note that the above arguments also explain why travel restrictions do not help much once the epidemic is noticeable in community 1. At this point, the outbreak in community 1 is already of order $n^{1-\alpha}$ (possibly slightly larger, by a factor of $\ln n$), large enough to make the epidemic grow in community 2 until it has reached its final size. Travel restrictions might reduce the additional seeds after that time a little, but at this point, the growth in community 2 is itself exponential, while the community 1 cannot grow by more than a constant factor. The effect on delaying the onset coming from the reduction of seeds is of lower order than the leading order $\ln n$. To understand our final statement, the fact that the overall outbreak size is only influenced minimally, we note that as long as the seed of an infection is sublinear, the final size of an epidemic barely depends (to the first order) on the seed size, a fact that is straightforward to establish for the standard SIR process on $G(n, p)$.

To make the above considerations rigorous, we proceed in three steps to prove Theorem 6.3.2, which we now explain. After that, we also give a brief overview of the proofs of Theorems 6.3.6 and 6.3.7.

One community lower bounds. We first establish a lower bound on the size of the outbreak in community 1 by only considering individuals who never travel. More precisely, since we expect the total time of the epidemic to be of order $\ln n$, we exclude everyone who travels up to time $T = \ln^2 n$. Since the number of travelers up to this time is bounded in probability by $\tilde{O}(n^{1-\alpha})$, this gives a lower bound on the size of the epidemic in community 1 by that on $G(n', p)$ with $n' = n - \tilde{O}(n^{1-\alpha})$. Asymptotically, these bounds are indistinguishable from those for $G(n, p)$, giving us an upper bound on the probability of a noticeable outbreak, and the timings in such an outbreak, including $\tau_1(\epsilon)$ and the time it takes to grow to, say, size

$n^\alpha \ln n$ in community 1. We may now leverage these observations to show that, conditional on a large outbreak, $I_2(t) > 0$ whp whenever $t > \frac{1}{\lambda} \ln(n^\alpha \ln n)$, giving an upper bound on $\tau_{1 \rightarrow 2}$. Finally, by considering a different coupling to an SIR process on a static Erdős-Rényi network with multiple (super-constant number of) seed infections, we can bound the subsequent growth in community 2 from below, getting the desired upper bound on $\tau_2(\epsilon)$, and a lower bound on the final sizes in both communities.

2-type branching process upper bounds. Next, we use the well known fact that when we replace the dynamic infecting only susceptible vertices by a dynamic where we do not decrease the number of susceptible vertices, we can describe the resulting dynamic by a branching process that gives an upper bound on the number of infected and recovered vertices for all times. It is less clear whether this is still possible for dynamic graphs, since one might worry that speeding up earlier parts of the time evolution might lead to infections which are now possible or not possible anymore because individuals one would have infected may now be traveling, leading to an unexpected slowdown instead of a speedup. But it turns out that a careful coupling of (a) the exploration of the random graph over the two communities, (b) the travel dynamic, and (c) the dynamics of the epidemic does give a branching process which gives upper bounds valid for all times. The analysis of this branching process will lead us to analyze the rate matrix and a renewal process, which we then treat using perturbation theory around the dynamics without travel. While the branching process upper bound eventually becomes trivial as it exceeds the population size n , it turns out that for the purpose of determining the stopping times $\tau_{1 \rightarrow 2}$, $\tau_1(\epsilon)$ and $\tau_2(\epsilon)$, the relevant bounds are off by only an additive constant, which is good enough for us since these do not show up in our leading $\ln n$ behavior.

Herd immunity and the final size of the epidemic. While our branching process upper bounds are tight enough to establish the asymptotics of the stopping times $\tau_{1 \rightarrow 2}$, $\tau_1(\epsilon)$ and $\tau_2(\epsilon)$, they become non-informative as we move past these times and do not provide any upper bounds on the final size of the epidemic. To deal with this fact, we loosely decompose the total number of infections into two parts: (1) travelers who get infected, and (2) the spread contained within each of the two non-traveling populations (possibly started with a transmission from a traveler). The contribution of the first part is limited since the total number of travelers is of order $\tilde{O}(n^{1-\alpha})$ up to time $T = \ln^2 n$. With travelers isolated out, the spread within the non-traveling parts can be coupled to an SIR process within a single community, except that we may have multiple initial seeds due to contact with travelers. It is known that, for a single community, the infection tree(s) can be studied as connected components (the giant) in a percolation on the contact network, and the largest component has a fractional size tending to r_∞ while the rest are of finite average size. Since we expect at most $\tilde{O}(n^{1-\alpha})$ seeds up to time T (the number of travelers is of this order and each has, in expectation, a constant number of neighbors), the contribution from the non-traveling

population is dominated by the giant, which has the size we claimed, while the remaining components are of lower order in total.

To wrap up this argument, we only need to show that the infection dies out by time $T = \ln^2 n$. This relies on the intuition of herd immunity towards the end of the epidemic: By time $C \ln n < T$ for some sufficiently large constant $C < \infty$, the level of infection in both communities should be arbitrarily low (as our lower and upper bounds are converging at later times), and the remaining fraction of susceptible individuals is too small to sustain further growth. Using a subcritical branching process to bound the dynamics in the remaining time, we show that the number of actively infected individuals decreases exponentially and thus drops to zero in another $O(\ln n)$ time, ending the epidemic before T as claimed.

Travel bans. To prove Theorem 6.3.6, we first note that our lower bound never used the additional seeds after $\tau_2(\epsilon)$, in fact, we entirely ignored travels beyond the time $\frac{1}{\lambda}(\alpha + o(1)) \ln n$ (which is before $\tau_1(\epsilon)$ whp) in the proof anyway.

For the upper bound on the sizes of outbreak in both communities, the analysis is analogous to the proof in the presence of travel. A branching process with growth rate λ serves as an upper bound until the outbreak reaches a linear size; the upper bound of r_∞ on the final fractional sizes of outbreak in both communities up to time T admits the same proof as before; and finally a subcritical branching process again provides an upper bound on the elapsed time of the outbreak using the fact that we have reached herd immunity.

Social distancing. Suppose that we implement social distancing in community 2 at time $\tau_1(\epsilon)$. Since the infection in community 1 dies out in time $O(\ln n)$, only $\tilde{O}(n^{1-\alpha})$ vertices from community 1 travel to community 2 before the infection dies out in community 1. However, since effective R_0 drops below 1 in community 2 due to social distancing, any newly seeded infection tree dies out exponentially quickly and only infects a constant number of individuals in community 2. Thus, only $\tilde{O}(n^{1-\alpha})$ individuals in community 2 will ever be infected and the final size of the infection in community 2 is sublinear in the number of individuals. The analysis in this case is highly analogous to our discussion above on herd immunity and sub-critical branching processes.

6.5 Discussion

Our results demonstrate that common intervention policies such as travel bans and social distancing have varied levels of efficacy in limiting epidemic impact across populations with weak ties. While travel bans are often perceived as a natural first response to emerging outbreaks, our analysis shows that they have minimal effect on both the final size of the outbreak and its timing. Even when implemented early, travel restrictions tend not to substantially delay the onset of infection in the second community: the excess delay is of lower order compared to the natural delay (without interventions). This suggests that travel bans offer limited utility as a standalone measure.

In contrast, social distancing, modeled as a reduction in transmission rate, proves highly effective. When enacted in the second community at the moment the outbreak has reached moderate levels in the first, it dramatically reduces the final number of infections. Notably, the reduction is not just in constant factors: the number of infections becomes sublinear in the population size, demonstrating a strong containment effect even under persistent travel.

Generalizations and open directions. While our model is based on a simple Erdős–Rényi network, the insights likely extend to more general and realistic contact structures, such as configuration models, clustered networks, or household models, provided travel remains rare and mixing within communities is sufficient.

With more than two communities, we expect qualitatively similar dynamics, with delays in outbreak timing depending on the inter-community travel densities, and infectious spread following the “most densely connected” paths. For networks with stochastic block structure, our methods should carry over directly, though analyzing multiple types may require a more detailed branching process. Configuration models, while also expected to exhibit similar outcomes, would necessitate further technical work to account for heterogeneous degrees and correlation structures.

Figure 6.2 shows a simulation on a configuration model contact network with $n = 10^6$ individuals in each community: an underlying network with geometrically distributed degrees is sampled, and then the individuals are randomly assigned into the two types. We choose the average degree such that the basic reproduction number $R_0 = 2$. We normalize $\rho_H = 1$ and set $\rho_T = n^{-1/2}$. Rates of infection and recovery are 1.5 and 3, respectively. We assume that travel ban and social distancing are implemented when a 1% outbreak is detected in the first community,³ and social distancing has the effect of reducing the rate of infection to $\beta' = \beta/3$. The message remains clear: social distancing will successfully prevent a noticeable outbreak in the second community, while a travel ban is ineffective in this respect and only delays the peak by a minor amount.

One could also consider the regime of very sparse travel (for example, when the travel rate scales as $\Theta(1/n)$), which is similar in spirit to the setting studied in [31]. This regime differs slightly from our current model in that even without intervention, infections may not always reach the second community, due to the infrequent travel. We conjecture that a travel ban could modestly reduce the likelihood of cross-community transmission, though not eliminate it completely. Nevertheless, we expect that at a high level the core insights persist: when infections are seeded in a second community, interventions that reduce within-community transmission remain substantially more effective than travel bans.

³Although the 1% intervention threshold may appear conservative, it is expressed in terms of the *true* cumulative infections, which typically exceed the cases that are detected or reported and grow exponentially during the early phase. Even under ideal surveillance—where cases are perfectly identified and reported and a travel ban is imposed promptly—the ban’s effectiveness still hinges on whether the earliest infectious individuals have already traveled. By contrast, social-distancing measures remain robust: once introduced, they limit transmission regardless of when they start (except for the infections that have already occurred).




Figure 6.2: Simulation on two communities of 1 million individuals each, with an underlying network generated from a configuration model with geometrically distributed degrees (parameter chosen such the basic reproduction number $R_0 = 2$). Other parameters are: $\beta = 1.5$, $\gamma = 3$, $\rho_T = 10^{-3}$, and $\rho_H = 1$.

Chapter 7

Power of Two Choices in Token Systems

In this final chapter, we apply further ideas from stochastic analysis to the design of decentralized exchange systems. In particular, we study token-based coordination mechanisms that enable agents, such as hospitals participating in a national kidney exchange, to self-organize without requiring centralized matching or monetary transfers. Using Lyapunov techniques and fluid approximations, we show that a family of extremely simple provider selection rules can be remarkably effective at maintaining stable and sustainable system operation, even with thin exogenous supplies.

7.1 Introduction

Token systems have been introduced as a market solution to economies in which monetary transfers are undesirable or repugnant. Examples include trading favors in babysitting co-operatives [195], exchanging resources in peer-to-peer systems [204], and distributing food to food banks [175]. The use of such artificial currency is intended to sustain cooperation, alleviate free riding, and increase efficiency. The incentive for agents to provide service (or resources) is to earn tokens and the ability to spend them in future exchanges. This ability relies on the liquidity of agents' availability to provide service upon request. In this chapter, we study the behavior of token systems in thin marketplaces, where demand exceeds supply, and agents' availability for transactions is sparse.

One motivation for this study arises from kidney exchange. Many patients in need for a kidney transplant have a willing but incompatible living donor, which has led to the emergence of platforms that arrange swaps in order to find a compatible donor [183]. Efficiency of these platforms relies on the thickness of the pool, as a large fraction of these pools include pairs that are hard-to-match [23]. However, free riding is a common behavior; in the US, hospitals often submit to national platforms the pairs they cannot match internally [9],¹

¹Agarwal et al. [9] document that more than half of the exchanges in the US being arranged within

and similar behavior was observed by countries in international collaborations.² To alleviate free riding, token systems based on hospitals' contribution to the pool have been adopted in an ad-hoc manner, most notably by the National Kidney Registry in the US.³ A major concern is whether such systems can reward back contributing hospitals successfully, given that liquidity is low due to the sparsity of the kidney exchange pool. In this work, we ask if a token system can be effective in such sparse marketplaces and sustain cooperation between agents.

To address this question, we study a stylized model in which agents request and provide service over time, and tokens are used as artificial currency to pay for service provision. We consider an infinite horizon model with finitely many agents, where each agent has initially 0 tokens, and agents are allowed to have negative number of tokens. At each time period, one randomly chosen agent requests a service, and a random subset of agents of size d become available to provide the service. One of these agents is selected to provide the requested service, and the service requester pays one token to the service provider. The number of tokens each agent has represents the difference between the amount of service provisions and service requests by the agent. As balancing the number of tokens across agents is key to sustain cooperation, the agent with the fewest tokens among the available agents is selected to provide the service.

We are interested in the case in which d is a small constant; in the context of kidney exchange, this aims to capture few match opportunities for a given patient-donor pair. Our model is based on Johnson, Simchi-Levi, and Sun [115] and Kash, Friedman, and Halpern [117]. Their models, however, assume that at each time period, a constant fraction of agents are available to provide service, i.e., the service availability is not minimal. Intuitively, the larger the number of agents who are available to provide service, the easier it is to balance the amount of service provisions among agents.

Cooperation is important for token systems. For example, the Capitol Hill Babysitting Co-op, which aimed to exchange babysitting hours between families has crashed, since tokens' values depreciated [195]. So in a healthy market, agents should not accumulate too many tokens (which can lead to unraveling), or accumulate a large debt (which leads to free riding).

Motivated by these potential frictions, we analyze the token distribution in our model, and identify conditions under which the token system is *stable*, in the sense that the Markov chain describing the process admits a stationary probability distribution (the formal definition is given in Section 7.2). Informally, stability is described by two desired conditions. The first condition is a uniform boundedness condition; the number of tokens each agent has does not deviate much from its initial state, with high probability. The second condition is a fairness condition; the number of tokens each agent has oscillates in such a way that agents return to their initial endowment in finite expected time. In other words, the first condition

hospitals.

²See [22] for discussion about collaborations and mergers.

³Numerous platforms count the number of altruistic donors (who have no intended donor) and end altruistic donor chains with patients in these hospitals. Some platforms attempt to equalize the number of donors and patients from each "player" matched in each exchange [42, 155].

ensures that agents will not accumulate or lose too many tokens. Thus, agents will not lose their incentive to cooperate. The second condition implies that the market clears in finite expected time, which alleviates free riding and balances the number of requests and provisions for each agent continuously over time. In the context of the classical model of mean-variance preferences (e.g., see Markowitz [149] and Tobin [199]), our stability conditions offer an attractive solution for strategic agents, where stable systems reduce agents' risk (for further discussion on strategic considerations, see Section 7.5).

Overview of results. In the baseline case ($d = 1$), only one agent is available to provide service. Thus at each time period, service requester and provider are chosen independently to exchange service for tokens. In this case the system is unstable, since the number of tokens each agent has behaves like a divergent or null recurrent random walk.

When $d > 1$, the service provider who has the minimal number of tokens is chosen among all available agents. In our main finding, we show that already when $d = 2$, the system is stable. We show that the long-run probability of agents having more than M tokens (or less than $-M$ tokens) is at most $O(1/M)$. We further show, using Lyapunov arguments, that this probability is bounded by $O(a^M)$ for any $a > 2/3$; here $O(\cdot)$ includes a constant that may depend on n . We further study stability in a large market, and show that as n grows large, this probability is bounded by $(1/2)^M$. The large market analysis offers intuitive insights into the token dynamics.

Techniques. This chapter is inspired and borrows from the literature on the power of two choices (see, e.g., Mitzenmacher [156], Vvedenskaya, Dobrushin, and Karpelevich [205], and Azar et al. [28]). The main finding of this literature is that in load balancing problems, minimal choice can significantly reduce congestion both in dynamic and static settings. In our model, it is simple to analyze the system directly using a birth-death process when $n = 2$. For $n \geq 3$, the system does not seem amenable to a complete analytical solution, but softer techniques allow us to show that it is stable.

For finite n , we use the Lyapunov method to study a suitable exponential Lyapunov function's negative drift to establish tail bounds on the number of tokens agents have in the long-run. We further analyze the token distribution, and we explicitly characterize the tail bound as n grows large using Kurtz's theorem on *density dependent Markov chains* [131]; the same techniques are used in Mitzenmacher [156] and Vvedenskaya, Dobrushin, and Karpelevich [205] to analyze various load balancing problems including the supermarket model. While there are subtle differences between our model and the load balancing problems in the literature (as we explain below), our work can be viewed as another application of the power of two choices paradigm. Kurtz's theorem provides conditions under which a stochastic process can be approximated by a deterministic process in the limit. This allows us to show that the token distribution is balanced, and the number of tokens each agent has is unlikely to be far from its initial state.

Related literature. Numerous papers study exchange economies using tokens or models for exchanging favors [157, 87, 97, 2, 118, 117, 115]. This literature is concerned with whether cooperation can be sustained in equilibrium, and whether efficiency can be achieved. The main finding of this literature is that if agents are sufficiently patient, token mechanisms may lead to efficient outcomes. Closely related are Johnson, Simchi-Levi, and Sun [115], Kash, Friedman, and Halpern [117] and Bo, Dawande, and Janakiraman [43]. These papers study token systems as a strategic game in an infinite horizon with discounting, but with similar dynamics. The key difference is that these papers assume that either all or a constant fraction of agents are available to provide service at each time period. Kash, Friedman, and Halpern [117] study a model in which the service provider is chosen independently from the token distribution, and show the existence of an equilibrium, in which agents provide service when their tokens is below some threshold. Johnson, Simchi-Levi, and Sun [115] study the same model and show that under the minimum token selection rule, agents always provide service in equilibrium when punishments are feasible. Bo, Dawande, and Janakiraman [43] extend their findings without using punishments. Our work contributes to this literature, by studying whether stability can be achieved with low liquidity (i.e., low availability of service).

Also related is the literature on the power of two choices in load balancing problems [27, 28, 156, 205]. In this classic problem, n balls are sequentially thrown into n bins. The key finding is that if two bins are selected randomly and the ball is thrown to the bin with the lower load, then the fullest bin has exponentially fewer balls than if only one bin is chosen randomly in the throwing process. Mitzenmacher [156] and Vvedenskaya, Dobrushin, and Karpelevich [205] find a similar result for the supermarket model, which is a dynamic queueing system, where the longest queue is much longer if customers choose randomly among all queues rather than choosing intelligently between two random queues. The state space of the process we are interested in can be obtained by truncating the state space of the supermarket model. We describe in detail the difference between our model and the supermarket model in Remark 7.4.7 in Section 7.4.3. In short, using a queueing language, Mitzenmacher [156] focuses on the effect of the parameter d on the expected time customers spend in the system and the length of the longest queue in the long-run. We are interested, however, in the overall distribution of all queue-lengths, and provide a more detailed characterization for the system, which is needed for our analysis. In fact, Azar et al. [28] also study an infinite horizon process in which at each time period, one random ball is removed from the bins, and one ball arrives which is assigned to one of the two random bins intelligently. Our stochastic process has subtle differences (using this language, instead of removing one random ball, we first pick a random bin, and then remove a ball), making the machinery of Azar et al. [28] inapplicable.

Notation. We use $\mathbb{Z}_{\geq 0}$ and $\mathbb{Z}_+$ to denote the set of non-negative integers and the set of strictly positive integers, respectively. We use $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_+$ to denote the set of non-negative real numbers and the set of strictly positive real numbers, respectively. We write $\mathbb{E}_\pi[\cdot]$ and

$\mathbb{P}_\pi(\cdot)$ to write the expectation and the probability with respect to a given distribution π , respectively.

7.2 Model

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

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

We refer the tuple (n, P, Q, d) as the token system. We will analyze the behavior of the token system under a natural matching policy called the *minimum token selection rule* (see, e.g., Johnson, Simchi-Levi, and Sun [115]). This policy, at each time period, selects the available provider with the lowest number of tokens as the *service provider* (ties are broken by choosing uniformly at random). At each time t , if agent i is the service requester and agent j is the service provider, i pays one token to j . Note that an agent can provide service to herself.⁴ In this case, $s^{t+1} = s^t$. Otherwise, $s_i^{t+1} = s_i^t - 1$, $s_j^{t+1} = s_j^t + 1$, and $s_k^{t+1} = s_k^t$ for all $k \in \mathcal{A} \setminus \{i, j\}$. In either case, we have $\sum_{i \in \mathcal{A}} s_i^t = 0$ for all $t \geq 0$. The case $d = 1$ is the degenerate case, where the system simply selects one service requester and one service provider independently at random.

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

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

- **(C1)** There is a uniformly small probability that the number of tokens owned or owed by any agent is large. Formally, there is a function $f: \mathbb{Z}_{\geq 0} \rightarrow [0, 1]$ such that $\lim_{M \rightarrow \infty} f(M) = 0$, and for all times t large enough and all agents $i \in \mathcal{A}$ it holds that $\mathbb{P}(|s_i^t| > M) < f(M)$.

⁴This technical assumption is useful for the proofs. Moreover, it is motivated by the kidney exchange setting, where patient-donor pairs from the same hospital can be matched internally in a centralized setting when hospitals merge their patient-donor pools. Our results hold with minor modifications if an agent cannot serve herself.

- **(C2)** The expected time for the token system to clear is finite. Formally, let T_0 be the first time the system returns to 0, i.e., $T_0 = \min\{t > 0 : s^t = 0\}$. Then $\mathbb{E}(T_0)$ is finite. Note that by the Markov property, this is also the expected time to return to 0 after a later visit to 0.

While we do not explicitly incorporate strategic considerations in our model, we view stability as a necessary condition to sustain cooperation in an appropriately defined strategic game.⁵ Since the chain is irreducible and aperiodic, stability implies that the stationary distribution is unique, and that over time the distribution of s^t will converge to it. Conversely, if there is no stationary distribution, the distribution of s^t will become more and more “spread out” as t increases, with some agents either owning or owing a large number of tokens. We thus interpret stability as a necessary condition for the prevention of unraveling and free riding in a token system. In Section 7.5, we offer a microeconomic justification for stability via mean-variance preferences, which is a simple model of risk aversion used in economics.

The key novel feature of our model is that the service availability density d is small, rather than being a constant fraction of n . The larger the service availability density d is, the easier it is to achieve stability, as weakly more agents can provide service at any time period, which provides more flexibility to balance service provisions among agents. Therefore unless stated otherwise, we focus on the case when the service availability is minimal, i.e., $d = 2$.

7.2.1 The Case $d = 1$

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

An immediate corollary is that the token system is also not stable under the *random tie-breaking selection rule*, where the rule selects the service provider, even if all providers are available ($d = n$), uniformly at random (see Corollary F.1.2 and all missing details in Appendix F.1).

7.3 Two Agents

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

⁵For example, in order to prove the existence of an equilibrium, where all agents play a threshold strategy, [117] first determine the token distribution of agents in the long-run.

Proposition 7.3.1. *For $d \geq 2$, the token system with 2 agents is stable if and only if $q_1^d < p_1$, $q_2^d < p_2$. In this case, let π be the steady-state distribution of the Markov chain $(s^t : t \geq 0)$. Then for all $M \in \mathbb{Z}_+$, we have*

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

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

The proof is straightforward and given in Appendix F.2.⁶ One implication of Proposition 7.3.1 is that the probability of owning or owing a large number of tokens decays exponentially, i.e., $f(M) = O(a^M)$, where the constant $a \in (0, 1)$ can be found following Remark F.2.1 in Appendix F.2.

Proposition 7.3.1 identifies the level of asymmetry that can be tolerated between service request and service provision (within and across agents) rates. Moving forward, we focus on the symmetric case. Note that for symmetric agents and $d \geq 2$, Proposition 7.3.1 implies that the system is stable. When furthermore $d = 2$, we get that

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

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

7.4 The General Case

We analyze here the symmetric case with any number of agents.⁷ The results for the general case can be summarized as follows:

1. In Section 7.4.1, we show that stability holds for all $n \geq 2$ and $d \geq 2$. We further show that (C1) holds with $f(M) = 5/M$ for all $n \geq 2$ (Theorem 7.4.1). In particular, the probability that a given agent has a large number of tokens decays to zero in a rate that does not depend on the number of agents.

⁶We note that similar results hold even with less amount of service availability in the following sense. Suppose that at each time period, at most 2 agents are available to provide service with probability $0 < \beta < 1$, and only one agent is available to provide service with probability $1 - \beta$, independently. The analysis of this system is very similar and given in Appendix F.2. This suggests that even with having a weaker power of tie-breaking, two agents can still trade favors while keeping the token system stable.

⁷See Ashlagi et al. [24] for more discussion on the asymmetric case.

2. In Section 7.4.2, we show that the tail bound in (C1) can be strengthened with $f(M) = O((\frac{2}{3} + \delta)^M)$ for any $\delta > 0$, where $O(\cdot)$ hides an implicit constant that may depend on n (Theorem 7.4.2). The proof relies on Lyapunov arguments and is given in Appendix F.4.
3. In Section 7.4.3, we show that as n grows large, (C1) holds with $f(M) = (1/2)^M$ (Theorem 7.4.4). The proof organization for Theorem 7.4.4 is as follows. We first present *density dependent Markov chains* and Kurtz's theorem. We then model our system as a density dependent Markov chain, and use Kurtz's theorem to characterize our system in the limit via a system of ordinary differential equations. Finally, we study the solution of this system to prove Theorem 7.4.4.
4. We conjecture that (C1) holds for any $n \geq 2$ with $f(M) = a^M$, where $a \in [1/3, 1/2]$ (Conjecture 7.4.9).

7.4.1 Stability

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

Theorem 7.4.1. *The token system is stable for any $d \geq 2$ when the agents are symmetric. Furthermore, (C1) holds with $f(M) = 5/M$.*

The proof of Theorem 7.4.1 is deferred to Appendix F.3. By analyzing the expected one-step transition difference of the potential function $\sum_{i=1}^n \mathbb{E}[(s_i^t)^2]$, we show that the Markov chain $(s^t : t \geq 0)$ has a stationary distribution, and $\mathbb{E}[|s_i^t|] \leq 5$ for all $i \in \mathcal{A}$ and for all t large enough by exploiting a recursion on the expected one-step transition difference of the potential function. Theorem 7.4.1 then follows from Markov's inequality.

7.4.2 Exponential Decay in Tail Mass

We next offer a refined result, showing that under symmetry conditions, the tail bound on token profiles decays exponentially. This improves upon the polynomial bound previously established in Theorem 7.4.1.

Theorem 7.4.2. *For the token system with $d \geq 2$ and symmetric agents, (C1) holds with $f(M) = O(a^M)$ for any $a > 2/3$, where $O(\cdot)$ hides implicit constant factors possibly depending on n .*

Unlike in Theorem 7.4.1, the constant in the bound is implicit and may depend on the size n of the system. Since we are unable to explicitly characterize the implicit constant, it is not clear under which regimes the exponential bound dominates the polynomial bound

in Theorem 7.4.1 and vice versa. The proof of Theorem 7.4.2 requires the following key observation. For notational convenience, we write $Z_\lambda^t := \sum_{i=1}^n e^{\lambda s_i^t}$ and $E_\lambda^t := \mathbb{E}[Z_\lambda^t]$.

Lemma 7.4.3. *For $\lambda > 0$, if both $\limsup_{t \rightarrow \infty} E_\lambda^t < \infty$ and $\limsup_{t \rightarrow \infty} E_{-\lambda}^t < \infty$, then (C1) holds with $f(M) = O(e^{-\lambda M})$.*

Proof. Note that whenever $|s_j^t| > M$ for some $j \in \mathcal{A}$, we must have $Z_\lambda^t + Z_{-\lambda}^t = \sum_{i=1}^n e^{\lambda |s_i^t|} + e^{-\lambda |s_i^t|} > e^{\lambda M}$. By Markov's inequality,

$$\mathbb{P}(|s_i^t| > M) \leq \mathbb{P}\left(\sum_{i=1}^n e^{\lambda |s_i^t|} + e^{-\lambda |s_i^t|} > e^{\lambda M}\right) \leq e^{-\lambda M} (E_\lambda^t + E_{-\lambda}^t).$$

The quantity $E_\lambda^t + E_{-\lambda}^t$ is uniformly bounded for all t sufficiently large, finishing our proof. $\square$

Lemma 7.4.3 prepares the ground for us to apply the Lyapunov method with E_λ^t as a family of potential functions; it will suffice to establish the asymptotic finiteness of E_λ^t for λ in the interval $(-\log 1.5, \log 1.5)$, where $\log$ is the natural logarithm. Our proof will proceed in two steps. First, we show that $\limsup_{t \rightarrow \infty} E_\lambda^t < \infty$ for λ sufficiently close to zero (possibly depending on n). Then we remove this dependency and widen the range of λ to cover the desired interval using a fixed point argument, exploiting the log-convexity of the mapping $\lambda \mapsto E_\lambda^t$. The proof of Theorem 7.4.2 is deferred to Appendix F.4.

7.4.3 Stability in the Large

In this section, we study the large market setting to offer intuition regarding the dynamics of the market by focusing on the densities of agents with a given number of tokens. What we establish next suggests that stability is “well-behaved” in the large; as n grows large, (C1) holds with $f(M) = (\frac{1}{2})^M$.

Let π be the steady-state distribution of the Markov chain $(s^t : t \geq 0)$ granted by Theorem 7.4.1. For any $M \in \mathbb{Z}_+$, define $p_{n,M} := \mathbb{P}_\pi(|s_1^t| \leq M)$ when there are n symmetric agents.

Theorem 7.4.4. $\lim_{n \rightarrow \infty} p_{n,M} \geq 1 - (1/2)^M$ for all $M \in \mathbb{Z}_+$.

We will use Kurtz's theorem on *density dependent Markov chains*, which allows us to analyze the stochastic process as a deterministic process in the limit. This analysis helps us to characterize the steady-state distribution of the system as the number of agents grows large and thus, proving Theorem 7.4.4.

Density dependent Markov chains. We begin with the definition of density dependent Markov chains, which is given in Kurtz [131] for finite dimensional systems and extended by Mitzenmacher [156] to countably infinite dimensional systems. Let $\mathbb{Z}^*$ be either $\mathbb{Z}^m$ for some finite dimension m , or $\mathbb{Z}^{\mathbb{N}}$, and similarly define $\mathbb{R}^*$. Let $L \subseteq \mathbb{Z}^*$ be the set of possible non-zero transitions of the system. For each $\vec{l} \in L$, define a nonnegative function $\beta_{\vec{l}}: \mathbb{R}^* \rightarrow [0, 1]$.

Definition 7.4.5. A sequence (indexed by n) of continuous time Markov chains $(X_n(t) : t \geq 0)$ on the state spaces $S_n = \{\vec{k}/n : \vec{k} \in \mathbb{Z}^*\}$ is a density dependent Markov chain if there exists a $\beta_{\vec{l}} : \mathbb{R}^* \rightarrow [0, 1]$ such that for all n the transition rate of X_n is given by $q_{x,y}^{(n)} = n\beta_{n(y-x)}(x)$, $x, y \in S_n$.

In Definition 7.4.5, the index n can be interpreted as the total population or volume of the system, and the components of $\vec{k}/n$ can be interpreted as the densities of different types present in the system. The $\beta_{\vec{l}}(x)$ can be interpreted as the probability of transition $\vec{l}$ from $x \in S_n$ to $y \in S_n$, where $nx + \vec{l} = ny$. Given a density dependent Markov chain X_n with transition rates $q_{x,y}^{(n)} = q_{\vec{k}, \vec{k}+\vec{l}}^{(n)} = n\beta_{\vec{l}}(\vec{k}/n)$, define $F(x) = \sum_{\vec{l} \in L} \vec{l}\beta_{\vec{l}}(x)$. The following theorem is key in our analysis:

Theorem 7.4.6 (Kurtz's theorem [131, 156]). Suppose we have a density dependent Markov chain X_n (of possibly countably infinite dimension) satisfying the Lipschitz condition $|F(x) - F(y)| \leq M|x - y|$ for some constant M . Further suppose $\lim_{n \rightarrow \infty} X_n(0) = x_0$, and let X be the deterministic process:

$$X(t) = x_0 + \int_0^t F(X(u))du, \quad t \geq 0. \quad (7.1)$$

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

$$\sum_{\vec{l} \in L} |\vec{l}| \sup_{x \in K} \beta_{\vec{l}}(x) < \infty. \quad (7.2)$$

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

The Lipschitz condition ensures the uniqueness of the solution for the differential equation $\dot{X} = F(X)$, which follows by taking the derivative of (7.1) with respect to t .⁸ Condition (7.2) ensures that the jump rate is bounded in the process. Kurtz's theorem implies that as $n \rightarrow \infty$, the behavior of a density dependent Markov chain can be characterized by the deterministic process given in (7.1), where the convergence holds on a finite time interval $[0, T]$ for an arbitrary T . We next model and study our system as a density dependent Markov chain, which we refer to as the *finite model*.

The finite and infinite models. Let us model the system with n symmetric agents as a density dependent Markov chain and denote it by $(X_n(t), t \geq 0)$. Note that in Definition 7.4.5, the $\beta_{\vec{l}}$'s are independent of n , and the transition rates are linear in n . In order to fit our system to this definition, we assume that each agent has an exponential clock with rate 1. The ticking of agent i 's clock corresponds to a service request by i , and the service

⁸For example, see Lemma 4.1.6 in [7].

provider is selected immediately using the minimum token selection rule. Note that because of the memoryless property and the continuity of the distribution that governs the clocks, agents request service uniformly, and exactly one agent requests service at a time. As a slight abuse of notation, let $n_i(t)$ be the number of agents with i tokens at time t , $m_i(t)$ be the number of agents with at least i tokens at time t , and $z_i(t) := m_i(t)/n$ be the fraction of agents with at least i tokens at time t . Let $\vec{z}(t) = (\dots, z_{-2}(t), z_{-1}(t), z_0(t), z_1(t), z_2(t), \dots)$, and we drop the time index t when the meaning is clear. We represent the state of X_n by $\vec{z} = \vec{k}/n \in \mathbb{Z}^{\mathbb{N}}/n$. Let us call this process the *finite model*. Note that the initial state of X_n is $\vec{z}(0) = (\dots, 1, 1, 1, 0, 0, \dots)$, where $z_i(0) = 1$ for all $i \leq 0$, and $z_i(0) = 0$ for all $i \geq 1$.

Next, we describe the transition probabilities $\beta_{\vec{l}}$'s. The set of possible non-zero transitions from $\vec{k} = n\vec{z}$ is $L = \{e_{ij} : i, j \in \mathbb{Z}, i \neq j\}$, where e_{ij} is an infinite dimensional vector of all zeros except the i 'th index (which corresponds to the index of z_i) is -1 and the j 'th index (which corresponds to the index of z_j) is 1 . Note that after transition e_{ij} occurs, nz_i decreases by 1 and nz_j increases by 1 simultaneously. Hence, the transition e_{ij} corresponds to the event when an agent with i many tokens requests service and an agent with $j - 1$ many tokens provides service. Since the probability that the service requester has i many tokens is $z_i - z_{i+1}$, and the probability that the service provider has j many tokens is $z_j^d - z_{j+1}^d$, we have $\beta_{e_{ij}}(\vec{z}) = (z_i - z_{i+1})(z_{j-1}^d - z_j^d)$.⁹ Denote the *infinite model* by X , which is the limit of the finite model X_n , i.e., $X = \lim_{n \rightarrow \infty} X_n$. Since X is characterized by the deterministic process (7.1), we need to analyze the components of $F(x)$. Note that the i 'th component of $F(x) = \sum_{\vec{l} \in L} \vec{l} \beta_{\vec{l}}(x)$ (which corresponds to z_i) is $\sum_{j \in \mathbb{A} \setminus \{i\}} (z_j - z_{j+1})(z_{i-1}^d - z_i^d) - \sum_{j \in \mathbb{A} \setminus \{i\}} (z_i - z_{i+1})(z_{j-1}^d - z_j^d)$, which simplifies to

$$(1 - z_i + z_{i+1})(z_{i-1}^d - z_i^d) - (1 - z_{i-1}^d + z_i^d)(z_i - z_{i+1}) = (z_{i-1}^d - z_i^d) - (z_i - z_{i+1}). \quad (7.3)$$

Remark 7.4.7. *In the well-known supermarket model, customers arrive according to a Poisson process with rate λn , $\lambda < 1$, where there are n servers that serve according to the FIFO (first in, first out) rule. An arriving customer considers only a constant number (d) of servers independently and uniformly at random from the n servers with replacement, and she joins the queue that contains fewest customers (any ties are broken arbitrarily). The service time for each customer is distributed exponentially with rate 1. Using this language, our model corresponds to this dynamic system, where the total number of customers is fixed throughout the process, which restricts the state space of the underlying process drastically. Our system belongs to the family of dynamic systems referred as the closed model in [156]. Our technical contributions here are twofold. First, the main focus of this literature is on the expected longest-queue length and its dependence on the service availability parameter (d), whereas we refine this goal and further study tail-bounds on queue-lengths. Second, as stated in [156], the steady-state probability that a certain queue has more than M customers can be found by solving certain constraints that are derived from Theorem 7.4.6, and in general,*

⁹The probability that all available agents have at least j many tokens is z_j^d , and we subtract the probability that all available agents have at least $j + 1$ many tokens, which is z_{j+1}^d .

solution to these constraints does not appear to have a closed form. As we argue next, we will be able to bound the steady-state probabilities without obtaining a closed form solution.

Towards the proof of Theorem 7.4.4. Now that we have represented our system using the finite model, we are ready to prove Theorem 7.4.4. The proof is organized as follows. We first show that the conditions of Theorem 7.4.6 hold. Then using Theorem 7.4.6, we obtain the system of ordinary differential equations that characterize the infinite model. This characterization lets us represent the probability of interest in (C1) as n grows large using π_0 (the fraction of agents that have at least 0 tokens in the long-run). Finally, we find lower and upper bounds for π_0 to conclude.

Condition (7.2) is clearly satisfied since the magnitude of any jump is bounded, and the jump rate is bounded above by 1 for any state. We also show in Appendix F.5 that the Lipschitz condition of Theorem 7.4.6 holds with $M = 2 + 2d$. By differentiating (7.1) with respect to t and using (7.3), we get the following system of ordinary differential equations that characterizes the infinite model:

$$\frac{dz_i}{dt} = (z_{i-1}^d - z_i^d) - (z_i - z_{i+1}) \text{ for all } i \in \mathbb{Z}. \quad (7.4)$$

Intuitively, (7.4) can be interpreted as follows. Let us consider the expected change in m_i (the number of agents with at least i many tokens) over a small time interval dt . First note that a transition occurs whenever one of the agent's exponential clock ticks, which happens with rate ndt . Under such transition, m_i increases by 1 if an agent with $i - 1$ many tokens is selected as the service provider, which happens with probability $z_{i-1}^d - z_i^d$. m_i decreases by 1 if an agent with i many tokens is selected as the service requester, which happens with probability $z_i - z_{i+1}$. Hence, the expected increase in m_i is $(z_{i-1}^d - z_i^d)ndt$, and the expected decrease in m_i is $(z_i - z_{i+1})ndt$, which gives $dm_i = (z_{i-1}^d - z_i^d)ndt - (z_i - z_{i+1})ndt$, and since $m_i/n = z_i$, dividing both sides by ndt gives (7.4).

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

$$\sum_{i \in \mathbb{Z}} i \cdot \frac{n_i}{n} = \sum_{i \geq 1} i \cdot \frac{n_i}{n} + \sum_{i \leq 0} i \cdot \frac{n_i}{n} = \sum_{i \geq 1} \frac{m_i}{n} - \sum_{i \leq 0} \frac{n - m_i}{n} = \sum_{i \geq 1} z_i - \sum_{i \leq 0} (1 - z_i) = 0. \quad (7.5)$$

Using (7.4) and (7.5), $\vec{\pi}$ can be found by solving the following system of equations:

$$(\pi_{i-1}^2 - \pi_i^2) - (\pi_i - \pi_{i+1}) = 0 \text{ for all } i \in \mathbb{Z}, \quad (7.6)$$

$$\sum_{i \geq 1} \pi_i - \sum_{i \leq 0} (1 - \pi_i) = 0. \quad (7.7)$$

Note that (7.6) implies $\pi_{i+1} - \pi_i^2 = \pi_0 - \pi_{-1}^2$ for all $i \in \mathbb{Z}$. Since $\lim_{i \rightarrow \infty} \pi_i = 0$, we have $\pi_0 = \pi_{-1}^2$, and inductively we have the following relation:

$$\pi_{i+1} = \pi_i^2 \text{ for all } i \in \mathbb{Z}. \quad (7.8)$$

Using (7.8), (7.7) becomes $\sum_{i \geq 1} \pi_0^{2^i} - \sum_{i \geq 0} (1 - \pi_0^{2^{-i}}) = 0$. Such series are known as *lacunary series*, where the function has no analytic continuation across its disc of convergence (see Hadamard's Gap Theorem). There is no closed form expression for such series to the best of our knowledge and thus, we are unable to find the equilibrium point explicitly. Note that in the long-run, the probability that $|s_i^t| \leq M$ for any $i \in \mathcal{A}$ is equal to $\pi_{-M} - \pi_{M+1}$. Using (7.8), proving that for all $M \in \mathbb{Z}_+$, $\lim_{n \rightarrow \infty} p_{n,M} \geq 1 - a^M$ for some $a \in (0, 1)$, is equivalent to proving that the following inequalities hold:

$$\pi_0^{2^{-M}} - \pi_0^{2^{M+1}} \geq 1 - a^M \text{ for all } M \in \mathbb{Z}_+. \quad (7.9)$$

Lemma 7.4.8. *We have $\frac{1}{2} < \pi_0 < \frac{3}{4}$.*

The proof of Lemma 7.4.8 is given in Appendix F.5, and now we use it to prove Theorem 7.4.4.

Proof of Theorem 7.4.4. We will show that $g(M) := \pi_0^{2^{-M}} - \pi_0^{2^{M+1}} - 1 + 2^{-M} \geq 0$ for all positive integers M , which implies that (7.9) is satisfied with $a = \frac{1}{2}$. Note that $\lim_{M \rightarrow \infty} g(M) = 0$. Hence, we will show that $g(M) \geq g(M+1)$ for all $M \in \mathbb{Z}_+$. It is easy to check that $g(M) \geq 0$ for all $M \leq 6$ using Lemma 7.4.8. The derivative of g with respect to M is $\frac{dg(M)}{dM} = \log(\pi_0) \cdot \log(0.5) \cdot \pi_0^{2^{-M}} \cdot 2^{-M} + \log(\pi_0) \cdot \log(0.5) \cdot \pi_0^{2^{M+1}} \cdot 2^{2M+1} \cdot 2^{-M} - \log(2) \cdot 2^{-M}$, where $\log$ is the natural logarithm. By Lemma 7.4.8, we have $\frac{1}{2} < \pi_0 < \frac{3}{4}$, and thus $0.19 < \log \pi_0 \cdot \log(0.5) < 0.5$. Since $\pi_0^{2^{-M}} \leq 1$, the first term in $\frac{dg(M)}{dM}$ is upper bounded by $\frac{1}{2} \cdot 2^{-M}$. For the second term, note that $\pi_0^8 \cdot 2 < \frac{1}{3}$. Since $2^{M+1} > 8(2M+1)$ for all $M \geq 6$, the second term is upper bounded by $\frac{1}{2} \cdot \frac{1}{3} \cdot 2^{-M}$, and $\frac{dg(M)}{dM}$ is upper bounded by $\frac{1}{2} \cdot 2^{-M} + \frac{1}{6} \cdot 2^{-M} - \log(2) \cdot 2^{-M}$,

which is negative since $\log(2) > 2/3$. Hence, we have shown that $g(M)$ is a decreasing function on $[6, \infty]$, which concludes the proof. $\square$

We close this section with the following conjecture, which improves upon Theorem 7.4.2 by removing the implicit dependency on n and also upon Theorem 7.4.4 by providing finite-size guarantees.

Conjecture 7.4.9. *Assume that the agents are symmetric. Then for all $n \geq 2$, (C1) holds with $f(M) = a^M$ for all $M \in \mathbb{Z}_+$, for some $a \in [1/3, 1/2]$.*

Numerical experiments offer compelling evidence for the conjecture. See discussions in Ashlagi et al. [24] for additional details.

7.5 Strategic Considerations

We now offer a microeconomic justification for the notion of stability defined by conditions (C1) and (C2). Suppose that agents gain a payoff $\pi > 0$ whenever they receive service, and incur a cost $c \in [0, \pi)$ whenever they provide service. We assume that $c < \pi$, since otherwise providing service is inefficient, and the optimal mechanism is to never provide service.

We consider a token system (n, P, Q, d) under the minimum token selection rule. For simplicity, we focus on the symmetric case of $p_i = q_i = \frac{1}{n}$ for all $i \in \mathcal{A}$, so that all agents are equally likely to request service and be available to provide service. We suppose that at time $t = 0$, agents can either opt out of the token system, or else commit to participate for some fixed number of periods T . Assuming all n agents participate, each agent i serves as a requester some (random) number of times during the process, which we denote by A_i , and each agent serves as a provider B_i times. Thus, the total payoff to agent i is $U_i := \pi A_i - c B_i$. Note that by the symmetry assumption, the expectation of U_i is

$$\mathbb{E}[U_i] = \frac{T}{n}(\pi - c).$$

In particular, this is strictly larger than zero, so if agents maximize their expected payoffs, then they will opt in. Note that this holds independent of d (in particular, even in the unstable case of $d = 1$).

The advantage of stability is not that it provides agents with higher expected payoff, but rather that it lowers their risk. A simple model of risk aversion used in economics is that of mean-variance preferences.¹⁰ Under these preferences, which are parametrized by $\kappa > 0$, agents desire a higher expected payoff but also a lower variance, and evaluate a risky prospect by subtracting from its expectation κ times its variance. In particular, in our case agent i will choose to opt in if $\mathbb{E}[U_i] - \kappa \text{Var}(U_i) > 0$.

When $d = 1$, the variance of U_i is linear in T , since A_i and B_i are independent. Thus, agents will only opt in if κ is low enough. Intuitively, this unstable system is risky because agents are not unlikely to be left with many tokens at time T , which would mean that they provided service more often than they received it, resulting in a negative payoff.

In contrast, when $d \geq 2$, stability implies that the variance of U_i is bounded from above, independent of T . Since the expectation of U_i increases linearly, it follows that even agents with high κ will opt in, assuming T is large enough. This observation provides a foundation for why stable token systems are attractive to strategic economic agents: stable systems reduce agents' risk.

Relating to the literature. We further discuss our stability notion in the context of the literature that considers dynamic favor exchange games with discounting. Several papers

¹⁰We choose this classical model [e.g., see 149, 199] for its simplicity, but a similar conclusion applies for other, more sophisticated models, such as CARA preferences over monetary gambles. See [174] for a recent axiomatization of mean-variance preferences.

study settings with payoff π and cost $c < \pi$ per service as above, but (instead of choosing whether to opt in) agents' strategies depend on histories of exchanges. So an agent can refuse to provide service at any time period. In these models, each period one agent is selected at random to request service, and another (available) agent is selected to provide service according to some selection rule.

Since $c < \pi$, always trading (provisioning a service request) maximizes social welfare. However, demand is not always met, and there are two sources for this efficiency loss. First, an agent may decline to provide service when her number of tokens exceeds some threshold. This indeed happens when the service provider is selected uniformly at random [87, 117]. The incentive to decline service is alleviated when there is a steady-state [115, 43], since an agent will be rewarded for her service in finite time in expectation. Each of the conditions (C1) and (C2) implies that the token system has a steady-state.

While the token system under the minimum token selection rule has a steady-state, an agent may still prefer to decline service provision when her number of tokens exceeds some threshold. Indeed, Johnson, Simchi-Levi, and Sun [115] and Bo, Dawande, and Janakiraman [43] show the existence of an ϵ -Nash equilibrium, thus bounding agents' benefit from declining service.¹¹ Our results on condition (C1) suggest how concentrated the token profiles are, and in particular, the likelihood of different level of thresholds being reached. So, for example, even if punishments are feasible, how concentrated the token profile is (i.e., the structure of f in (C1)) can guide a social planner whether punishments are needed.

Second, inefficiency can arise when an agent demands a service, but has zero tokens. [115] show that (when service is always provided when feasible) the minimum token selection rule minimizes the probability of this event. Bo, Dawande, and Janakiraman [43] show that the likelihood of this event vanishes when n grows large and a constant fraction of agents are available to provide service for each request. In our setting, instead of zero, a social planner may decide to ban an agent from requesting service if her tokens are below some (possibly negative) threshold. Condition (C1) implies that an agent is unlikely to have a large deficit of tokens and explicitly quantifies the probability of this event. Finally, (C2) complements (C1) so that agents benefit from providing service in expectation, ensuring average positive utility in finite time.

7.6 Conclusion

This chapter adapts methodologies from stochastic processes and ideas from the power of two choices literature to illustrate that token systems based on the minimum token selection rule are likely to behave well even in thin marketplaces, where there can be very little availability of supply. We identified settings, under which the token system is stable when only few agents are available to provide service for any service request. Our analysis further provides

¹¹Johnson, Simchi-Levi, and Sun [115] use punishments to attain this, whereas Bo, Dawande, and Janakiraman [43] assume punishments are not feasible, but they require a large market, in which many agents are available to provide service.

concentration bounds for the token distribution. We further discuss why a stable token system is attractive for strategic agents interested in low participation risk.

In the context of kidney exchange, multi-hospital exchange platforms increase the chance for hospitals to match their hard-to-match pairs. To alleviate free riding by hospitals, token systems have been proposed and applied to incentivize participation by hospitals by accounting their contribution to the platform [9]. Our findings suggest that the possibility of breaking ties in the matching process, even among few hospitals, enables the stability of the token system. Breaking ties based on token balances allows platforms to reward hospitals with their contribution, assuring the difference between hospitals' tokens remains small. Numerical experiments reveal that tie-breakings are likely to happen in kidney exchange, and token systems that account for hospitals' contributions can ensure cooperation between hospitals. In practice, platforms use ad-hoc rules to increase participation, and several platforms seek to equalize the number of donors and patients of each "player" transplanted in every match [42, 155]. Our findings suggest that an arguably simpler (and less restrictive) token-based system is likely to behave well. It is interesting to expand this work to identify conditions under which dynamic token systems can be sustainable when the market exhibits more heterogeneity.

Bibliography

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. “Improved algorithms for linear stochastic bandits”. In: *Advances in neural information processing systems* 24 (2011).
- [2] Atila Abdulkadiroglu and Kyle Bagwell. *The optimal chips mechanism in a model of favors*. Tech. rep. Working paper, 2012.
- [3] Atila Abdulkadiroglu, Parag Pathak, and Alvin Roth. “The New York City high school match”. In: *American Economic Review* 95.2 (May 2005), pp. 364–367.
- [4] Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. “Strategy-proofness versus efficiency in matching with indifference: Redesigning the NYC high school match”. In: *American Economic Review* 99.5 (2009), pp. 1954–78.
- [5] Atila Abdulkadiroğlu and Tayfun Sönmez. “School choice: A mechanism design approach”. In: *American Economic Review* 93.3 (2003), pp. 729–747.
- [6] Atila Abdulkadiroğlu et al. “The Boston public school match”. In: *American Economic Review* 95.2 (2005), pp. 368–371.
- [7] Ralph Abraham, Jerrold E Marsden, and Tudor Ratiu. *Manifolds, tensor analysis, and applications*. Vol. 75. Springer Science & Business Media, 2012.
- [8] Nikhil Agarwal and Paulo Somaini. “Demand analysis using strategic reports: An application to a school choice mechanism”. In: *Econometrica* 86.2 (2018), pp. 391–444.
- [9] Nikhil Agarwal et al. “Market failure in kidney exchange”. In: *American Economic Review* 109.11 (2019), pp. 4026–4070.
- [10] José Alcalde and Antonio Romero-Medina. “Fair Student Placement”. In: *Theory and Decision* 83.2 (2017), pp. 293–307. DOI: 10.1007/s11238-017-9598-8.
- [11] David Aldous and Russell Lyons. “Processes on unimodular random networks”. In: *Electronic Journal of Probability* 12 (2007), pp. 1454–1508.

- [12] David Aldous and J. Michael Steele. “The Objective Method: Probabilistic Combinatorial Optimization and Local Weak Convergence”. In: *Probability on Discrete Structures*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 1–72. ISBN: 978-3-662-09444-0. DOI: 10.1007/978-3-662-09444-0_1. URL: https://doi.org/10.1007/978-3-662-09444-0_1.
- [13] Yeganeh Alimohammadi, Christian Borgs, and Amin Saberi. “Algorithms Using Local Graph Features to Predict Epidemics”. In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2022, pp. 3430–3451.
- [14] Yeganeh Alimohammadi, Christian Borgs, and Amin Saberi. “Locality of Random Digraphs on Expanders”. In: *arXiv preprint arXiv:2103.09952* (2021).
- [15] Noga Alon, Itai Benjamini, and Alan Stacey. “Percolation on finite graphs and isoperimetric inequalities”. In: *The Annals of Probability* 32.3 (2004), pp. 1727–1745.
- [16] Samson Alva and Vikram Manjunath. *Stable-dominating rules*. Tech. rep. Working paper, University of Ottawa, 2019.
- [17] Reid Andersen et al. “Local computation of PageRank contributions”. In: *Workshop on Algorithms and Models for the Web-Graph (WAW)*. Vol. 4863. Springer. 2007, pp. 150–165.
- [18] Dana Angluin. “Local and global properties in networks of processors”. In: *Proceedings of the twelfth annual ACM symposium on Theory of computing*. 1980, pp. 82–93.
- [19] Itai Ashlagi, Mark Braverman, and Avinatan Hassidim. “Stability in large matching markets with complementarities”. In: *Operations Research* 62.4 (2014), pp. 713–732.
- [20] Itai Ashlagi, Mark Braverman, and Geng Zhao. “Welfare Distribution in Two-sided Random Matching Markets”. In: *Proceedings of the 24th ACM Conference on Economics and Computation*. EC ’23. London, United Kingdom: Association for Computing Machinery, 2023, p. 122. ISBN: 9798400701047. DOI: 10.1145/3580507.3597730.
- [21] Itai Ashlagi, Yash Kanoria, and Jacob D Leshno. “Unbalanced random matching markets: The stark effect of competition”. In: *Journal of Political Economy* 125.1 (2017), pp. 69–98.
- [22] Itai Ashlagi and Alvin E Roth. “Kidney exchange: An operations perspective”. In: *Management Science* 67.9 (2021), pp. 5455–5478.
- [23] Itai Ashlagi et al. “On matching and thickness in heterogeneous dynamic markets”. In: *Operations Research* 67.4 (2019), pp. 927–949.
- [24] Itai Ashlagi et al. “The Power of Two in Token Systems”. In: *arXiv preprint arXiv:2405.12414* (2025).
- [25] Itai Ashlagi et al. “Tiered random matching markets: Rank is proportional to popularity”. In: *12th Innovations in Theoretical Computer Science Conference, ITCS 2021*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing. 2021, p. 46.

- [26] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. “Finite-time analysis of the multiarmed bandit problem”. In: *Machine learning* 47.2 (2002), pp. 235–256.
- [27] Yossi Azar, Andrei Z Broder, and Anna R Karlin. “On-line load balancing”. In: *Theoretical Computer Science* 130.1 (1994), pp. 73–84.
- [28] Yossi Azar et al. “Balanced allocations”. In: *SIAM journal on computing* 29.1 (1999), pp. 180–200.
- [29] Yu Bai et al. “Sample-efficient learning of Stackelberg equilibria in general-sum games”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 25799–25811.
- [30] Frank Ball and Tom Britton. “Epidemics on networks with preventive rewiring”. In: *Random Structures and Algorithms* (2020). arXiv: 2008.06375 [math.PR]. URL: <https://arxiv.org/abs/2008.06375>.
- [31] Frank Ball, Tom Britton, and Peter Neal. “An epidemic model on a network having two group structures with tunable overlap”. In: *arXiv preprint arXiv:2410.06696* (2024).
- [32] Frank Ball, David Sirl, and Pieter Trapman. “Analysis of a stochastic SIR epidemic on a random network incorporating household structure”. In: *Mathematical Biosciences* 224.2 (2010), pp. 53–73. DOI: 10.1016/j.mbs.2009.12.001.
- [33] Frank Ball, David Sirl, and Pieter Trapman. “SIR epidemics in populations with large sub-communities”. In: *The Annals of Applied Probability* 34.5 (2024), pp. 4408–4454.
- [34] Frank Ball, David Sirl, and Pieter Trapman. “Threshold behaviour and final outcome of an epidemic on a random network with household structure”. In: *Advances in Applied Probability* 41.3 (2009), pp. 765–796. DOI: 10.1239/aap/1253281063.
- [35] Keisuke Bando. “On the existence of a strictly strong Nash equilibrium under the student-optimal deferred acceptance algorithm”. In: *Games and Economic Behavior* 87 (2014), pp. 269–287.
- [36] Albert-László Barabási and Réka Albert. “Emergence of scaling in random networks”. In: *Science* 286.5439 (1999), pp. 509–512.
- [37] A. Barbour and G. Reinert. “Approximating the epidemic curve”. In: *Electronic Journal of Probability* (2013).
- [38] Marc Barthélemy, Claude Godreche, and Jean-Marc Luck. “Fluctuation effects in metapopulation models: percolation and pandemic threshold”. In: *Journal of theoretical biology* 267.4 (2010), pp. 554–564.
- [39] Itai Benjamini, Asaf Nachmias, and Yuval Peres. “Is the critical percolation probability local?”. In: *Probability theory and related fields* 149 (2011), pp. 261–269.
- [40] Itai Benjamini and Oded Schramm. “Recurrence of distributional limits of finite planar graphs”. In: *Selected Works of Oded Schramm* (2011), pp. 533–545.

- [41] J. van den Berg and Michael Keane. “On the continuity of the percolation probability function”. In: *Contemp. Math* 26 (1984), pp. 61–65.
- [42] Péter Biró et al. “Modelling and optimisation in European Kidney Exchange Programmes”. In: *European Journal of Operational Research* (2019).
- [43] Yang Bo, Milind Dawande, and Ganesh Janakiraman. “Analysis of Scrip Systems: On an Open Question in Johnson et al.(2014)”. In: *Operations Research* (2018).
- [44] T. Bohman and M. Piccollelli. “SIR epidemics on random graphs with a fixed degree sequence”. In: *Random Structures & Algorithms* 41.2 (2012), pp. 179–214. DOI: 10.1002/rsa.20386.
- [45] Béla Bollobás, Svante Janson, and Oliver Riordan. “The phase transition in inhomogeneous random graphs”. In: *Random Structures & Algorithms* 31.1 (2007), pp. 3–122.
- [46] Martin C. J. Bootsma and Neil M. Ferguson. “The effect of public health measures on the 1918 influenza pandemic in U.S. cities”. In: *Proceedings of the National Academy of Sciences* 104.18 (2007), pp. 7588–7593. DOI: 10.1073/pnas.0611071104. URL: <https://www.pnas.org/doi/epdf/10.1073/pnas.0611071104>.
- [47] Christian Borgs, Karissa Huang, and Christian Ikeokwu. *A Law of Large Numbers for SIR on the Stochastic Block Model: A Proof via Herd Immunity*. 2024. arXiv: 2410.07097 [math.PR]. URL: <https://arxiv.org/abs/2410.07097>.
- [48] Christian Borgs and Geng Zhao. “Locality via global ties: Stability of the 2-core against misspecification”. In: *arXiv preprint arXiv:2304.06170* (2023).
- [49] T. Britton et al. “Epidemics on random graphs with tunable clustering”. In: *Journal of Applied Probability* 45.3 (2008), pp. 743–756.
- [50] Tom Britton, Maria Deijfen, and Fredrik Liljeros. “A stochastic SIR network epidemic model with preventive dropping of edges”. In: *Journal of Mathematical Biology* 78.7 (2019), pp. 2125–2159. DOI: 10.1007/s00285-019-01380-6.
- [51] Linda Cai and Clayton Thomas. “The short-side advantage in random matching markets”. In: *SIAM Symposium on Simplicity in Algorithms (SOSA)*. 2022.
- [52] Claudia Cerrone, Yoan Hermstrüwer, and Onur Kesten. “School choice with consent: An experiment”. In: *The Economic Journal* (Jan. 2024).
- [53] Yeon-Koo Che and Fuhito Kojima. “Asymptotic equivalence of probabilistic serial and random priority mechanisms”. In: *Econometrica* 78.5 (2010), pp. 1625–1672.
- [54] Yeon-Koo Che and Olivier Tercieux. “Efficiency and stability in large matching markets”. In: *Journal of Political Economy* 127.5 (2019), pp. 2301–2342.
- [55] Yeon-Koo Che and Olivier Tercieux. “Payoff equivalence of efficient mechanisms in large matching markets”. In: *Theoretical Economics* 13.1 (2018), pp. 239–271.

- [56] Yiqiu Chen and Markus Möller. “Regret-free truth-telling in school choice with consent”. In: *Theoretical Economics* (2023).
- [57] Matteo Chinazzi et al. “The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak”. In: *Science* 368.6489 (2020), pp. 395–400. DOI: 10.1126/science.aba9757. eprint: <https://www.science.org/doi/pdf/10.1126/science.aba9757>. URL: <https://www.science.org/doi/abs/10.1126/science.aba9757>.
- [58] Juniper Cocomello and Kavita Ramanan. “Exact description of limiting SIR and SEIR dynamics on locally tree-like graphs”. In: *arXiv preprint arXiv:2309.08829* (2023). URL: <https://arxiv.org/abs/2309.08829>.
- [59] Vittoria Colizza and Romualdo Pastor-Satorras. “Reaction-diffusion processes and metapopulation models in heterogeneous networks”. In: *Nature Physics* 3 (2007), pp. 276–282. DOI: 10.1038/nphys560.
- [60] Vittoria Colizza and Alessandro Vespignani. “Epidemic modeling in metapopulation systems with heterogeneous coupling pattern: Theory and simulations”. In: *Journal of Theoretical Biology* 251.3 (2008), pp. 450–467. DOI: 10.1016/j.jtbi.2007.11.028.
- [61] Vittoria Colizza and Alessandro Vespignani. “Invasion threshold in heterogeneous metapopulation networks”. In: *Physical Review Letters* 99.14 (2007), p. 148701. DOI: 10.1103/PhysRevLett.99.148701.
- [62] Vittoria Colizza et al. “Modeling the worldwide spread of pandemic influenza: baseline case and containment interventions”. In: *PLoS medicine* 4.1 (2007), e13.
- [63] Vincent Conitzer and Tuomas Sandholm. “Computing the optimal strategy to commit to”. In: *Proceedings of the 7th ACM Conference on Electronic Commerce*. 2006, pp. 82–90.
- [64] Kenny S Crump and Charles J Mode. “A general age-dependent branching process. I”. In: *Journal of mathematical analysis and applications* 24.3 (1968), pp. 494–508.
- [65] John K Dagsvik. “Aggregation in matching markets”. In: *International Economic Review* 41.1 (2000), pp. 27–58.
- [66] L. Decreusefond et al. “Large graph limit for an SIR process in random network with heterogeneous connectivity”. In: *Annals of Applied Probability* 22.2 (2012), pp. 541–575. DOI: 10.1214/10-AAP745.
- [67] Souvik Dhara, Debankur Mukherjee, and Subhabrata Sen. “Phase transitions of extremal cuts for the configuration model”. In: *Electronic Journal of Probability* 22 (2017), pp. 1–29.
- [68] Franz Diebold and Martin Bichler. “Matching with indifference: A comparison of algorithms in the context of course allocation”. In: *European Journal of Operational Research* 260.1 (2017), pp. 268–282.

- [69] Odo Diekmann, J. A. P. Heesterbeek, and J. A. J. Metz. “On the definition and the computation of the basic reproduction ratio R_0 in models for infectious diseases in heterogeneous populations”. In: *Journal of Mathematical Biology* 28.4 (1990), pp. 365–382.
- [70] Benjamin Doerr. “Probabilistic tools for the analysis of randomized optimization heuristics”. In: *Theory of evolutionary computation*. Springer, 2020, pp. 1–87.
- [71] Battal Doğan and Lars Ehlers. “Minimally unstable Pareto improvements over deferred acceptance”. In: *Theoretical Economics* 16.4 (2021), pp. 1249–1279.
- [72] Jinshuo Dong et al. “Strategic classification from revealed preferences”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 2018, pp. 55–70.
- [73] Kefan Dong, Jiaqi Yang, and Tengyu Ma. “Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 26168–26182.
- [74] Umut Dur, A Arda Gitmez, and Özgür Yılmaz. “School choice under partial fairness”. In: *Theoretical Economics* 14.4 (2019), pp. 1309–1346.
- [75] Rick Durrett and Dong Yao. “Susceptible–infected epidemics on evolving graphs”. In: *Electronic Journal of Probability* 27.110 (2022), pp. 1–66. DOI: 10.1214/22-EJP828. URL: <https://doi.org/10.1214/22-EJP828>.
- [76] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. “Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator”. In: *The Annals of Mathematical Statistics* (1956), pp. 642–669.
- [77] Lars Ehlers and Thayer Morrill. “(Il) legal assignments in school choice”. In: *The Review of Economic Studies* 87.4 (2020), pp. 1837–1875.
- [78] Aytek Erdil. “Strategy-proof stochastic assignment”. In: *Journal of Economic Theory* 151 (2014), pp. 146–162.
- [79] Pál Erdős and Alfréd Rényi. “On a classical problem of probability theory”. In: *A Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei* 6.1-2 (1961), pp. 215–220.
- [80] Paul Erdős, Alfréd Rényi, et al. “On the evolution of random graphs”. In: *Publ. Math. Inst. Hung. Acad. Sci* 5.1 (1960), pp. 17–60.
- [81] William Feller. *An introduction to probability theory and its applications, Volume 2*. Vol. 81. John Wiley & Sons, 1971.
- [82] Jacques Ferber and Gerhard Weiss. *Multi-agent systems: an introduction to distributed artificial intelligence*. Vol. 1. Addison-wesley Reading, 1999.
- [83] Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.

- [84] Antoine Flahault and Alain-Jacques Valleron. “A method for assessing the global spread of HIV-1 infection based on air travel”. In: *Mathematical Population Studies* 3.3 (1992), pp. 161–171.
- [85] Dylan J Foster et al. “The statistical complexity of interactive decision making”. In: *arXiv preprint arXiv:2112.13487* (2021).
- [86] Nikolaos Fountoulakis. “Percolation on sparse random graphs with given degree sequence”. In: *Internet Mathematics* 4.4 (2007), pp. 329–356.
- [87] Eric J Friedman, Joseph Y Halpern, and Ian Kash. “Efficiency and Nash equilibria in a scrip system for P2P networks”. In: *Proceedings of the 7th ACM conference on Electronic commerce*. ACM. 2006, pp. 140–149.
- [88] David Gale and Lloyd S Shapley. “College admissions and the stability of marriage”. In: *The American Mathematical Monthly* 69.1 (1962), pp. 9–15.
- [89] Matthias Gerstgrasser and David C Parkes. “Oracles & Followers: Stackelberg Equilibria in Deep Multi-Agent Reinforcement Learning”. In: *arXiv preprint arXiv:2210.11942* (2022).
- [90] Mohsen Ghaffari. “Local computation of maximal independent set”. In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 438–449.
- [91] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. “On the complexity of local distributed graph problems”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 784–797.
- [92] Hugo Gimbert, Claire Mathieu, and Simon Mauras. “Two-sided matching markets with correlated random preferences have few stable pairs”. In: *arXiv preprint arXiv:1904.03890* (2019).
- [93] Michael A Goodrich, Alan C Schultz, et al. “Human–robot interaction: a survey”. In: *Foundations and Trends® in Human–Computer Interaction* 1.3 (2008), pp. 203–275.
- [94] R.F. Grais, J. Hugh Ellis, and G.E. Glass. “Assessing the impact of airline travel on the geographic spread of pandemic influenza”. In: *European Journal of Epidemiology* 18 (2003), pp. 1065–1072. URL: <https://doi.org/10.1023/A:1026140019146>.
- [95] Rebecca F. Grais, James H. Ellis, Adam Kress, et al. “Modeling the Spread of Annual Influenza Epidemics in the U.S.: The Potential Role of Air Travel”. In: *Health Care Management Science* 7.2 (2004), pp. 127–134. DOI: 10.1023/B:HCMS.0000020652.38181.da.
- [96] Moritz Hardt et al. “Strategic classification”. In: *Proceedings of the 2016 ACM conference on Innovations in Theoretical Computer Science*. 2016, pp. 111–122.
- [97] Christine Hauser and Hugo Hopenhayn. “Trading favors: Optimal exchange and forgiveness”. In: *Collegio Carlo Alberto Carlo Alberto Notebooks* 88 (2008).

- [98] Herbert W Hethcote. “An immunization model for a heterogeneous population”. In: *Theoretical population biology* 14.3 (1978), pp. 338–349.
- [99] Günter J Hitsch, Ali Hortaçsu, and Dan Ariely. “Matching and sorting in online dating”. In: *American Economic Review* 100.1 (2010), pp. 130–163.
- [100] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. “Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems”. In: *Proceedings of the fifteenth ACM conference on Economics and computation*. 2014, pp. 359–376.
- [101] Remco van der Hofstad. “The giant in random graphs is almost local”. In: *arXiv preprint arXiv:2103.11733* (2021).
- [102] Remco van der Hofstad, Johan SH van Leeuwen, and Clara Stegehuis. “Hierarchical configuration model”. In: *Internet Mathematics* (2016), p. 1214.
- [103] T. House. “Modelling epidemics on networks”. In: *Contemp. Phys.* 53.3 (2012), pp. 213–225.
- [104] Baihe Huang et al. “Optimal gradient-based algorithms for non-concave bandit optimization”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 29101–29115.
- [105] Yuanfei Huang and Adrian Röllin. “The SIR epidemic on a dynamic Erdős–Rényi random graph”. In: *arXiv preprint arXiv:2404.12566* (2024).
- [106] Lars Hufnagel, Dirk Brockmann, and Theo Geisel. “Forecast and control of epidemics in a globalized world”. In: *Proceedings of the National Academy of Sciences* 101.42 (Oct. 2004), pp. 15124–15129. DOI: 10.1073/pnas.0308344101.
- [107] Nicole Immorlica and Mohammad Mahdian. “Incentives in large random two-sided markets”. In: *ACM Transactions on Economics and Computation (TEAC)* 3.3 (2015), pp. 1–25.
- [108] Nicole Immorlica and Mohammad Mahdian. “Marriage, honesty, and stability”. In: *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA ’05. 2005, pp. 53–62.
- [109] Transforma Insights. *Number of Internet of Things (IoT) connected devices worldwide from 2022 to 2023, with forecasts from 2024 to 2033 (in billions)*. Available at <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>. Sept. 2024. URL: <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/> (visited on 01/02/2025).
- [110] Peter Jagers. “Convergence of general branching processes and functionals thereof”. In: *Journal of applied probability* 11.3 (1974), pp. 471–478.
- [111] S. Janson, M. Luczak, and P. Windridge. “Law of large numbers for the SIR epidemic on a random graph with given degrees”. In: *Random Structures & Algorithms* 45.4 (2014), pp. 726–763.

- [112] Svante Janson. “Tail bounds for sums of geometric and exponential variables”. In: *Statistics & Probability Letters* 135 (2018), pp. 1–6.
- [113] Svante Janson and Malwina J Luczak. “A simple solution to the k -core problem”. In: *Random Structures & Algorithms* 30.1-2 (2007), pp. 50–62.
- [114] Jiacheng Jiang et al. “SIR epidemics on evolving graphs”. In: *Electronic Journal of Probability* 24 (2019), pp. 1–28. DOI: 10.1214/19-EJP320. URL: <https://projecteuclid.org/euclid.ejp/1565747743>.
- [115] Kris Johnson, David Simchi-Levi, and Peng Sun. “Analyzing scrip systems”. In: *Operations Research* 62.3 (2014), pp. 524–534.
- [116] Hsu Kao, Chen-Yu Wei, and Vijay Subramanian. “Decentralized cooperative reinforcement learning with hierarchical information structure”. In: *International Conference on Algorithmic Learning Theory*. PMLR. 2022, pp. 573–605.
- [117] Ian A Kash, Eric J Friedman, and Joseph Y Halpern. “An equilibrium analysis of scrip systems”. In: *ACM Transactions on Economics and Computation* 3.3 (2015), p. 13.
- [118] Ian A Kash, Eric J Friedman, and Joseph Y Halpern. “Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion”. In: *Distributed Computing* 25.5 (2012), pp. 335–357.
- [119] William Ogilvy Kermack and Anderson G. McKendrick. “A Contribution to the Mathematical Theory of Epidemics”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115.772 (1927), pp. 700–721. DOI: 10.1098/rspa.1927.0118.
- [120] Onur Kesten. “School choice with consent”. In: *The Quarterly Journal of Economics* 125.3 (2010), pp. 1297–1348.
- [121] Robert Kleinberg and Tom Leighton. “The value of knowing a demand curve: Bounds on regret for online posted-price auctions”. In: *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. IEEE. 2003, pp. 594–605.
- [122] Taylor Knipe and Josue Ortega. “Improvable Students in School Choice”. In: *arXiv preprint arXiv:2504.12871* (2025).
- [123] Donald Knuth. *Mariages stables et leurs relations avec d’autres problèmes combinatoires: introduction à l’analyse mathématique des algorithmes*. Presses de l’Université de Montréal, 1976. ISBN: 9780840503428.
- [124] Donald E Knuth, Rajeev Motwani, and Boris Pittel. “Stable husbands”. In: *Random Structures & Algorithms* 1.1 (1990), pp. 1–14.
- [125] Donald Ervin Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*. Vol. 10. American Mathematical Soc., 1997.

- [126] Jens Kober, J Andrew Bagnell, and Jan Peters. “Reinforcement learning in robotics: A survey”. In: *The International Journal of Robotics Research* 32.11 (2013), pp. 1238–1274.
- [127] Fuhito Kojima and Parag A. Pathak. “Incentives and stability in large two-sided matching markets”. In: *American Economic Review* 99.3 (June 2009), pp. 608–27. DOI: 10.1257/aer.99.3.608.
- [128] Michael Krivelevich, Eyal Lubetzky, and Benny Sudakov. “Asymptotics in percolation on high-girth expanders”. In: *Random Structures & Algorithms* 56.4 (2020), pp. 927–947. DOI: <https://doi.org/10.1002/rsa.20903>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/rsa.20903>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rsa.20903>.
- [129] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. “The price of being near-sighted”. In: *SODA 2006: 17th ACM-SIAM Symposium on Discrete Algorithms, Miami, Florida, USA*. 2006.
- [130] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. “What cannot be computed locally!” In: *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*. 2004, pp. 300–309.
- [131] Thomas G Kurtz. *Approximation of population processes*. Vol. 36. SIAM, 1981.
- [132] Sohee Kwon et al. “Association of social distancing and face mask use with risk of COVID-19”. In: *Nature Communications* 12 (2021). DOI: 10.1038/s41467-021-24115-7. URL: <https://www.nature.com/articles/s41467-021-24115-7>.
- [133] John Langford and Tong Zhang. “The epoch-greedy algorithm for contextual multi-armed bandits”. In: *Advances in neural information processing systems* 20.1 (2007), pp. 96–1.
- [134] Anwar Lashari, Pieter Trapman, and David Sirl. “A branching process approach for epidemics in dynamic partnership networks”. In: *Journal of Mathematical Biology* 76.5 (2018), pp. 1231–1269. DOI: 10.1007/s00285-017-1175-1.
- [135] Niklas Lauffer et al. “No-Regret Learning in Dynamic Stackelberg Games”. In: *arXiv preprint arXiv:2202.04786* (2022).
- [136] SangMok Lee. “Incentive compatibility of large centralized matching markets”. In: *The Review of Economic Studies* 84.1 (2016), pp. 444–463.
- [137] SangMok Lee and Leeat Yariv. “On the efficiency of stable matchings in large markets”. In: *Available at SSRN 2464401* (2014).
- [138] Reut Levi, Dana Ron, and Ronitt Rubinfeld. “Local algorithms for sparse spanning graphs”. In: *Algorithmica* 82.4 (2020), pp. 747–786.
- [139] Don R Lick and Arthur T White. “ k -Degenerate graphs”. In: *Canadian Journal of Mathematics* 22.5 (1970), pp. 1082–1096.

- [140] Timothy P Lillicrap et al. “Continuous control with deep reinforcement learning”. In: *arXiv preprint arXiv:1509.02971* (2015).
- [141] Nathan Linial. “Locality in distributed graph algorithms”. In: *SIAM Journal on computing* 21.1 (1992), pp. 193–201.
- [142] Qingmin Liu and Marek Pycia. “Ordinal efficiency, fairness, and incentives in large markets”. In: *SSRN preprint* (2016).
- [143] Yang Liu and Yiling Chen. “A bandit framework for strategic regression”. In: *Advances in Neural Information Processing Systems* 29 (2016).
- [144] Ira M. Longini. “A mathematical model for predicting the geographic spread of new infectious agents”. In: *Mathematical Biosciences* 90.1 (1988), pp. 367–383. ISSN: 0025-5564. DOI: [https://doi.org/10.1016/0025-5564\(88\)90075-2](https://doi.org/10.1016/0025-5564(88)90075-2). URL: <https://www.sciencedirect.com/science/article/pii/0025556488900752>.
- [145] Peter Macgregor and He Sun. “Local algorithms for finding densely connected clusters”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 7268–7278.
- [146] Yishay Mansour and Shai Vardi. “A local computation approximation scheme to maximum matching”. In: *16th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) 2013 and 17th International Workshop on Randomization and Computation (RANDOM) 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*. Springer. 2013, pp. 260–273.
- [147] Janusz Marecki, Gerry Tesauro, and Richard Segal. “Playing repeated Stackelberg games with unknown opponents”. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems- Volume 2*. 2012, pp. 821–828.
- [148] Sharon Marko and Dana Ron. “Distance approximation in bounded-degree and general sparse graphs”. In: *9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) 2006 and 10th International Workshop on Randomization and Computation (RANDOM) 2006, Barcelona, Spain, August 28-30 2006. Proceedings*. Springer. 2006, pp. 475–486.
- [149] Harry Markowitz. “Portfolio selection”. In: *The Journal of Finance* 7.1 (1952), pp. 77–91.
- [150] Ana L P Mateus et al. “Effectiveness of travel restrictions in the rapid containment of human influenza: a systematic review”. In: *Bulletin of the World Health Organization* 92.12 (2014), pp. 868–880D. DOI: 10.2471/BLT.14.135590. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4264390/>.
- [151] Simon Muras. “Two-Sided Random Matching Markets: Ex-Ante Equivalence of the Deferred Acceptance Procedures”. In: *ACM Transactions on Economics and Computation* 9.4 (2021), pp. 1–14.

- [152] Peter McCullagh. “An asymptotic approximation for the permanent of a doubly stochastic matrix”. In: *Journal of Statistical Computation and Simulation* 84.2 (2014), pp. 404–414.
- [153] Konrad Menzel. “Large matching markets as two-sided demand systems”. In: *Econometrica* 83.3 (2015), pp. 897–941.
- [154] Meta. *Earnings Presentation Q3 2024*. Available at https://s21.q4cdn.com/399680738/files/doc_financials/2024/q3/Earnings-Presentation-Q3-2024.pdf. Oct. 2024. URL: https://s21.q4cdn.com/399680738/files/doc_financials/2024/q3/Earnings-Presentation-Q3-2024.pdf (visited on 01/02/2025).
- [155] Radu-Stefan Mincu et al. “IP solutions for international kidney exchange programmes”. In: *Central European Journal of Operations Research* (2020), pp. 1–21.
- [156] Michael David Mitzenmacher. “The power of two random choices in randomized load balancing”. PhD thesis. PhD thesis, Graduate Division of the University of California at Berkley, 1996.
- [157] Markus Möbius. “Trading favors”. In: *Mimeo, Microsoft Research* (2001).
- [158] Moni Naor and Larry Stockmeyer. “What can be computed locally?”. In: *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*. 1993, pp. 184–193.
- [159] Peter Neal. “SIR epidemics on a Bernoulli random graph”. In: *Journal of applied probability* 40.3 (2003), pp. 779–782.
- [160] VB Nevzorov. “Representations of order statistics, based on exponential variables with different scaling parameters”. In: *Journal of Soviet Mathematics* 33.1 (1986), pp. 797–798.
- [161] Donald J Newman and Lawrence Shepp. “The double dixie cup problem”. In: *The American Mathematical Monthly* 67.1 (1960), pp. 58–61.
- [162] Huy N Nguyen and Krzysztof Onak. “Constant-time approximation algorithms via local improvements”. In: *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE. 2008, pp. 327–336.
- [163] Afshin Nikzad. “Rank-optimal assignments in uniform markets”. In: *Theoretical Economics* 17.1 (2022), pp. 25–55.
- [164] Josué Ortega and Thilo Klein. “The cost of strategy-proofness in school choice”. In: *Games and Economic Behavior* 141 (2023), pp. 515–528.
- [165] Josué Ortega et al. “Identifying and Quantifying (Un) Improvable Students”. In: *arXiv preprint arXiv:2407.19831* (2025).
- [166] Cory Palmer and Dömötör Pálvölgyi. “At most $3.55n$ stable matchings”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 217–227.

- [167] Michal Parnas and Dana Ron. “Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms”. In: *Theoretical Computer Science* 381.1-3 (2007), pp. 183–196.
- [168] Merav Parter et al. “Local computation algorithms for spanners”. In: *arXiv preprint arXiv:1902.08266* (2019).
- [169] Marcin Peski. “Large roommate problem with non-transferable random utility”. In: *J. Econ. Theory* 168 (2017), pp. 432–471.
- [170] Boris Pittel. “On likely solutions of a stable marriage problem”. In: *The Annals of Applied Probability* 2.2 (1992), pp. 358–401.
- [171] Boris Pittel. “On likely solutions of the stable matching problem with unequal numbers of men and women”. In: *Mathematics of Operations Research* 44.1 (2019), pp. 122–146.
- [172] Boris Pittel. “The average number of stable matchings”. In: *SIAM Journal on Discrete Mathematics* 2.4 (1989), pp. 530–549.
- [173] Boris Pittel, Joel Spencer, and Nicholas Wormald. “Sudden emergence of a giant k -core in a random graph”. In: *Journal of Combinatorial Theory, Series B* 67.1 (1996), pp. 111–151.
- [174] Luciano Pomatto, Philipp Strack, and Omer Tamuz. “Stochastic dominance under independent noise”. In: *Journal of Political Economy* 128.5 (2020), pp. 1877–1900.
- [175] Canice Prendergast. “The Allocation of Food to Food Banks.” In: *EAI Endorsed Trans. Serious Games* 3.10 (2016), e4.
- [176] Derek J de Solla Price. “Networks of scientific papers: The pattern of bibliographic references indicates the nature of the scientific research front”. In: *Science* 149.3683 (1965), pp. 510–515.
- [177] Marek Pycia. “Evaluating with statistics: Which outcome measures differentiate among matching mechanisms?” In: *Unpublished paper, University of Zurich* (2019).
- [178] Filippo Radicchi et al. “Defining and identifying communities in networks”. In: *Proceedings of the national academy of sciences* 101.9 (2004), pp. 2658–2663.
- [179] Nived Rajaraman et al. “On the Value of Interaction and Function Approximation in Imitation Learning”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 1325–1336.
- [180] Philip J Reny. “Efficient matching in the school choice problem”. In: *American Economic Review* 112.6 (2022), pp. 2025–43.
- [181] Oliver Riordan. “The k -core and branching processes”. In: *Combinatorics, Probability and Computing* 17.1 (2008), pp. 111–136.
- [182] Amit Ronen et al. “Stable Marriage: Loyalty vs. Competition”. In: *arXiv preprint arXiv:2501.18442* (2025).

- [183] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. “Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences”. In: *American Economic Review* 97.3 (2007), pp. 828–851.
- [184] Ronitt Rubinfeld et al. “Fast local computation algorithms”. In: *arXiv preprint arXiv:1104.1377* (2011).
- [185] Daniel Russo and Benjamin Van Roy. “Eluder dimension and the sample complexity of optimistic exploration”. In: *Advances in Neural Information Processing Systems* 26 (2013).
- [186] Leonid A Rvachev and Ira M Longini Jr. “A mathematical model for the global spread of influenza”. In: *Mathematical biosciences* 75.1 (1985), pp. 3–22.
- [187] Ahmad EL Sallab et al. “Deep reinforcement learning framework for autonomous driving”. In: *Electronic Imaging* 2017.19 (2017), pp. 70–76.
- [188] Sourav Sarkar. “A note on the local weak limit of a sequence of expander graphs”. In: *Electronic Communications in Probability* 26 (2021), pp. 1–6.
- [189] Johannes Schneider and Roger Wattenhofer. “A log-star distributed maximal independent set algorithm for growth-bounded graphs”. In: *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*. 2008, pp. 35–44.
- [190] Pál Schweitzer. “Egy egyenlőtlenség az aritmetikai középértékről (An inequality concerning the arithmetic mean)”. In: *Math. Phys. Lapok* 23 (1914), pp. 257–261.
- [191] Stephen B Seidman. “Network structure and minimum degree”. In: *Social networks* 5.3 (1983), pp. 269–287.
- [192] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. “Safe, multi-agent, reinforcement learning for autonomous driving”. In: *arXiv preprint arXiv:1610.03295* (2016).
- [193] Richard Sinkhorn. “A relationship between arbitrary positive matrices and doubly stochastic matrices”. In: *The annals of mathematical statistics* 35.2 (1964), pp. 876–879.
- [194] Heinrich von Stackelberg. *Market Structure and Equilibrium*. Springer Science & Business Media, 2010.
- [195] Joan Sweeney and Richard James Sweeney. “Monetary theory and the great Capitol Hill Baby Sitting Co-op crisis: comment”. In: *Journal of Money, Credit and Banking* 9.1 (1977), pp. 86–89.
- [196] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
- [197] Qianfeng Tang and Jingsheng Yu. “A new perspective on Kesten’s school choice with consent idea”. In: *Journal of Economic Theory* 154 (2014), pp. 543–561.

- [198] Qianfeng Tang and Yongchao Zhang. “Weak stability and Pareto efficiency in school choice”. In: *Economic Theory* 71 (2021), pp. 533–552.
- [199] James Tobin. “Liquidity preference as behavior towards risk”. In: *The Review of Economic Studies* 25.2 (1958), pp. 65–86.
- [200] Pieter Trapman. “On analytical approaches to epidemics on networks”. In: *Theoretical Population Biology* 71.2 (2007), pp. 160–173. ISSN: 0040-5809. DOI: <https://doi.org/10.1016/j.tpb.2006.11.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0040580906001626>.
- [201] Peter Troyan, David Delacrétaz, and Andrew Kloosterman. “Essentially stable matchings”. In: *Games and Economic Behavior* 120 (2020), pp. 370–390.
- [202] Peter Troyan and Thayer Morrill. “Obvious manipulations”. In: *Journal of Economic Theory* 185 (2020), p. 104970.
- [203] Charles Van Loan. “The sensitivity of the matrix exponential”. In: *SIAM Journal on Numerical Analysis* 14.6 (1977), pp. 971–981.
- [204] Vivek Vishnumurthy, Sangeeth Chandrakumar, and Emin Gun Sirer. “Karma: A secure economic framework for peer-to-peer resource sharing”. In: *Workshop on Economics of Peer-to-Peer Systems*. Vol. 35. 2003.
- [205] Nikita Dmitrievna Vvedenskaya, Roland L’vovich Dobrushin, and Fridrikh Izrailevich Karpelevich. “Queueing system with selection of the shortest of two queues: An asymptotic approach”. In: *Problemy Peredachi Informatsii* 32.1 (1996), pp. 20–34.
- [206] Martin J Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Vol. 48. Cambridge University Press, 2019.
- [207] Chih-Chun Wang, Sanjeev R Kulkarni, and H Vincent Poor. “Bandit problems with side observations”. In: *IEEE Transactions on Automatic Control* 50.3 (2005), pp. 338–355.
- [208] LB Wilson. “An analysis of the stable marriage assignment algorithm”. In: *BIT Numerical Mathematics* 12.4 (1972), pp. 569–575.
- [209] Michael Wooldridge. *An introduction to multiagent systems*. John wiley & sons, 2009.
- [210] Annie Xie et al. “Learning latent representations to influence multi-agent interaction”. In: *Conference on robot learning*. PMLR. 2021, pp. 575–588.
- [211] Boling Yang et al. “Stackelberg MADDPG: Learning Emergent Behaviors via Information Asymmetry in Competitive Games”. In: (2022).
- [212] Yaolong Yu, Haifeng Xu, and Haipeng Chen. “Learning Correlated Stackelberg Equilibrium in General-Sum Multi-Leader-Single-Follower Games”. In: *arXiv preprint arXiv:2210.12470* (2022).
- [213] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. “Multi-agent reinforcement learning: A selective overview of theories and algorithms”. In: *Handbook of Reinforcement Learning and Control* (2021), pp. 321–384.

- [214] Geng Zhao et al. “Online learning in stackelberg games with an omniscient follower”. In: *International Conference on Machine Learning*. PMLR. 2023, pp. 42304–42316.
- [215] Han Zhong et al. “Can Reinforcement Learning Find Stackelberg-Nash Equilibria in General-Sum Markov Games with Myopic Followers?” In: *arXiv preprint arXiv:2112.13521* (2021).
- [216] Banghua Zhu et al. “The Sample Complexity of Online Contract Design”. In: *arXiv preprint arXiv:2211.05732* (2022).

Appendix A

Appendix to Chapter 2

A.1 Proofs in Section 2.3

A.1.1 Proof of Theorem 2.3.1

Proof. Consider Example 2.3.2. The expected reward is given by

$$h_\theta(a, b) := \theta \cdot \phi(a, b) = (1 - b)\theta_{-d} \cdot a + b(1 - \Delta), \quad (\text{A.1})$$

Optimizing over $b \in [0, 1]$ yields

$$\bar{h}_\theta(a) = \max\{1 - \Delta, \theta_{-d} \cdot a\}. \quad (\text{A.2})$$

Note that for any $a \in \mathcal{A}$ such that $\theta_{-d} \cdot a < 1 - \Delta$, the best response of the follower is $b = 1$, yielding an expected reward of $1 - \Delta$; for any $a \in \mathcal{A}$ such that $\theta_{-d} \cdot a \geq 1 - \Delta$, the best response of the follower is $b = 0$, yielding an expected reward of $\theta_{-d} \cdot a$. The optimal joint response $a = \theta_{-d}$ and $b = 0$ achieves the optimal expected reward of $\|\theta_{-d}\| = 1 > 1 - \Delta$. From the leader's perspective, this now reduces to the problem of a ReLU bandit considered in Dong, Yang, and Ma [73], since the response provides no information until the average regret falls below Δ .¹ Thus we have

$$\inf_{\hat{\pi}} \sup_{\theta \in \Theta} \mathcal{R}(T) \geq \Omega(T^{1 - \frac{2}{d-2}}).$$

□

¹Same as in Dong, Yang, and Ma [73], we allow the reward and response observations to be noiseless. We believe, however, the proof of Dong, Yang, and Ma [73, Theorem 5.1] has a small gap, where the packing number should be computed for radius $\sqrt{\epsilon}$ instead of ϵ . This lower bound can be further improved if we assume noisy observations.

A.1.2 Proof of Theorem 2.3.3

Proof. Let $\mathcal{H}(\epsilon)$ be a minimal ϵ -covering of $\mathcal{H}$ under the metric $\|\cdot\|_\infty$. Let

$$\mathcal{A}(\epsilon) = \left\{ \operatorname{argmax}_{a \in \mathcal{A}} \max_{b \in \mathcal{B}} h(a, b) \mid h \in \mathcal{H}(\epsilon) \right\},$$

where we break ties arbitrarily when the optimal action is non-unique. Note that we have $|\mathcal{A}(\epsilon)| \leq |\mathcal{H}(\epsilon)| \leq N(\epsilon)$. Let h^* be the true reward function. By the definition of a covering, there exists some $h_\epsilon \in \mathcal{H}(\epsilon)$ such that $\|h^* - h_\epsilon\|_\infty \leq \epsilon$. Thus we have

$$\mathcal{R}(T) = \sum_{t=1}^T \mathbb{E}[\bar{h}^*(a^*) - \bar{h}^*(a_t)] \leq \epsilon T + \sum_{t=1}^T \mathbb{E}[\bar{h}_\epsilon^*(a^*) - \bar{h}_\epsilon^*(a_t)].$$

We know that the optimal action for $\bar{h}_\epsilon$ must be inside the set $\mathcal{A}(\epsilon)$. Thus any worst-case optimal no-regret algorithm on the set $\mathcal{A}(\epsilon)$ gives a regret of $\sqrt{|\mathcal{A}(\epsilon)|T} \leq \sqrt{N(\epsilon)T}$. This gives that

$$\mathcal{R}(T) \leq \epsilon T + \sqrt{N(\epsilon)T}.$$

Taking infimum over ϵ finishes the proof. $\square$

A.2 Proofs in Section 2.4

A.2.1 Proof of Lemma 2.4.5

Proof. Recall the notation from Example 2.4.4: let $\theta_t^{(b)} = \Pi_{\mathcal{A}}(\hat{\theta}_t)$ for $t \geq 2$, with $\hat{\theta}_t := \frac{1}{t-1} \sum_{i=1}^{t-1} \hat{b}_i$. The first round incurs at most a constant regret and can be ignored. It suffices to show that, with probability at least $1 - \delta$,

$$\|\theta - \theta_t^{(b)}\| \leq \frac{\alpha_t}{\sqrt{t}} \tag{A.3}$$

for $\alpha_t = \Theta(\sigma_b \sqrt{d + \log \frac{T}{\delta}})$.

First, we bound the distance between $\hat{\theta}_t$ and θ . By our assumption,

$$\|\hat{\theta}_t - \theta\| = \left\| \frac{1}{t-1} \sum_{i=1}^{t-1} w_i \right\|,$$

where $w_1, \dots, w_t$ are i.i.d. zero-mean σ_b -sub-Gaussian. We proceed using a covering argument. Construct $U \subseteq \mathbb{S}^{d-1}$ such that

$$\inf_{v \in \mathbb{S}^{d-1}} \sup_{u \in U} u \cdot v \geq \frac{1}{2}. \tag{A.4}$$

Note that $\|u - v\| = \sqrt{2 - 2u \cdot v}$ for $u, v \in \mathbb{S}^{d-1}$. Hence, equivalently, we may choose U as a minimal 1-covering of $\mathbb{S}^{d-1}$ in Euclidean metric. Then

$$\log |U| \leq \log N^{\text{int}}(\mathbb{S}^{d-1}, 1, \|\cdot\|) \leq \log M(\mathbb{B}^d, 1, \|\cdot\|) = \Theta(d), \quad (\text{A.5})$$

where N^{int} and M denote the internal covering number and the packing number of the space under a given metric. The choice of U ensures that

$$\|w\| \leq 2 \sup_{u \in U} u \cdot w \quad (\text{A.6})$$

for all $w \in \mathbb{R}^d$, and ignoring the constant factor, we may focus on upper bounding the supremum $\sup_{u \in U} \sum_{i=1}^{t-1} u \cdot w_i$.

For each choice of $u \in U$, let $Z_{u,i} = u \cdot w_i$, so that $Z_{u,1}, \dots, Z_{u,t-1}$ are i.i.d. zero-mean σ_b -sub-Gaussian by definition of sub-Gaussian random vectors. By Hoeffding's inequality for sub-Gaussian random variables, we have

$$\mathbb{P}\left(\sum_{i=1}^t Z_{u,i} > x\right) \leq \exp\left(-\frac{x^2}{2t\sigma_b^2}\right) \quad (\text{A.7})$$

for all $x > 0$. Applying union bound over U and using (A.6) gives

$$\mathbb{P}\left(\left\|\sum_{i=1}^t w_i\right\| \geq 2x\right) \leq \mathbb{P}\left(\sup_{u \in U} \sum_{i=1}^t Z_{u,i} \geq x\right) \leq |U| \exp\left(-\frac{x^2}{2t\sigma_b^2}\right). \quad (\text{A.8})$$

Choosing $x = \sigma_b \sqrt{2t \log(|U|T)} \lesssim \sigma_b \sqrt{t(d + \log \frac{T}{\delta})}$ ensures that, by another union bound over $t \in [T]$,

$$\|\hat{\theta}_t - \theta\| \lesssim \sigma_b \sqrt{t^{-1}(d + \log \frac{T}{\delta})} \quad (\text{A.9})$$

with probability at least $1 - \delta$. By the triangle inequality and the definition of projection,

$$\|\theta_t^{(b)} - \theta\| \leq \|\theta_t^{(b)} - \hat{\theta}_t\| + \|\hat{\theta}_t - \theta\| \leq 2\|\hat{\theta}_t - \theta\| \lesssim \sigma_b \sqrt{t^{-1}(d + \log \frac{T}{\delta})} \quad (\text{A.10})$$

with the same probability. This gives (A.3) and completes the proof. $\square$

A.2.2 Proof of Proposition 2.4.6

Proof. We will condition upon the validity of the confidence sets, which holds with probability at least $1 - \delta$ per our choice of $\{\alpha_t\}_{t \in [T]}$.

UCB always chooses a_t in the confidence set Θ_t , with radius of order $O(\sigma_b \sqrt{t^{-1}(d + \log \frac{T}{\delta})})$.

When $\theta^* \in \Theta_t$, we have $\|a_t - \theta^*\| \lesssim \sigma_b \sqrt{t^{-1}(d + \log \frac{T}{\delta})}$. Since both a_t and θ^* are unit vectors,

we have

$$\begin{aligned}\mathcal{R}_{UCB}(T) &\leq 2\delta T + \sum_{t=1}^T (1 - \theta^* \cdot a_t) = 2\delta T + 2 + \frac{1}{2} \sum_{t=1}^T \|\theta^* - a_t\|^2 \\ &\lesssim 2\delta T + \sum_{t=2}^T \frac{\sigma_b^2}{t} \left(d + \log \frac{T}{\delta}\right) = O\left(\delta T + \sigma_b^2 \log T \cdot \left(d + \log \frac{T}{\delta}\right)\right),\end{aligned}$$

where the term $2\delta T$ bounds the contribution of the event that the confidence sets fails to be all valid. Choosing $\delta = 1/T$ gives our desired bound. $\square$

A.2.3 Proof of Proposition 2.4.10

Proof. After the first round, the leader's task reduces to a linear bandit with action space $\mathcal{A}_1$: only actions within $\mathcal{A}_1$ will be played, and the reward is linear in this region. As is well known for linear bandit (e.g., [185]), with probability $1 - \delta$, the regret in this linear stage (i.e., excluding the first round) is upper bounded by

$$2\delta T + O\left(\sqrt{d \log T \cdot (d \log T + \log \delta^{-1}) \cdot T}\right).$$

The first round adds at most a constant to this and can be ignored. By choosing $\delta = T^{-1}$, we have

$$\mathcal{R}_{UCB}(T) \leq \tilde{O}(d\sqrt{T}). \quad (\text{A.11})$$

$\square$

A.2.4 Proof of Proposition 2.4.11

Proof. Let $\Theta_1 = \{\theta_a \in \mathbb{S}^{d-1} | \theta_a \cdot b_1 \geq \zeta\} \times \{b_1\}$, and denote the true parameter by $\theta^* = (\theta_a^*, \theta_b^*)$. By our assumption on the problem structure, we have $\theta_a^* \in \Theta^{(b)}$.

As in the proof of Theorem 2.3.3, let $\Theta(\epsilon)$ be a minimal ϵ -covering of Θ_1 in Euclidean metric, with $\epsilon > 0$ to be specified later. In particular, there is some $\tilde{\theta}_a \in \Theta_1$ with $\|\tilde{\theta}_a - \theta_a^*\| \leq \epsilon$. Let $\mathcal{A}(\epsilon) = \{\arg\max_{a \in \mathcal{A}} \text{ReLU}(\theta_a \cdot a - \Delta) \mid \theta_a \in \Theta(\epsilon)\}$, where we break tie arbitrarily when the optimal action is non-unique. Note that $|\mathcal{A}(\epsilon)| \leq |\Theta(\epsilon)| = N(\Theta_1, \epsilon, \|\cdot\|)$.

Now, let the leader play UCB on the discrete action set $\mathcal{A}(\epsilon)$ after the first round. The regret satisfies

$$\mathcal{R}(T) \leq 1 + \sum_{t=2}^T \mathbb{E}[\bar{h}^*(a^*) - \bar{h}^*(a_t)] \leq 1 + T \cdot \mathbb{E}[\bar{h}^*(a^*) - \bar{h}^*(\tilde{a}^*)] + \sum_{t=1}^T \mathbb{E}[\bar{h}^*(\tilde{a}^*) - \bar{h}^*(a_t)], \quad (\text{A.12})$$

where $a^* = \theta_a^*$ and $\tilde{a}^* \in \arg\max_{a \in \mathcal{A}(\epsilon)} \bar{h}^*(a)$. Note that $\bar{h}^*(\tilde{a}^*) \geq \bar{h}^*(\tilde{\theta}_a) \geq \bar{h}^*(a^*) - \epsilon$ by our choice of $\tilde{\theta}_a$ and $\mathcal{A}(\epsilon)$, the second term in (A.12) is at most ϵT . The third term, the regret of UCB on $\mathcal{A}(\epsilon)$, is bounded by $O(\sqrt{N(\Theta_1, \epsilon, \|\cdot\|) \cdot T})$ in expectation.

It remains to bound $N(\Theta_1, \epsilon, \|\cdot\|)$. Note that for any $\theta_a, \theta'_a \in \Theta_1$, we have

$$\begin{aligned}\theta_a \cdot \theta'_a &= (\theta_a \cdot b_1)(\theta'_a \cdot b_1) + (\theta_a - (\theta_a \cdot b_1)b_1) \cdot (\theta'_a - (\theta'_a \cdot b_1)b_1) \\ &\geq \zeta^2 - \|\theta_a - (\theta_a \cdot b_1)b_1\| \|\theta'_a - (\theta'_a \cdot b_1)b_1\| \\ &\geq \zeta^2 - (1 - \zeta^2) = 2\zeta^2 - 1.\end{aligned}$$

Equivalently, $\|\theta_a - \theta'_a\| = \sqrt{2 - 2\theta_a \cdot \theta'_a} \leq 2\sqrt{1 - \zeta^2} = 2C_\zeta$. Thus, the covering number of Θ_1 is upper bounded by $\left(\frac{KC_\zeta^d}{\epsilon}\right)$ for some absolute constant K , which yields a regret bound of $1 + \epsilon T + O(\sqrt{K^d C_\zeta^d T / \epsilon^d})$. Choosing $\epsilon \asymp (KC_\zeta)^{\frac{d}{d+2}} T^{-\frac{1}{d+2}}$ reduces this upper bound to $O\left(C_\zeta^{\frac{d}{d+2}} T^{\frac{d+1}{d+2}}\right)$ as desired. $\square$

A.3 Proofs in Section 2.5

A.3.1 Proof of Proposition 2.5.2

Proof. Let the leader run the phased elimination algorithm Huang et al. [104, Algorithm 6] using the response $b_\theta^*(a_t)$ as the proxy reward to maximize. This proxy reward, in expectation, is a homogeneous polynomial of degree $2k - 1$. By Corollary 3.16 in Huang et al. [104], the algorithm achieves

$$\widehat{\mathcal{R}}(T) \leq \widetilde{O}(\sqrt{d^{2k-1}T}), \quad (\text{A.13})$$

where $\widehat{\mathcal{R}}(T) = \sum_{t=1}^T 1 - b_\theta^*(a_t)$ is the proxy regret measured based on the the proxy reward (i.e., absolute response). Note that the reward is maximized exactly when the proxy reward is maximized. Thus, the Lipschitz property (2.19) suggests that

$$\mathcal{R}(T) \leq \frac{2k}{2k-1} \widehat{\mathcal{R}}(T) \leq \widetilde{O}(\sqrt{d^{2k-1}T}). \quad (\text{A.14})$$

$\square$

Appendix B

Appendix to Chapter 3

B.1 Some Concentration Inequalities

In this section, we will establish concentration inequalities that will serve as key technical tools in many of our subsequent proofs.

B.1.1 Concentration for Independent Non-Identically Distributed Exponential Random Variables

We first establish some large-deviation-type concentration bounds for exponential random variables.

Lemma B.1.1. *Let $\mathbf{u} \in \mathbb{R}_+^n$ be a vector with $\|\mathbf{u}\|_1 = n$ and let $\mathbf{Z}$ be a random vector with independent $\text{Exp}(1)$ components. Then for any $t \in [0, 1)$, we have*

$$\mathbb{P}(\mathbf{u} \cdot \mathbf{Z} \leq tn) \leq (te^{1-t})^n \prod_{i=1}^n u_i^{-1} \leq (te)^n \prod_{i=1}^n u_i^{-1}. \quad (\text{B.1})$$

In fact, the upper bound given by the second inequality holds trivially when $t \geq 1$ and is invariant under simultaneous scaling of $\mathbf{u}$ and t .

Further, when $1/K \leq u_i \leq K$ for some constant $K \geq 1$, we have

$$\mathbb{P}(\mathbf{u} \cdot \mathbf{Z} \leq tn) \geq e^{-O(n^{2/3})} (te^{1-Kt})^n \prod_{i=1}^n u_i^{-1}. \quad (\text{B.2})$$

For $t = o(1)$, this lower bound becomes

$$e^{-O(n^{2/3}) + (1-K)tn} (te^{1-t})^n \prod_{i=1}^n u_i^{-1} = e^{o(n)} (te)^n \prod_{i=1}^n u_i^{-1},$$

indicating that the upper bound is tight up to a factor of $e^{o(n)}$.¹

Proof. First we establish the upper bound. Directly applying Chernoff's method on $\mathbf{u} \cdot \mathbf{Z}$, we have

$$\mathbb{P}(\mathbf{u} \cdot \mathbf{Z} \leq tn) \leq \inf_{\lambda \geq 0} \frac{\mathbb{E}[\exp(-\lambda \mathbf{u} \cdot \mathbf{Z})]}{\exp(-\lambda tn)} = \inf_{\lambda \geq 0} e^{\lambda tn} \prod_{i=1}^n \frac{1}{1 + \lambda u_i}. \quad (\text{B.3})$$

Taking $\lambda = 1/t - 1$ (which is the minimizer when $\mathbf{u} = \mathbf{1}$) gives

$$\mathbb{P}(\mathbf{u} \cdot \mathbf{X} \leq tn) \leq e^{n-tn} \prod_{i=1}^n \frac{t}{t + (1-t)u_i}. \quad (\text{B.4})$$

Notice that $\mathbf{u} \mapsto \sum_{i=1}^n \log u_i$ is a concave function on $\mathbb{R}_+^n$, and hence

$$\prod_{i=1}^n (t + (1-t)u_i) \geq \left(\prod_{i=1}^n u_i \right)^{1-t} \geq \prod_{i=1}^n u_i$$

since $\prod_{i=1}^n u_i \leq (n^{-1} \sum_{i=1}^n u_i)^n = 1$. Plugging the above inequality into (B.4) gives the desired upper bound.

Now we establish the tightness of the bound under the additional assumption that $1/K \leq u_i \leq K$ for all $i \in [n]$. Consider independent random variables $W_i \sim \text{Exp}(u_i R/t)$ for $i = 1, \dots, n$ with $R = 1 + n^{-1/3}$, so that by Chebyshev's inequality

$$q_n := \mathbb{P}(\mathbf{u} \cdot \mathbf{W} \leq tn) = \mathbb{P}_{T \sim \Gamma(n,1)}(T \leq nR) \geq 1 - n^{-1/3}.$$

For convenience, we similarly write

$$p_n := \mathbb{P}(\mathbf{u} \cdot \mathbf{Z} \leq tn)$$

and write the (joint) distributions of $\mathbf{Z}$ and $\mathbf{W}$ as $P_n = \text{Exp}(1)^{\otimes n}$ and $Q_n = \bigotimes_{i=1}^n \text{Exp}(u_i R/t)$, respectively. Applying the data processing inequality to the channel $\mathcal{C}$ that maps $\zeta \in \mathbb{R}^n$ to $\mathbb{1}\{\mathbf{u} \cdot \zeta \leq tn\}$ gives

$$\begin{aligned} D(Q_n \| P_n) &\geq D(\mathcal{C}(Q_n) \| \mathcal{C}(P_n)) = q_n \log \frac{q_n}{p_n} + (1 - q_n) \log \frac{1 - q_n}{1 - p_n} \\ &= -q_n \log p_n + (1 - q_n) \log(1 - p_n) + (q_n \log q_n + (1 - q_n) \log(1 - q_n)), \end{aligned} \quad (\text{B.5})$$

where $D(\cdot \| \cdot)$ denotes the Kullback-Leibler (KL) divergence between two probability distributions. A direct computation gives

$$\begin{aligned} D(Q_n \| P_n) &= \sum_{i=1}^n \left(\frac{t}{Ru_i} - 1 - \log \frac{t}{Ru_i} \right) \leq \sum_{i=1}^n \left(\frac{Kt}{R} - 1 - \log \frac{t}{Ru_i} \right) \\ &= -n + R^{-1}Ktn - n \log t + n \log R + \sum_{i=1}^n \log u_i \leq -n + Ktn - n \log t + n^{2/3} + \sum_{i=1}^n \log u_i. \end{aligned} \quad (\text{B.6})$$

¹The tail bounds here are analogous to the bounds in [112]. Our upper bound is sharper when entries of $\mathbf{u}$ show “moderate heterogeneity,” and our lower bound is new.

Combining this with (B.5) and letting $n \rightarrow \infty$ gives

$$-n + Ktn - n \log t + n^{2/3} + \sum_{i=1}^n \log u_i \geq -(1 - O(n^{-1/3})) \log p_n + o(1), \quad (\text{B.7})$$

where we used the fact that $\log(1 - p_n) \rightarrow 0$ (due to our upper bound). Exponentiating both sides gives the desired lower bound for p_n . $\square$

As a consequence, we have the following lemma.

Lemma B.1.2. *Let $\mathbf{u}, \mathbf{v} \in \mathbb{R}_+^n$ be two vectors with bounded components, namely $\|\mathbf{u}\|_1 = \|\mathbf{v}\|_1 = n$ and $1/K \leq u_i, v_i \leq K$ for some fixed $K \geq 1$. For independent $Z_1, \dots, Z_n \sim \text{Exp}(1)$, we have*

$$\mathbb{P}\left(\left|\frac{\mathbf{u} \cdot \mathbf{Z}}{t\mathbf{u} \cdot \mathbf{v}^{-1}} - 1\right| > \zeta \mid \mathbf{v} \cdot \mathbf{Z} < tn\right) \leq \exp(-\Theta(n\zeta^2)) \quad (\text{B.8})$$

for $t = o(1)$ and any fixed constant $\zeta > 0$, where $\mathbf{v}^{-1}$ denotes the component-wise inverse of vector $\mathbf{v}$.²

Proof. We first prove the concentration bound for the lower tail. Writing

$$\begin{aligned} \mathbb{P}(u \cdot x < (1 - \zeta)tu \cdot v^{-1} \mid v \cdot x < tn) &= \frac{\mathbb{P}(u \cdot x < (1 - \zeta)tu \cdot v^{-1}, v \cdot x < tn)}{\mathbb{P}(v \cdot x < tn)} \\ &\leq \frac{\mathbb{P}((\lambda u + (1 - \lambda)v) \cdot x < (1 - \zeta)\lambda tu \cdot v^{-1} + (1 - \lambda)tn)}{\mathbb{P}(v \cdot x < tn)} \end{aligned}$$

for some $\lambda > 0$ to be determined later, the previous Lemma bounds the numerator by

$$t^n \left(1 - \lambda + \frac{(1 - \zeta)\lambda u \cdot v^{-1}}{n}\right)^n e^{n - (1 - \zeta)\lambda tu \cdot v^{-1} - (1 - \lambda)tn} \prod_{i=1}^n \frac{1}{\lambda u_i + (1 - \lambda)v_i}.$$

For lower bounding the denominator $\mathbb{P}(v \cdot x < tn)$, Lemma B.1.1 indicates that for $t = o(n)$, the denominator is well approximated by $t^n e^{n - tn} \prod_i v_i^{-1}$, up to an error of $e^{o(n)}$. Taking the ratio between the two quantities gives

$$\left(1 - \lambda + \frac{(1 - \zeta)\lambda u \cdot v^{-1}}{n}\right)^n e^{\lambda tn - (1 - \zeta)\lambda tu \cdot v^{-1}} \prod_{i=1}^n \frac{1}{\lambda u_i v_i^{-1} + 1 - \lambda}. \quad (\text{B.9})$$

²Note that this result is invariant under simultaneous scaling of vector $\mathbf{u}$, $\mathbf{v}$, and t . Essentially, we only need $tn/\|\mathbf{v}\|_1 = o(1)$ and bounded ratios between among the entries of $\mathbf{u}$ and $\mathbf{v}$. Further, the result remains unchanged if $Z_i \sim \text{Exp}(c_i)$ independently with c_i 's bounded on some $[1/K', K']$; the c_i 's can simply be absorbed into $\mathbf{u}$ and $\mathbf{v}$.

Focus on the quantity $\prod_{i=1}^n (\lambda u_i v_i^{-1} + 1 - \lambda)$. Using the AM-GM-HM inequality and the fact that $e^{a-1} \geq a$ for all $a \in \mathbb{R}$, we obtain

$$\prod_{i=1}^n \frac{1}{\lambda u_i v_i^{-1} + 1 - \lambda} \leq \left(1 - \lambda + \frac{\lambda u \cdot v^{-1}}{n}\right)^{-n} \exp\left(-n + \sum_{i=1}^n \frac{1 - \lambda + \lambda u \cdot v^{-1}/n}{1 - \lambda + \lambda u_i v_i^{-1}}\right).$$

Hence, the conditional probability of interest is upper bounded, up to $e^{o(n)}$, by the following expression

$$\begin{aligned} \left(1 - \lambda + \frac{(1 - \zeta)\lambda u \cdot v^{-1}}{n}\right)^n e^{\lambda t n - (1 - \zeta)\lambda t u \cdot v^{-1}} \left(1 - \lambda + \frac{\lambda u \cdot v^{-1}}{n}\right)^{-n} \\ \exp\left(-n + \sum_{i=1}^n \frac{1 - \lambda + \lambda u \cdot v^{-1}/n}{1 - \lambda + \lambda u_i v_i^{-1}}\right). \end{aligned}$$

Denote the negative logarithm of the n -th root of the quantity above by $\underline{\chi}(\lambda)$. That is,

$$\mathbb{P}(u \cdot x < (1 - \zeta)t u \cdot v^{-1} | v \cdot x < t n) \leq \inf_{\lambda > 0} e^{o(n) - n \underline{\chi}(\lambda)} = \exp\left(o(n) - n \sup_{\lambda > 0} \underline{\chi}(\lambda)\right)$$

for any $\lambda > 0$ with

$$\begin{aligned} \underline{\chi}(\lambda) := -\log\left(1 - \lambda + \frac{(1 - \zeta)\lambda u \cdot v^{-1}}{n}\right) - \lambda t + \frac{(1 - \zeta)\lambda t u \cdot v^{-1}}{n} + \log\left(1 - \lambda + \frac{\lambda u \cdot v^{-1}}{n}\right) \\ + 1 - \frac{1}{n} \sum_{i=1}^n \frac{1 - \lambda + \lambda u \cdot v^{-1}/n}{1 - \lambda + \lambda u_i v_i^{-1}}. \quad (\text{B.10}) \end{aligned}$$

The $o(n)$ factor is of lower order, and it suffices to show that there exists some λ such that $\underline{\chi}(\lambda) = \Theta(\zeta^2)$. For λ sufficiently small (e.g., $\lambda \leq K^{-2}/2$, recalling that $u_i, v_i \in [1/K, K]$), we may approximate the logarithm function near its zero and obtain

$$\log\left(1 - \lambda + \frac{\lambda u \cdot v^{-1}}{n}\right) \geq \frac{\lambda u \cdot v^{-1}}{n} - \lambda - \left(\frac{\lambda u \cdot v^{-1}}{n} - \lambda\right)^2.$$

Then the two log terms in (B.10) combined can be bounded below by

$$\lambda - \frac{(1 - \zeta)\lambda u \cdot v^{-1}}{n} + \frac{\lambda u \cdot v^{-1}}{n} - \lambda - \left(\frac{\lambda u \cdot v^{-1}}{n} - \lambda\right)^2 = \zeta \lambda \frac{u \cdot v^{-1}}{n} - \lambda^2 \left(\frac{u \cdot v^{-1}}{n} - 1\right)^2.$$

With $u_i, v_i \in [1/K, K]$, a naive lower bound is the following

$$\underline{\chi}(\lambda) \geq \zeta \lambda K^{-2} - \lambda^2 K^4 - \lambda t + (1 - \zeta)\lambda t K^{-2} + 1 - \frac{(2 - 2\lambda + K^2\lambda + K^{-2}\lambda)^2}{4(1 - \lambda + K^2\lambda)(1 - \lambda + K^{-2}\lambda)}, \quad (\text{B.11})$$

where the summation at the end of (B.10) is bounded using Schweitzer's inequality [190] for the ratio between arithmetic and harmonic means, stating

$$\frac{1}{n} \sum_{i=1}^n \frac{\bar{z}}{z_i} \leq \frac{(a+b)^2}{4ab}$$

for $z \in \mathbb{R}^n$ with bounded components $0 < a \leq z_i \leq b$. Further, we observe

$$1 - \frac{(2 - 2\lambda + K^2\lambda + K^{-2}\lambda)^2}{4(1 - \lambda + K^2\lambda)(1 - \lambda + K^{-2}\lambda)} = -\frac{(4K^2(K^{-2} - 1)^2 + (K - K^{-1})^4)}{4(1 - \lambda + K^2\lambda)(1 - \lambda + K^{-2}\lambda)}\lambda^2 \geq -3K^4\lambda^2.$$

Taking $\lambda = \zeta K^{-6}/8$ in (B.11) yields

$$\chi\left(\frac{1}{8}\zeta K^{-6}\right) \geq \frac{1}{16}\zeta^2 K^{-8} - \frac{1}{8}t\zeta K^{-6} \geq \Theta(\zeta^2),$$

hence finishing our proof for the lower tail.

The proof for the upper tail follows a similar structure. Writing

$$\begin{aligned} \mathbb{P}(u \cdot x > (1 + \zeta)tu \cdot v^{-1} | v \cdot x < tn) &= \frac{\mathbb{P}(u \cdot x > (1 + \zeta)tu \cdot v^{-1}, v \cdot x < tn)}{\mathbb{P}(v \cdot x < tn)} \\ &\leq \frac{\mathbb{P}((- \lambda u + (1 + \lambda)v) \cdot x < -(1 + \zeta)\lambda tu \cdot v^{-1} + (1 + \lambda)tn)}{\mathbb{P}(v \cdot x < tn)} \end{aligned}$$

for some $0 < \lambda < K^{-2}$ (so that $- \lambda u + (1 + \lambda)v \in \mathbb{R}_+^n$) to be determined later, Lemma B.1.1 implies that the ratio is, up to $e^{o(n)}$,

$$\left(1 + \lambda - \frac{(1 + \zeta)\lambda u \cdot v^{-1}}{n}\right)^n e^{(1 + \zeta)\lambda tu \cdot v^{-1} - \lambda tn} \prod_{i=1}^n \frac{1}{1 + \lambda - \lambda u_i v_i^{-1}}. \quad (\text{B.12})$$

As in the proof of lower tail bound, the product term can be bounded by

$$\prod_{i=1}^n (1 + \lambda - \lambda u_i v_i^{-1}) \leq \left(1 + \lambda - \frac{\lambda u \cdot v^{-1}}{n}\right)^{-n} \exp\left(-n + \sum_{i=1}^n \frac{1 + \lambda - \lambda u \cdot v^{-1}/n}{1 + \lambda - \lambda u_i v_i^{-1}}\right),$$

giving an upper bound, again up to $e^{o(n)}$, of

$$\begin{aligned} \left(1 + \lambda - \frac{(1 + \zeta)\lambda u \cdot v^{-1}}{n}\right)^n e^{(1 + \zeta)\lambda tu \cdot v^{-1} - \lambda tn} \left(1 + \lambda - \frac{\lambda u \cdot v^{-1}}{n}\right)^{-n} \\ \exp\left(-n + \sum_{i=1}^n \frac{1 + \lambda - \lambda u \cdot v^{-1}/n}{1 + \lambda - \lambda u_i v_i^{-1}}\right) \end{aligned}$$

for the conditional probability of interest.

Denote the negative logarithm of the n -th root of the quantity above by $\bar{\chi}(\lambda)$. That is,

$$\mathbb{P}(u \cdot x > (1 + \zeta)tu \cdot v^{-1} | v \cdot x < tn) \leq \inf_{0 < \lambda < K^{-2}} e^{o(n) - n\bar{\chi}(\lambda)}$$

for any $\lambda \in (0, K^{-2})$ with

$$\begin{aligned} \bar{\chi}(\lambda) := & -\log \left(1 + \lambda - \frac{(1 + \zeta)\lambda u \cdot v^{-1}}{n} \right) + \lambda t - \frac{(1 + \zeta)\lambda t u \cdot v^{-1}}{n} + \log \left(1 + \lambda - \frac{\lambda u \cdot v^{-1}}{n} \right) \\ & + 1 - \frac{1}{n} \sum_{i=1}^n \frac{1 + \lambda - \lambda u \cdot v^{-1}/n}{1 + \lambda - \lambda u_i v_i^{-1}}. \end{aligned}$$

Again, it suffices to prove that for some choice of λ we have $\bar{\chi}(\lambda) = \Theta(\zeta^2)$. With similar arithmetic as in the proof for the lower tail, we observe that for λ sufficiently small (e.g., $\lambda \leq K^{-2}/2$)

$$\bar{\chi}(\lambda) \geq \zeta \lambda K^{-2} + \lambda t - (1 + \zeta)\lambda t K^2 - 4\lambda^2 K^4. \quad (\text{B.13})$$

Again, taking $\lambda = \zeta K^{-6}/8$ in (B.13) gives the desired $\Theta(\zeta^2)$ lower bound for $\sup_{0 < \lambda < K^{-2}} \bar{\chi}(\lambda)$ and thus finishes our proof. $\square$

B.1.2 A Generalized DKW Inequality for Independent and Nearly Identically Distributed Random Variables

Lemma B.1.3. *Let X_i , $i = 1, \dots, n$ be independent random variables each with (non-identical) distribution function G_i , and assume that there exists a constant $\delta > 0$ and a distribution F such that $\|G_i - F\|_\infty \leq \delta$ uniformly across all $i = 1, \dots, n$. Let $\hat{G}$ be the empirical distribution function of $\{X_i\}_{i=1}^n$. Then*

$$\mathbb{P}(\|\hat{G} - F\|_\infty > 2\delta + \epsilon) < 4 \exp(-2n\epsilon^2/9). \quad (\text{B.14})$$

Proof. Let $U_i = G_i(X_i)$ so that $U_1, \dots, U_n$ are i.i.d. uniform on $[0, 1]$, and denote their empirical distribution function by $\hat{J}$. Let $Y_i = F^{-1}(U_i) = F^{-1}(G_i(X_i))$ so that $Y_1, \dots, Y_n$ are i.i.d. each with distribution function F , and denote their empirical distribution function by $\hat{F}$. Notice that

$$\begin{aligned} \|\hat{G} - F\|_\infty &= \sup_{x \in \mathbb{R}} \left| n^{-1} \sum_{i=1}^n I_{(-\infty, x)}(X_i) - F(x) \right| \\ &= \sup_{x \in \mathbb{R}} \left| n^{-1} \sum_{i=1}^n I_{(-\infty, x)}(Y_i) - F(x) + n^{-1} \sum_{i=1}^n (I_{(-\infty, x)}(Y_i) - I_{(-\infty, x)}(X_i)) \right| \\ &\leq \sup_{x \in \mathbb{R}} \left| n^{-1} \sum_{i=1}^n I_{(-\infty, x)}(Y_i) - F(x) \right| + \sup_{x \in \mathbb{R}} \left(n^{-1} \sum_{i=1}^n |I_{(-\infty, x)}(Y_i) - I_{(-\infty, x)}(X_i)| \right) \\ &= \|\hat{F} - F\|_\infty + \sup_{x \in \mathbb{R}} A(x). \end{aligned}$$

From the classic result of DKW inequality [76] applied to $\hat{F}$ and F , we know that

$$\mathbb{P}(\|\hat{F} - F\|_\infty > \epsilon/3) < 2\exp(-2n\epsilon^2/9). \quad (\text{B.15})$$

For the second supremum of in the sum above, we now consider U_i , $i = 1, \dots, n$ as the underlying random variables. Each term in the summation in A contributes 1 to the sum if and only if

$$F^{-1}(U_i) = Y_i < x \leq X_i = G_i^{-1}(U_i) \quad \text{or} \quad G_i^{-1}(U_i) = X_i < x \leq Y_i = F^{-1}(U_i),$$

or alternatively

$$F(x) \wedge G_i(x) \leq U_i \leq F(x) \vee G_i(x),^3$$

where $\wedge$ and $\vee$ denote the operators of taking the minimum and maximum, respectively. Hence,

$$\begin{aligned} A(x) &= n^{-1} \sum_{i=1}^n |I_{(-\infty, x)}(Y_i) - I_{(-\infty, x)}(X_i)| \\ &= n^{-1} \sum_{i=1}^n I_{(F(x) \wedge G_i(x), F(x) \vee G_i(x))}(U_i) \\ &\leq n^{-1} \sum_{i=1}^n I_{(\bigwedge_j G_j(x) \wedge F(x), \bigvee_j G_j(x) \vee F(x))}(U_i) \\ &= \hat{J}(M(x)) - \hat{J}(m(x)), \end{aligned}$$

where M and m denote the maximum and minimum across F and G_i , $i = 1, \dots, n$, respectively. By our assumption that $\|G_i - F\|_\infty \leq \delta$ across all i , we have that

$$0 \leq M(x) - m(x) \leq 2\delta$$

for all $x \in \mathbb{R}$. Noticing that the true distribution function J of U_i , $i = 1, \dots, n$ is the identity function on $[0, 1]$, we have

$$\begin{aligned} A(x) &= \hat{J}(M(x)) - \hat{J}(m(x)) \\ &\leq \left| \hat{J}(M(x)) - J(M(x)) \right| + \left| \hat{J}(m(x)) - J(m(x)) \right| + |J(M(x)) - J(m(x))| \\ &\leq 2\|\hat{J} - J\|_\infty + 2\delta \end{aligned}$$

on $\mathbb{R}$ uniformly. Therefore, applying DKW inequality again, we see that

$$\mathbb{P}(\sup A > 2\delta + 2\epsilon/3) \leq \mathbb{P}(\|\hat{J} - J\|_\infty > \epsilon/3) < 2\exp(-2n\epsilon^2/9). \quad (\text{B.16})$$

Combining (B.15) and (B.16) yields the desired bound in (B.14). $\square$

³We may safely ignore the case where the two sides are equal, as it happens with probability zero.

B.2 Typical Behaviors of Values in Stable Matching

In this section, we will characterize the typical behaviors of values in stable matchings, including first and second moment bounds and concentration. These results will ultimately lead to results in Section 3.6.3 such as Propositions 3.6.1 and 3.6.2.

Recall that for a matching μ with (latent) value vectors $\mathbf{X}(\mu)$ and $\mathbf{Y}(\mu)$, we define $\mathbf{U}(\mu) = F(\hat{\mathbf{X}}(\mu))$ and $\mathbf{V}(\mu) = f(\hat{\mathbf{X}}(\mu))$ with the standard exponential CDF $F(z) = 1 - e^{-z}$ applied component-wise to the renormalized values $\hat{X}_i(\mu) = X_i(\mu)a_{i,\mu(i)}$ and $\hat{Y}_i(\mu) = Y_i(\mu)b_{i,\mu^{-1}(i)}$ for $i \in [n]$, so that $U_i, V_i \sim \text{Unif}([0, 1])$ and are mutually independent due to the way the score matrices are generated.

Lemma B.2.1. *For any $c \in (0, 1/2)$, there exists a constant $\theta_1 > 0$ (depending on c and C) such that in a random instance of the market, at least $\theta_1 n \ln n$ proposals are made during man-proposing deferred acceptance with probability $1 - \exp(-n^c)$.*

Proof. This result follows from a standard analysis of the deferred acceptance algorithm executed as in [25, Section 3]. In [25], the preference model involves tiers, where fitness values among different tiers differ by at most a constant factor. It turns out that this bounded ratio of fitness is the only thing used in the proofs, and is also satisfied by our matching market under Assumption 3.2.2. The main steps of analysis are as follows.

1. Consider (man-proposing) deferred acceptance *with re-proposals*, where each time when man i proposes, his proposal will go to woman i with probability proportional to a_{ij} , independent of all previous proposals (and their acceptance/rejection). The total number T of proposals in this process is equal in distribution to the number of draws in the coupon collector problem, and from standard concentration bounds we can show $\mathbb{P}(T \geq kn^{1+c}) \leq e^{-n^{c-2}}$ for k sufficiently large and $\mathbb{P}(T < \alpha n \ln n) \leq \exp(-n^c - 2)$ for $\alpha > 0$ sufficiently small (see also [70]). The details mirror Appendices A and B in [25].
2. Analogous to Lemmas 3.5 and 3.6 in [25], we can show that, with probability $1 - \exp(-n^c - 1)$, no single man makes more than ℓn^{2c} proposals during deferred acceptance with re-proposal for ℓ sufficiently large.
3. Conditional on $T \geq \alpha n \ln n$ and the maximum number of proposals any man makes being at most ℓn^{2c} , the fraction of re-proposals should be no greater than $C\ell n^{2c-1}$ in expectation, since each proposal will be a duplicate of a previous proposal independently with probability at most $\frac{C\ell n^{2c}}{n}$. It follows immediately from a binomial concentration that the (conditional) probability that the number of repeated proposals exceeds $T/2$ is exponentially small. Hence, the actual number of proposals during deferred acceptance (without re-proposals) is at least $\frac{\alpha}{2} n \ln n$ with probability $1 - \exp(-n^c)$.

□

Lemma B.2.2. *For any $c \in (0, 1/2)$, there exists a constant $\theta_1 > 0$ (depending on c and C) such that, with probability $1 - \exp(-n^c)$, there exist no stable matchings μ where $\|\mathbf{U}(\mu)\|_1 \leq \theta_2 \ln n$.*

Proof. Note that $U_i(\mu) = 1 - e^{-a_{i,\mu(i)}X_i(\mu)} \in [C^{-1}F(X_i(\mu)), CF(X_i(\mu))]$ since $a_{i,\mu(i)} \in [1/C, C]$. For any stable matching μ , $\mathbf{X}(\mu) \succeq \mathbf{X}(\mu_{\text{MOSM}})$ where μ_{MOSM} is the man-optimal stable matching obtained from the man-proposing deferred acceptance, where all the n men are matched with their optimal possible stable partner (and hence achieves best value) simultaneously, and hence

$$\|\mathbf{U}(\mu)\|_1 \geq \frac{1}{C} \|F(\mathbf{X}(\mu))\|_1 \geq \frac{1}{C} \|F(\mathbf{X}(\mu_{\text{MOSM}}))\|_1 \geq \frac{1}{C^2} \|\mathbf{U}(\mu_{\text{MOSM}})\|_1. \quad (\text{B.17})$$

Thus, it suffices to consider the event where $\|\mathbf{U}(\mu_{\text{MOSM}})\|_1 \leq \theta_2 C^2 \ln n$. Without loss of generality, we may assume that μ_{MOSM} matches man i with woman i for $i \in [n]$. Denote by $R_i \in [n]$ and $X_i = X_{ii} \in \mathbb{R}_+$ the (random) rank of partner and the latent value in μ_{MOSM} for man $i \in [n]$, and let $U_i = F(a_{ii}X_i)$. By definition, $X_i \sim \text{Exp}(a_{ii})$ and $R_i = \sum_{j \neq i} \mathbb{1}\{X_{ij}/a_{ij} < X_i\}$.

We condition on a specific execution of the man-proposing deferred acceptance algorithm, i.e., on the sequence of proposals, which specifies the rank R_i and an ordering over his top R_i most preferred partners for each man. Notice that the specific value X_{ij} only affect the execution through the ordering of proposals, and hence conditional on a particular ordering, the values of the men are independent. Further, the value X_i conditional on R_i and an ordering $w_{j_1} \succeq_{m_i} w_{j_2} \succeq_{m_i} w_{j_1} \succeq \dots \succeq_{m_i} w_{j_{R_i}}$ is equal in distribution to $\tilde{X}_{(R_i)}$, i.e., the R_i -th order statistic of $(\tilde{X}_1, \dots, \tilde{X}_n)$ with $\tilde{X}_j \sim \text{Exp}(a_{ij})$, conditional on $\tilde{X}_{(k)} = \tilde{X}_{j_k}$ for all $k \in [R_i]$. By the representation of exponential order statistics given in [160], under such conditions,

$$\tilde{X}_{(R_i)} \stackrel{d}{=} \sum_{t=1}^{R_i} \frac{Z_{i,t}}{\sum_{k=t}^n a_{i,j_k}}, \quad (\text{B.18})$$

where $Z_{i,t} \sim \text{Exp}(1)$ are independently sampled for $t \in [n]$. Conditional on R_i and the sequence $j_1, \dots, j_{R_i}$, we have

$$\begin{aligned} U_i = F(a_{ii}X_i) &\stackrel{d}{=} F\left(\sum_{t=1}^{R_i} \frac{a_{ii}Z_{i,t}}{\sum_{k=t}^n a_{i,j_k}}\right) \geq F\left(\frac{1}{n} \sum_{t=1}^{R_i} \frac{Z_{i,t}}{C^2}\right) \\ &\geq \frac{1}{n} \sum_{t=1}^{R_i} F(Z_{i,t}/C^2) \geq \frac{1}{C^2 n} \sum_{t=1}^{R_i} F(Z_{i,t}) = \frac{1}{C^2 n} \sum_{t=1}^{R_i} W_{i,t}, \end{aligned} \quad (\text{B.19})$$

where the second last inequality follows from Jensen's inequality (applied to the concave function F), and $W_{i,t} = F(Z_{i,t}) \sim \text{Unif}([0, 1])$ independently. Thus, conditional on $\mathbf{R}$, we have $C^2 n \|\mathbf{U}\|_1 \succeq \sum_{i=1}^n \sum_{t=1}^{R_i} W_{i,t}$, independent of the specific ordering (and the identity) of

the proposals made. Therefore, we may marginalize over this ordering to get

$$\mathbb{P}\left(\|\mathbf{U}\|_1 \leq \theta_2 C^2 \ln n \mid \mathbf{R}\right) \leq \mathbb{P}\left(\sum_{i=1}^n \sum_{t=1}^{R_i} W_{i,t} \leq \theta_2 C^4 n \ln n \mid \mathbf{R}\right). \quad (\text{B.20})$$

Whenever $\|\mathbf{R}\|_1 \geq \theta_1 n \ln n$, the probability above is at most $\exp(-\Theta(n \ln n)) \ll \exp(-n^c)$ by Hoeffding's inequality granted that we choose $\theta_2 < \frac{\theta_1}{2C^4}$, and our proof is complete as we marginalize over all possible realizations of $\mathbf{R}$ with $\|\mathbf{R}\|_1 \geq \theta_1 n \ln n$. $\square$

Lemma B.2.3. *For any $\kappa \geq 0$ and $\theta_3 > 0$, the expected number of stable matchings μ with $\|\mathbf{U}(\mu)\|_1 \|\mathbf{V}(\mu)\|_1 \geq \theta_3 n (\ln n)^{1/8}$ is upper bounded by $\exp(-\kappa n)$. In particular, with high probability, no such stable matchings exists.*

Proof. It suffices to show that the probability that any fixed matching μ is stable and satisfies $\|\mathbf{U}(\mu)\|_1 \|\mathbf{V}(\mu)\|_1 \geq \theta_3 n (\ln n)^{1/8}$ is upper bounded by $o(e^{-\kappa n}/n!)$ for any $\theta_3 > 0$. Write $I = [0, 1]$ for the unit interval. Let

$$\Omega = \{(\mathbf{u}, \mathbf{v}) \in I^n \times I^n : \|\mathbf{u}\|_1 \|\mathbf{v}\|_1 > \theta_3 n (\ln n)^{1/8}\}$$

and let

$$P := \mathbb{P}(\mu \in \mathcal{S}, (\mathbf{U}(\mu), \mathbf{V}(\mu)) \in \Omega) = \int_{\mathbb{R}_+^n \times \mathbb{R}_+^n} p(\mathbf{x}, \mathbf{y}) \cdot \mathbb{1}_\Omega(F(\hat{\mathbf{x}}), F(\hat{\mathbf{y}})) \cdot \prod_{i=1}^n f(\hat{x}_i) f(\hat{y}_i) d\hat{\mathbf{x}} d\hat{\mathbf{y}}, \quad (\text{B.21})$$

where $f(t) = e^{-t}$ denotes the standard exponential density function. We apply the simple bound (3.10) on $p(\mathbf{x}, \mathbf{y})$ to obtain

$$\begin{aligned} P &\leq \int_{\mathbb{R}_+^n \times \mathbb{R}_+^n} \prod_{i \neq j} \left(1 - (1 - e^{-\hat{x}_i/C^2})(1 - e^{-\hat{y}_j/C^2})\right) \cdot \mathbb{1}_\Omega(F(\hat{\mathbf{x}}), F(\hat{\mathbf{y}})) \cdot \prod_{i=1}^n f(\hat{x}_i) f(\hat{y}_i) d\hat{\mathbf{x}} d\hat{\mathbf{y}} \\ &= \int_{I^n \times I^n} \prod_{i \neq j} \left(1 - (1 - (1 - u_i)^{1/C^2})(1 - (1 - v_j)^{1/C^2})\right) \cdot \mathbb{1}_\Omega(\mathbf{u}, \mathbf{v}) d\mathbf{u} d\mathbf{v} \\ &\leq \int_{I^n \times I^n} \prod_{i \neq j} \left(1 - \frac{1}{C^4} u_i v_j\right) \cdot \mathbb{1}_\Omega(\mathbf{u}, \mathbf{v}) d\mathbf{u} d\mathbf{v} \\ &\leq \int_{I^n \times I^n} \exp\left(-\frac{1}{C^4} (\|\mathbf{u}\|_1 \|\mathbf{v}\|_1 - \mathbf{u} \cdot \mathbf{v})\right) \cdot \mathbb{1}_\Omega(\mathbf{u}, \mathbf{v}) d\mathbf{u} d\mathbf{v} \\ &\leq \int_{I^n \times I^n} \exp\left(\frac{1}{C^4} (n - \|\mathbf{u}\|_1 \|\mathbf{v}\|_1)\right) \cdot \mathbb{1}_\Omega(\mathbf{u}, \mathbf{v}) d\mathbf{u} d\mathbf{v}, \end{aligned} \quad (\text{B.22})$$

where we use the basic facts that $1 - z^{1/C^2} \geq (1 - z)/C^2$ for all $z \in [0, 1]$ and $C \geq 1$ and that $1 + z \leq e^z$ for all $z \in \mathbb{R}$.

Let $s = \|\mathbf{u}\|_1$ and $t = \|\mathbf{v}\|_1$. It is well known (e.g., [81, Ch. I, Sec. 9]) that the probability density of $S = \|\mathbf{U}\|_1$ with $U_1, \dots, U_n$ independent samples from $\text{Unif}(I)$ is bounded above by $\frac{s^{n-1}}{(n-1)!}$. Hence,

$$P \leq \int_{s,t \in [0,n] : st \geq \theta_3 n (\ln n)^{1/8}} e^{(n-st)/C^4} \cdot \frac{n^2 (st)^{n-1}}{(n!)^2} ds dt. \quad (\text{B.23})$$

Note that when $st \geq n(\ln n)^2$, we have $\exp((n-st)/C^4) \leq \exp((n-n(\ln n)^2)/C^4) = o(\exp(-\kappa n))/n!$ and therefore the region $\{s, t \in [0, n] : st \geq n(\ln n)^2\}$ contributes a negligible amount to the integral. Hence,

$$\begin{aligned} P &\leq \frac{o(e^{-\kappa n})}{n!} + \int_{s,t \in [0,n] : \theta_3 n (\ln n)^{1/8} \leq st \leq n(\ln n)^2} e^{(n-st)/C^4} \cdot \frac{n^2 (st)^{n-1}}{(n!)^2} ds dt \\ &\stackrel{(i)}{\leq} \frac{o(e^{-\kappa n})}{n!} + n^2 \cdot e^{(n-\theta_3 n (\ln n)^{1/8})/C^4} \cdot \frac{n^2 (n(\ln n)^2)^{n-1}}{(n!)^2} \\ &\stackrel{(ii)}{\leq} \frac{o(e^{-\kappa n})}{n!} + \frac{1}{n!} \cdot \exp\left(\frac{1}{C^4}(n - \theta_3 n (\ln n)^{1/8}) + 3 \ln n + 2(n-1) \ln \ln n + n\right) = \frac{o(e^{-\kappa n})}{n!}, \end{aligned} \quad (\text{B.24})$$

where in step (i) we upper bound the integral by the product of the Lebesgue measure of its domain (bounded by n^2) and the supremum of its integrand, and in step (ii) we invoke Stirling's approximation. $\square$

Corollary B.2.4. *For any constant $c \in (0, 1/2)$, there exists a constant $\theta_4 > 0$ such that, with probability $1 - \exp(-n^c)$, there exist no stable matchings μ with $\|\mathbf{U}(\mu)\|_1 \geq \frac{\theta_4 n}{(\ln n)^{7/8}}$.*

Proof. This follows immediately from Lemma B.2.2 and B.2.3. $\square$

Proposition B.2.5. *Let $\Omega \subset I^n$ be a region in the n -dimensional hypercube. For $k \in \mathbb{Z}_+$, define interval $I_k = (2^{-k}n, 2^{-k+1}n]$. If*

$$\mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k) \otimes n}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) \leq \frac{g(n)}{e^{6n} C^{8n}} \quad (\text{B.25})$$

for some function $g(n)$ and uniformly for all $k \in \mathbb{Z}_+$, we can guarantee that the expected number of stable matchings μ with value vector $\mathbf{X}(\mu) \in \mathbb{R}^n$ satisfying $\mathbf{U}(\mu) \in \Omega$ is upper bounded by $g(n) + e^{-\Theta(n^2)}$; in particular, with high probability, no such stable matchings exist if $g(n) = o(1)$.

Proof. We focus on a fixed matching μ and by union bound, it suffices to show that

$$\mathbb{P}(\mu \in \mathcal{S}, \mathbf{U}(\mu) \in \Omega) = \frac{g(n) + e^{-\Theta(n^2)}}{n!} \quad (\text{B.26})$$

under the condition of (B.25). The same chain of reasoning as in (B.22) gives

$$\begin{aligned} P &:= \mathbb{P}(\mu \in \mathcal{S}, \mathbf{U}(\mu) \in \Omega) \leq \int_{I^n \times I^n} \exp\left(\frac{1}{C^4}(n - \|\mathbf{u}\|_1 \|\mathbf{v}\|_1)\right) \cdot \mathbb{1}_\Omega(\mathbf{u}) \, d\mathbf{u} \, d\mathbf{v} \\ &\leq e^n \int_{I^n \times I^n} \exp(-\|\mathbf{u}\|_1 \|\mathbf{v}\|_1 / C^4) \cdot \mathbb{1}_\Omega(\mathbf{u}) \, d\mathbf{u} \, d\mathbf{v} \end{aligned} \quad (\text{B.27})$$

since $C \geq 1$. Observing that $\|\mathbf{U}(\mu)\|_1 = 0$ and $\|\mathbf{V}(\mu)\|_1 = 0$ are both probability zero events, we may split the domain of the integral above into sub-regions according to which intervals $\|\mathbf{U}(\mu)\|_1$ and $\|\mathbf{V}(\mu)\|_1$ fall into, and then bound the value of the integral within each sub-region. That is, with the help of the monotone convergence theorem to interchange summation with integral,

$$\begin{aligned} P &\leq e^n \sum_{k,\ell=1}^{\infty} \int_{I^n \times I^n} \exp(-\|\mathbf{u}\|_1 \|\mathbf{v}\|_1 / C^4) \cdot \mathbb{1}_\Omega(\mathbf{u}) \mathbb{1}_{I_k}(\|\mathbf{u}\|_1) \mathbb{1}_{I_\ell}(\|\mathbf{v}\|_1) \, d\mathbf{u} \, d\mathbf{v} \\ &\leq e^n \sum_{k,\ell=1}^{\infty} \int_{I^n \times I^n} \exp(-2^{-k-\ell} n^2 / C^4) \cdot \mathbb{1}_\Omega(\mathbf{u}) \mathbb{1}_{I_k}(\|\mathbf{u}\|_1) \mathbb{1}_{I_\ell}(\|\mathbf{v}\|_1) \, d\mathbf{u} \, d\mathbf{v}. \end{aligned} \quad (\text{B.28})$$

For all $k \in \mathbb{Z}_+$, we have $e^{2^n} \exp(-2^k \|\mathbf{u}\|_1) \geq 1$ whenever $\mathbf{u} \in I^n$ with $\|\mathbf{u}\|_1 \in I_k$. Thus,

$$\begin{aligned} P &\leq e^{5n} \sum_{k,\ell=1}^{\infty} \int_{I^n \times I^n} \exp(-2^{-k-\ell} n^2 / C^4) \cdot \mathbb{1}_\Omega(\mathbf{u}) \mathbb{1}_{I_k}(\|\mathbf{u}\|_1) \mathbb{1}_{I_\ell}(\|\mathbf{v}\|_1) \\ &\quad \cdot \exp(-2^k \|\mathbf{u}\|_1 - 2^\ell \|\mathbf{v}\|_1) \, d\mathbf{u} \, d\mathbf{v} \\ &= e^{5n} \sum_{k,\ell=1}^{\infty} 2^{-(k+\ell)n} \exp(-2^{-k-\ell} n^2 / C^4) \cdot \\ &\quad \mathbb{E}_{\mathbf{U} \sim \text{Exp}(2^k)^{\otimes n}, \mathbf{V} \sim \text{Exp}(2^\ell)^{\otimes n}} [\mathbb{1}_\Omega(\mathbf{U}) \mathbb{1}_{I^n}(\mathbf{V}) \mathbb{1}_{I_k}(\|\mathbf{U}\|_1) \mathbb{1}_{I_\ell}(\|\mathbf{V}\|_1)]. \end{aligned} \quad (\text{B.29})$$

Observe that all the terms with $k + \ell \geq n$ combined contribute at most

$$\begin{aligned} e^{5n} \sum_{k+\ell \geq n} 2^{-(k+\ell)n} &= e^{5n} \left(\sum_{k=1}^{n-1} 2^{-kn} \sum_{\ell=n-k}^{\infty} 2^{-\ell n} + \sum_{k=n}^{\infty} 2^{-kn} \sum_{\ell=1}^{\infty} 2^{-\ell n} \right) \\ &= e^{5n} \cdot 2^{-n^2-n} \cdot \frac{n(1 - 2^{-n}) + 1}{(1 - 2^{-n})^2} = \frac{e^{-\Theta(n^2)}}{n!}, \end{aligned} \quad (\text{B.30})$$

which is negligible. Therefore, we only need to consider $O(n^2)$ terms and

$$\begin{aligned} P &\leq \frac{e^{-\Theta(n^2)}}{n!} + n^2 e^{5n} \max_{k,\ell \in \mathbb{Z}_+} \frac{\mathbb{E}_{\mathbf{U} \sim \text{Exp}(2^k)^{\otimes n}, \mathbf{V} \sim \text{Exp}(2^\ell)^{\otimes n}} [\mathbb{1}_\Omega(\mathbf{U}) \mathbb{1}_{I^n}(\mathbf{V}) \mathbb{1}_{I_k}(\|\mathbf{U}\|_1) \mathbb{1}_{I_\ell}(\|\mathbf{V}\|_1)]}{2^{(k+\ell)n} \exp(2^{-k-\ell} n^2 / C^4)} \\ &\leq \frac{e^{-\Theta(n^2)}}{n!} + e^{6n} \max_{k \in \mathbb{Z}_+} \frac{\mathbb{P}_{\mathbf{U} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{U} \in \Omega, \|\mathbf{U}\|_1 \in I_k)}{2^{(k+\ell)n} \exp(2^{-k-\ell} n^2 / C^4)}. \end{aligned} \quad (\text{B.31})$$

Let $\alpha = 2^{-(k+\ell)}$. When $\alpha \leq C^8/n$, the denominator of the second term can be bounded as $\alpha^{-n} e^{\alpha n^2/C^4} \geq n^n/C^{8n} \geq n!/C^{8n}$; when $\alpha > C^8/n$, let $\tau = n\alpha > C^8$ and we have $\alpha^{-n} e^{\alpha n^2/C^4} = \frac{n^n}{\tau^n} e^{\tau n/C^4} = n^n \exp(n(\tau/C^4 - \ln \tau)) \geq n^n \exp(n(C^4 - 4 \ln C)) \geq n!$. Thus, the denominator is bounded below by $n!/C^{8n}$ and

$$P \leq \frac{e^{-\Theta(n^2)}}{n!} + \frac{e^{\delta n C^{8n}}}{n!} \max_{k \in \mathbb{Z}_+} \mathbb{P}_{\mathbf{U} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{U} \in \Omega, \|\mathbf{U}\|_1 \in I_k) \leq \frac{1}{n!} (e^{-\Theta(n^2)} + g(n)). \quad (\text{B.32})$$

The claim follows immediately. $\square$

Lemma B.2.6. *For any fixed $\delta > 0$ and $\kappa > 0$, there exists an absolute constant θ_5 (depending on δ, κ , and C) such that, with probability $1 - \exp(-\kappa n)$, there exist no stable matchings μ with $\|\mathbf{X}_\delta(\mu)\|_\infty \leq \theta_5$. (Recall that $\mathbf{X}_\delta(\mu)$ is the value vector of the $(1-\delta)$ -partial matching obtained from μ that excludes the least happy $\delta/2$ fraction of men and women.)*

Proof. Since $\hat{\mathbf{X}}$ and $\mathbf{X}$ differ by at most a factor of C component-wise, we have

$$\|\mathbf{X}_\delta(\mu)\|_\infty \leq X_{(n-\lfloor \delta n/2 \rfloor)}(\mu) \leq C \hat{X}_{(n-\lfloor \delta n/2 \rfloor)}(\mu) = -C \log(1 - U_{(n-\lfloor \delta n/2 \rfloor)}(\mu)). \quad (\text{B.33})$$

Thus, it suffices to bound the upper $\delta/2$ quantile $U_{(n-\lfloor \delta n/2 \rfloor)}(\mu)$ away from 1.

Let $\Omega = \{\mathbf{u} \in I^n : \mathbf{u}_{(n-\lfloor \delta n/2 \rfloor)} > 1 - e^{-s}\}$ for some $s \geq 1$ that we will specified later. Then $\mathbf{W} \in \Omega$ implies that $\sum_{i=1}^n \mathbb{1}_{(1-e^{-s}, 1]}(W_i) \geq \delta n/2$. For $W_i \sim \text{Exp}(2^k)$, we have

$$\mathbb{P}(W_i \in (1 - e^{-s}, 1]) = \int_{1-e^{-s}}^1 2^k e^{-2^k t} dt \leq e^{-s} \cdot 2^k e^{-2^{k-1}} \leq e^{-s}.$$

Thus, $\sum_{i=1}^n \mathbb{1}_{(1-e^{-s}, 1]}(W_i)$ is stochastically dominated by a $\text{Bin}(n, e^{-s})$ random variable, and as a result

$$\begin{aligned} \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) &\leq \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega) \\ &\leq \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}\left(\sum_{i=1}^n \mathbb{1}_{(1-e^{-s}, 1]}(W_i) \geq \frac{\delta n}{2}\right) \\ &\leq \mathbb{P}_{Z \sim \text{Bin}(n, e^{-s})}\left(Z \geq \frac{\delta n}{2}\right) \leq \exp(-nD(\delta/2 \| e^{-s})). \end{aligned}$$

Since $D(\delta/2 \| z) \rightarrow \infty$ as $z \rightarrow 0$, it suffices to take s sufficiently large to ensure $D(\delta/2 \| e^{-s}) > \kappa + 6 + 8 \log C$. Proposition B.2.5 then guarantees with probability $1 - \exp(-\kappa n)$ that no stable matchings μ have $\mathbf{U}(\mu) \in \Omega$, and as a result of (B.33), with at least the desired probability, no stable matchings μ should have $\|\mathbf{X}_\delta(\mu)\|_\infty > \theta_5 := Cs$. $\square$

Lemma B.2.7. *For any $\delta > 0$ and $\kappa > 0$, there exists an absolute constant θ_6 (again depending on δ, κ , and C) such that, with probability $1 - \exp(-\kappa n)$, there exist no stable matchings μ with $\|\mathbf{X}_\delta(\mu)\|_1 \geq \theta_6 \|\mathbf{U}(\mu)\|_1$.*

Proof. Take $\theta_6 = C \sup_{0 < x \leq C\theta_5} \frac{x}{F(x)} = \frac{C^2\theta_5}{1-e^{-C\theta_5}}$, where θ_5 is the constant in Lemma B.2.6. Assume that $\|\mathbf{X}_\delta(\mu)\|_\infty \leq \theta_5$ for all stable matchings μ since the probability otherwise is at most $\exp(-\kappa n)$ as desired. Note that for each i in the support of $\mathbf{X}_\delta(\mu)$ (i.e., $(X_\delta)_i(\mu) > 0$), we have

$$\hat{X}_i(\mu) \leq CX_i(\mu) = C(X_\delta)_i(\mu) \leq C\theta_5, \quad (\text{B.34})$$

and subsequently

$$U_i(\mu) = F(\hat{X}_i(\mu)) \geq \frac{C\hat{X}_i(\mu)}{\theta_6} \geq \frac{X_i(\mu)}{\theta_6} = \frac{(X_\delta)_i(\mu)}{\theta_6}, \quad (\text{B.35})$$

and this final inequality is trivial for any i not in the support of $\mathbf{X}_\delta(\mu)$. The claim then follows immediately. $\square$

Lemma B.2.8. *For any fixed $\delta > 0$ and $\kappa > 0$, there exists an absolute constant θ_7 (depending on δ, κ , and C) such that, with probability $1 - \exp(-\kappa n)$, there exist no stable matchings μ with $\|\mathbf{X}_\delta(\mu)\|_2^2 \geq \theta_7 \|\mathbf{U}(\mu)\|_1^2/n$.*

Proof. Taking advantage of Lemma B.2.6, let us assume that $X_{(n-\lfloor \delta n/2 \rfloor)}(\mu) \leq \theta_5$ is satisfied simultaneously by all stable matchings μ (see the proof, in particular (B.33), for more details); the event otherwise has probability bounded by $\exp(-\kappa n)$.

Notice that

$$\|\mathbf{X}_\delta(\mu)\|_2^2 \leq \sum_{i=1}^{n-\lfloor \delta n/2 \rfloor} X_{(i)}(\mu)^2 \leq C^2 \sum_{i=1}^{n-\lfloor \delta n/2 \rfloor} \hat{X}_{(i)}(\mu)^2 \leq \theta'_6 \sum_{i=1}^{n-\lfloor \delta n/2 \rfloor} U_{(i)}(\mu)^2,$$

where the (i) subscript denotes the i -th (lower) order statistics (and in particular, $\hat{X}_{(i)}(\mu)$ is the i -th smallest entry of $\hat{\mathbf{X}}(\mu)$) with $\theta'_6 = C^2 \left(\frac{\theta_5}{F(\theta_5)} \right)^2$. Now it suffices to compare $\sum_{i=1}^{n-\lfloor \delta n/2 \rfloor} U_{(i)}(\mu)^2$ with $\|\mathbf{U}(\mu)\|_1^2/n$.

Consider $\Omega := \{\mathbf{w} \in I^n : \sum_{i=1}^{n-\lfloor \delta n/2 \rfloor} w_{(i)}^2 \geq \gamma \|\mathbf{w}\|_1^2/n\}$ for some $\gamma \in \mathbb{R}_+$ to be specified. By Proposition B.2.5, it suffices to show that for some appropriate value of γ we have

$$\mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) \leq e^{-(\kappa+6)n} C^{-8n}$$

for all $k \in \mathbb{Z}_+$. Observe that

$$\begin{aligned}
& \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) \\
&= \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}} \left(\sum_{i=1}^{n - \lfloor \delta n/2 \rfloor} W_{(i)}^2 \geq \frac{\gamma \|\mathbf{W}\|_1^2}{n}, 2^{-k}n < \|\mathbf{W}\|_1 \leq 2^{-k+1}n \right) \\
&\leq \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}} \left(\sum_{i=1}^{n - \lfloor \delta n/2 \rfloor} W_{(i)}^2 \geq \gamma 2^{-2k}n \right) \\
&\leq \mathbb{P}_{\mathbf{W} \sim \text{Exp}(1)^{\otimes n}} \left(\sum_{i=1}^{n - \lfloor \delta n/2 \rfloor} W_{(i)}^2 \geq \gamma n \right) \\
&\leq \mathbb{P}_{\mathbf{W} \sim \text{Exp}(1)^{\otimes n}} (W_{(n - \lfloor \delta n/2 \rfloor)} \geq \sqrt{\gamma}) \\
&\leq \mathbb{P}_{Z \sim \text{Bin}(n, e^{-\sqrt{\gamma}})} \left(Z \geq \frac{\delta n}{2} \right).
\end{aligned}$$

By the large deviation bound for binomial distribution, choosing γ sufficiently large such that $D(\delta/2 \|e^{-\sqrt{\gamma}}) > \kappa + 6 + 8 \log C$ ensures that this probability is $o(e^{-(\kappa+6)n} C^{-8n})$. This finishes the proof with the choice of $\theta_7 = \theta'_6 \gamma$. $\square$

Remark B.2.9. *This is the only part of our analysis that relies on the δ -truncation of values. Without the truncation, $\frac{1}{n} \|\mathbf{W}\|_2^2$ would concentrate poorly—in fact not even having a finite mean—for $\mathbf{W} \sim \text{Exp}(1)^{\otimes n}$.*

Corollary B.2.10. *For any constant $c \in (0, 1/2)$, there exists a constant $\theta_8 > 0$ such that, with probability $1 - \exp(-n^c)$, there is no stable matching μ with $\sum_{i=1}^n (X_\delta)_i \mathbb{1}_{[2/C, \infty)}((X_\delta)_i) \geq \theta_8 \|\mathbf{U}\|_1 / (\ln n)^{7/8}$.*

Proof. Notice that

$$\|\mathbf{X}_\delta\|_2^2 \geq \sum_{i=1}^n (X_\delta)_i^2 \mathbb{1}_{[2/C, \infty)}((X_\delta)_i) \geq \frac{2}{C} \sum_{i=1}^n (X_\delta)_i \mathbb{1}_{[2/C, \infty)}((X_\delta)_i).$$

The statement then follows from Lemma B.2.8 and Corollary B.2.4. $\square$

Lemma B.2.11. *For any fixed $\delta > 0$ and $\kappa > 0$, there exists an absolute constant θ_9 (depending on δ, κ , and C) such that, with probability $1 - \exp(-\kappa n)$, there exist no stable matchings μ with $\|\mathbf{X}_\delta(\mu)\|_1 \leq \theta_9 \|\mathbf{U}(\mu)\|_1$.*

Proof. Since $\mathbf{X} \succeq \mathbf{U}$ component-wise, we have $\|\mathbf{X}_\delta(\mu)\|_1 \geq \|\mathbf{U}_\delta(\mu)\|_1$. Thus, it suffices to consider the condition $\|\mathbf{U}_\delta(\mu)\|_1 \leq \theta_9 \|\mathbf{U}(\mu)\|_1$.

Consider $\Omega := \{\mathbf{w} \in I^n : \exists S \subseteq [n], |S| = n - \lfloor \delta n \rfloor, \sum_{i \in S} w_i \leq \alpha \|\mathbf{w}\|_1\}$ for some $\alpha \in \mathbb{R}_+$ to be specified. By union bound, for any $k \in \mathbb{Z}_+$,

$$\begin{aligned} \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) &\leq \binom{n}{\lfloor \delta n \rfloor} \mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}\left(\sum_{i=1}^{n-\lfloor \delta n \rfloor} w_i \leq \alpha \|\mathbf{W}\|_1 \leq 2^{-k+1} \alpha n\right) \\ &= \binom{n}{\lfloor \delta n \rfloor} \mathbb{P}_{\mathbf{W} \sim \text{Exp}(1)^{\otimes n}}\left(\sum_{i=1}^{n-\lfloor \delta n \rfloor} w_i \leq 2\alpha n\right) \\ &= \exp(-h(\delta)n + o(n)) \cdot \left(\frac{2\alpha e}{1-\delta}\right)^{n-\lfloor \delta n \rfloor}, \end{aligned}$$

where in the last step we use Stirling's approximation to bound the first factor and standard (lower) concentration of $\text{Exp}(1)$ to bound the probability term (e.g., see Lemma B.1.1). For α sufficiently small, e.g., $\alpha < \exp\left(\frac{1}{1-\delta}(h(\delta) - \kappa - 6 - 8 \ln C) - h(\delta)\right)$, we have

$$\mathbb{P}_{\mathbf{W} \sim \text{Exp}(2^k)^{\otimes n}}(\mathbf{W} \in \Omega, \|\mathbf{W}\|_1 \in I_k) \leq e^{-(\kappa+6)n} C^{-8n}$$

for all $k \in \mathbb{Z}_+$. Invoking Lemma B.2.5 concludes the proof with $\theta_9 = \alpha$. $\square$

Corollary B.2.12. *For any constant $c \in (0, 1/2)$, there exists a constant $\theta_{10} > 0$ such that, with probability $1 - \exp(-n^c)$, there exist no stable matchings μ with $\|\mathbf{X}_\delta(\mu)\|_1 \leq \theta_{10} \ln n$.*

Proof. This follows from Lemmas B.2.2 and B.2.11, with $\theta_{10} = \theta_2 \theta_9$. $\square$

The following Corollary combines all the previous into the typical behavior of value vectors in stable matchings.

Corollary B.2.13. *Define $\mathcal{R}^*(\mu) \subseteq \mathbb{R}_+^n \times \mathbb{R}_+^n$, in the context of a matching μ , to be the set of all pairs of vectors $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n$ that satisfy all of the following conditions:*

$$\theta_2 \ln n \leq \|\mathbf{u}\|_1, \|\mathbf{v}\|_1 \leq \frac{\theta_4 n}{(\ln n)^{7/8}}, \quad (\text{B.36})$$

$$\|\mathbf{u}\|_1 \|\mathbf{v}\|_1 \leq \theta_3 n (\ln n)^{1/8}, \quad (\text{B.37})$$

$$\|\mathbf{x}_\delta\|_1 \leq \theta_6 \|\mathbf{u}\|_1 \text{ and } \|\mathbf{y}_\delta\|_1 \leq \theta_6 \|\mathbf{v}\|_1, \quad (\text{B.38})$$

$$\|\mathbf{x}_\delta\|_2^2 \leq \frac{\theta_7 \|\mathbf{u}\|_1^2}{n} \text{ and } \|\mathbf{y}_\delta\|_2^2 \leq \frac{\theta_7 \|\mathbf{v}\|_1^2}{n}, \quad (\text{B.39})$$

$$\sum_{i=1}^n (x_\delta)_i \mathbb{1}_{[2/C, \infty)}((x_\delta)_i) \leq \frac{\theta_8 \|\mathbf{u}\|_1}{(\ln n)^{7/8}} \text{ and } \sum_{i=1}^n (y_\delta)_i \mathbb{1}_{[2/C, \infty)}((y_\delta)_i) \leq \frac{\theta_8 \|\mathbf{v}\|_1}{(\ln n)^{7/8}}, \quad (\text{B.40})$$

$$\|\mathbf{x}_\delta\|_1, \|\mathbf{y}_\delta\|_1 \geq \theta_{10} \ln n, \quad (\text{B.41})$$

where $u_i = F(a_{i, \mu(i)} x_i)$ and $v_j = F(b_{j, \mu^{-1}(j)} y_j)$ for $i, j \in [n]$; $\mathbf{x}_\delta$ and $\mathbf{y}_\delta$ denote the truncated version of $\mathbf{x}$ and $\mathbf{y}$; $\theta_2, \theta_3, \theta_6, \theta_7, \theta_4, \theta_8, \theta_{10} \in \mathbb{R}_+$ are absolute constants (independent of μ)

chosen appropriately as in Lemmas B.2.2, B.2.3, B.2.7, B.2.8, and Corollaries B.2.4, B.2.10, and B.2.12. Then, for any $c \in (0, 1/2)$, with probability $1 - \exp(-n^c)$, $(\mathbf{X}(\mu), \mathbf{Y}(\mu)) \in \mathcal{R}^*(\mu)$ for all stable matchings μ .

The proof simply summarizes the aforementioned Lemmas and Corollaries and shall be omitted.

B.2.1 Proof of Propositions 3.6.1 and 3.6.2

Proof of Proposition 3.6.1. Note that $1 - e^{-tx} \geq \left(tx - \frac{t^2x}{2}\right)$ for all $x, t \geq 0$. In particular, $1 - e^{-tx} \geq \left(tx - \frac{t^2x}{2}\right) \mathbb{1}_{[0, 2/t]}(x) \geq 0$. Using this to approximate $p(\mathbf{x}, \mathbf{y})$ gives

$$\begin{aligned} p(\mathbf{x}, \mathbf{y}) &= \prod_{i \neq j} (1 - (1 - e^{-a_{ij}x_i})(1 - e^{-b_{ji}y_j})) \\ &\leq \prod_{i \neq j} \left(1 - \mathbb{1}_{[0, 2/a_{ij}]}(x_i) \mathbb{1}_{[0, 2/b_{ji}]}(y_j) \left(a_{ij}x_i - \frac{a_{ij}^2}{2}x_i^2\right) \left(b_{ji}y_j - \frac{b_{ji}^2}{2}y_j^2\right)\right) \\ &\leq \exp \left(- \sum_{i \neq j} \mathbb{1}_{[0, 2/C]}(x_i) \mathbb{1}_{[0, 2/C]}(y_j) \left(a_{ij}x_i - \frac{a_{ij}^2}{2}x_i^2\right) \left(b_{ji}y_j - \frac{b_{ji}^2}{2}y_j^2\right)\right). \end{aligned} \quad (\text{B.42})$$

Taking logarithm for simplicity and expanding the expression above gives

$$\begin{aligned} \ln p(\mathbf{x}, \mathbf{y}) &\leq - \sum_{i \neq j} \left(a_{ij}b_{ji}x_iy_j - (\mathbb{1}_{(2/C, \infty)}(x_i) + \mathbb{1}_{(2/C, \infty)}(y_j))a_{ij}b_{ji}x_iy_j \right. \\ &\quad \left. - \mathbb{1}_{[0, 2/C]}(x_i) \mathbb{1}_{[0, 2/C]}(y_j) \left(a_{ij}^2b_{ji}x_i^2y_j + a_{ij}b_{ji}^2x_iy_j^2\right) \right) \\ &\leq - \sum_{i, j=1}^n a_{ij}b_{ji}x_iy_j + \sum_{i=1}^n C^2x_iy_i \\ &\quad + \sum_{i, j=1}^n \left(C^2(\mathbb{1}_{(2/C, \infty)}(x_i) + \mathbb{1}_{(2/C, \infty)}(y_j))x_iy_j + C^3 \left(x_i^2y_j + x_iy_j^2\right) \right). \end{aligned} \quad (\text{B.43})$$

Notice that $-\ln q(\mathbf{x}, \mathbf{y}) = \sum_{i, j=1}^n a_{ij}b_{ji} \frac{x_i}{a_{ii}} \frac{y_j}{b_{jj}}$. Thus,

$$\begin{aligned} \ln \frac{p(\mathbf{x}, \mathbf{y})}{q(\mathbf{x}, \mathbf{y})} &\leq C^2 \mathbf{x}^\top \mathbf{y} + C^2 \left(\|\mathbf{x}\|_1 \sum_{i=1}^n \mathbb{1}_{(2/C, \infty)}(y_j)y_j + \|\mathbf{y}\|_1 \sum_{i=1}^n \mathbb{1}_{(2/C, \infty)}(x_i)x_i \right) \\ &\quad + C^3 (\|\mathbf{x}\|_2^2 \|\mathbf{y}\|_1 + \|\mathbf{x}\|_1 \|\mathbf{y}\|_2^2). \end{aligned} \quad (\text{B.44})$$

In light of Corollary B.2.13, it suffices to upper bound $\ln \frac{p(\mathbf{x}_\delta, \mathbf{y}_\delta)}{q(\mathbf{x}_\delta, \mathbf{y}_\delta)}$ by $cn/(\ln n)^{1/2}$ for all $(\mathbf{x}, \mathbf{y}) \in \mathcal{R}^*(\mu)$ and for all μ . To simplify notation, we will make the dependency on μ implicit in the rest of the proof. By Cauchy-Schwarz inequality, the first term in (B.44), up to a factor of C^2 , is at most

$$\|\mathbf{x}_\delta\|_2 \|\mathbf{y}_\delta\|_2 \leq \frac{\theta_7 \|\mathbf{u}\|_1 \|\mathbf{v}\|_1}{n} \leq \theta_3 \theta_7 (\ln n)^{1/8} = o\left(\frac{n}{(\ln n)^{1/2}}\right)$$

by (B.39) and (B.37). The middle term in (B.44), up to a factor of $2C^2$, is at most

$$\|\mathbf{x}_\delta\|_1 \sum_{i=1}^n \mathbb{1}_{(2/C, \infty)}((y_\delta)_i) (y_\delta)_i \leq \theta_6 \|\mathbf{u}\|_1 \frac{\|\mathbf{v}\|_1}{(\ln n)^{7/8}} \leq \theta_3 \theta_6 \frac{n}{(\ln n)^{3/4}}$$

by (B.38), (B.40), and (B.37). Finally, the last term, up to a factor of $2C^2$, is upper bounded by

$$\begin{aligned} \|\mathbf{x}_\delta\|_2^2 \|\mathbf{y}_\delta\|_1 &= \frac{\|\mathbf{x}_\delta\|_2^2 \|\mathbf{y}_\delta\|_1}{\|\mathbf{u}\|_1^2 \|\mathbf{v}\|_1} \frac{1}{\|\mathbf{v}\|_1} (\|\mathbf{u}\|_1 \|\mathbf{v}\|_1)^2 \\ &\leq \frac{\theta_7}{n} \cdot \theta_6 \cdot \frac{1}{\theta_2 \ln n} \cdot \theta_3^2 n^2 (\ln n)^{1/4} = \frac{\theta_7 \theta_6 \theta_3^2}{\theta_2} \frac{n}{(\ln n)^{3/4}} \end{aligned}$$

due to (B.39), (B.38), (B.36), and (B.37). Putting these together gives the proposition. $\square$

Proposition 3.6.2 is also a direct consequence of Corollary B.2.13. In fact, the conditions (3.12) and (3.13) follow directly from the conditions (B.36)–(B.38) and (B.41) in Corollary B.2.13 once we note that $\mathbf{x}_\delta^\top \mathbf{M} \mathbf{y}_\delta \leq \Theta(\|\mathbf{x}_\delta\|_1 \|\mathbf{y}_\delta\|_1)$ due to Assumption 3.2.2 on the boundedness of the entries of $\mathbf{M}$. The proof is omitted.

For convenience, we will express the conditions (3.12) and (3.13) as a likely region $\mathcal{R} \subseteq \mathbb{R}_+^n \times \mathbb{R}_+^n$ for the value vectors of the two sides. Formally, define

$$\mathcal{R}_1 = \{\mathbf{x} : \underline{c}_1 \log n \leq \|\mathbf{x}\|_1 \leq \bar{c}_1 n (\log n)^{-7/8}\} \quad (\text{B.45})$$

and

$$\mathcal{R}_2 = \{(\mathbf{x}, \mathbf{y}) : \mathbf{x}^\top \mathbf{M} \mathbf{y} \leq c_2 (\log n)^{1/8}\}, \quad (\text{B.46})$$

$$\mathcal{R} := \{(\mathbf{x}, \mathbf{y}) \in \mathcal{R}_2 : \mathbf{x}, \mathbf{y} \in \mathcal{R}_1\}, \quad (\text{B.47})$$

where $\underline{c}_1, \bar{c}_1$ and c_2 are appropriately chosen constants independent of n hidden by the $O(\cdot)$ and $\Omega(\cdot)$ notations in the Proposition statement. Then, Proposition 3.6.2 can be equivalently rephrased as

$$\mathbb{P}(\exists \mu \in \mathcal{S}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \mathcal{R}) \lesssim e^{-n^c} \quad \text{as } n \rightarrow \infty. \quad (\text{B.48})$$

B.2.2 A Key Reduction Lemma

Proposition 3.6.1 allows us to study high probability behaviors in stable partial matchings obtained from truncating stable (full) matchings pretending that the conditional probability of stability were given by $q(\mathbf{x}, \mathbf{y})$. Concretely, consider a fixed constant $\delta > 0$ and a region $\Omega \subseteq \mathbb{R}_+^n \times \mathbb{R}_+^n$ that defines an event on the (truncated) value vectors. If there exists a stable matching μ whose truncated value vectors $(\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \Omega$, the induced partial matching μ_δ of size $n - \lfloor \delta n \rfloor$ between $\mathcal{M}'_{\mu, \delta}$ and $\mu(\mathcal{M}'_{\mu, \delta})$ must also be stable with value vectors $\mathbf{X}_{\mathcal{M}'_{\mu, \delta}}(\mu_\delta) = \mathbf{X}_\delta(\mu)$ and $\mathbf{Y}_{\mu(\mathcal{M}'_{\mu, \delta})}(\mu_\delta) = \mathbf{Y}_\delta(\mu)$. Thus, we will end up either having a stable matching whose truncated value vectors fall outside $\mathcal{R}$ (defined above in (B.47)), or a stable partial matching of size $n - \lfloor \delta n \rfloor$ whose value vectors (already truncated) lie in $\Omega \cap \mathcal{R}$. By Proposition 3.6.1, the former event happens with probability $o(1)$. Therefore, we may focus on the second event, where a stable partial matching of size $n - \lfloor \delta n \rfloor$ exists with value vectors in $\Omega \cap \mathcal{R}$. This is summarized by the following Lemma, which will be a major tool in the remainder of the proof.

Lemma B.2.14. *Let $\delta > 0$, $c \in (0, 1/2)$, and $\Omega \subseteq \mathbb{R}_+^n \times \mathbb{R}_+^n$. Then,*

$$\mathbb{P}(\exists \mu \text{ stable}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \Omega) \leq e^{-n^c} + \exp\left(\Theta\left(\frac{n}{(\log n)^{1/2}}\right)\right) \cdot \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu'))]. \quad (\text{B.49})$$

Proof. Note that

$$\begin{aligned} \mathbb{P}(\exists \mu \text{ stable}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \Omega) &\leq \mathbb{P}(\exists \mu \text{ stable}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \mathcal{R}) \\ &\quad + \mathbb{P}(\exists \mu \text{ stable}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \mathcal{R} \cap \Omega), \end{aligned} \quad (\text{B.50})$$

where the first term is e^{-n^c} by Proposition 3.6.2. Let $\mathcal{E}$ denote the event that there exists a *stable* partial matching μ' between subsets $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$ with $|\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor$ where $(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \in \mathcal{R} \cap \Omega$. Clearly, the existence of a stable matching μ with $(\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \mathcal{R} \cap \Omega$ implies $\mathcal{E}$. Thus, the second term in (B.50) is bounded by

$$\mathbb{P}(\mathcal{E}) \leq \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} \mathbb{P}(\mu' \text{ stable}, (\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \in \mathcal{R} \cap \Omega). \quad (\text{B.51})$$

using union bound. For each $\mathcal{M}', \mathcal{W}'$, and μ' in the summation, we compute the above

probability through conditioning on $\mathbf{X}_{\mathcal{M}'}(\mu')$ and $\mathbf{Y}_{\mathcal{W}'}(\mu')$ as

$$\begin{aligned}
& \mathbb{P}(\mu' \text{ stable}, (\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \in \mathcal{R} \cap \Omega) \\
&= \mathbb{E}[\mathbb{P}(\mu' \text{ stable} | \mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')); (\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \in \mathcal{R} \cap \Omega] \\
&= \mathbb{E}[p_{\mu'}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu'))] \\
&\leq \mathbb{E}\left[\exp\left(\Theta\left(\frac{n}{(\log n)^{1/2}}\right)\right) q(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu'))\right] \\
&\leq \exp\left(\Theta\left(\frac{n}{(\log n)^{1/2}}\right)\right) \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu'))],
\end{aligned} \tag{B.52}$$

which completes the proof. $\square$

Remark B.2.15. *The choice of the constant $c \in (0, 1/2)$ affects the implicit constants in defining $\mathcal{R}$. Once we have a target convergence of e^{-n^c} for some $c \in (0, 1/2)$, we will assume that c is fixed in the rest of our discussion unless otherwise mentioned.*

Lemma B.2.14 will be a key tool in the proof to establish further likely behaviors of value (and rank) vectors. It will be a recurring theme where we first identify a likely region Ω_{likely} for truncated value vectors of stable (full) matchings to fall in ($\mathcal{R}$ to start with), then rule out a bad event Ω_{bad} within Ω_{likely} by showing $\Omega_{\text{likely}} \cap \Omega_{\text{bad}}$ is unlikely for value vectors of any stable *partial* matching, and apply Lemma B.2.14 to conclude that Ω_{bad} is unlikely for truncated value vectors of any stable matching and that $\Omega_{\text{likely}} \cap \Omega_{\text{bad}}$ can be used as the likely region moving forward.

Based on Lemma B.2.14, in order to rule out some unlikely event represented by a region Ω of value vectors, it suffices to study partial matchings μ' of size $n - \lfloor \delta n \rfloor$ between $\mathcal{M}'$ and $\mathcal{W}'$ and upper bound $\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{X}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\Omega}(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{X}_{\mathcal{W}'}(\mu'))]$. From now on, we fix $\mathcal{M}'$, $\mathcal{W}'$, and μ' , and make the dependency of $\mathbf{X}_{\mathcal{M}'}$ and $\mathbf{Y}_{\mathcal{W}'}$ on μ' implicit when the context is clear.

B.3 Number of Stable Matchings: Proof of Proposition 3.6.3

In this section, we prove Proposition 3.6.3. The proof relies on the following technical lemma, whose proof is involved and is deferred to Appendix B.5. Recall from Appendix B.2.1 the definitions of likely regions $\mathcal{R}_1$ and $\mathcal{R}_2$ for value vectors ((B.45) and (B.46)).

Lemma B.3.1. *For a fixed a partial matching μ' on $\mathcal{M}'$ and $\mathcal{W}'$ of size $n - \lfloor \delta n \rfloor$,*

$$\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})] \leq e^{o(n) + o_{\delta}(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}. \tag{B.53}$$

Assuming Lemma B.3.1, we now present the proof of Proposition 3.6.3.

Proof of Proposition 3.6.3. We begin by reducing the bound on the total number of stable matchings N to a bound on the number N_δ of stable partial matchings of size $n - \lfloor \delta n \rfloor$. We claim that

$$N \leq e^{O(\delta n)} \cdot N_\delta = \exp(o_\delta(n)) \cdot N_\delta$$

with probability 1. To see this, fix any stable partial matching μ' defined on subsets $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$, both of size $n - \lfloor \delta n \rfloor$. Any stable full matching μ that agrees with μ' on $\mathcal{M}'$ must also be stable on the remaining sub-market $(\mathcal{M} \setminus \mathcal{M}', \mathcal{W} \setminus \mathcal{W}')$ of size $\lfloor \delta n \rfloor$. Therefore, the number of full matchings consistent with a given μ' is at most the number of stable matchings in this smaller sub-market. As is shown by Palmer and Pálvölgyi [166], the number of stable matchings in any market of size $\lfloor \delta n \rfloor$ is at most $e^{3.55\delta n}$. It follows that each stable partial matching can extend to at most $e^{O(\delta n)}$ full stable matchings. Hence, multiplying by N_δ yields the claimed bound.

We now show the high-probability bound on N_δ . Observe that for each $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$ with $|\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor$ and partial matching μ' between $\mathcal{M}'$ and $\mathcal{W}'$,

$$\begin{aligned} \mathbb{P}(\mu' \text{ is stable and satisfies } \mathcal{R}) &\leq e^{o(n)} \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \leq \\ &e^{o(n)} \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})] \leq e^{o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i} \end{aligned}$$

by Proposition B.3.1. Summing over $\mathcal{M}'$, $\mathcal{W}'$, and μ' bounds the expected number of such stable partial matchings above by

$$\begin{aligned} \mathbb{E}[N_\delta] &\leq \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} e^{o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i} \\ &\stackrel{(i)}{=} \frac{1}{\lfloor \delta n \rfloor!} \sum_{\substack{\mu: \mathcal{M} \rightarrow \mathcal{W} \\ \text{bijection}}} \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M} \\ |\mathcal{M}'| = n - \lfloor \delta n \rfloor}} e^{o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu(i)} b_{\mu(i), i} \\ &\stackrel{(ii)}{\leq} \sum_{\substack{\mu: \mathcal{M} \rightarrow \mathcal{W} \\ \text{bijection}}} \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M} \\ |\mathcal{M}'| = n - \lfloor \delta n \rfloor}} e^{o_\delta(n)} \frac{1}{n!} \cdot C^{2\lfloor \delta n \rfloor} \prod_{i \in \mathcal{M}} a_{i, \mu(i)} b_{\mu(i), i} \\ &\stackrel{(iii)}{\leq} e^{o_\delta(n)} \binom{n}{\lfloor \delta n \rfloor} \cdot \frac{1}{n!} \text{Perm}(\mathbf{A} \circ \mathbf{B}^\top) \\ &\stackrel{(iv)}{\leq} e^{o_\delta(n)}, \end{aligned}$$

where in (i) we use an alternative counting of partial matchings by counting sub-matchings of size $n - \lfloor \delta n \rfloor$ in full matchings and then deduplicate by a factor of $\lfloor \delta n \rfloor!$; in (ii) we use the boundedness assumption on the components of $\mathbf{A}$ and $\mathbf{B}$; in (iii) we merge $C^{2\lfloor \delta n \rfloor}$ into $e^{o_\delta(n)}$;

and finally in (iv) we merge $\binom{n}{\lfloor \delta n \rfloor} = \exp(h(\delta)n + o(n))$ into $e^{o_\delta(n)}$ and bound the permanent term by $n^n \text{Perm}(\mathbf{M}) \leq \Theta(n!)$ using the moderate deviation property of $\mathbf{M}$ [152, Sec. 3]. $\square$

B.4 Proof of Proposition 3.7.1

In this section, we prove Proposition 3.7.1, which will be an immediate corollary to the following Lemma, whose proof is similar to that of Lemma B.2.14.

Lemma B.4.1. *Let μ' be a partial matching of size $n - \lfloor \delta n \rfloor$ on $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$. Fix any $\zeta > 0$, and let*

$$\Omega_{\text{eig}}(\zeta) := \left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}^n \times \mathbb{R}^n : \exists t \in \mathbb{R}_+, \sum_{i=1}^n \mathbb{1} \left\{ |(\mathbf{M}\mathbf{y})_i - t| \geq \sqrt{\zeta}t \right\} \leq \sqrt{\zeta}n \right\}.$$

Then

$$\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \setminus \Omega_{\text{eig}}(\Theta(\delta) + \zeta)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \leq \exp(o_\delta(n) - \Theta(\zeta^2 n)) \cdot \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}, \quad (\text{B.54})$$

with the implicit constants uniform over all $\mathcal{M}', \mathcal{W}'$, and μ' .

Proof. Denote the expectation in (B.54) by E , and express it as

$$\begin{aligned} E &= \int_0^\infty \mathbb{P}(q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \setminus \Omega_{\text{eig}}(\zeta)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) > s) ds \\ &= \int_0^1 \mathbb{P}(\exp(-n \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} > s, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \mathcal{R} \setminus \Omega_{\text{eig}}(\zeta)) ds \\ &= \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \mathcal{R}_2 \setminus \Omega_{\text{eig}}(\zeta), \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot n e^{-nt} dt \\ &= \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_{\text{eig}}(\zeta), \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot n e^{-nt} dt, \end{aligned} \quad (\text{B.55})$$

where $\bar{t} := t \wedge (c_2(\log n)^{1/8})$. If we can find two families of regions $\Omega_1(\zeta; s), \Omega_2(\zeta; s) \subseteq \mathbb{R}_+^n \times \mathbb{R}_+^n$ such that $\Omega_{\text{eig}}(\zeta) \supseteq \Omega_1(\Theta(\zeta); s) \cap \Omega_2(\Theta(\zeta); s)$ for all $0 < s < c_2(\log n)^{1/8}$, by union bound

and relaxing the requirement that $\mathbf{X}_{\mathcal{M}'}$ (resp. $\mathbf{Y}_{\mathcal{W}'}$) is in $\mathcal{R}_1$, we will obtain

$$\begin{aligned}
E &\leq \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_1(\Theta(\zeta); \bar{t}), \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt \\
&\quad + \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_2(\Theta(\zeta); \bar{t}), \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt \\
&\leq \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_1(\Theta(\zeta); \bar{t}), \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt \\
&\quad + \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_2(\Theta(\zeta); \bar{t}), \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt.
\end{aligned}$$

Rewriting the probabilities through conditioning and further relaxing the requirement gives

$$\begin{aligned}
E &\leq \int_0^\infty \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_1(\Theta(\zeta); \bar{t}) | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1) \\
&\quad \cdot \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt \\
&\quad + \int_0^\infty \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_2(\Theta(\zeta); \bar{t}) | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \\
&\quad \cdot \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt.
\end{aligned} \tag{B.56}$$

Due to the symmetry between the two integrals, it then suffices to bound one of the two integrals (e.g., the latter) by showing

$$\sup_{0 < t < c_2(\log n)^{1/8}} \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \notin \Omega_2(\Theta(\zeta); t) | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t, \mathbf{X}_{\mathcal{M}'} \in \mathcal{R}_1) \leq \exp(-\Theta(\zeta^2 n)) \tag{B.57}$$

and

$$\int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt \leq e^{o(n) + o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}, \tag{B.58}$$

from which the desired upper bound immediately follows. Recognizing that

$$\int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot ne^{-nt} dt = \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})),$$

we reduce (B.58) to Proposition B.3.1. Our road map is to first find the desirable choices for Ω_1 and Ω_2 and establish (B.57), and then prove Proposition B.3.1 in Appendix B.5. Note that Proposition B.3.1 is in fact independent of our choice of Ω_1 and Ω_2 , but we will develop useful intermediate results to prepare for its proof.

Concretely, we consider events Ω_1 and Ω_2 as follows:

$$\Omega_1(\zeta; t) := \left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \max_{i \in [n]} \left| (1 - \delta) \frac{\mathbf{M}_{i, \cdot} \cdot \mathbf{y}}{t \mathbf{M}_{i, \cdot} \cdot (\mathbf{M}^\top \mathbf{x})_{\mathcal{W}'}^{-1}} - 1 \right| > \zeta \right\}, \tag{B.59}$$

$$\Omega_2(\zeta; t) := \left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \max_{j \in [n]} \left| (1 - \delta) \frac{\mathbf{M}_{\cdot, j} \cdot \mathbf{x}}{t \mathbf{M}_{\cdot, j} \cdot (\mathbf{M} \mathbf{y})_{\mathcal{M}'}}^{-1} - 1 \right| > \zeta \right\}, \quad (\text{B.60})$$

where $\mathbf{M}_{i, \cdot}$ and $\mathbf{M}_{\cdot, j}$ denote the i -th row and the j -th column of $\mathbf{M}$, respectively; inverse is applied coordinate-wise on vectors; and $\mathbf{v}_S$ denotes the n -dimensional vector obtained by zeroing out the i -th component v_i of $\mathbf{v} \in \mathbb{R}_+^n$ for all $i \in [n] \setminus S$ (with this operation performed after coordinate-wise inverse). We first verify the following lemma.

Lemma B.4.2. *There exist absolute constants $\zeta_0, \delta_0 > 0$ and $k_1, k_2 > 0$ such that for all $\zeta \in (0, \zeta_0)$, $\delta \in (0, \delta_0)$, and $t > 0$ we have*

$$\Omega_1(\zeta; t) \cap \Omega_2(\zeta; t) \subseteq \Omega_{\text{eig}}(k_1 \delta + k_2 \zeta).$$

Proof. Let $\mathbf{d} = \mathbf{M}^\top \mathbf{x}$ and $\mathbf{e} = \mathbf{M} \mathbf{y}$. Under the event that $(\mathbf{x}, \mathbf{y}) \in \Omega_1(\zeta; t) \cap \Omega_2(\zeta; t)$, we have

$$\begin{aligned} \frac{1}{d_j} &= \frac{1}{\mathbf{M}_{\cdot, j} \cdot \mathbf{x}} \stackrel{(i)}{\leq} \frac{1 - \delta}{(1 - \zeta) t \mathbf{M}_{\cdot, j} \cdot \mathbf{e}_{\mathcal{M}'}}^{-1} \stackrel{(ii)}{\leq} \frac{1 - \delta}{(1 - \zeta) t (1 + 2C^2 \delta) \mathbf{M}_{\cdot, j} \cdot \mathbf{e}^{-1}} \\ &\stackrel{(iii)}{\leq} \frac{1 - \delta}{(1 - \zeta) t (1 + 2C^2 \delta)} \sum_{i=1}^n m_{ij} e_i = \frac{1 - \delta}{(1 - \zeta) t (1 + 2C^2 \delta)} \sum_{i=1}^n m_{ij} \mathbf{M}_{i, \cdot} \cdot \mathbf{y} \\ &\stackrel{(iv)}{\leq} \frac{1 - \delta}{(1 - \zeta) t (1 + 2C^2 \delta)} \sum_{i=1}^n m_{ij} \left(\frac{1 + \zeta}{1 - \delta} t \mathbf{M}_{i, \cdot} \cdot \mathbf{d}_{\mathcal{W}'}^{-1} \right) \leq \frac{1 + \zeta}{(1 - \zeta) (1 + 2C^2 \delta)} (\mathbf{M}^\top \mathbf{M} \mathbf{d}^{-1})_j, \end{aligned} \quad (\text{B.61})$$

where (i) uses the definition of $\Omega_2(\zeta; t)$; (ii) uses the fact that $\mathbf{M}$ and $\mathbf{e}$ both have bounded ratios (at most C) among their entries and assumed $\delta < 1/2$; (iii) is due to Jensen's inequality (or equivalently, harmonic-mean-arithmetic-mean inequality); and (iv) uses the definition of $\Omega_1(\zeta; t)$.

Recall our assumption that $\mathbf{M}$ has entries bounded on $[1/(Cn), C/n]$. It is straightforward to verify that for any vector $\mathbf{v} \in \mathbb{R}_+^n$ with $\bar{v} = \frac{1}{n} \sum_{i=1}^n v_i$, we have $\max_{i \in [n]} (\mathbf{M} \mathbf{v})_i - \bar{v} \leq \max_{i \in [n]} v_i - \frac{1}{Cn} \cdot n(\max_{i \in [n]} v_i - \bar{v}) - \bar{v} = (1 - C^{-1})(\max_{i \in [n]} v_i - \bar{v})$. In the case of $\mathbf{v} = \mathbf{d}^{-1}$, this implies that

$$\begin{aligned} (1 - C^{-1})^2 (d_{i^*}^{-1} - \bar{d}_{(H)}^{-1}) &\geq \max_{i \in [n]} (\mathbf{M}^\top \mathbf{M} \mathbf{d}^{-1})_i - \bar{d}_{(H)}^{-1} \\ &\geq (\mathbf{M}^\top \mathbf{M} \mathbf{d}^{-1})_{i^*} - \bar{d}_{(H)}^{-1} \geq (1 + 2C^2 \delta) \frac{1 - \zeta}{1 + \zeta} d_{i^*}^{-1} - \bar{d}_{(H)}^{-1}, \end{aligned} \quad (\text{B.62})$$

where $i^* = \operatorname{argmin}_{i \in [n]} d_i$ and $\bar{d}_H = (n^{-1} \sum_{i=1}^n d_i^{-1})^{-1}$ is the harmonic mean of $d_1, \dots, d_n$. Solving (B.62) gives

$$\frac{d_{i^*}^{-1} - \bar{d}_{(H)}^{-1}}{\bar{d}_{(H)}^{-1}} \leq \Theta(\delta) + \frac{2\zeta}{1 - \zeta - (1 - C^{-2})^2(1 + \zeta)} \leq \Theta(\delta + \zeta)$$

with hidden constants independent of δ and ζ , granted that ζ is sufficiently small. Hence, for all but $\sqrt{\delta + \zeta}n$ indices $i \in [n]$, we have $1 - \Theta(\sqrt{\delta + \zeta}) \leq \frac{d_{(H)}}{d_i} \leq 1 + \Theta(\delta + \zeta)$, implying that $(\mathbf{x}, \mathbf{y}) \in \Omega_{\text{eig}}(\Theta(\delta + \zeta))$. $\square$

Let $\mathbf{D} = \mathbf{M}^\top \mathbf{X}_{\mathcal{M}'}$ and $\mathbf{E} = \mathbf{M} \mathbf{Y}_{\mathcal{W}'}$. Note that $\mathbf{D}$ and $\mathbf{E}$ both have bounded ratios among their components due to the bounded ratio assumption on $\mathbf{M}$, and in addition $\|\mathbf{D}\|_1 = \|\mathbf{X}_{\mathcal{M}'}\|_1$ and $\|\mathbf{E}\|_1 = \|\mathbf{Y}_{\mathcal{W}'}\|_1$. By Lemma B.1.2, whenever $t \leq c_2(\log n)^{1/8}$, we have for each column $\mathbf{M}_{\cdot j}$ of $\mathbf{M}$, $j = 1, \dots, n$,

$$\mathbb{P} \left(\left| (1 - \delta) \frac{\mathbf{M}_{\cdot j} \cdot \mathbf{X}_{\mathcal{M}'}}{t \mathbf{M}_{\cdot j} \cdot \mathbf{E}_{\mathcal{M}'}} - 1 \right| > \zeta \mid \mathbf{X}_{\mathcal{M}'} \cdot \mathbf{E} < t, \|\mathbf{E}\|_1 \geq \underline{c}_1 \log n \right) \leq \exp(-\Theta(n\zeta^2)),$$

where we note that the effective dimension of $\mathbf{X}_{\mathcal{M}'}$ is $n - \lfloor \delta n \rfloor$ instead of n . By a union bound over $j \in [n]$, this gives

$$\mathbb{P} \left(\max_{j \in [n]} \left| (1 - \delta) \frac{\mathbf{M}_{\cdot j} \cdot \mathbf{X}_{\mathcal{M}'}}{t \mathbf{M}_{\cdot j} \cdot \mathbf{E}_{\mathcal{M}'}} - 1 \right| > \zeta, \|\mathbf{E}\|_1 \geq \underline{c}_1 \log n \mid \mathbf{X}_{\mathcal{M}'} \cdot \mathbf{E} < t \right) \leq \exp(-\Theta(n\zeta^2)),$$

which is simply (B.57). $\square$

We now prove Proposition 3.7.1.

Proof. Summing over all partial matchings with size $n - \lfloor \delta n \rfloor$ gives

$$\begin{aligned} & \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}(\mu'), \mathbf{Y}_{\mathcal{W}'}(\mu')) \cdot \mathbb{1}_{\mathcal{R} \setminus \Omega_{\text{eig}}(\zeta)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \\ & \leq \exp(o_\delta(n) - \Theta(\zeta^2 n)) \cdot \frac{(\delta n)!}{n!} \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} \prod_{i \in \mathcal{M}'} (nm_{i, \mu'(i)}). \quad (\text{B.63}) \end{aligned}$$

To bound the summation, notice that

$$\begin{aligned} & \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M}, \mathcal{W}' \subseteq \mathcal{W} \\ |\mathcal{M}'| = |\mathcal{W}'| = n - \lfloor \delta n \rfloor}} \sum_{\substack{\mu': \mathcal{M}' \rightarrow \mathcal{W}' \\ \text{bijection}}} \prod_{i \in \mathcal{M}'} (nm_{i, \mu'(i)}) = \frac{1}{(\lfloor \delta n \rfloor)!} \sum_{\substack{\mu: \mathcal{M} \rightarrow \mathcal{W} \\ \text{bijection}}} \sum_{\substack{\mathcal{M}' \subseteq \mathcal{M} \\ |\mathcal{M}'| = n - \lfloor \delta n \rfloor}} \prod_{i \in \mathcal{M}'} (nm_{i, \mu'(i)}) \\ & \leq \frac{1}{(\lfloor \delta n \rfloor)!} \sum_{\substack{\mu: \mathcal{M} \rightarrow \mathcal{W} \\ \text{bijection}}} \binom{n}{\lfloor \delta n \rfloor} C^{\lfloor \delta n \rfloor} \prod_{i \in \mathcal{M}} (nm_{i, \mu(i)}) \\ & = \frac{e^{o_\delta(n)}}{(\delta n)!} \cdot n^n \text{Perm}(\mathbf{M}). \quad (\text{B.64}) \end{aligned}$$

Under the assumption that the bistochastic matrix $\mathbf{M}$ is of moderate deviation (cf. [152, Section 3]), we know that $n^n \text{Perm}(\mathbf{M}) = O(n!)$. Hence, the quantity in (B.63) is bounded by $\exp(o_\delta(n) - \Theta(\zeta^2 n))$. Invoking Lemma B.2.14 finishes the proof. $\square$

B.5 Proof of Proposition B.3.1

In this section, we present the proof of Proposition B.3.1, restated below.

Lemma B.3.1. *For a fixed a partial matching μ' on $\mathcal{M}'$ and $\mathcal{W}'$ of size $n - \lfloor \delta n \rfloor$,*

$$\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})] \leq e^{o(n)+o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}. \quad (\text{B.53})$$

Denote the target expectation by E and express it as an integral of tail probability

$$E = \int_0^\infty \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1) \cdot n e^{-nt} dt, \quad (\text{B.65})$$

where $\bar{t} = t \wedge (c_2(\log n)^{1/8})$. It suffices to upper bound probabilities of the form $\mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < \bar{t}, \mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1)$ for all $t \in (0, c_2(\log n)^{1/8})$. We will go one step further and prove a stronger result by relaxing the $\mathbf{Y}_{\mathcal{W}'} \in \mathcal{R}_1$ condition, which will eventually translate to a bound on $\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})]$.

Lemma B.5.1. *There exists a positive constant β such that, for any $t \in (0, c_2(\log n)^{1/8})$,*

$$\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn}, \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn} \mid \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t\right) \leq 0.1$$

for n sufficiently large.

Proof. Let p denote the target probability. We have

$$\begin{aligned} p &= \mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn} \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}, \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t\right) \\ &\quad \cdot \mathbb{P}\left(\|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn} \mid \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t\right) \\ &\leq \mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn} \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}, \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t\right) \\ &= \frac{\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn}, \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}\right)}{\mathbb{P}\left(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}\right)} \\ &\leq \frac{\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn} \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}\right)}{\mathbb{P}\left(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \beta\sqrt{tn}\right)} \\ &\leq \frac{\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn}\right)}{\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \sqrt{tn}/(C\beta)\right)} = \mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn} \mid \|\mathbf{X}_{\mathcal{M}'}\|_1 \leq (C\beta)^{-1}\sqrt{tn}\right), \quad (\text{B.66}) \end{aligned}$$

where the last inequality follows from the independence between $\mathbf{X}$ and $\mathbf{Y}$ and the fact that $n\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} \leq C\|\mathbf{X}_{\mathcal{M}'}\|_1 \|\mathbf{Y}_{\mathcal{W}'}\|_1$. By choosing $\beta = (2C)^{-1/2}$, the upper bound becomes

$$\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \beta\sqrt{tn} \mid \|\mathbf{X}_{\mathcal{M}'}\|_1 \leq 2\beta\sqrt{tn}\right).$$

A direct invocation of Lemma B.1.2 implies an $\exp(-\Theta(n))$ upper bound for this probability. $\square$

Lemma B.5.2. *There exists a positive constant γ such that, for any $t \in (0, c_2(\log n)^{1/8})$,*

$$\mathbb{P}(\|\mathbf{Y}_{\mathcal{W}'}\|_1 \geq \gamma n(\log n)^{-7/8} | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t) \leq 0.1 \quad (\text{B.67})$$

Proof. To free ourselves from always carrying the notation for the partial matching, let us observe that, once we relinquish the condition on the bistochasticity of $\mathbf{M}$ (i.e., we will only use the boundedness of the entries in the matrix), it becomes irrelevant that μ' is a partial matching between $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$ (instead of a complete one between $\mathcal{M}$ and $\mathcal{W}$), since the difference in the market size $|\mathcal{M}| = |\mathcal{W}| = n$ and $|\mathcal{M}'| = |\mathcal{W}'| = n - \delta n$ does not affect the final asymptotics in the Lemma. Hence, it suffices to establish a version of (B.67) with $\mathbf{X}_{\mathcal{M}'}$ and $\mathbf{Y}_{\mathcal{W}'}$ replaced by non-truncated value vectors $\mathbf{X}$ and $\mathbf{Y}$ in a complete (instead of partial) matching μ , as long as we do not rely on bistochasticity of $\mathbf{M}$.

In the simplified notation, let $\mathbf{Z} = \mathbf{a} \circ \mathbf{X}$ and $\mathbf{W} = \mathbf{b} \circ \mathbf{Y}$ with $\mathbf{a} = (a_{i,\mu(i)})_{i \in [n]}$ and $\mathbf{b} = (b_{j,\mu^{-1}(j)})_{j \in [n]}$, so that $\mathbf{Z}, \mathbf{W} \sim \text{Exp}(1)^n$ and are independent. Moreover, let $R = \|\mathbf{Z}\|_1$ and $\mathbf{U} = R^{-1}\mathbf{Z}$ so that, as is well known, $R \sim \Gamma(n, 1)$, $\mathbf{U} \sim \text{Unif}(\Delta_{n-1})$, and R and $\mathbf{U}$ are independent. Similarly, let $S = \|\mathbf{W}\|_1$, and $\mathbf{V} = S^{-1}\mathbf{W}$. Then

$$\mathbf{X}^\top \mathbf{M} \mathbf{Y} = \mathbf{Z}^\top \text{diag}(\mathbf{a}^{-1}) \mathbf{M} \text{diag}(\mathbf{b}^{-1}) \mathbf{W} = R S \mathbf{U}^\top \tilde{\mathbf{M}} \mathbf{V},$$

where $\tilde{\mathbf{M}} := \text{diag}(\mathbf{a}^{-1}) \mathbf{M} \text{diag}(\mathbf{b}^{-1})$ again has entries bounded on $[1/(C^2 n), C^2/n]$. Since $\|\mathbf{Y}\|_1 = \Theta(S)$, it suffice to find a positive constant γ such that

$$\mathbb{P}(S \geq \gamma n(\log n)^{-7/8} | R S \mathbf{U}^\top \tilde{\mathbf{M}} \mathbf{V} < t) < 0.1$$

for all n sufficiently large and $t \in (0, c_2(\log n)^{1/8})$. Note that $1/(C^2 n) \leq \mathbf{U}^\top \tilde{\mathbf{M}} \mathbf{V} \leq C^2/n$ a.s. By conditional on all possible values of $\mathbf{U}^\top \tilde{\mathbf{M}} \mathbf{V}$, it suffices to show that for all $t' \in (0, c_2 C^2(\log n)^{1/8})$

$$\mathbb{P}(S \geq \gamma n(\log n)^{-7/8} | R S < t' n) < 0.1 \quad (\text{B.68})$$

asymptotically.

First, we write $\mathbb{P}(S \geq \gamma n(\log n)^{-7/8}, R S < t' n)$ explicitly as

$$\mathbb{P}(S \geq \gamma n(\log n)^{-7/8}, R S < t' n) = \int_{\gamma n(\log n)^{-7/8}}^{\infty} G(t' n/s) g(s) ds,$$

where $g(x) = \frac{x^{n-1} e^{-x}}{(n-1)!}$ is the probability density function of $\Gamma(n, 1)$ and G is the corresponding CDF. Since $t' n/s \ll n$, we may use Lemma B.1.1 to upper bound $G(t' n/s)$, giving

$$\begin{aligned} \mathbb{P}(S \geq \gamma n(\log n)^{-7/8}, R S < t' n) &\leq \int_{\gamma n(\log n)^{-7/8}}^{\infty} \left(\frac{t' e}{s} \right)^n \frac{s^{n-1} e^{-s}}{(n-1)!} ds \\ &= \frac{(t' e)^n}{(n-1)!} \int_{\gamma n(\log n)^{-7/8}}^{\infty} \frac{e^{-s}}{s} ds \\ &\leq \frac{(t' e)^n}{(n-1)!} e^{-\gamma n(\log n)^{-7/8}}. \end{aligned} \quad (\text{B.69})$$

Next, we lower bound $\mathbb{P}(RS < t'n)$ by

$$\mathbb{P}(RS < t'n) \geq \mathbb{P}(n^{1/2} \leq S \leq n^{2/3}, RS < t'n) = \int_{n^{1/2}}^{n^{2/3}} G(t'n/s)g(s)ds.$$

Note that $t'n/s \leq O(n^{2/3})$ for all $s \in [n^{1/2}, n^{2/3}]$. Using the lower bound in Lemma B.1.1, we have

$$\begin{aligned} \mathbb{P}(n^{1/2} \leq S \leq n^{2/3}, RS < t'n) &\geq e^{-O(n^{2/3})} \int_{n^{1/2}}^{n^{2/3}} \left(\frac{t'e}{s}\right)^n \frac{s^{n-1}e^{-s}}{(n-1)!} ds \\ &= e^{-O(n^{2/3})} \frac{(t'e)^n}{(n-1)!} \int_{n^{1/2}}^{n^{2/3}} \frac{e^{-s}}{s} ds \\ &\geq \frac{(t'e)^n}{(n-1)!} n^{-2/3} e^{-O(n^{2/3})}. \end{aligned} \quad (\text{B.70})$$

Comparing (B.69) with (B.70), we see that $\mathbb{P}(S \geq \gamma n(\log n)^{-7/8}, RS < t'n) / \mathbb{P}(RS < t'n) \rightarrow 0$ as $n \rightarrow \infty$, which establishes (B.68) and hence finishes the proof. $\square$

Remark B.5.3. Note that this lemma should be treated only as a technical result about the typical behavior of $\mathbf{X}_{\mathcal{M}'}$ and $\mathbf{Y}_{\mathcal{W}'}$ when $q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) = \exp(-n\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'})$ is large, and should not be confused with any attempt to bound the number of stable (partial) matchings with women's total values in a certain range. For example, one might hope to replace $\gamma n(\log n)^{-7/8}$ with $\gamma n(\log n)^{-1}$ in the proof to conclude that stable matchings with $\|\mathbf{Y}_\delta\|_1 \in [n^{1/2}, n^{2/3}]$ are over $e^{n^{2/3}}$ times more common than those with $\|\mathbf{Y}_\delta\|_1 \geq \Omega(n(\log n)^{-1})$. This, however, is generally not true as we know in the classic case with uniformly random preferences. To see why this fact is not contradictory to our proof, recall from Proposition 3.6.1 (see Section 3.5 and Appendix B.2) that $q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})$ is only a good approximation to $p_{\mu'}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})$ when, among other conditions, $\|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \Theta(n(\log n)^{-7/8})$; even then, the approximation is only valid up to an $e^{o(n)}$ factor. As the ratio between (B.69) and (B.70) is only $e^{o(n)}$, the quality of approximation is insufficient for us to rule out the possibility for a (partial) stable matching to have $\|\mathbf{Y}_{\mathcal{W}'}\|_1 \geq \Theta(n(\log n)^{-1})$: the man-optimal stable matching obtained from the man-proposing deferred acceptance algorithm will be such an example.

Corollary B.5.4. There exists a positive constant γ' such that, for any $t \in (0, c_2(\log n)^{1/8})$,

$$\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \gamma' t(\log n)^{7/8}, \|\mathbf{Y}_{\mathcal{W}'}\|_1 \geq \beta \sqrt{tn} \mid \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t\right) \leq 0.2 \quad (\text{B.71})$$

for n sufficiently large, where β is the constant appearing in Lemma B.5.1.

Proof. Note that $\|\mathbf{M} \mathbf{Y}_{\mathcal{W}'}\|_1 = \|\mathbf{Y}_{\mathcal{W}'}\|_1 \gtrsim \sqrt{tn} \gg t$ for $t \lesssim (\log n)^{1/8}$. For any $\mathbf{y}$ supported on coordinates indexed by $\mathcal{W}'$ with $t \ll \|\mathbf{y}\|_1 \leq \gamma n(\log n)^{-7/8}$, Lemma B.1.2 implies

$$\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq 0.9 \frac{t}{n - \lfloor \delta n \rfloor} \mid \left\|(\mathbf{M} \mathbf{y})_{\mathcal{M}'}^{-1}\right\|_1 \mid \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{Y}_{\mathcal{W}'} < t, \mathbf{Y}_{\mathcal{W}'} = \mathbf{y}\right) \leq e^{-\Theta(n)}. \quad (\text{B.72})$$

Plugging in $\|(\mathbf{M}\mathbf{y})_{\mathcal{M}'}^{-1}\|_1 \geq (n - \lfloor \delta n \rfloor) \frac{n}{C\|\mathbf{y}\|_1} \geq \frac{n - \lfloor \delta n \rfloor}{C\gamma} (\log n)^{7/8}$ gives

$$\mathbb{P}(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq 0.9(C\gamma)^{-1}t(\log n)^{7/8} | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, \mathbf{Y}_{\mathcal{W}'} = \mathbf{y}) \leq e^{-\Theta(n)}. \quad (\text{B.73})$$

Marginalizing over all relevant values of $\mathbf{y}$ implies

$$\mathbb{P}(\|\mathbf{X}_{\mathcal{M}'}\|_1 \leq \gamma' t (\log n)^{7/8}, \beta \sqrt{tn} \leq \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \gamma n (\log n)^{-7/8} | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t) \leq e^{-\Theta(n)} \quad (\text{B.74})$$

with $\gamma' = 0.9(C\gamma)^{-1}$. Combining this with Lemma B.5.2 through a union bound completes the proof. $\square$

Corollary B.5.5. *For any $t \in (0, c_2(\log n)^{1/8})$ and n sufficiently large,*

$$\mathbb{P}(\|\mathbf{X}_{\mathcal{M}'}\|_1 \wedge \|\mathbf{Y}_{\mathcal{W}'}\|_1 \geq \gamma' t (\log n)^{7/8}, \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t) \geq \frac{1}{2} \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t). \quad (\text{B.75})$$

Proof. Note that $\gamma' t (\log n)^{7/8} \ll \beta \sqrt{tn}$ for $t < c_2(\log n)^{1/8}$. The Corollary follows directly from Lemma B.5.1 and Corollary B.5.4 (and its symmetric version with the roles of $\mathbf{X}$ and $\mathbf{Y}$ swapped). $\square$

We are now ready to state the proof of Lemma B.3.1.

Proof of Lemma B.3.1. Define events

$$\begin{aligned} A_1(t) : \|\mathbf{X}_{\mathcal{M}'}\|_1 &\geq \gamma' t (\log n)^{1/8}, & A_2(t) : \|\mathbf{Y}_{\mathcal{W}'}\|_1 &\geq \gamma' t (\log n)^{1/8}, \\ B_1(t) : (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) &\in \Omega_1(\zeta; t), & B_2(t) : (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) &\in \Omega_2(\zeta; t), \end{aligned}$$

where Ω_1 and Ω_2 are defined in (B.59) and (B.60) in Appendix B.4, and $\zeta = o_\delta(1)$ is to be specified later. From Corollary B.5.5, we have

$$\begin{aligned} \frac{1}{2} \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t) &\leq \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), A_2(t)) \\ &\leq \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), A_2(t), B_1(t), B_2(t)) \\ &\quad + \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), B_1(t)^c) \\ &\quad + \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_2(t), B_2(t)^c) \\ &\leq \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), A_2(t), B_1(t), B_2(t)) \\ &\quad + \mathbb{P}(B_1(t)^c | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t)) \cdot \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t) \\ &\quad + \mathbb{P}(B_2(t)^c | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_2(t)) \cdot \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t). \quad (\text{B.76}) \end{aligned}$$

For any fixed $\delta > 0$ and $\zeta = o_\delta(1)$, $\mathbb{P}(B_1(t)^c | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t)) \rightarrow 0$ by Lemma B.1.2; in particular, for n sufficiently large, this term (and by symmetry, the $\mathbb{P}(B_2(t)^c | \mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_2(t))$ term) is at most $1/8$. Thus,

$$\mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t) \leq 4 \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), A_2(t), B_1(t), B_2(t)). \quad (\text{B.77})$$

By Lemma B.4.2, $B_1(t)$ and $B_2(t)$ together imply that

$$n\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} = (1 + o_\delta(1))\|\mathbf{X}_{\mathcal{M}'}\|_1\|\mathbf{Y}_{\mathcal{W}'}\|_1.$$

Further, along with the events $\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t$, $A_1(t)$, and $A_2(t)$, they imply $t(\log n)^{1/8} \lesssim \|\mathbf{Y}_{\mathcal{W}'}\|_1 \lesssim n(\log n)^{-1/8}$. Hence,

$$\begin{aligned} & \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M}\mathbf{Y}_{\mathcal{W}'} < t, A_1(t), A_2(t), B_1(t), B_2(t)) \\ & \leq \mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1\|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \frac{nt}{1 + o_\delta(1)}, t(\log n)^{1/8} \leq \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \Theta(n(\log n)^{-1/8})\right) \\ & \leq \mathbb{E}\left[\mathbb{P}\left(\|\mathbf{X}_{\mathcal{M}'}\|_1\|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \frac{nt}{1 + o_\delta(1)}, t(\log n)^{1/8} \leq \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \Theta(n(\log n)^{-1/8}) \mid \|\mathbf{Y}_{\mathcal{W}'}\|_1\right)\right] \\ & \leq e^{o_\delta(n)} \left(\frac{ent}{n - \lfloor \delta n \rfloor}\right)^{n - \lfloor \delta n \rfloor} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} \\ & \quad \cdot \mathbb{E}\left[\|\mathbf{Y}_{\mathcal{W}'}\|_1^{-n + \lfloor \delta n \rfloor}; t(\log n)^{1/8} \leq \|\mathbf{Y}_{\mathcal{W}'}\|_1 \leq \Theta(n(\log n)^{-1/8})\right]. \end{aligned} \quad (\text{B.78})$$

It is straightforward, albeit a bit tedious, to explicitly bound the expectation term in (B.78) by

$$(\Theta(\log n) - \log t)e^{n - \lfloor \delta n \rfloor} \prod_{i \in \mathcal{M}'} b_{\mu'(i), i},$$

again using Lemma B.1.1. Carrying out the integral over t in (B.65) finishes the proof. $\square$

B.6 Empirical Distribution of Values: Proof of Theorem 3.8.1

We now present the proof of Theorem 3.8.1.

Again, the key to the proof is to establish the following technical lemma. The proof of this lemma formalizes the intuition that, conditional on stability of μ' and $\mathbf{Y}_{\mathcal{W}'} = \mathbf{y}$, the value X_i for $i \in \mathcal{M}'$ should behave approximately as $\text{Exp}(\lambda_i)$ for some $\lambda_i = (1 + \Theta(\delta + \sqrt{\zeta}))\|\mathbf{y}\|_1$.

Lemma B.6.1. *Fix any $\delta, \zeta > 0$. Let μ' be a partial matching of size $n - \lfloor \delta n \rfloor$ on $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$. For any $\epsilon > 0$, consider*

$$\Omega_{\text{emp}}(\epsilon) := \left\{(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \exists \lambda \in \mathbb{R}_+, \|\hat{\mathcal{F}}(\mathbf{x}) - F_\lambda\|_\infty \leq \epsilon + \Theta(\delta + \sqrt{\zeta})\right\}. \quad (\text{B.79})$$

Then

$$\mathbb{E}\left[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})\right] \leq \exp(o_\delta(n) - \Theta(\epsilon^2 n)) \cdot \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}, \quad (\text{B.80})$$

where again the implicit constants are uniform over all $\mathcal{M}'$, $\mathcal{W}'$, and μ' .

Proof. In light of Proposition B.3.1, it suffices to show that for all $\mathbf{y} \in \text{Proj}_y(\mathcal{R} \cap \Omega_{\text{eig}}(\zeta))$, we have

$$\frac{\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'} = \mathbf{y})]}{\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'} = \mathbf{y})]} \leq \exp(-\Theta(\epsilon^2 n)). \quad (\text{B.81})$$

It then follows that

$$\begin{aligned} & \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \\ &= \mathbb{E}[\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'})]] \\ &\leq \mathbb{E}[\exp(-\Theta(\epsilon^2 n)) \cdot \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\text{Proj}_y(\mathcal{R} \cap \Omega_{\text{eig}}(\zeta))}(\mathbf{Y}_{\mathcal{W}'})]] \\ &\leq \exp(-\Theta(\epsilon^2 n)) \cdot \mathbb{E}[\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})]] \\ &\leq \exp(-\Theta(\epsilon^2 n)) \cdot \mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_2}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'} \mid \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R}_1}(\mathbf{Y}_{\mathcal{W}'})], \end{aligned} \quad (\text{B.82})$$

and Proposition B.3.1 immediately implies the desired bound.

To show (B.81), we rewrite the numerator and denominator on its left-hand side as

$$\mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)) \quad \text{and} \quad \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R}_2),$$

respectively, under the distribution

$$(\mathbf{X}_{\mathcal{M}'})_i = \begin{cases} X_i \sim \text{Exp}(a_{i, \mu'(i)} + n(\mathbf{M}\mathbf{y})_i) & \text{for } i \in \mathcal{M}' \text{ independently,} \\ 0 & \text{for } i \notin \mathcal{M}'. \end{cases}$$

We will express this compactly as $\mathbf{X} \sim \prod_{i=1}^n \text{Exp}(a_{i, \mu'(i)} + n(\mathbf{M}\mathbf{y})_i)$, assuming arbitrary behavior of μ' outside $\mathcal{M}'$ —they will be zeroed out when we consider $\mathbf{X}_{\mathcal{M}'}$. This will be the distribution we work with throughout the rest of this proof.

By definition of $\Omega_{\text{eig}}(\zeta)$, for any $\mathbf{y} \in \text{Proj}_y(\mathcal{R} \cap \Omega_{\text{eig}}(\zeta))$, there exists $\hat{y} \in \mathbb{R}_+$ such that for all but at most $\sqrt{\zeta}n$ indices $i \in \mathcal{M}'$ we have $|(\mathbf{M}\mathbf{y})_i - \hat{y}| \leq \sqrt{\zeta}\hat{y}$. Note that $n\hat{y}\mathbf{X}_{\mathcal{M}'} \sim \prod_{i=1}^n \text{Exp}(\lambda_i)$ with

$$\lambda_i := \frac{a_{i, \mu'(i)}}{n\hat{y}} + \frac{(\mathbf{M}\mathbf{y})_i}{\hat{y}}$$

for all $i \in \mathcal{M}'$ (and again the behavior for $i \notin \mathcal{M}'$ will not matter). For those indices $i \in \mathcal{M}'$ with $|(\mathbf{M}\mathbf{y})_i - \hat{y}| \leq \sqrt{\zeta}\hat{y}$, we further have

$$\lambda_i = 1 + \Theta(\sqrt{\zeta}) + \Theta(1/\log n),$$

where we used the fact that $n\hat{y} \geq \Theta(\|\mathbf{y}\|_1) \geq \Theta(\log n)$ as implied by $\mathbf{y} \in \text{Proj}_y(\mathcal{R} \cap \Omega_{\text{eig}}(\zeta))$. In other words, the majority of the non-trivial components of $\mathbf{X}_{\mathcal{M}'}$ are distributed like nearly i.i.d. $\text{Exp}(\hat{y})$ samples. The generalized Dvoretzky–Kiefer–Wolfowitz (DKW) inequality (Lemma B.1.3) for independent and nearly-identically distributed random variables implies the upper bound

$$\mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \Omega_{\text{emp}}(\epsilon)) \leq \mathbb{P}(\|\hat{\mathcal{F}}(\mathbf{x}) - F_{n\hat{y}}\|_\infty > \epsilon + \Theta(\delta + \sqrt{\zeta})) \leq \exp(-\Theta(\epsilon^2 n)). \quad (\text{B.83})$$

To finish our proof, it remains to lower bound the denominator

$$\mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R}_2) = \mathbb{P}(\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{y} \leq c_2 (\ln n)^{1/8}).$$

This turns out to require only a first-order argument: under our distribution,

$$\mathbb{E}[\mathbf{X}_{\mathcal{M}'}^\top \mathbf{M} \mathbf{y}] = \sum_{i \in \mathcal{M}'} \frac{(\mathbf{M} \mathbf{y})_i}{a_{i, \mu(i)} + n(\mathbf{M} \mathbf{y})_i} \leq \frac{|\mathcal{M}'|}{n} \leq 1.$$

By Markov's inequality, $\mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R}_2) \geq 1 - \frac{1}{c_2 (\ln n)^{1/8}}$. Combining this with (B.83) finishes the proof. $\square$

The proof implies a choice of $\lambda = n\hat{y}$ in the definition (B.79) of $\Omega_{\text{emp}}(\epsilon)$. By the boundedness assumption on the entries of $\mathbf{M}$, we can show with some standard albeit tedious arithmetic that $n\hat{y} = (1 + O(\sqrt{\zeta}))\|\mathbf{y}\|_1$. Hence, instead of looking for the *optimal* λ that minimizes $\|\hat{\mathcal{F}}(\mathbf{X}_{\mathcal{M}'} - F_\lambda)\|_\infty$, we may settle for a slightly sub-optimal choice of $\lambda = \|\mathbf{y}\|_1$. Let

$$\tilde{\Omega}_{\text{emp}}(\epsilon) := \left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \|\hat{\mathcal{F}}(\mathbf{x}) - \mathcal{F}(\text{Exp}(\|\mathbf{y}\|_1))\|_\infty \leq \epsilon + \Theta(\delta + \sqrt{\zeta}) \right\},$$

except for allowing for a slightly worse implicit constant in $\Theta(\delta + \sqrt{\zeta})$ term than in $\Omega_{\text{emp}}(\epsilon)$, the same conclusion holds as in Lemma B.6.1 with $\Omega_{\text{emp}}(\epsilon)$ replaced by $\tilde{\Omega}_{\text{emp}}(\epsilon)$; that is,

$$\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \cap \Omega_{\text{eig}}(\zeta) \setminus \tilde{\Omega}_{\text{emp}}(\epsilon)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \leq \exp(o_\delta(n) - \Theta(\epsilon^2 n)) \cdot \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}. \quad (\text{B.84})$$

Using Lemma B.6.1, we now prove our main Theorem 3.8.1 regarding the uniform limit of empirical distribution of men's (or women's) value in stable matchings.

Proof of Theorem 3.8.1. Plugging (B.84) into Lemma B.2.14 and repeating the same arithmetic as in (B.63) and (B.64) immediately give

$$\mathbb{P}(\exists \mu \in \mathcal{S}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \in \Omega_{\text{eig}}(\zeta) \setminus \tilde{\Omega}_{\text{emp}}(\epsilon)) \leq e^{-n^c} + \exp((o_\delta(1) - \Theta(\epsilon^2))n) \lesssim e^{-n^c} \quad (\text{B.85})$$

granted that we achieve a negative rate in the exponential term above: this requires $\epsilon \geq \epsilon_0(\delta)$ with some function $\epsilon_0(\delta)$ that tends to zero as $\delta \rightarrow 0$. By Proposition 3.7.1, with probability at least $1 - \Theta(e^{-n^c})$, there exists no stable matching μ with $(\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \Omega_{\text{eig}}(\zeta)$; combined with (B.85) through a union bound, this implies

$$\mathbb{P}(\exists \mu \in \mathcal{S}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \tilde{\Omega}_{\text{emp}}(\epsilon/3)) \lesssim e^{-n^c} \quad (\text{B.86})$$

granted that $\epsilon > 3\epsilon_0(\delta)$. By also requiring ϵ to be large enough relative to δ (and also $\zeta = \zeta(\delta)$), we can ensure that the $\Theta(\delta + \sqrt{\zeta})$ term in the definition of $\tilde{\Omega}_{\text{emp}}$ is at most $\epsilon/3$. Thus,

$$\mathbb{P}(\exists \mu \in \mathcal{S}, \|\hat{\mathcal{F}}(\mathbf{X}_\delta(\mu)) - \mathcal{F}(\text{Exp}(\|\mathbf{Y}_\delta(\mu)\|_1))\|_\infty \geq 2\epsilon/3) \lesssim e^{-n^c}. \quad (\text{B.87})$$

Similar, for δ sufficiently small compared to ϵ , we may absorb the difference caused by the δ -truncation on $\mathbf{X}(\mu)$ into an extra term of $\Theta(\delta) \leq \epsilon/3$, since $\|\hat{\mathcal{F}}(\mathbf{X}_\delta(\mu)) - \hat{\mathcal{F}}(\mathbf{X}(\mu))\|_\infty \leq \delta$. Combining this with (B.87) through a triangle inequality gives the Theorem immediately. $\square$

With essentially the same analysis as in Lemma B.6.1 and Theorem 3.8.1, except for replacing the DKW inequality with Bernstein's inequality for empirical average, we can also establish the “law of hyperbola” as stated in Proposition 3.8.2. The proof is omitted.

The effect of the δ -truncation is nontrivial to remove because the sum of values can be sensitive to outliers, in particular given the heavy tail of the exponential distribution. We believe, however, that a refined analysis should suggest that $\sup_{\mu \in \mathcal{S}} |n^{-1} \|\mathbf{X}(\mu)\|_1 \|\mathbf{Y}(\mu)\|_1 - 1| \xrightarrow{p} 0$. This is the analogue of the “law of hyperbola” in Pittel [170].

B.7 Ranks Scale (Approximately) With Values: Proof of Theorem 3.8.3

Heuristically, we would expect the rank R_i for a man to be proportional to his value $X_i(\mu)$. We will see below that this is approximately the case when $x_i \ll 1$. There are, however, going to be some x_i of constant order, making it hard for us to say anything exact about R_i . But as we will soon see, for all but a tiny $o(1)$ fraction of the n men, we will indeed have $x_i = o(1)$. As we are concerned with the empirical distribution, such small fraction becomes negligible in the limit and can be simply ignored. This heuristic is formalized in the next Lemma.

Lemma B.7.1. *Fix any $\delta > 0$. Let μ' be a partial matching of size $n - \lfloor \delta n \rfloor$ on $\mathcal{M}' \subseteq \mathcal{M}$ and $\mathcal{W}' \subseteq \mathcal{W}$. For any $0 < \xi < \rho < 1$, consider $\Omega_{\text{tail}}(\xi, \rho)$ defined as*

$$\left\{ (\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \sum_{i=1}^n \mathbb{1} \left\{ nx_i(\mathbf{M}\mathbf{y})_i \notin (F^{-1}(\xi/2)/2, F^{-1}(1-\xi/2)) \right\} \leq \lfloor \delta n \rfloor + \rho(n - \lfloor \delta n \rfloor) \right\}. \quad (\text{B.88})$$

Then

$$\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \setminus \Omega_{\text{tail}}(\xi, \rho)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})] \leq \exp(o_\delta(n) - \Theta(D(\rho \parallel \xi)n)) \cdot \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}, \quad (\text{B.89})$$

where $D(q \parallel p)$ denotes the KL divergence from $\text{Bern}(p)$ to $\text{Bern}(q)$.

Proof. The proof entirely mirrors that of Lemma B.6.1. It suffices to show that for all $\mathbf{y} \in \text{Proj}_{\mathcal{y}}(\mathcal{R})$ we have

$$\frac{\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \cdot \mathbb{1}_{\mathcal{R} \setminus \Omega_{\text{tail}}(\xi, \rho)}(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \mid \mathbf{Y}_{\mathcal{W}'} = \mathbf{y}]}{\mathbb{E}[q(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \mid \mathbf{Y}_{\mathcal{W}'} = \mathbf{y}]} \leq \exp(-\Theta(D(\rho \parallel \xi)n)). \quad (\text{B.90})$$

The quotient is simply

$$\mathbb{P}_{\mathbf{X} \sim \prod_{i=1}^n \text{Exp}(a_{i,\mu'(i)} + n(\mathbf{M}\mathbf{y})_i)}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \in \mathcal{R} \setminus \Omega_{\text{tail}}(\xi, \rho)) \leq \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \notin \Omega_{\text{tail}}(\xi, \rho)), \quad (\text{B.91})$$

where we will have $\mathbf{X} \sim \prod_{i=1}^n \text{Exp}(a_{i,\mu'(i)} + n(\mathbf{M}\mathbf{y})_i)$ for the rest of this proof. Note that under this specified distribution, $\left((a_{i,\mu'(i)} + n(\mathbf{M}\mathbf{y})_i)X_i\right)_{i \in \mathcal{M}'}$ are $n - \lfloor \delta n \rfloor$ i.i.d. samples from $\text{Exp}(1)$, each falling outside the interval $(F^{-1}(\xi/2), F^{-1}(1 - \xi/2))$ with probability precisely ξ . Hence,

$$\begin{aligned} \mathbb{P}\left(\sum_{i \in \mathcal{M}'} \mathbb{1}\left\{X_i(a_{i,\mu'(i)} + n(\mathbf{M}\mathbf{y})_i) \notin (F^{-1}(\xi/2), F^{-1}(1 - \xi/2))\right\} \leq \rho(n - \lfloor \delta n \rfloor)\right) \\ = \mathbb{P}_{Z \sim \text{Bin}(n - \lfloor \delta n \rfloor, \xi)}(Z > \rho(n - \lfloor \delta n \rfloor)) \leq \exp(-D(\rho \parallel \xi)(n - \lfloor \delta n \rfloor)) \end{aligned} \quad (\text{B.92})$$

by the Hoeffding bound for binomial distribution. Since $n(\mathbf{M}\mathbf{y})_i \leq a_{i,\mu'(i)} + n(\mathbf{M}\mathbf{y})_i \leq 2n(\mathbf{M}\mathbf{y})_i$ across all $i \in \mathcal{M}'$ for $\mathbf{y} \in \mathcal{R}_1$ and n sufficiently large, the probability (B.92) upper bounds $\mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{y}) \notin \Omega_{\text{tail}}(\xi, \rho))$. This establishes (B.90) and concludes the proof. $\square$

By fixing some small δ, ρ and choosing ξ sufficiently small, we can make $D(\rho \parallel \xi)$ arbitrarily large and obtain the following Corollary.

Corollary B.7.2. *For any $0 < \delta, \rho < 1/2$, there exists a choice of $\xi > 0$ such that*

$$\mathbb{P}(\exists \mu \in \mathcal{S}, (\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \tilde{\Omega}_{\text{tail}}(\xi, \rho)) \lesssim e^{-n^c} \quad (\text{B.93})$$

asymptotically as $n \rightarrow \infty$, where $\tilde{\Omega}_{\text{tail}}(\xi, \rho)$ is defined as

$$\left\{(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_+^n \times \mathbb{R}_+^n : \sum_{i=1}^n \mathbb{1}\left\{x_i \notin \left(F^{-1}(\xi/2) \frac{(\log n)^{7/8}}{C^2 \bar{c}_1 n}, 2F^{-1}(1 - \xi/2) \frac{C^2}{\underline{c}_1 \log n}\right)\right\}\right\} \leq (\delta + \rho)n \right\}. \quad (\text{B.94})$$

That is, with high probability, no stable matchings μ have more than $\delta + \rho$ fraction of the men's post-truncation values outside an interval $(\Theta(n^{-1}(\log n)^{7/8}), \Theta(1/\log n))$.

Proof. Observe that $\mathcal{R} \setminus \tilde{\Omega}_{\text{tail}}(\xi, \rho) \subseteq \mathcal{R} \setminus \Omega_{\text{tail}}(\xi, \rho)$ by our definition of $\mathcal{R}$ and the boundedness assumption on the entries of $\mathbf{M}$. Again, invoking Lemma B.2.14 using inequality (B.89) in Lemma B.7.1 yields the $\Theta(e^{-n^c})$ upper bound on $\mathbb{P}((\mathbf{X}_\delta(\mu), \mathbf{Y}_\delta(\mu)) \notin \Omega_{\text{tail}}(\xi, \rho))$. $\square$

Remark B.7.3. Recall that in a market with uniform preferences, the man-optimal (and woman-pessimal) stable matching realizes an average rank of $\Theta(\log n)$ for men and $\Theta(n/\log n)$ for women. Under the heuristic that values multiplied by n roughly correspond to ranks (which we will formalize below), Lemma B.7.1 nicely matches our expectation that even in the most extreme cases, few individuals will strike a rank better (smaller) than $\Theta((\log n)^{7/8})$ or worse (larger) than $\Theta(n/\log n)$. The lower bound can be refined to $\Theta(\log n / \log \log n)$ with a more careful analysis.

Now let us consider a specific partial matching μ' of size $n - \lfloor \delta n \rfloor$ between $\mathcal{M}'$ and $\mathcal{W}'$ and condition on μ' being stable with value vectors $(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) = (\mathbf{x}, \mathbf{y}) \in \tilde{\Omega}_{\text{tail}}(\xi, \rho)$. That is, there exists a subset $\bar{\mathcal{M}}' \subseteq \mathcal{M}'$ with $|\bar{\mathcal{M}}'| \geq (1 - \delta - \rho)n$ such that $\Theta(n^{-1}(\log n)^{7/8}) \leq (X_{\mathcal{M}'})_i \leq \Theta(1/\log n)$ for all $i \in \bar{\mathcal{M}}'$. By symmetry, we may further assume that there exists $\bar{\mathcal{W}}' \subseteq \mathcal{W}'$ with $|\bar{\mathcal{W}}'| \geq (1 - \delta - \rho)n$ such that $\Theta(n^{-1}(\log n)^{7/8}) \leq (Y_{\mathcal{W}'})_j \leq \Theta(1/\log n)$ for all $j \in \bar{\mathcal{W}}'$. We want to show that, for $i \in \bar{\mathcal{M}}'$, the *pre-truncation* rank R_i of man m_i (i.e., over the entire market, including the $\lfloor \delta n \rfloor$ women outside $\mathcal{M}'$) is well characterized by his value $X_{i, u'(i)}$ in the matching, up to some proper scaling. From now on, we will consider some $i \in \mathcal{M}'$ with value $X_{i, u'(i)} = x_i$, and write

$$R_i = 1 + \sum_{j \neq \mu'(i)} \mathbb{1}_{[0, x_i]}(X_{ij}). \quad (\text{B.95})$$

The condition that μ' is stable requires $(X_{ij}, Y_{ji}) \notin [0, x_i] \times [0, y_j]$ for all $j \in \mathcal{W}' \setminus \{\mu'(i)\}$. Thus, for all $j \in \mathcal{W}' \setminus \{\mu'(i)\}$,

$$\begin{aligned} \mathbb{P}(X_{ij} \leq x_i | \mu' \text{ stable}, (\mathbf{X}_{\mathcal{M}'})_i = x_i, (\mathbf{Y}_{\mathcal{W}'})_j = y_j) &= \frac{\mathbb{P}(X_{ij} \leq x_i, Y_{ji} > y_j)}{1 - \mathbb{P}(X_{ij} \leq x_i, Y_{ji} \leq y_j)} \\ &= \frac{(1 - e^{-a_{ij}x_i})e^{-b_{ji}y_j}}{1 - (1 - e^{-a_{ij}x_i})(1 - e^{-b_{ji}y_j})}, \end{aligned} \quad (\text{B.96})$$

and for all $j \in \mathcal{W} \setminus \mathcal{W}'$ (so $(\mathbf{Y}_{\mathcal{W}'})_j = 0$),

$$\mathbb{P}(X_{ij} \leq x_i | \mu' \text{ stable}, (\mathbf{X}_{\mathcal{M}'})_i = x_i) = 1 - e^{-a_{ij}x_i}. \quad (\text{B.97})$$

Define

$$p_{ij} = \begin{cases} 1 & \text{when } j = \mu'(i), \\ \frac{(1 - e^{-a_{ij}x_i})e^{-b_{ji}y_j}}{1 - (1 - e^{-a_{ij}x_i})(1 - e^{-b_{ji}y_j})} & \text{when } j \in \mathcal{W}' \setminus \{\mu'(i)\}, \\ 1 - e^{-a_{ij}x_i} & \text{when } j \in \mathcal{W} \setminus \mathcal{W}', \end{cases}$$

and $I_{ij} \sim \text{Bern}(p_{ij})$ independently for $i \in [n]$ so that $R_i = \sum_{j=1}^n I_{ij}$ conditional on $(\mathbf{X}_{\mathcal{M}'})_i = X_i = x_i$. Note that for any $j \neq \mu'(i)$ and $j \notin \bar{\mathcal{W}}'$, we always have

$$p_{ij} \leq 1 - e^{-a_{ij}x_i} \leq a_{ij}x_i \quad (\text{B.98})$$

and

$$\begin{aligned} p_{ij} &\geq \frac{(1 - e^{-a_{ij}x_i})e^{-b_{ji}y_j}}{1 - (1 - e^{-a_{ij}x_i})(1 - e^{-b_{ji}y_j})} \geq (1 - e^{-a_{ij}x_i})e^{-b_{ji}y_j} \\ &\geq e^{-\Theta(\frac{1}{\log n})} \left(1 - \Theta\left(\frac{1}{\log n}\right) \right) a_{ij}x_i = (1 - o(1))a_{ij}x_i. \end{aligned} \quad (\text{B.99})$$

For $j \in \bar{\mathcal{W}}' \setminus \{\mu'(i)\}$, p_{ij} admits the same upper bound (B.98) and the trivial lower bound of zero. Hence, conditional on

$$\mu' \text{ stable and } (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) = (\mathbf{x}, \mathbf{y}) \in \tilde{\Omega}_{\text{tail}}(\xi, \rho) \cap \tilde{\Omega}_{\text{emp}}(\epsilon) \cap \mathcal{R} \quad (\dagger)$$

for any (fixed) $\xi, \rho, \epsilon > 0$, we have the stochastic dominance

$$1 + \sum_{j \notin \bar{\mathcal{W}}' \cup \{\mu'(i)\}} \underline{I}_{ij} \preceq R_i \preceq 1 + \sum_{j \neq \mu'(i)} \bar{I}_{ij}, \quad (\text{B.100})$$

where $\underline{I}_{ij} \sim \text{Bern}((1 - o(1))a_{ij}x_i)$ and $\bar{I}_{ij} \sim \text{Bern}(a_{ij}x_i)$. Since $i \in \bar{\mathcal{M}}'$ by our assumption and thus $\Theta((\log n)^{7/8}/n) \leq x_i \leq \Theta(1/\log n)$, the expectation of R_i/x_i can be upper bounded by

$$\mathbb{E} \left[\frac{R_i}{x_i} \middle| (\dagger) \right] \leq \frac{1}{x_i} + \sum_{j \neq \mu'(i)} a_{ij} = (1 + o(1)) \sum_{j=1}^n a_{ij} \quad (\text{B.101})$$

and lower bounded by

$$\mathbb{E} \left[\frac{R_i}{x_i} \middle| (\dagger) \right] \geq (1 - o(1)) \sum_{j \neq \mu'(i)} a_{ij} = (1 - \Theta(\delta)) \sum_{j=1}^n a_{ij}. \quad (\text{B.102})$$

Similarly, we may bound the variance of R_i/x_i by

$$\text{Var} \left(\frac{R_i}{x_i} \middle| (\dagger) \right) \leq \sum_{j \neq \mu'(i)} a_{ij}(1 - a_{ij}x_i) \leq \sum_{j=1}^n a_{ij}. \quad (\text{B.103})$$

Hence, we have

$$1 - \Theta(\delta) \leq \mathbb{E} \left[\frac{R_i}{x_i \sum_{j=1}^n a_{ij}} \middle| (\dagger) \right] \leq 1 + o(1) \quad \text{and} \quad \text{Var} \left(\frac{R_i}{x_i \sum_{j=1}^n a_{ij}} \middle| (\dagger) \right) \leq \Theta(n^{-1}), \quad (\text{B.104})$$

with these quantities conditionally independent for all $i \in \bar{\mathcal{M}}'$ and the hidden constants depending only on C , implying concentration of R_i around $x_i \sum_{j=1}^n a_{ij}$ in the following sense.

Proposition B.7.4. *Conditional on $(\dagger)$, for any fixed $\theta > 0$ and $\delta, \rho, \gamma \in (0, 1/2)$, we have*

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^n \mathbb{1}_{(\theta + \Theta(\delta), \infty)} \left(\left| \frac{R_i}{x_i \sum_{j=1}^n a_{ij}} - 1 \right| \right) \geq (\delta + \rho + \gamma)n \middle| (\dagger) \right) \\ \lesssim \mathbb{P}_{N \sim \text{Poi}(\Theta(\theta^{-2}))}(N \geq \gamma n) \leq e^{-\Theta(\gamma n)}. \end{aligned} \quad (\text{B.105})$$

Proof. By Chebyshev's inequality and (B.104), $\mathbb{P}(|\frac{R_i}{x_i \sum_{j=1}^n a_{ij}} - \mathbb{E}[\frac{R_i}{x_i \sum_{j=1}^n a_{ij}} | (\dagger)]| \geq \theta | (\dagger)) \leq \Theta((n\theta^2)^{-1})$ for all $i \in \bar{\mathcal{M}}'$. Hence, by conditional independence of the ranks,

$$\sum_{i \in \bar{\mathcal{M}}'} \mathbb{1}_{(\theta + \Theta(\delta), \infty)} \left(\left| \frac{R_i}{x_i \sum_{j=1}^n a_{ij}} - 1 \right| \right) \preceq \text{Bin}(n, \Theta((n\theta^2)^{-1}))$$

, which converges to $\text{Poi}(\Theta(\theta^{-2}))$ in total variance. The Proposition follows from the well known tail bound for $N \sim \text{Poi}(\lambda)$ that $\mathbb{P}(N \geq \lambda + t) \leq \exp\left(-\frac{t^2}{2(\lambda+t)}\right)$, which implies $\mathbb{P}_{N \sim \text{Poi}(\Theta(\theta^{-2}))}(N \geq \gamma n) \leq \exp\left(-\frac{(\gamma n - \Theta(\theta^{-2}))^2}{2\gamma n}\right) \lesssim \exp(-\gamma n/2)$. $\square$

We are now ready to proof Theorem 3.8.3.

Proof. Note that in Corollary B.7.2, ρ can be chosen arbitrarily small once δ is fixed. In particular, we may always guarantee $\rho \leq \delta$. Similarly, we may assume $\theta \leq \delta$ in Proposition B.7.4. Thus,

$$\mathbb{P}\left(\sum_{i=1}^n \mathbb{1}_{(\Theta(\delta), \infty)}\left(\left|\frac{R_i}{x_i w_i} - 1\right|\right) \geq (2\delta + \gamma)n \middle| (\dagger)\right) \leq e^{-\Theta(\gamma n)}, \quad (\text{B.106})$$

where $w_i = \sum_{j=1}^n a_{ij}$ is the fitness value of man m_i . Marginalizing over all pairs of relevant value vectors $(\mathbf{x}, \mathbf{y}) \in \tilde{\Omega}_{\text{tail}}(\xi, \rho) \cap \tilde{\Omega}_{\text{emp}}(\epsilon) \cap \mathcal{R}$ in the condition $(\dagger)$, we obtain

$$\mathbb{P}\left(\mathcal{E}_{\text{ratio}}(\delta, \gamma) \middle| (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'}) \in \tilde{\Omega}_{\text{tail}}(\xi, \rho) \cap \tilde{\Omega}_{\text{emp}}(\epsilon) \cap \mathcal{R}\right) \leq e^{-\Theta(\gamma n)}, \quad (\text{B.107})$$

where $\mathcal{E}_{\text{ratio}}(\delta, \gamma)$ denotes the undesirable event that $\sum_{i=1}^n \mathbb{1}_{(\Theta(\delta), \infty)}\left(\left|\frac{R_i}{(\mathbf{X}_{\mathcal{M}'})_i w_i} - 1\right|\right) \geq (2\delta + \gamma)n$ for the partial matching μ' . By Proposition B.3.1,

$$\begin{aligned} \mathbb{P}((\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \tilde{\Omega}_{\text{tail}}(\xi, \rho) \cap \Omega_{\text{emp}}(\epsilon) \cap \mathcal{R}) &\leq \mathbb{P}(\mu' \text{ stable}, (\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \mathcal{R}) \\ &\leq e^{o(n) + o_\delta(n)} \frac{(\delta n)!}{n!} \prod_{i \in \mathcal{M}'} a_{i, \mu'(i)} b_{\mu'(i), i}. \end{aligned} \quad (\text{B.108})$$

By choosing $\gamma = \gamma(\delta) = o_\delta(1)$ sufficiently large (relative to δ) and following a similar computation as in Lemma B.2.14 and Corollary 3.6.3, we can ensure that with probability $1 - \Theta(e^{-n^c})$ there exists no stable partial matching μ' where both $\mathcal{E}_{\text{ratio}}(\delta, \gamma)$ and $(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \tilde{\Omega}_{\text{emp}}(\epsilon_0(\delta)) \cap \mathcal{R}$ happen, where the function ϵ_0 is defined in the proof of 3.8.1. Notice that by repeated uses of the triangle inequality, $(\mathbf{X}_{\mathcal{M}'}, \mathbf{Y}_{\mathcal{W}'})) \in \tilde{\Omega}_{\text{emp}}(\epsilon_0(\delta)), \mathcal{E}_{\text{ratio}}(\delta, \gamma)^c$ implies that

$$\|\hat{\mathcal{F}}(\mathbf{w}^{-1} \circ \mathbf{R}(\mu')) - F_\lambda\|_\infty \leq \Theta(\delta) + \gamma(\delta) + \epsilon_0(\delta) = o_\delta(1)$$

for the choice of $\lambda = \|\mathbf{Y}_{\mathcal{W}'}\|_1$. Combining this with (B.86) and Proposition 3.6.2, we conclude that with probability $1 - \Theta(e^{-n^c})$, all stable matchings $\mu \in \mathcal{S}$ induce δ -truncated stable partial matchings μ_δ with $\|\hat{\mathcal{F}}(\mathbf{w}^{-1} \circ \mathbf{R}(\mu_\delta)) - F_{\lambda(\mu)}\|_\infty = o_\delta(1)$, where $\lambda(\mu) = \|\mathbf{Y}_\delta(\mu)\|_1$. The δ -truncation affects the distance by at most δ , which can be absorbed into the $o_\delta(1)$ upper bound. Thus, by choosing δ sufficiently small relative to any fixed $\epsilon > 0$, we complete our proof of Theorem 3.8.3. $\square$

B.7.1 Approximate Stability: Proof of Theorem 3.4.5

Proof of Theorem 3.4.5. There are $\binom{n}{\alpha n}^2 = \exp(2h_b(\alpha)n + O(\ln n))$ sub-markets of size at least $(1 - \alpha)n$, where $h_b(p) = -p \log p - (1 - p) \log(1 - p)$ is the binary entropy function. Under Assumption 3.2.2 for the whole market, each of such sub-markets also satisfies Assumption 3.2.2. Fix any $\epsilon > 0$. By Theorem 3.8.1, each of such sub-markets can only contain a stable matching with men's empirical distribution of value deviating from the family of exponential distributions by at least $\epsilon/2$ in Kolmogorov-Smirnov distance with probability at most $1 - \exp(-n^c)$ for any fixed $c \in (0, 1/2)$. Whenever $\alpha < n^{-\eta}$ for some $\eta > 1/2$, we have $h_b(\alpha)n < n^{1-\eta}$. Choosing $c \in (1 - \eta, 1/2)$ and applying union bound over all relevant sub-markets gives the first part of (3.6), since the additional α fraction of the market affects the empirical distribution by at most $\alpha \rightarrow 0$. The second part follows analogously from Theorem 3.8.3. $\square$

Appendix C

Appendix to Chapter 4

C.1 Many-to-One Matching

In the main text, we have shown that a giant unique SCC appears in $G^{\text{DA}(P_n)}$ if there are as many schools as students. We relax such assumption here, allowing for a school to be able to admit many students, as follows.

Theorem C.1.1 (Giant SCC in Many-to-One Matching). *Consider a random school choice problem P_n with n schools, where each school has the same fixed quota q , and nq students. Fix an arbitrary $\epsilon > 0$. With high probability, the envy digraph $G^{\text{DA}(P_n)}$ has a unique giant strongly connected component containing at least $(1 - \epsilon)nq$ students.*

Proof. We follow a strategy similar to Theorem 4.4.1, adapting it to the many-to-one setting. First, we establish bounds on the number of applications in the DA algorithm, then analyze partitions of the resulting envy digraph.

Step 1: Bounds on applications. We first establish the following bound on the number of applications a typical student makes, similar to Lemma 4.4.2:

Lemma C.1.2. *Fix $\epsilon > 0$. With high probability, all but at most ϵnq students each make at least $\sqrt{\log n}$ applications.*

Proof. Again, it suffices to show that for a fixed student i , the number of applications she makes is at least $\sqrt{\log n}$ with high probability; the lemma follows as a consequence of Markov's inequality applied to the number of students making fewer applications than this, same as we saw in the proof of Lemma 4.4.2.

The standard approach for one-to-one matchings using the coupon collector problem as an approximation can be adapted to this case with slight modifications. In particular, one can establish that the total number of applications throughout DA is of order $\Theta(n \log n)$. Since the order of applications in DA is irrelevant, one can hold off i and run DA on the rest of the students $I \setminus \{i\}$, and when this first stage terminates (i.e., when all the students

except for i are stably matched) there should be already $\Theta(n \log n)$ applications. At this point, all but one school has their capacity filled; Without loss of generality, let it be school s_n . Using terminologies from [25], we say that this intermediate state of DA is *smooth* if (1) at most $O(n \log n)$ applications are made until this point and (2) all but at most n^a schools have received applications from $\Omega(\log n)$ distinct students with constant $a < 1$. With nearly identical analysis as in Ashlagi et al. [25], we can show that a smooth state is achieved with high probability; the details are repetitive and hence omitted.¹ We now let i start making applications until the full DA terminates (i.e., some one applies to this last unfilled school s_n). Note that, by symmetry, an application to a school who has received $m \geq q$ applications will be accepted with probability $q/(m+1)$. Denote the number of applications to school s_k at the beginning of this second stage by M_k for $k = 1, \dots, n-1$. Under a smooth state, we may assume without loss of generality that $M_k \geq \Theta(\log n)$ for $k = 1, \dots, \lfloor n - n^a \rfloor$; as the second stage of DA unfolds, the number of applications to each school can only increase. The probability that an application from i is accepted is at most

$$\frac{1}{n} \left(1 + \sum_{k=1}^{n-1} \frac{q}{M_k + 1} \right) \leq \frac{n^a}{n} + (n - n^a) \frac{q}{\Theta(\log n)} \leq \Theta(q/\log n).$$

Thus, conditional on the likely smooth state before i starts to apply, the number of applications that i makes until receiving a (tentative) acceptance is stochastically lower bounded by a geometric distribution with success probability $\Theta(q/\log n)$, and therefore, the conditional probability (under a smooth state) that i makes fewer than $\sqrt{\log n}$ applications is at most $1 - (1 - \Theta(q/\log n))^{\sqrt{\log n}} \rightarrow 0$. Since a smooth state is achieved with high probability, the marginal probability of i making less than $\sqrt{\log n}$ applications vanishes as $n \rightarrow \infty$, finishing our proof. $\square$

Step 2: Partition Argument. Suppose $G^{\text{DA}(P_n)}$ has two disjoint subsets $I_0, I_1 \subseteq I$, each of size $\geq \epsilon n q$, with no edges from I_0 to I_1 . Let $S_1 = \{s \in S : \exists j \in I_1 \text{ with } \text{DA}_j(P_n) = s\}$, so $|S_1| \geq \epsilon n$. We now show that for any fixed subsets I_0 and S_1 , it is extremely unlikely that no one in I_0 ever applies to schools in S_1 .

Lemma C.1.3. *For any fixed $I_0 \subseteq I$ and $S_1 \subseteq S$ with $|I_0| \geq \epsilon n q$ and $|S_1| \geq \epsilon n$, the probability that students in I_0 make at least $n\sqrt{\log n}$ applications yet none apply to S_1 is at most $(1 - \epsilon)^{n\sqrt{\log n}}$.*

¹The first bound on the number of applications can be derived from classic results on the generalized coupon collector problem; see, e.g., [161, 79]. The second bound on the number of applications received by the schools is slightly more involved. To derive it, we run the amnesiac DA [123], and using concentration arguments on the number of draws in (generalized) coupon collector problem, one can show that the number of applications, counting duplicates, is of order $\Omega(n \log n)$ up to this stage; further, duplicates are rare among them, so the total number of applications in the classic DA is again $\Omega(n \log n)$ with high probability prior to the second stage. The number of applications each school receives can then be bounded using a Poissonization argument or with direct Chernoff-type bounds. See, e.g., Ashlagi et al. [25, Proposition 4.3] for a formal statement and analysis in the one-to-one case.

Proof. We use the similar idea as in the proof of Lemma 4.4.2. In the amnesiac DA model, each application independently targets a school in S_1 with probability $\geq \epsilon$. Thus, the probability of the event in the lemma is at most $(1 - \epsilon)^{n\sqrt{\log n}}$ as claimed. $\square$

Step 3: Union Bound over Partitions. The event that the giant SCC contains strictly less than $(1 - 2\epsilon)nq$ students implies the existence of linear-sized $I_0, I_1 \subseteq I$ with no edges going from I_0 to I_1 , and thus no one in I_0 ever applies to the set of schools $S_1 := \{s \in S : \exists j \in I_1, \text{DA}_j(P) = s\}$. Denote such an event by $\mathcal{E}_{I_0, S_1}$. By a similar union bound argument as in the proof of Theorem 4.4.1, the probability p the giant SCC contains $< (1 - 2\epsilon)nq$ students is bounded by:

$$p \leq \sum_{|I_0| \geq \epsilon nq, |S_1| \geq \epsilon n} \mathbb{P}(\mathcal{E}_{I_0, S_1} \cap \mathcal{A}_{I_0}) + \mathbb{P}\left(\bigcup_{|I_0| \geq \epsilon nq} \mathcal{A}_{I_0}^c\right), \quad (\text{C.1})$$

where $\mathcal{A}_{I_0}$ denotes the event that the subset of students I_0 make at least $n\sqrt{\log n}$ applications in total.

Note that $\mathcal{E}_{I_0, S_1} \cap \mathcal{A}_{I_0}$ is precisely the kind of events characterized in Lemma C.1.3, and hence the first summation term (which has fewer than $(2^{nq})^2$ terms) is at most

$$2^{2nq} \cdot (1 - \epsilon)^{n\sqrt{\log n}} = \exp\left(2nq \ln 2 + n\sqrt{\log n} \log(1 - \epsilon)\right) \rightarrow 0 \quad (\text{C.2})$$

as $n \rightarrow \infty$. Lemma C.1.2 further ensures $\mathbb{P}(\bigcup_{|I_0| \geq \epsilon nq} \mathcal{A}_{I_0}^c) \rightarrow 0$ as $n \rightarrow \infty$. Thus, the giant SCC contains $\geq (1 - \epsilon)nq$ students with high probability. $\square$

Despite the different structure of the many-to-one market with uniform quotas, the fundamental conclusion remains: in large random markets, a unique giant SCC emerges, and thus most students are improvable. However, we emphasize that the rate of convergence in the many-to-one setting differs from the one-to-one case, and depends on the size of the constant q : As q increases, each school forms larger “envy-free clusters” (students assigned to the same school), potentially slowing convergence.

In Figures C.1a and C.1b, we observe exactly that: the unique giant SCC emerges but the fraction of nodes in it for fixed n becomes smaller for larger constant q .

Note that although the convergence appears slow, scaling like $1/\log n$ and consistent with our analysis, the main driver behind this slow rate is the highly dispersed distribution of applications submitted by students. In particular, an $O(1/\log n)$ fraction of students submit only one application and thus envy no one. Empirical simulations suggest that, aside from these trivially unimprovable students, the giant strongly connected component (SCC) contains nearly all of the remaining students.

Equivalence in Many-to-One Matching. The cycle packing result in Theorem 4.4.6 also extends naturally to the many-to-one setting with fixed quota q for each school. We present this generalization as follows:



(a) $n = [10 : 10 : 200]$, $q = 5$.(b) $n = [5 : 5 : 100]$, $q = 10$.

Figure C.1: Average fraction of nodes in the largest (blue) and second-largest (red) SCCs (average over 1,000 random problems for each n), preferences and priorities uniform i.i.d.

Theorem C.1.4 (Cycle Packing in Many-to-One Matching). *Consider a random school choice problem P_n with n schools, where each school has the same fixed quota q , and nq students. Fix an arbitrary $\epsilon > 0$. With high probability, every cycle packing of the envy digraph $G^{\text{DA}}(P_n)$ covers at least $(1 - \epsilon)nq$ students.*

Proof. The proof follows the approach of Theorem 4.4.6, adapted to the many-to-one setting. Let $H \subseteq G^{\text{DA}}(P_n)$ be a cycle packing, and consider the induced subgraph $G^{\text{DA}}(P_n) \setminus H = G^{\text{DA}}(P_n)[I \setminus V(H)]$, which contains no cycles by definition.

Suppose, for contradiction, that $|V(H)| < (1 - \epsilon)nq$. Then, the directed acyclic subgraph $G^{\text{DA}}(P_n) \setminus H$ contains $N = |I \setminus V(H)| > \epsilon nq$ vertices. We can find a topological ordering $i_1, \dots, i_N$ of vertices in $I \setminus V(H)$ such that no directed edge from i_k to i_ℓ for $1 \leq k < \ell \leq N$ is present in $G^{\text{DA}}(P_n) \setminus H$.

Partitioning these nodes into $I_0 = \{i_1, \dots, i_{\lfloor N/2 \rfloor}\}$ and $I_1 = \{i_{\lfloor N/2 \rfloor + 1}, \dots, i_N\}$ yields two subsets, each with size at least $\lfloor \epsilon nq/2 \rfloor$, with no edges going from I_0 to I_1 in $G^{\text{DA}}(P_n)$.

This scenario implies that no student in I_0 has applied to any school matched to students in I_1 during the execution of DA. By Lemma C.1.3, such a partition is extremely unlikely for sufficiently large n , with probability vanishing as n grows.

Thus, we must have $|V(H)| \geq (1 - \epsilon)nq$ with high probability. $\square$

These findings collectively demonstrate that our main conclusions—both about the pervasiveness of improvable students and the asymptotic equivalence of mechanisms that Pareto dominate DA—are robust to varying market structures, including the practically important case of schools with multiple seats.

C.2 Correlation in Preferences

In the main text, we assume that students' preferences are independently drawn. Here we relax such assumption by dividing schools into two tiers of equal size, and assume that each

school in Tier 1 is more preferred than any school in Tier 2. This tiered model aims to capture the empirical observation that some schools are known to be of high quality and preferred by most students. Within tiers, preferences are drawn uniformly at random.

Interestingly, in this scenario we observe the emergence of two (rather than one) giant SCCs. Simulations suggest that the fraction of nodes in the two giant SCCs combined converges to one as the market grows, as follows. The emergence of two giant SCCs is




Figure C.2: Average fraction of nodes in the largest (blue) and second-largest (red) SCCs in a two-tiered model (average over 1 000 random envy digraphs for each n).

expected. Every student matched with a school in Tier 2 is not envied by any student in Tier 1, so therefore any SCCs can contain either only Tier 1 students (i.e. matched to schools from Tier 1) or Tier 2 students, which implies no SCC includes more than half of the students. But, since each subgraph is dense enough, we find that a giant SCC obtains for each tier of the market, and thus almost all students in each tier are improvable via within-tier exchanges. Based on these observations, we conjecture that in a model that allows for a constant c number of tiers, c giant SCCs grow and they contain asymptotically almost all nodes as the market grows.

Appendix D

Appendix to Chapter 5

D.1 Local Algorithm for the Size of 2-Core

We offer an explicit implementation of Algorithm 2 in pseudocode. The algorithm consumes a graph G , a locality parameter K , and the number T of repetitions as inputs, draws T i.i.d. vertices from $V(G)$, runs a local exploration subroutine on each of the samples, and returns the average across the T samples.

Algorithm 3 details the local exploration subroutine using breadth-first search. It decides whether a vertex is potentially in the (giant) 2-core of G based on the local neighborhood.

D.2 Size of the Giant Component

D.2.1 Sprinkling and the Size of the Giant Component

The proof of Theorem 5.2.3 is analogous to that of Proposition 3.5 in [14]. Recall from [14] that we say G is an (ϵ, α) -large-set expander if

$$\phi(G, \epsilon) := \min_{A \subseteq V, \epsilon n \leq |A| \leq (1-\epsilon)n} \frac{|E(A, V \setminus A)|}{|A|} \geq \alpha.$$

The quantity $\phi(G, \epsilon)$ is referred to as the ϵ -large-set expansion of G . Our definition of a (structurally) robust sequence simply requires the ϵ -large-set expansion to be asymptotically bounded away from zero for any $\epsilon > 0$ (with $\delta = \epsilon\phi(G, \epsilon)$ in our definition).

The following lemma restates Lemma 3.3 in [14] and captures the gist of the sprinkling argument. For the ease of reference, we state the proof here, which is identical to the one in [14].

Lemma D.2.1 (Sprinkling. Lemma 3.3, [14]). *Let $G = (V, E)$ be an (ϵ, α) -expander on n vertices with average degree at most $\bar{d}$. Fix real constants $\kappa > 0$ and $\beta, \epsilon \in (0, 1)$. Let S be a collection of disjoint connected subgraphs of G each of size at least κ . Consider the event that*

Algorithm 3: A local exploration subroutine for 2-core candidacy

Input: Graph $G = (V, E)$, locality parameter K **Output:** Indicators of potential membership of $C^{(2)}(G)$ and $C^{(2)}(C_{\max}(G))$

```

1 Sample  $v \in V$  uniformly at random
2 if  $\deg(v) \geq K$  then                                     // a high degree vertex
3   return  $I^{(2)} = I^{(2,\infty)} = 1$ 
4 Set  $r = 0, S = 1$  // initialize radius of exploration, and size of the neighborhood
5 Initialize  $q = \text{Queue}(v)$ 
6 for  $u \in \text{Neighbors}(v)$  do
7   Set  $r_u = r + 1, p_u = u$  // store the distance to  $v$  and the ancestor towards  $v$ 
8   Set  $\text{reach}_u = r_u, \text{cycle}_u = \text{False}$ 
9    $S = S + 1$ 
10   $q.\text{enqueue}(u)$ 
11 while  $q \neq \emptyset$  do
12    $u = q.\text{pop}()$ 
13    $r = r_u$ 
14   for  $w \in \text{Neighbors}(u)$  do
15     if  $p_w = \text{undefined}$  then
16       Set  $r_w = r + 1, p_w = p_u$  // store the distance to  $v$  and the ancestor
17       towards  $v$ 
18       Set  $\text{reach}_{p_u} = r_w$ 
19        $S = S + 1$ 
20       if  $S > K$  then
21         break for and while
22        $q.\text{enqueue}(w)$ 
23     else
24       Set  $\text{cycle}_{p_u} = \text{True}$ 
25 if  $S \geq K$  then // connected component of size at least  $K$ 
26   return  $I^{(2)} = I^{(2,\infty)} = \mathbb{1}\{\sum_{u \in \text{Neighbors}(v)} \mathbb{1}\{\text{cycle}_u \text{ or } \text{reach}_u = r\} \geq 2\}$ 
27 else // connected component of size less than  $K$ 
28   return  $I^{(2)} = \mathbb{1}\{\sum_{u \in \text{Neighbors}(v)} \mathbb{1}\{\text{cycle}_u\} \geq 2\}$  and  $I^{(2,\infty)} = 0$ 

```

there exists a partition of S into S_0 and $S \setminus S_0$ each with size at least ϵn that are disconnected in $G(\beta)$ is at most $\exp\left(\frac{n}{\kappa} - \beta^{\bar{d}/\alpha\epsilon} \frac{\alpha\epsilon n}{2}\right)$.

Proof. It suffices to consider S_0 consisting of some complete pieces of subgraphs in S ; otherwise S_0 and $S \setminus S_0$ are clearly connected. We hope to lower bound the number of (disjoint) paths joining vertices in S_0 and those in $S \setminus S_0$. This is the max-flow between the two subsets of vertices, and is equal to the min-cut. By ϵ -large-set expansion of G , with $|S_0|, |S \setminus S_0| \geq \epsilon n$, the min-cut is at least $\alpha\epsilon n$, leading to at least $\alpha\epsilon n$ disjoint paths between the two sets. Since the total number of edges cannot exceed $n\bar{d}/2$, at least half of the paths have length bounded above by $\frac{\bar{d}}{\alpha\epsilon}$. Hence, the probability that S_0 is disconnected from $S \setminus S_0$ in $G(\beta)$ is at most $(1 - \beta^{\bar{d}/\alpha\epsilon})^{\alpha\epsilon n/2}$. There are at most $2^{n/\kappa}$ possible choices of S_0 , and hence the overall probability is bounded above by $2^{n/\kappa}(1 - \beta^{\bar{d}/\alpha\epsilon})^{\alpha\epsilon n/2} \leq \exp(n/\kappa - \beta^{\bar{d}/\alpha\epsilon} \alpha\epsilon n/2)$ as desired. $\square$

Note that the sprinkling argument is crafted for any fixed $\epsilon > 0$. Thus, when proving the uniqueness and locality of the giant component, it suffices to consider a fixed $\epsilon > 0$ and use the ϵ -large-set expansion (which depends on ϵ). With this in mind, we consider a locally converging robust sequence $\{G_n\}_{n \in \mathbb{N}}$. The exact same analysis as in [14] carries over and suggests the following proposition.

Proposition D.2.2. *For any $\epsilon > 0$ and any continuity point $p \in [0, 1]$ of ζ ,*

$$\mathbb{P}\left(\left|\frac{1}{n}|C_{\max}(G_n(p))| - \zeta(p)\right| \geq \epsilon\right) \rightarrow 0. \quad (\text{D.1})$$

D.2.2 Continuity of the Percolation Function ζ

Given Proposition D.2.2, the convergence of the relative size of $C_{\max}(G_n(p))$ as stated in Theorem 5.2.3 only requires us to establish the continuity of the percolation function ζ , which, again, nearly follows from the analysis in [14]. The slight generalization of large-set expansion into structural robustness leaves the proof unaffected other than the special attention to pay when we decide the expansion parameter.

First, we observe that the local limit of a robust sequence is ergodic.

Lemma D.2.3 (Cf. Lemma A.1, [14]). *Let $\{G_n\}_{n \in \mathbb{N}}$ be a robust sequence of (possibly random) graphs with bounded average degree converging locally to $(G, o) \in \mathcal{G}^*$ with law μ . Then μ is ergodic, and thus extremal among unimodular probability measures on $\mathcal{G}^*$.*

Proof sketch. The proof is almost exactly the same as the proof of Lemma A.1 in [14], except for the following minor difference: Once we fix $\epsilon \leq p_0/8$, the ϵ -large-set expansion of our robust sequence will be asymptotically $\alpha = \alpha(\epsilon)$, instead of some absolute quantity (independent of ϵ) as in the original proof. This, however, causes no issue for us as the rest of the proof relies on the fixed values of ϵ and α . $\square$

In particular, if μ is the local limit of a robust sequence, then the critical probability of percolation on $(G, o) \sim \mu$ is constant μ -a.s. Let $p_c = p_c(\mu)$ denote this critical probability, i.e.,

$$\inf\{p : \mathbb{P}(|C(o; G(p))| = \infty | (G, o)) > 0\} = p_c \quad \text{a.s.} \quad (\text{D.2})$$

By Theorem 6.7 in [11], for all p_1, p_2 satisfying $p_c < p_1 < p_2 \leq 1$, every infinite cluster in $G(p_2)$ contains an infinite cluster in $G(p_1)$ (under the standard coupling of percolation) μ -a.s. This allows us to immediately conclude that ζ is continuous at all $p \neq p_c$.

Proposition D.2.4 (Continuity of ζ . Cf. [14, Corollary 2.2]). *Let μ be the local limit of a robust sequence, and $\zeta : [0, 1] \rightarrow [0, 1]$ be the corresponding percolation function. Then ζ is continuous at all $p \in [0, 1]$ except possibly at $p = p_c$.*

Theorem 5.2.3 now follows immediately from Propositions D.2.2 and D.2.4. Further, a slightly more involved analysis suggests that, for all $p \in [0, 1]$ (including $p = p_c$), the second largest component $C_2(G_n(p))$ has sublinear size whp (cf. [14, Appendix C]).

Appendix E

Appendix to Chapter 6

E.1 Preliminaries

E.1.1 Single Community Results

To state the needed single community results, we recall our notation $R_0 = \frac{c\beta}{\beta+\gamma}$ for the basic reproduction number and $\lambda = (c-1)\beta - \gamma$ for the exponential growth rate. We will use $\mathcal{T}_{\text{Poi}(c)}$ to denote a random tree where each node has a number of offsprings sampled independently from the Poisson distribution with rate c , and $\text{BP}(c)(t)$ to denote the size of the infection tree at time t of an SIR infection on $\mathcal{T}_{\text{Poi}(c)}$ starting at the root.

Next, we note that if we start an infection with at least ϵn infected vertices on a single $G(n, c/n)$ network, then the SIR dynamic obeys a law of large numbers which can be expressed in terms of the solution of the following system of differential equations

$$\frac{ds}{dt} = -\beta xs, \quad \frac{dx}{dt} = -(\beta + \gamma)x + \beta cxs, \quad \frac{di}{dt} = \beta cx - \gamma i, \quad \frac{dr}{dt} = \gamma i \quad (\text{E.1})$$

where s , i and r represent the limit of $\frac{1}{n}S(t)$, $\frac{1}{n}I(t)$ and $\frac{1}{n}R(t)$ as $n \rightarrow \infty$, and x is the limit of $\frac{1}{n}X(t)$, where $X(t)$ captures the force of the infection, i.e., the rate at which a random vertex in S gets infected, see, e.g., [30, 47] for the derivation of this equations and the proof of the law of large numbers.

Finally, we define two stopping times: the time $\tau(\epsilon)$ when the infection has reached relative ϵ where $0 < \epsilon < r_\infty$,

$$\tau(\epsilon) := \inf\{t > 0 : S(t) < (1 - \epsilon)n\}, \quad (\text{E.2})$$

and the time τ_∞ when the infection dies out,

$$\tau_\infty = \sup\{t > 0 : I(t) > 0\}, \quad (\text{E.3})$$

both for an infection starting from a single infected individual at time 0. We call $R(\infty)$ the *size of the outbreak*, and say it is a *large outbreak* if $R(\infty) \geq \ln n$.

Theorem E.1.1. *Consider the SIR epidemic spreading on an Erdős–Rényi random graph $G(n, c/n)$. Assume there is one initially infected vertex unless stated explicitly otherwise. Then the following hold:*

- (a) **(Probability of outbreak)** *The probability of a large outbreak of at least $\ln n$ vertices in the graph converges to π . That is*

$$\mathbb{P}(R(\infty) \geq \ln n) \rightarrow \pi \text{ as } n \rightarrow \infty. \quad (\text{E.4})$$

Further, for all $\epsilon \in (0, r_\infty)$

$$\lim_{n \rightarrow \infty} \mathbb{P}(R(\infty) \geq \epsilon n \mid R(\infty) \geq \ln n) = 1. \quad (\text{E.5})$$

- (b) **(Initial growth)** *The SIR epidemic on $G(n, c/n)$ and the branching processes $\text{BP}(c)$ and $\text{BP}(c(1 - \epsilon))(t)$ can be coupled in such a way that the following holds whenever $N_n = o(\sqrt{n})$.*

$$\mathbb{P}(|\text{BP}(c)(t)| = (I(t) + R(t)) \text{ for all } t \text{ s.t. } |\text{BP}(c)(t)| \leq N_n) \rightarrow 1 \text{ as } n \rightarrow \infty.$$

Furthermore, for any $a \in (0, 1)$, conditioned on a large outbreak, we have

$$\frac{\tau(n^{a-1})}{\lambda^{-1} \ln n} \xrightarrow{p} a.$$

- (c) **(Final size)** *Conditioned on not having a large outbreak, $R(\infty)$ is bounded in probability, and conditioned on a large outbreak*

$$\frac{R(\infty)}{n} \xrightarrow{p} r_\infty.$$

- (d) **(Continuity of r_∞ and π)** *Both $r_\infty = r_\infty(R_0)$ and $\pi = \pi(c, \beta, \gamma)$ are continuous functions, and both are strictly positive if and only if $R_0 > 1$.*

- (e) **(Duration of the epidemic)** *Conditioned on not having a large outbreak, τ_∞ is bounded in probability, and conditioned on a large outbreak, $(\ln n)^{-1} \tau_\infty$ is bounded in probability.*

- (f) **(Growth of a linear-sized infection)** *There exist constants $\epsilon_0 > 0$ and $0 < C_1 < C_2 < \infty$ such that for $0 < \epsilon \leq \epsilon_0$, conditioned on a large outbreak, whp $C_1 \epsilon \leq X(\tau(\epsilon)) \leq C_2 \epsilon$. Furthermore, conditioned on the state of the infection at time $\tau(\epsilon)$, the evolution of the SIR dynamic is well described by set of equations (E.1), in the sense that*

$$\sup_{t \geq \tau(\epsilon)} \left| \frac{1}{n} (S(t), X(t), I(t)) - (s(t), x(t), i(t)) \right| \xrightarrow{p} 0.$$

- (g) (**Concentration for τ**) For $0 < \epsilon' < \epsilon_0$ and $\epsilon' < \epsilon < r_\infty$, there exists $K = K(\epsilon', \epsilon)$ such that, conditioned on a large outbreak, whp $\tau(\epsilon) - \tau(\epsilon') < K$. In particular, combined with (b), this implies that conditioned on a large outbreak

$$\frac{\tau(\epsilon)}{\lambda^{-1} \ln n} \xrightarrow{p} 1.$$

- (h) (**Many initially infected vertices**) Let $0 < \epsilon < r_\infty$, and assume $1 \ll I(0) \ll n$, with the initially infected vertices chosen independently from the vertex set of $G(n, c/n)$. Then

$$\frac{R(\infty)}{n} \xrightarrow{p} r_\infty$$

Proof. Except for the last statements in (b) and (e) and both statements in (g), the proofs of the rest of the statements are standard and can be found in the literature, e.g. in Neal [159], Ball and Britton [30], and Borgs, Huang, and Ikeokwu [47].

The proof of the last statement of (b) follows from standard theory of branching processes once we observe that we can upper bound $I(t) + R(t)$ by the size of an infection tree on the random tree $\mathcal{T}_{\text{Bin}(n,p)}$ with $\text{Bin}(n, p)$ offspring distribution, and similarly lower bound it by the size of an infection tree on $\mathcal{T}_{\text{Bin}(n',p)}$ up to the stopping time when $I(t) + R(t)$ first exceeds $n - n'$. The proof of (g) follows by combining these arguments with the law of large numbers behavior of the process. Specifically, combining the just mentioned upper and lower bounds on $I(t) + R(t)$ with standard branching process theory, we get that for $0 < \delta < 1$, whp,

$$(1 - \delta)\lambda^{-1} \leq \frac{\tau(\epsilon')}{\ln n} \leq \frac{1 + \delta}{\lambda - c\beta\epsilon'}.$$

Finally, using the bound $\epsilon' C_1 < X(\tau(\epsilon')) < \epsilon' C_2$ from (f) and the law of large numbers for $n^{-1}S(t)$ and $n^{-1}X(t)$, it is straightforward to prove that whp the epidemic will grow from size (measured in $I(t) + S(t)$) $\epsilon'n$ to ϵn in finite time, i.e., in time at most $K(\epsilon', \epsilon) < \infty$. Sending first n to infinity and then δ and ϵ' to zero, one gets that

$$\frac{\tau(\epsilon)}{\lambda^{-1} \ln n} \xrightarrow{p} 1.$$

The last statement of (e) is easily proven by first noting that (g) implies that it takes time of order $\ln n$ to pass the herd immunity threshold and then using the arguments of, e.g., Borgs, Huang, and Ikeokwu [47] that after passing the herd immunity threshold $X(t)$ and $I(t)$ decay exponentially in t . $\square$

E.1.2 Useful Couplings for SIR on Random Graphs

Next, we recall several couplings for the SIR dynamics which will be useful in this chapter.

1. **(Infection digraph)** This coupling is defined for an SIR infection on an arbitrary finite graph G , and involves the construction of an infection digraph, to be denoted $\mathcal{D}_G^{SIR}$. It is defined by drawing recovery times $L_v \sim \text{Exp}(\gamma)$ i.i.d. for all vertices v as well infection times $W_{uv} \sim \text{Exp}(\beta)$ i.i.d. for all oriented edges uv such that $\{u, v\}$ is an edge in G . The infection digraph $\mathcal{D}_G^{SIR}$ is then defined by drawing an edge from i to j whenever $W_{ij} \leq L_i$, giving edges (i, j) a weight or length W_{ij} . If we infect a single vertex v at time 0, the time the infection has reached vertex w is then equal to the length $t_{v \rightarrow w}$ of the shortest oriented path from v to w in $\mathcal{D}_G^{SIR}$. More generally, in a setting where we infect several vertices, w will get infected at time $t = \min_i (t_i + t_{v_i \rightarrow w})$ where the minimum goes over all seeds, and t_i is the time vertex v_i gets infected (from an external source).
2. When G is random itself, it is often useful to couple the randomness of G and that of G in a way which more tightly represents the course of the infection. In the setting of this section, where $G \sim G(n, p)$, this can be done in several ways.
 - a) **(Poisson coupling)** Here one first defines a multi-graph analog of $G(n, p)$ where edges have a multiplicity given by a Poisson random variable $\text{Poi}(c/n)$. Next, when a vertex gets first infected, one attaches $d \sim \text{Poi}(c)$ half-edges to the infected vertex, each running an infection clock at rate β , until the vertex recovers. When an infection clock on a half-edge clicks, we stop the Poisson clock, chose an endpoint for the half-edge uniformly among all vertices, and if this endpoint happens to be susceptible, infect the endpoint. After establishing all desired results for the Poisson random graph, one finally conditions onto the Poisson graph being simple to get the desired results for $G(n, p)$. This approach is used, e.g., when deriving the differential equations (E.1), with $X(t)$ representing the number of infection clocks which are running at time t ; see Ball and Britton [30] and Borgs, Huang, and Ikeokwu [47] for more details.
 - b) **(Binomial coupling to susceptible neighbors)** In our context, the Poisson coupling is less useful, since it does not generalize well in the presence of travel clocks. Here the most natural coupling proceeds by drawing the neighbors of an infected vertex when it first gets infected. More precisely, we will draw the neighbors from all currently susceptible vertices, including each of them independently with probability c/n , then draw the recovery time L_v for the infected vertex, and run infection clocks at rate β along the edges to susceptible vertices until the vertex recovers. Removing edges that used to point to susceptible vertices as they get infected, active edges, i.e., edges with a running β clock, will always point to susceptible vertices, and when one of them clicks, we infect the new vertex and iterate. Note that in this setting, when a vertex gets infected at time t , it starts with $d \sim \text{Bin}(S(t), c/n)$ active edges, and when the first edge of an infected vertex clicks, its endpoint is uniform among the susceptible individuals at that time.

- c) **(Binomial coupling to all vertices)** A variant of the last coupling draws a larger set of edges when a vertex first gets infected, namely i.i.d. with probability c/n from all n vertices including thus the possibility of a self-loop to the infected vertex itself. Infections along active edges should then be considered infection attempts, and will be ignored if the endpoint of such an infection is already infected. An advantage of this coupling is that we could first draw the number d of active half-edges according to $d \sim \text{Bin}(n, c/n)$, and only commit which edges we choose once the first infection clock clicks, at which point we choose d endpoints uniformly without replacement. In this way, the first infection attempt will be uniformly among all n vertices, which turns out to be useful for some of our bounds.

E.1.3 Single Community with Multiple Seeds

The next lemma gives upper and lower bounds on the spread of the SIR epidemic in the case of multiple initial “seed” infections appearing at times t_v that are not necessarily all equal (as they were in the setting of Statement (h) of the last theorem). It will be a key ingredient we use to tie the development in the two communities together.

Lemma E.1.2 (Outbreak in a single community with multiple seeds). *Consider an SIR epidemic over $G \sim G(n, c/n)$ with infection rate β and recovery rate γ , assume $R_0 > 1$, and let $\delta > 0$ and $\epsilon \in (0, r_\infty)$ be fixed, i.e., independent of n . Assume that the SIR dynamic evolves as before, except that at times $0 \leq t_1 < t_2 < \dots$ we seed additional infections at some vertices $v_1, v_2, \dots$ in $V(G)$.*

1. *If a set of possible seeds $\mathcal{S}$ of size $|\mathcal{S}| = o(n)$ is chosen independently of the draw $G \sim G(n, c/n)$ and before the start of the epidemic and, all seeds v_i are in $\mathcal{S}$, then whp*

$$\frac{R(\infty)}{n} \leq r_\infty + \delta.$$

2. *If for a subset $\mathcal{S}' \subseteq \mathcal{S}$ with $|\mathcal{S}'| \gg 1$, the seeds within $\mathcal{S}'$ are chosen by time t_0 and i.i.d. uniformly at random in the n vertices, then then whp*

$$\tau(\epsilon) \leq t_0 + (1 + \delta) \frac{\ln n}{\lambda} \quad \text{and} \quad \frac{R(\infty)}{n} \geq r_\infty - \delta.$$

Proof. The first statement is an easy consequence of the representation of the epidemic in term of the infection digraph $\mathcal{D}_G^{\text{SIR}}$ and the observation that in this representation, the final set of infected vertices is just the set of vertices that can be reached from the set of seeds. Since this set does not depend on the infection times t_v , the statement follows from statement (h) of the previous theorem.

For the next part, it suffices to prove that for any given $\delta' > 0$ and any M sufficiently large, the probability of $\tau(\epsilon) \leq t_0 + (1 + \delta) \ln n / \lambda$ is at least $1 - \delta'$. Let $\mathcal{A}_i$ be the event that

the SIR epidemics with the single seed v_i at time t_i would satisfy the second statements in the lemma, and let $\mathcal{B}_i$ be the event that the SIR epidemics with seed v_i would terminate before infecting L individuals for some $L < \infty$ to be specified. Note that in the above coupling where we pre-sample the infection and recovery clocks, we can ensure that $\bigcup_{i=1}^M \mathcal{A}_i$ implies the desired events: if any seed infection is capable of producing a large outbreak, then having additional seeds of infection only expedites the spread. Thus, it suffices to bound $\mathbb{P}(\bigcap_{i=1}^N \mathcal{A}_i^c)$ to be at most δ' for all $\delta' > 0$.

Using inclusion-exclusion, we first bound

$$\begin{aligned} \mathbb{P}\left(\bigcap_{i=1}^M \mathcal{A}_i^c\right) &= \mathbb{P}\left((\mathcal{B}_1 \cap \mathcal{A}_1^c) \bigcap_{i=2}^M \mathcal{A}_i^c\right) + \mathbb{P}\left((\mathcal{B}_1^c \cap \mathcal{A}_1^c) \bigcap_{i=2}^M \mathcal{A}_i^c\right) \leq \mathbb{P}\left(\mathcal{B}_1 \cap \bigcap_{i=2}^M \mathcal{A}_i^c\right) + \mathbb{P}\left(\mathcal{A}_1^c \setminus \mathcal{B}_1\right) \\ &\leq \mathbb{P}\left(\bigcap_{i=1}^M \mathcal{B}_i\right) + \sum_{i=1}^M \mathbb{P}\left(\mathcal{A}_i^c \setminus \mathcal{B}_i\right), \end{aligned} \tag{E.6}$$

and then note that by statements (c), (e) and (g) of the previous theorem, we can make $\mathbb{P}(\mathcal{A}_i^c \setminus \mathcal{B}_i)$ as small as desired by choosing L and n sufficiently large. Thus given M , for L and n sufficiently large, we can bound the second term on the right by at most $\delta'/2$.

To bound the first term, we first observe that for the first seed, the bad event $\mathcal{B}_1$ is just the event that $R(\infty) < L$ on the graph $G(n, c/n)$ with a single random seed. By the fact that conditioned on not having a large outbreak, $R(\infty)$ is bounded in probability, and the fact that the probability of a large outbreak converges to π in probability, we have that for L and n sufficiently large, $\mathbb{P}(R(\infty) < L)$ can be bounded by, say, $1 - \pi/2$. It turns out we need slightly more, namely that

$$\mathbb{P}(R(\infty) < L) \leq 1 - \frac{\pi}{2} \quad \text{on } G(n', c/n) \quad \text{whenever } n' \geq n - \ln n \tag{E.7}$$

and L and n are sufficiently large. To see that this holds, we first note that the probability $\mathbb{P}(R(\infty) < L)$ is decreasing in n' , showing that $n' = n - \ln n$ is the worst case. The proof now follows once we use that $G(n', c/n) = G(n', c'/n')$ with $c' = cn'/n$, together with the fact that $\pi(c', \beta, \gamma)$ is continuous in c' .

To prove an upper bound on the first term in (E.6), we now use the following coupling in which we expose the set of vertices $C_+(v_i)$ that can be reached from the seed v_i in $\mathcal{D}_G^{SIR}$ by the usual breath first search (BFS) starting from v_i . In this process, whenever a vertex v gets infected, we first draw its recovery clock, then draw the edges to the remaining susceptible vertices i.i.d. with probability c/n , and then draw the infection clocks for these edges, keeping only those edges for which the infection clock clicks before the recover clock, and declaring the end point infected. If at any point in this process, we have discovered at least L infected vertices, we declare success, and if the process dies before infecting L vertices, we conclude that $|C_+(v_i)| < L$, declare the vertices in $C_+(v_i)$ used, and move on to the next seed v_{i+1} . Furthermore, since we are interested in an upper bound on the probability of failure event $\mathcal{B} = \bigcap_{i=1}^M \mathcal{B}_i$, we will additionally add $k = L - |C_+(v_i)|$ random vertices to the

set of used vertices, and declare failure if the new seed v_{i+1} falls into the set of used vertices. We furthermore declare failure if the BFS process starting from v_{i+1} discovers one of the used vertices. In this way, the failure probability in step i (assuming failure in all previous steps) can be bounded from above by $\frac{(i-1)L}{n}$ plus the probability that a BFS exploration on $\mathcal{D}_{G(n-(i-1)L, c/n)}^{SIR}$ fails before reaching L vertices, which in turn is nothing but the probability that $R(\infty) < L$ on $G(n - (i-1)L, c/n)$. Combined with the bound (E.7), we see that

$$\mathbb{P}\left(\bigcap_{i=1}^M \mathcal{B}_i\right) \leq \prod_{i=1}^M \left(1 - \frac{\pi}{2} + \frac{(i-1)L}{n}\right) \leq \exp\left(-\sum_{i=1}^M \left(\frac{\pi}{2} - (i-1)\frac{L}{n}\right)\right) \leq e^{-M\pi/3}$$

provided $\frac{LM}{n} \leq \pi/6$. Choosing first M large enough so that the right hand side is smaller than $\delta'/2$, and then L and n large enough to guarantee that the bound (E.7) holds and the second term in (E.6) is smaller than $\delta'/2$ and finally increasing n if needed to make sure that $\frac{LM}{n} \leq \pi/6$ and $ML \leq \ln n$, we obtain the desired bound $\mathbb{P}(\bigcap_{i=1}^N \mathcal{A}_i^c) \leq \delta'$. $\square$

Lemma E.1.3. *Consider an SIR epidemic over $G \sim G(n, c/n)$ with infection rate β and recovery rate γ , assume $R_0 > 1$, let $a \in (0, 1)$, and let $0 < \delta < a$ and let x be an arbitrary positive number. Then with probability tending to π , $\tau(n^{a-1}) < \infty$ and among these first $\lceil n^a \rceil$ infected nodes, at least $n^{a-\delta}$ will have infectious periods of at least x (i.e., having a recovery time at least x). With probability tending to $1 - \pi$, the process dies out before reaching n^a vertices (hence $\tau(n^{a-1}) = \infty$).*

Proof. Note that the second part is a direct consequence of Theorem E.1.1 (Statement (a) in particular).

It is helpful to consider an augmented process where a fresh new node will be infected as a new seed whenever the SIR process dies out (and this gets repeated). More precisely, whenever the number $I(t)$ of actively infectious vertices drops to zero, if $S(t) > 0$, we pick a vertex v among the remaining susceptible ones uniformly at random, and infect it at time $t + 1$. This way the process runs until the entire population have been infected once, and we can be sure to reach n^a number of infections eventually. The true SIR process is simply the stopped process on the augmented process upon first extinction (i.e., number of active infections dropping to zero).

In the context of the augmented process, denote by $\tilde{\tau}(\epsilon)$ the stopping time for an ϵ -fraction infection, analogous to our $\tau(\epsilon)$ in the original process. Note that $\tilde{\tau}(n^{a-1}) < \infty$ a.s. Denote by $\tilde{\tau}_0$ the first time infections drop to zero. Per our observation above, our original SIR process is simply the augmented process stopped at $\tilde{\tau}_0$. Further, the coupling between the two process is exact until $\tilde{\tau}_0$; in particular, whenever $\tilde{\tau}(n^{a-1}) < \tilde{\tau}_0$ (that is, when the standard SIR process leads to a large outbreak), $\tilde{\tau}(n^{a-1}) = \tau(n^{a-1})$ and those nodes infected by then share identical lives in these two version.

It now suffices to show the following event $\tilde{\mathcal{E}}$ happens whp in the augmented process: among the first $\lceil n^a \rceil$ infected nodes, at least $n^{a-\delta}$ will have infectious periods at least x . To derive our target statement in the lemma, the event $\tilde{\mathcal{E}}$ along with $\tilde{\tau}(n^{a-1}) < \tilde{\tau}_0$ implies

that $\tau(n^{a-1}) = \tilde{\tau}(n^{a-1}) < \infty$ and, as these $\lceil n^a \rceil$ nodes have the same infection periods in the original SIR as in the augmented process, at least $n^{a-\delta}$ of them have infectious periods at least x in the original SIR process. Thus, the probability of our original event is lower bounded by $\mathbb{P}(\tilde{\tau}(n^{a-1}) < \tilde{\tau}_0, \tilde{\mathcal{E}}) \geq \mathbb{P}(\tilde{\tau}(n^{a-1}) < \tilde{\tau}_0) - \mathbb{P}(\tilde{\mathcal{E}}^c)$, which has a limit π once we show $\mathbb{P}(\tilde{\mathcal{E}}^c) \rightarrow 0$. (On the other hand, π as an upper bound is trivial, since we know that the original SIR process dies out with asymptotic probability $1 - \pi$.)

It remains to understand the event $\tilde{\mathcal{E}}$. By our definition of the infection process, each node has an $\text{Exp}(\gamma)$ recovery clock independent from all other clocks; in other words, when a node is infected, no matter what past event we condition on (including other infection clocks we have seen, past extinction events, etc.), its recovery time is $\text{Exp}(\gamma)$ distributed. Thus, the recovery times for the first $\lceil n^a \rceil$ infected nodes form a sequence of $\lceil n^a \rceil$ i.i.d. $\text{Exp}(\gamma)$ samples¹. By Hoeffding's inequality,

$$\mathbb{P}(\tilde{\mathcal{E}}^c) = \mathbb{P}_{N \sim \text{Bin}(\lceil n^a \rceil, e^{-\gamma x})}(N < n^{a-\delta}) \leq \exp\left(-\frac{2(\lceil n^a \rceil e^{-\gamma x} - n^{a-\delta})^2}{\lceil n^a \rceil}\right),$$

which tends to zero as $n \rightarrow \infty$ for any constant $x > 0$. This finishes our proof. $\square$

E.1.4 Useful Perturbation Bounds on Matrix Exponential

In later analysis, we need to understand the growth of the solutions to certain system of renewal equations (or inequalities). The solution will involve a matrix exponential of the form e^{tM} with some 16×16 matrix M in our case, whose leading behavior will be given by the dominant eigenvalue of M , i.e., the eigenvalue with the largest real part. Unfortunately, the actual matrix M is hard to diagonalize directly. To deal with this, we will write M as $M_0 + W$, where M_0 is obtained by setting all terms of order $\tilde{O}(n^{-\alpha})$ to zero, and then use the following lemma to control the difference between e^{tM} and e^{tM_0} .

Lemma E.1.4. *Let M_0 and W be real matrices, and let $\|\cdot\|$ be a submultiplicative matrix norm. Assume that $\lambda \geq 0$ is an eigenvalue of M_0 , and that all other eigenvalues of M_0 have real parts less than or equal to λ . Define*

$$C := \sup_{t \geq 0} e^{-\lambda t} \|e^{tM_0}\|.$$

Then for all $t \geq 0$, we have

$$\|e^{t(M_0+W)} - e^{tM_0}\| \leq C (e^{C\|W\|t} - 1) e^{\lambda t} = O(\|W\| e^{C\|W\|t}) e^{\lambda t}.$$

Proof. For normal matrices, a bound of the above form (with $C = 1$) follows from Theorem 4.2 in [203]. Unfortunately, without such the normality assumption, the results from [203]

¹Another way to see this is to presample an infinite sequence of i.i.d. recovery time. Each time a new vertex gets infected, it just chooses the next recovery time on the list, which means the sequence we see is an i.i.d. sequence.

only imply a bound of the form $O(\|W\|e^{C\|W\|t})e^{\kappa t}$, where $\kappa = \sup_{t>0} \frac{1}{t} \log \|e^{tM_0}\| > \lambda$, which is too weak for our purpose. We therefore include the (straightforward) proof here.

We use the Dyson expansion for the matrix exponential:

$$e^{t(M_0+W)} = e^{tM_0} + \sum_{k=1}^{\infty} \int_{\Delta_k(t)} e^{(t-s_1)M_0} W e^{(s_1-s_2)M_0} \dots W e^{s_k M_0} ds_1 \dots ds_k,$$

where $\Delta_k(t) = \{0 \leq s_k \leq \dots \leq s_1 \leq t\}$ is the standard k -simplex of time-ordered variables. By submultiplicativity of the norm and the assumption on C , we estimate the norm of the k -th term as

$$\|e^{(t-s_1)M_0} W \dots W e^{s_k M_0}\| \leq C^{k+1} \|W\|^k e^{\lambda t}$$

where we used that the sum of the increments $(t - s_1) + (s_1 - s_2) + \dots + s_k$ is t . Since the volume of $\Delta_k(t)$ is $\frac{t^k}{k!}$, each term in the series is bounded by $\frac{t^k}{k!} C^{k+1} \|W\|^k e^{\lambda t}$, giving the desired bound

$$\|e^{t(M_0+W)} - e^{tM_0}\| \leq C e^{\lambda t} \sum_{k=1}^{\infty} \frac{(C\|W\|t)^k}{k!} = C (e^{C\|W\|t} - 1) e^{\lambda t}.$$

□

E.2 Branching Process Upper Bounds

In addition to the properties of SIR on a single community random graph $G(n, p)$, we will rely on approximations of SIR processes with properly constructed branching process to derive upper bounds on the growth of the outbreak. The key intuition is as follows. When the epidemic initially starts to spread, all individual in $V_1 \cup V_2$ are susceptible and can be infected. As the epidemic evolves, less and less vertices are susceptible, and hence can be infected. One might thus expect that if we replace the number of susceptible vertices in each community by n (in other words, allowing vertices to get infected multiple times), we could get upper bounds on the speed at which the epidemic spreads. Thus, one might hope that the following process gives upper bounds.

Start with one infected individual in community 1 (who whp will not be traveling, and hence can be assume to be in V_1). Upon each new infection, we explore the neighbors of the newly infected in our random, two-community network, where by definition the intra-type and inter-type (forward) degrees should follow a distribution of $\text{Bin}(n-1, c/n)$ and $\text{Bin}(n, c/n)$, respectively. Since these neighbors may or may not get infected later on, we think of them as possible, or not yet released children in a branching process. For our branching process, as mentioned above, we allow infected (or already recovered) individuals to get re-infected, and record the number of (re-)infections with multiplicity to type- i individuals up to time t as $\tilde{Y}_i(t)$ for $i \in \{1, 2\}$. Clearly, $\tilde{Y}_i(t)$ stochastically dominates $Y_i(t)$ as the re-infections are simply ineffective and hence ignored in the actual SIR process.

A naive way to treat the processes $\tilde{Y}_1(t)$ and $\tilde{Y}_2(t)$ is to view them as a branching process with travel, where each vertex sees a random number of potential children of the same or different types, each vertex has independent travel trajectories, and the potential children are realized at times determined the infection and recovery clocks (or possibly never if recovery takes place before an infection can strike). This approach bears one issue: in the actual SIR process, the travel state of a neighbor to a newly infected node can depend on the history. For instance, consider a neighbor u of a newly infected v . It is possible that u is also a neighbor of a previously infected vertex w , who attempted to infect u but failed because u was traveling in a different community. Thus, to ensure a valid coupling to our original SIR process, we must incorporate past information on the vertices into the branching process, and sample the travel state conditional upon such information.

However, the complication above only adds to the complexity of describing the processes $\tilde{Y}_1(t)$ and $\tilde{Y}_2(t)$ ². Nevertheless, for our purposes, it suffices to obtain upper bounds on the *expectations* of these two quantities, in which case standard tools for studying generalized branching processes remain useful.

To formalize the idea described above, we introduce the following notation. First, we will use an augmented type (k, A) to describe a node v of type $k \in \{1, 2\}$ and infected while in travel state $A \in \{\mathbf{H}, \mathbf{T}\}$. For $k \in \{1, 2\}$ and $A \in \{\mathbf{H}, \mathbf{T}\}$, let $Y_{k,A}(t)$ denote the total number of nodes of type (k, A) that have been infected up to time t . Note that our previously defined $Y_k(t)$ is simply $Y_{k,\mathbf{H}}(t) + Y_{k,\mathbf{T}}(t)$. Also, define a stopping time τ to be the earliest time when the number of active travelers exceeds $2n^{1-\alpha} \ln n$. A standard supermartingale argument reveals the following bound; the proof is standard and hence omitted.

Lemma E.2.1. *Whp, $\tau > T = \ln^2 n$.*

Next, we define a dominating branching process with cloned travel trajectories, which will provide an almost sure upper bound on the SIR process. Like the original SIR process, we have a population of two types of nodes undergoing travel dynamics identical to that of our original SIR model, and our infection again starts from a single seed v_0 . We model replenishment of susceptible nodes as follows: Each time a vertex v in our original population gets infected, we make a *susceptible* clone u' for each already infected (or recovered) node u in the original $2n$ population such that u' shares the exact same type and travel trajectory as u ; edges between v and the clones and the still susceptible nodes are drawn independently with probability c/n ; for the potential children of u , i.e., susceptible nodes or clones of infected nodes that are linked to v , the infection clocks are set in the usual way (running whenever the two endpoints are in the same physical community and paused when they are not), and the recovery clock on v is also set upon infection. When a cloned vertex v gets infected, we make a susceptible clone u' for every node in the original population and run infections

²These are technically no longer branching processes in the classic sense due to the additional memory. It is possible to view them as a generalized branching process, or a Crump-Mode-Jagers process (see Crump and Mode [64] and Jagers [110]), which allows for dependencies between offsprings distributions and history of the process.

and recovery as usual. Observe that by construction, any of the $2n$ original nodes can only get infected by other original nodes, and hence by dropping the cloned copies, we recover our original SIR process with travel. Thus, the total number of infection events to (possibly cloned) nodes of type k in travel state A up to time t is an almost sure upper bound on $Y_{k,A}(t)$.

To allow for an easier analysis, we make one further adjustment to the branching process: when cloning infected nodes upon a newly infected node (original or cloned), we only allow cloning up to $n^{1-\alpha} \ln n$ currently traveling nodes of each type. Importantly, if our stopping time $\tau > t$, then this constraint makes no difference up to time t . In other words, under the event that $\tau > t$, the previously stated almost sure domination still holds. This branching process with constrained cloning of travelers will be our subject to study, and we will define $s\tilde{Y}_{k,A}(t)$ as the total number of infection events to (possibly cloned) nodes of type k in travel state A up to time t in this process. By Lemma E.2.1, we now have a coupling between the original processes $Y_{k,A}(t)$ and the constructed $\tilde{Y}_{k,A}(t)$ such that whp $\tilde{Y}_{k,A}(t) \geq Y_{k,A}(t)$ for time t up to $\ln^2 n$. It now suffices to bound the growth of $\tilde{Y}_{k,A}(t)$, which we will focus on next.

To analyze this branching process we constructed, the key is to bound the reproduction measure $\mu_v(s; k, A)$ for each newly infected vertex (or possibly a clone) v . It is defined as the expected number of type- (k, A) infections directly caused by v within period s of v 's infection, conditional on the history at the infection time b_v . Here, the history is defined with respect to the filtration capturing past infection and recovery events, travels, and neighbors discovered up to time t , including links between past infected vertices and v itself, but excluding yet to be discovered neighbors of v .

Lemma E.2.2. *Conditional on the history up to $t = b_v$, the reproduction measure μ_v is bounded as*

$$\mu_v(s; k, A) \leq \mu_{\text{type}(v), \sigma_v(t)}(s; k, A)$$

for all $s < T - t$, where $\sigma_v(t) \in \{\mathbf{H}, \mathbf{T}\}$ denotes the travel state of v at t and

$$\mu_{k,\mathbf{H}}(s; k, \mathbf{H}) = \mu_{k,\mathbf{T}}(s; k, \mathbf{H}) = \mu_{k,\mathbf{T}}(s; 3 - k, \mathbf{H}) = R_0(1 - e^{-(\gamma+\beta)t}) + \tilde{O}(\rho_{\mathbf{T}})$$

for $k = 1, 2$, and for all other combinations $\mu_{k',A'}(s; k, A) = \tilde{O}(\rho_{\mathbf{T}})$, where $\tilde{O}(\cdot)$ hides logarithmic dependency on n and R_0 is defined in (6.1).

Proof. The proof is a standard computation for continuous time Markov processes. The key is to notice that the quantity $\mu_v(s; k, A)$ depends only on the current states of all the vertices. By linearity of expectation, it suffices for us to understand the probability that a node u (either still susceptible, or a cloned copy of an infected one) of type k in state A gets infected directly by v within a period of s . To illustrate this, consider the case of $\text{type}(u) = \text{type}(v) = 1$ and $\sigma_u(t) = \sigma_v(t) = \mathbf{H}$. By construction, the edge between u and v is present with probability c/n . The joint state of v and u can be modeled as a continuous-time Markov chain on 7 states, with four transient states corresponding to different travel states

(σ_v, σ_u) of the two individuals while v remains infectious and u remains uninfected, and three absorbing states $\mathbf{R}$ for v 's recovery and $\mathbf{l}_1$ and $\mathbf{l}_2$ for u 's infection in community 1 and 2, respectively. Write $p(s)$ for the hitting probability of the $\mathbf{l}_1$ by time s . A straightforward computation (e.g., with Kolmogorov forward equations) yields that

$$p(s) = \int_0^s \mathbf{e}_1^\top \exp \left(s' \begin{pmatrix} -\beta - \gamma - 2\rho_\mathbf{T} & \rho_\mathbf{T} & \rho_\mathbf{T} & 0 \\ \rho_\mathbf{H} & -\gamma - \rho & 0 & \rho_\mathbf{T} \\ \rho_\mathbf{H} & 0 & -\gamma - \rho & \rho_\mathbf{T} \\ 0 & \rho_\mathbf{H} & \rho_\mathbf{H} & -\beta - \gamma - 2\rho_\mathbf{H} \end{pmatrix} \right) (\beta \mathbf{e}_1) ds', \quad (\text{E.8})$$

where $\rho = \rho_\mathbf{H} + \rho_\mathbf{T}$. Following the perturbation theory on matrix exponential, with Lemma E.1.4 as a key technical tool, we derive

$$p(s) \leq \frac{R_0}{n} (1 - e^{-(\beta+\gamma)s}) + O(\rho_\mathbf{T})$$

with $R_0 = c\beta/(\beta + \gamma)$. In the case of u currently traveling, we only need the trivial bound that $p(s) \leq 1$.

The number of non-traveling nodes (or clones) of each type is at most n ; the number of traveling nodes (or clones) of each type is at most $n^{1-\alpha} \ln n$ by construction. Invoking linearity of (conditional) expectation upon all the clones and original nodes give the claimed bounds. $\square$

The rest of the analysis of this branching process is morally straightforward but in practice rather tedious. Instead of providing the full details, we opt to give a high-level overview of the techniques. Other variants of analysis are likely possible and even easier to flesh out.

The first step is to construct a renewal inequality for the expected size of the branching process using the law of total expectation and a Grönwald-type bound. To do this, we express $Y_{k,A}(t)$ as

$$\tilde{Y}_{k,A}(t) = \mathbb{1}\{\text{type}(v_0) = k, \sigma_{v_0}(0) = A\} + \sum_{u: b_u < t} \xi_u(t - b_u; k, A), \quad (\text{E.9})$$

where the summation is over all nodes infected by time t , and $\xi_u(s; k, A)$ is the (stochastic) type- (k, A) reproduction process of a node u , i.e., the number of type- k state- A infections caused by u within time period s of its infection. Effectively, (E.9) aggregates the count of infections by their sources. Take expectations of (E.9), apply the tower property towards each summand (conditioned at the birth of the corresponding node), we obtain

$$\mathbb{E}[\tilde{Y}_{k,A}(t)] = \mathbb{1}\{\text{type}(v_0) = k, \sigma_{v_0}(0) = A\} + \mathbb{E} \left[\sum_{u: b_u < t} \mu_u(t - b_u; k, A) \right],$$

where we note that μ_u is simply the conditional expectation of ξ_u upon its birth. Using our bounds on μ_u and applying Fubini's Theorem to interchange integral and summations, we

get

$$\mathbb{E}[\tilde{Y}_{k,A}(t)] \leq \mathbb{1}\{\text{type}(v_0) = k, \sigma_{v_0}(0) = A\} + \sum_{k',A'} \int_0^t \mathbb{E}[\tilde{Y}_{k',A'}(t-s)] \mu_{k',A'}(ds; k, A), \quad (\text{E.10})$$

a renewal inequality.

It is convenient to vectorize this system of renewal inequalities into a 16-dimensional vector form. Standard techniques for renewal systems allow us to bound the growth as matrix exponential and we can derive a Malthusian upper bound on the growth rate, again using perturbation bounds such as Lemma E.1.4. The Malthusian parameter for this system turns out to converge rapidly to λ (defined in (6.2)) as $n \rightarrow \infty$. We leave such computation to the interested readers.

The following bounds are derived from solving the system (E.10) of renewal inequalities.

Lemma E.2.3. *Under our usual assumption that $\text{type}(v_0) = 1$, $\sigma_{v_0}(0) = \mathbf{H}$, and $\lambda > 1$, for $t < \ln^2 n$,*

$$\mathbb{E}[\tilde{Y}_{k,A}(t)] \leq \mathbb{1}\{k = 1, A = \mathbf{H}\} \frac{c\beta}{\lambda} e^{\lambda t} + \tilde{O}(n^{-\alpha} e^{\lambda t}).$$

As a consequence of Markov's inequality and the coupling between $Y_{k,A}$ and $\tilde{Y}_{k,A}$, we reach the following upper bounds.

Proposition E.2.4. *Whp, $Y_{1,\mathbf{H}}(t) \leq \tilde{Y}_{1,\mathbf{H}}(t) \leq \tilde{O}(e^{\lambda t})$ and $Y_{k,A}(t) \leq \tilde{Y}_{k,A}(t) \leq \tilde{O}(n^{-\alpha} e^{\lambda t})$ for all $(k, A) \neq (1, \mathbf{H})$. In particular, $Y_1(t) = \tilde{O}(e^{\lambda t})$ and $Y_2(t) = \tilde{O}(e^{-\alpha} e^{\lambda t})$.*

The proof simply follows by applying Markov's inequality to the expectation bounds from our previous lemma and noting that $\tilde{Y}_{k,A}$ is a high probability upper bound for $Y_{k,A}$.

E.2.1 Subcritical Processes and Late Stage of Epidemic

It is also useful for us to understand the subcritical case of $\lambda < 1$ for the branching process constructed above, in which case we expect an exponential decay in the actively infectious frontier. This is useful to analyze the late state of the epidemic, when the remaining number of susceptible nodes falls too low to sustain growth. We will describe our approach and the high level results, but again leaving out the full technical details.

Instead of analyzing $\tilde{Y}_{k,A}(t)$, we may also construct a dominating process $\tilde{I}_{k,A}(t)$ for the actively infectious population size $I_{k,A}(t)$ in the SIR process.³ When $\lambda < 0$, we can consider $e^{bt} \tilde{I}_{k,A}(t)$ for any $b \in (0, -\lambda)$, and similarly construct renewal inequalities to show that $\sup_{t \geq 0} e^{bt} \mathbb{E}[\tilde{I}_{k,A}(t)] < \infty$. This suggests an exponential decay in the actively infected population and as a consequence, the growth terminates in logarithm time (in the initial number of infections, which is at most n). Coupling such processes to the late stage, we obtain the following bound.

³We use the same coupling as before. Note that the cloned copies can be all removed without affecting the infections among the original $2n$ population, as long as we do not hit the stopping time τ .

Proposition E.2.5. *Assuming $R_0 > 1$, define a herd immunity threshold s_0 by $s_0 = 1/R_0$. Then, given $\delta > 0$, there exists a constant $C < \infty$ such that conditional on $S_1(t_0), S_2(t_0) \leq (s_0 - \delta)n$, whp, the SIR process dies out by time $t_0 + C \ln n$, i.e., $I_1(t_0 + C \ln n) = I_1(t_0 + C \ln n) = 0$.*

Proof. We provide a proof sketch, commenting on the key steps. First, roll out the original SIR process up to the target time t_0 and conditional on $S_1(t_0)$ and $S_2(t_0)$ both at most $(s_0 - \delta)n$. Note that for each pair of vertex $u \in I(t_0)$ and $v \in S(t_0)$, it is possible to resample the presence of an edge between them conditional on past trajectories (either there is no edge in the underlying graph G , or there is an edge but the infection clock on it has yet to tick), and with basic computation, the conditional probability of the presence of an edge is at most c/n , the original probability. Thus, it is straightforward to construct a dominating branching process with clones starting with the infectious nodes at time t_0 as before such that (1) $\tilde{I}_{k,A}(t)$ is a high probability upper bound on $I_{k,A}(t)$ up to time T , and (2) the reproduction measures for the nodes admit similar upper bounds as in Lemma E.2.2 except that c is now replaced with $c(s_0 - \delta)$ due to the now reduced number of susceptible nodes. This dominating process can be analyzed analogously using the same approach of renewal inequalities. Since $(s_0 - \delta)R_0 < 1$, the resulted Malthusian parameter for the process is $(\beta + \gamma)(R_0 - 1) < 0$ and the process is subcritical. A system renewal inequalities can be formulated for $e^{bt}\mathbb{E}[\tilde{I}_{k,A}(t)]$, which leads to the bound $\sup_{t \geq 0} e^{bt}\mathbb{E}[\tilde{I}_{k,A}(t)] < \infty$. The Proposition now follows from Markov's inequality applied to $\tilde{I}_{k,A}(t)$ and the high probability bound $\tilde{I}_{k,A}(t) \geq I_{k,A}(t)$ for $t < T$. $\square$

Note that from the definition of r_∞ and the concavity of $r \mapsto 1 - r - e^{-R_0 r}$, it is easy to check that $R_0(1 - r_\infty) < 1$ and thus $s_0 > 1 - r_\infty$.

E.3 Travelers and Non-Travelers

On an intuitive level, all our results should follow from the fact that we can divide the set of individuals into a large majority of non-travelers, with the epidemic growth mainly driven by the growth in the two resulting sub-communities, and a smaller population which does travel and transmits the epidemic between communities. It turns out that most of our proofs follow this intuition, except for some of our upper bounds, which require a more detailed understanding of the inter-community interactions and will be carried out with the help of suitable multi-type branching process approximations.

To translate this intuition into a rigorous proof, we consider the set of non-travelers up to time T at which point we expect the epidemic to have died with high probability. Specifically, we set $T := \ln^2 n$ and then define the random sets

$$\mathcal{T}_i = \{v \in V_i \mid v \text{ travels at least once during } [0, T]\},$$

and

$$V_i^0 = V_i \setminus \mathcal{T}_i, \quad V^0 = V_1^0 \cup V_2^0 \quad \text{and} \quad \mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2,$$

as well as the random variables $I_i^0(t)$ and $R_i^0(t)$, denoting the number of infected and recovered individuals in V_i^0 , respectively, writing

$$Y_i^0(t) = I_i^0(t) + R_i^0(t) \quad (\text{E.11})$$

for their sum. The following lemma says that whp, the set of travelers contains only a small proportion of all individuals.

Lemma E.3.1. *With high probability,*

$$|\mathcal{T}_i| = np_{\mathcal{T}}(1 + \rho_{\mathcal{H}} \ln^2 n) \left(1 + \tilde{O}(n^{-\tilde{\alpha}})\right) \leq \frac{1}{2} n^{1-\alpha} \ln^3 n \quad (\text{E.12})$$

where $\tilde{\alpha} = \min\{\alpha, \frac{1}{2}(1 - \alpha)\}$. Furthermore, conditioned on $v \in \mathcal{T}_i$, the probability that v travels at a given time $s \in (0, \ln^2 n)$ is independent of s and equal to

$$\tilde{p}_{\mathcal{T}} = \frac{p_{\mathcal{T}}}{p_{\mathcal{T}} + (1 - p_{\mathcal{T}})(1 - e^{-\rho_{\mathcal{T}} \ln^2 n})} = \frac{1}{1 + \rho_{\mathcal{H}} \ln^2 n} (1 + O(n^{-\alpha})). \quad (\text{E.13})$$

Proof. Each individual in $V_1 \cup V_2$ has probability

$$p_{\geq 1} = p_{\mathcal{T}} + (1 - p_{\mathcal{T}})(1 - e^{-\rho_{\mathcal{T}} \ln^2 n}) = p_{\mathcal{T}}(1 + \rho_{\mathcal{H}} \ln^2 n)(1 + \tilde{O}(n^{-\alpha})) \quad (\text{E.14})$$

of traveling at least once in the time interval $[0, T]$, leading to

$$\mathbb{E}[|\mathcal{T}_i|] = np_{\mathcal{T}}(1 + \rho_{\mathcal{H}} \ln^2 n)(1 + \tilde{O}(n^{-\alpha})).$$

The first statement follows from a multiplicative Chernoff bound.

To prove the second, we note that the unconditional probability of traveling at time s is $p_{\mathcal{T}}$. Since traveling at time $s < \ln^2 n$ means in particular that the first travel of v occurs before $T = \ln^2 n$, the conditional probability is obtained by dividing by $p_{\geq 1}$, proving the statement. $\square$

E.3.1 Early Stage of the Epidemic

Our first result concerns the early stages of the epidemic, which, as it turns out, is with high probability driven by the spread of the epidemic in the community of non-travelers of type 1.

Proposition E.3.2. *Let N_n be such that $N_n(n^{-1/2} + n^{-\alpha} \ln^3 n) \rightarrow 0$ as $n \rightarrow \infty$. Then the two community model with travel can be coupled to $\text{BP}(c)(t)$ such that whp*

$$I(t) + R(t) = I_1^0(t) + R_1^0(t) = |\text{BP}(c)(t)| \quad \text{for all } t \text{ such that } |\text{BP}(c)(t)| \leq N_n.$$

As a consequence, the probability of a large outbreak in community 1 converges to π , and conditioned on not having a large outbreak, the following statements hold:

1. *The time of an outbreak as well as its total size $R_1(\infty) + R_2(\infty)$ are bounded in probability;*
2. *Whp, $R_2(\infty) = 0$ and all vertices contributing to $R_1(\infty)$ recover before traveling.*

Proof of Proposition E.3.2. To prove the proposition, we couple the exploration of the two community network to the SIR dynamic as follows. First, we draw all the travel clocks, determining in particular the set $\mathcal{T}_1 \cup \mathcal{T}_2$ of travelers up to time $T = \ln^2 n$. Next, starting from a single infected individual v_1 in $V_1^H \cup V_2^T$ (which whp will be in V_1^0), whenever a new vertex v_j gets infected at soem time t (where we order the vertices according to their infection time so that $n - j = S(t)$), we sample its recovery time and explore the edges from v_j into the set of susceptible vertices. Note that for this newly infected vertex v_j , its “forward” degrees into the susceptible vertices in V_1 and V_2 have distributions $\text{Bin}(n_1, c/n)$ and $\text{Bin}(n_2, c/n)$ where n_1 and n_2 are the number of susceptible vertices in V_1 and V_2 , respectively. As a consequence, these forward degrees can be coupled to two independent $\text{Poi}(c)$ random variables in such a way that the binomial and Poisson version are equal with probability $1 - O(j/n)$. Finally, we start independent infection clocks with rate β on the newly discovered edges, pausing them whenever the two endpoints are not in the same physical community (and terminating them when v_j recovers).

Define a stopping time τ as the first time one of the following events happens

1. The number of “children” of v_j in the exploration process is different from $\text{Poi}(c)$ in either community,
2. $v_j \in \mathcal{T}_1 \cup \mathcal{T}_2$.

We see that up to the stopping time, the infection tree is equal to $\text{BP}(c)(t)$. Furthermore, if $k = o(\sqrt{n})$, whp none of the first type of events happens before the process reaches size k , and if $k = o(n^\alpha \ln^{-3} n)$, whp, none of the second type happens before the process reaches size k . This proves the first statement of the theorem. The next statement (on the probability of a large outbreak) follows by setting $N_n = \ln n$ and observing that $\mathbb{P}(|\text{BP}(c)(t)| \geq \ln n) \rightarrow \pi$, and the next statement follows from observing that conditioned on $\text{BP}(c)(t)$ not surviving, it lives for a time that is bounded in probability, and reaches a final size that is bounded in probability. Since this implies that with high probability the branching process dies before time T , we see that whp, all vertices that ever get infected (which whp are in V_1^0) recover before ever traveling, proving the last statement. $\square$

E.3.2 Early and Late Stage Bounds in the Two Communities

Next, we derive an upper bound on the size of the epidemic up to time T^4 . To do so, we now couple the epidemic on the full graph to the epidemics on the communities of non-travelers

⁴Note that the branching process upper bounds see exponential growth in time, and therefore quickly deteriorate to become uninformative for time beyond $O(\ln n)$.

as follows. First, we draw all travel clocks, determining in particular the sets V_1^0 and V_2^0 . Next we draw the two community network, noting that the induced graph on, say, V_1^0 is nothing but a copy of $G'_1 \sim G(n', c'/n')$, where $n' = |V_1^0|$ and c' is defined by $c'/n' = c/n$, and similarly for V_2^0 (with $n'' = |V_2^0|$). Finally, for each individual v , we draw a recovery clock T_v and for each directed edge (u, v) in the 2-community network we draw an infection clock $T_{u,v} \sim \text{Exp}(\gamma)$.⁵

Considering the epidemic as experience by individuals in V_1^0 , once the infection has reached V_1^0 , the epidemics will spread inside G'_1 , while injecting seeds into the external directed network along the edges from V_1^0 to $\mathcal{T}_1 \cup \mathcal{T}_2$ (note that there will be no active edges between the non-travelers in V_1^0 and V_2^0 until time T), and being infected via edges into the set of currently susceptible individuals in V_1^0 . The epidemic inside G'_2 proceeds by a similar process.

Proposition E.3.3. *Fix $\delta > 0$, and let $t_- = \frac{\alpha-\delta}{\lambda} \ln n$. Then whp, the following statements hold: $R_2(t_-) + I_2(t_-) = 0$, $R_1(T) + I_1(T) \leq (r_\infty + \delta)n$, and $R_2(T) + I_2(T) \leq (r_\infty + \delta)n$.*

Proof. To prove the last two statements, we note that with probability tending to 1, $R_i(T) + I_i(T) \leq Y_i^0(T) + n^{1-\alpha} \ln^3 n$ where Y_i^0 is defined in (E.11). As a consequence, it is enough to prove the upper bound for Y_i^0 , which can be obtained by considering an epidemic on $G'_i \sim G(n', c'/n')$ with external seeds emanating from infected vertices in $\mathcal{T}_1 \cup \mathcal{T}_2$. Since $I_i^0(T) + R_i^0(T)$ does not depend on the evolution of the SIR dynamics past time T , we can further assume that G'_i does not receive further seeds after $t = T$, which allows us to bound $I_i^0(T) + R_i^0(T)$ by $R(\infty)$ on G'_i . This puts us into the setting of Lemma E.1.2, with c and R_0 replaced by $c' = c \frac{n}{n'}$ and $R'_0 = R_0 \frac{n}{n'} \leq R_0$, implying in particular that $r'_\infty \leq r_\infty$. To apply the lemma, we note that each vertex in V_i^0 has $\text{Bin}(|\mathcal{T}_i|, c/n)$ edges joining it to $\mathcal{T}_i$, showing it has probability $q = 1 - (1 - c/n)^{|\mathcal{T}_i|} = \tilde{O}(n^{-\alpha})$ of being part of the seed set, which in turn shows that the seed set is $o(n)$ with high probability.

To prove the first statement, define a stopping time τ to be the first time an individual outside of V_1^0 gets infected. Up to the stopping time, the epidemic in V_1^0 grows at most as fast as the SIR epidemic on $G(n', p)$, showing that if we consider the time $\tilde{t}_- = \min\{\tau, t_-\}$ and an arbitrary $\delta' > 0$, then whp $I_1^0(\tilde{t}_-) + R_1^0(\tilde{t}_-) \leq e^{\lambda(\tilde{t}_- + \delta')}$. But each individual contribution to $I_1^0 + R_1^0$ has only $\text{Bin}(|\mathcal{T}|, c/n)$ neighbors in V_1^0 it can infect, showing that up to the time $\tilde{t}_-$, the probability of an infection leaving V_1^0 is bounded by $\tilde{O}(n^{-\alpha} \cdot e^{\lambda(\tilde{t}_- + \delta')})$, which tends to zero for δ' chosen small enough (e.g., $\delta' < \delta$). This completes the proof. $\square$

⁵The travel clocks introduce some subtleties here: When we pre-draw a time T_{uv} with either u or v in $\mathcal{T}_1 \cup \mathcal{T}_2$, we have to take into account that the original clock was only active when both endpoints were physically in the same community. In order to determine if an infection is passed from u to v and when, we therefore have to look at the whole travel history of the two vertices and the time t_u when u was initially infected. Specifically, taking T_{uv} as time-budget that we use whenever the edge uv is active (starting from the time t_u), we will replacing T_{uv} by the time t_{uv} we run out of the time budget (which by definition, will be a time when uv is active). If u has not recovered by the time t_{uv} , i.e., if $t_{uv} \leq t_u + T_u$, the edge transmits at time t_{uv} - otherwise it doesn't. The above process has to be done iteratively, starting by setting $t_u = 0$ for the first infected individual and $t_v = \infty$ for all others, updating the infection times as we carry out the above steps for the infected vertices, one by one as they get infected.

E.3.3 Lower Bounds in community 1

Let T , $\mathcal{T}$, V_i^0 , $I_i^0(t)$ and $R_i^0(t)$ as in the last section, and let

$$\tau_i^0(\epsilon) = \inf\{t \geq 0 : I_i^0(t) + R_i^0(t) \geq \epsilon n\}. \quad (\text{E.15})$$

In this section we prove the following proposition.

Proposition E.3.4. *Consider the setting of Theorem 2.2 in the main text, and assume that $R_0 > 1$. Fix $\epsilon \in (0, r_\infty)$, $\delta \in (0, 1)$ and let*

$$\tau^+ = \frac{1 + \delta}{\lambda} \ln n \quad \text{and} \quad \tau_{1 \rightarrow 2}^+ = \frac{\alpha + \delta}{\lambda} \ln n.$$

With probability at least $\pi - \delta$, the following holds: $\tau_1(\epsilon) \leq \tau_1^0(\epsilon) \leq \tau^+$ and

$$I_1(\tau_{1 \rightarrow 2}^+) + R_1(\tau_{1 \rightarrow 2}^+) \geq I_1^0(\tau_{1 \rightarrow 2}^+) + R_1^0(\tau_{1 \rightarrow 2}^+) \geq n^{\alpha + \delta/4}.$$

Before proving the proposition, we note that by choosing $\epsilon = r_\infty - \delta$, the proposition immediately implies that with probability at least $\pi - \delta$, $R_1(\infty) \geq (r_\infty - \delta)n$, complementing the upper bound from the previous proposition.

Proof. As in the proof of Proposition E.3.3, we couple the static subgraph process on $G'_1 \sim G(n', c/n')$, but this time, we will derive a *lower* by a coupling to what we call a *restricted epidemic process* on V_1^0 by disallowing any edges or infections involving $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2$ during $[0, T]$. To motivate this coupling, we note that throughout $[0, T]$, the epidemic on V_1^0 in the full travel-based setting is stochastically *at least as fast* as it is in the single-community setting on G'_1 , because:

- In the actual process, the nodes in V_1^0 can still gain exposure from infected travelers (type-1 or type-2) who move into community 1. Thus, the infection can only spread earlier or faster in the real process.
- Removing the travelers (i.e., ignoring any infections they cause or receive) can only reduce the outbreak size and slow the onset of infections, giving a lower bound for the real process.

Because edges among V_1^0 are exactly those in G'_1 , and these edges remain active for times $[0, T]$, the restricted process is precisely the single-community SIR on G'_1 . By monotonicity of SIR processes under edge removals, the actual epidemic on V_1 (and hence community 1) is stochastically at least as large—i.e., at any fixed time $t \leq T$, the number of infections in V_1^0 under the full model is at least the number of infections under the restricted single-community epidemic.

From here on the proof is straightforward: As before, let $n' = |V_1^0|$, and let

$$\mathcal{E}_{\text{travel}} = \{|\mathcal{T}| \leq n^{1-\alpha} \ln^3 n\}$$

be the high-probability event from Lemma E.3.1. Conditioned on $\mathcal{E}_{\text{travel}}$, we have $n' \geq n - n^{1-\alpha} \ln^3 n$. Since we assume $\alpha > 0$, the ratio $n'/n \xrightarrow{p} 1$ as $n \rightarrow \infty$. Furthermore, conditional on V_1^0 , the induced contact network on these non-traveling nodes is an Erdős–Rényi graph

$$G'_1 \sim G(n', c'/n')$$

with $c' = cn'/n$ since edges among these vertices remain active throughout $[0, T]$, and there are no new edges formed via travel into V_1^0 . The standard SIR epidemics over $G(n', c'/n')$ then has parameters $R'_0 = \frac{c'\beta}{\beta+\gamma} 1$ and $\lambda' = c'\beta - \beta - \gamma$. Noting that $R'_0 \xrightarrow{p} R_0$, $\lambda' \xrightarrow{p} \lambda$, we have that the infection probability and final size converge as well, $r'_\infty \xrightarrow{p} r_\infty$ and $\pi' \xrightarrow{p} \pi$ due to the continuity of r_∞ and π in R_∞ and c . By the properties of the standard SIR epidemics over $G(n', c'/n')$ summarized in Theorem E.1.1, the following statements then hold with probability at least $\pi' - \delta/2$ if n is sufficiently large:

1. Let $\epsilon' = \epsilon n/n'$ and let $\tau'_1(\epsilon')$ be the first time the infected set in G'_1 reaches $n'\epsilon' = n\epsilon$ many vertices. Then

$$\tau'_1(\epsilon') \leq (1 + \delta/4) \frac{\ln n'}{\lambda'} \leq \tau^+.$$

2. Similarly, setting $\alpha' = \alpha(1 + \delta/3)$, we have

$$\tau'_1((n')^{\alpha'-1}) \leq (1 + \delta/4) \frac{\alpha' \ln n'}{\lambda'} \leq (1 + 3\delta/4) \frac{\alpha \ln n'}{\lambda'} \leq \tau_{1 \rightarrow 2}^+.$$

and hence

$$I_1^0(\tau_{1 \rightarrow 2}^+) + R_1^0(\tau_{1 \rightarrow 2}^+) \geq (n')^{\alpha+\delta/3} \geq n^{\alpha(1+\delta/4)}.$$

Denote by $\mathcal{E}_1$ the event that the single-community SIR on G'_1 behaves “as expected”. We then combine the events $\mathcal{E}_{\text{travel}}$ and $\mathcal{E}_1$ via a union bound:

$$\mathbb{P}(\mathcal{E}_{\text{travel}} \cap \mathcal{E}_1) \geq 1 - \left[\mathbb{P}(\mathcal{E}_{\text{travel}}^c) + \mathbb{P}(\mathcal{E}_1^c) \right] \geq 1 - \delta - \pi.$$

On $\mathcal{E}_{\text{travel}}$, we have $n' \geq n - \tilde{O}(n^{1-\alpha})$ and thus $n'/n \rightarrow 1$. On $\mathcal{E}_1$, the single-community SIR in G'_1 infects at least $(r_\infty - \delta)n$ nodes by time T and reaches size ϵn by time at most $\tau^+ = (1 + \delta) \frac{\ln n}{\lambda}$.

By the coupling argument, the actual epidemic in V_1 must infect at least as many nodes in V_1^0 (and do so at least as fast). Hence: on the event $\mathcal{E}_{\text{travel}} \cap \mathcal{E}_1$

$$R_1(\infty) \geq R_1^0(T) \geq (r_\infty - \delta)n,$$

$$\tau_1(\epsilon) \leq \tau_1^0(\epsilon) \leq (1 + \delta) \frac{\ln n}{\lambda},$$

and

$$I_1(\tau_{1 \rightarrow 2}^+) + R_1(\tau_{1 \rightarrow 2}^+) \geq I_1^0(\tau_{1 \rightarrow 2}^+) + R_1^0(\tau_{1 \rightarrow 2}^+) \geq n^{\alpha+\delta/4},$$

all for n sufficiently large. This proves the proposition. $\square$

E.3.4 Lower Bounds in community 2

For the lower bounds in community 2, we will want to use Lemma E.1.2, which means we will need to understand the transmission of infections from community 1 to 2 via travelers. For this, it will be important to define a suitable coupling on the whole two-community network.

To define the coupling, we first predetermined the sets $\mathcal{T}_1$ and $\mathcal{T}_2$ of travelers, keeping track to the fact that when we draw the travel times for vertices in these sets that the first travel has to happen before time $T = \ln^2 n$. Assuming the first infected vertex is in V_1^0 , and event which happens with high probability, we then define a set of discovered vertices, $V_i^{disc}(t) \subset \mathcal{T}_i$ which is initially empty, and gets updated as part of our coupling as follows; starting with one initially infected individual chosen uniformly at random in v_1^0 , whenever an individual v gets infected at time b_v , we draw

1. Start a recovery clock which clicks at rate γ ;
2. Draw four forward degrees, $d_i^0(v) \sim \text{Bin}(|V_i^0|, c/n)$ and $d_i^{\mathcal{T}}(v) \sim \text{Bin}(|\mathcal{T}_i|, c/n)$ with $i = 1, 2$;
3. Draw half-edges without an end-point for the degrees d_i^0 into static vertices;
4. Draw vertices $v_{i,1}^{\mathcal{T}}, \dots, v_{i,d_i}^{\mathcal{T}} \in \mathcal{T}_i$ uniformly without replacement, and link them to v by an oriented edge pointing away from v ; if any of these vertices has not yet been discovered, add it to V_i^{disc} , draw its initial travel state from the conditional distribution, and start travel clocks with rates $\rho_{\mathcal{T}}$ and $\rho_{\mathcal{H}}$ to update their travel state going forward;
5. Start infection clocks for the edges and half-edges out of v (and note that we can determine which edges/half-edges are active, since the travel state of the vertices in V_i^0 is determined);
6. Once the first infection clock into V_i^0 clicks, we choose d_i^0 endpoints $v_{i,1}^0, \dots, v_{i,d_i^0}^0$ uniformly in V_i^0 without replacement, starting with the endpoint for the half-edge which just clicked. With a slight abuse of notation, we will say that at this point, the infection clock on the edge $(v, v_{i,1}^0)$ has clicked;
7. When an infection clock on an edge vw clicks, and w is still infected, we say that v tries to infect w ; if at this point in time w is still susceptible, we call the infection attempt successful and declare w infected.

We would like to couple this process to the following process, which is easier to analyze:

- Instead of Step 2, we set $d_i^{\mathcal{T}}(v) = 0$ whenever $v \notin V_1^0 \cup V_2^0$;
- When $v \in V_1^0 \cup V_2^0$, instead of choosing $v_{i,1}^{\mathcal{T}}, \dots, v_{i,d_i}^{\mathcal{T}} \in \mathcal{T}_i$ in Step 4, we create new vertices for all of them, and add them to V_i^{disc} , each of them with a fresh draw of travel states from the conditional distribution.

Lemma E.3.5. *Let $0 < \delta < \min\{\alpha, \frac{1-\alpha}{2}\}$. Then the above two processes can be coupled such that they are equal with high probability as long as*

$$I^0(t) + R^0(t) \leq n^{\alpha+\delta}. \quad (\text{E.16})$$

Under this condition, we furthermore have that whp, $d_i^T(v) \in \{0, 1\}$ for all $v \in V_1^0 \cup V_2^0$.

Proof. Let $V_{I+R}^0(t)$ be the set of infected vertices in $V^0 = V_1^0 \cup V_2^0$, and let $N(t) = |V_{I+R}^0(t)| = I^0(t) + R^0(t)$. Define τ to be the stopping time when the first infected vertex $v \in \mathcal{T}$ draws a degree $d_i^T > 0$. For $t < \tau$, the set of discovered vertices V^{disc} is just the set of vertices discovered by an infected vertex in V_{I+R}^0 , showing that

$$|V^{disc}| \leq \sum_{v \in V_{I+R}^0} (d_1^T(v) + d_2^T(v)).$$

This random variable has distribution $\text{Bin}(N(t)|\mathcal{T}|, c/n)$, which under the assumption (E.16) is bounded whp by, say $2c|\mathcal{T}|n^{\alpha+\delta-1}$. But this means that by this time, the total number of infections emanating from $\mathcal{T}$ is bounded by $2c|\mathcal{T}|n^{-(1-\alpha-\delta)}$, which shows that the probability of drawing a random variable $d_i^T(v) > 0$ is bounded by

$$2c|\mathcal{T}|n^{\alpha+\delta-1} \times |\mathcal{T}|c/n = O(n^{2-2\alpha}n^{\alpha+\delta-2}\ln^4 n) = O(n^{\delta-\alpha}\ln^4 n) = o(1).$$

This shows that with high probability, the stopping time τ happens after the condition (E.16) is violated. This allows us to couple Step 2.

Concerning Step 4, we note that under the assumption that $d_i^T = 0$ for all $v \in \mathcal{T}$, the discovery steps now involve just the edges from V_{I+R}^0 into $\mathcal{T}$. Since the total number of these edges is a random variable with distribution $\text{Bin}(N(t)|\mathcal{T}|, c/n)$, it is whp bounded by $2c|\mathcal{T}|n^{\alpha+\delta-1}$, where we again used the assumption (E.16). Since the endpoints of these edges are chosen i.i.d. uniformly at random in $\mathcal{T}$, the probability of a two of them hitting the same vertex is bounded by

$$\frac{1}{|\mathcal{T}|} (2c|\mathcal{T}|n^{\alpha+\delta-1})^2 = O(|\mathcal{T}|n^{2\alpha+2\delta-2}) = O(n^{2\delta-(1-\alpha)}) = o(1),$$

showing that whp there are not conflicts. Finally, we note that

$$\mathbb{P}(d_i^T(v) \geq 2) = O(|\mathcal{T}|^2 n^{-2}) = \tilde{O}(n^{-2\alpha}),$$

showing that the probability that there exist a $v \in V_{I+R}^0$ with $d_i^T \geq 2$ is at most $\tilde{O}(n^{\gamma-\alpha}) = o(1)$. This proves the last statement. $\square$

Proposition E.3.6. *In the setting of Proposition E.3.4, With probability at least $\pi - \delta$, the following holds*

$$\begin{aligned} I_2(\tau_{1 \rightarrow 2}^+) + R_2(\tau_{1 \rightarrow 2}^+) &\geq I_2^0(\tau_{1 \rightarrow 2}^+) + R_2^0(\tau_{1 \rightarrow 2}^+) \geq n^{\delta/8}, \\ R_2(\infty) &\geq R_2^0(T) \geq (r_\infty - \delta)n, \end{aligned}$$

and

$$\tau_2(\epsilon) \leq \tau_2^0(\epsilon) \leq \tau_{1 \rightarrow 2}^+ + \tau^+.$$

Proof. We use the coupling described at the beginning of this section. Furthermore, we note that the SIR epidemic on V_1^0 will only slow down if we neglect all re-infections into V_1^0 , and finally, we note that we can lower bound the number of seeds into V_2^0 by only considering the seeds which appear because a vertex in V_1^0 infects a vertex in $\mathcal{T}_2$ which is traveling and returns before its infection clock runs out. Keeping this in mind, we define a stopping time τ_0 as the first time the SIR epidemic on V_1^0 (with one internal seed, and no external seeds) reaches $n^{\alpha+\delta/4}$ many vertices, noting that by Proposition E.3.4 $\tau < \tau_{1 \rightarrow 2}^+$ with probability at least $\pi - \delta$. Define the set **GOOD** individuals as the set of individuals contributing to $Y_1^0(\tau)$ whose recovery clock lasts for at least one unit time, we then use Lemma E.1.3 to conclude that with probability at least $\pi - 2\delta$, $|\text{GOOD}| \geq n^{\alpha-\delta/8}$. Let $\tilde{p}_\tau$ be the probability that a vertex in $\mathcal{T}_2$ is traveling at a given time $s < \ln^2 n$ and consider a vertex $v \in \text{GOOD}$, born at some time $t \leq \tau_0$. With probability $(1 - e^{-\beta})\mathbb{P}(\text{Bin}(|\mathcal{T}_2|, c/n) > 0)\tilde{p}_\tau = \theta(n^{-\alpha}/\ln^2 n)$ the vertex v will create an infected vertex in $\mathcal{T}_2$ at some time $t' < t + 1 \leq \tau_0 + 1$. With probability at least $e^{-2\gamma}(1 - e^{-\rho_H}(1 - e^{-\beta})e^{-\rho_\tau \frac{1}{n}}|V_2^0|)$ this vertex will stay infected for at least 2 days, return home latest after one unit time and stay home for at least one unit time, and during this time send a seed towards V_2^0 . Thus with probability $q = \theta(n^{-\alpha}/\ln^2 n)$, a good vertex will lead to a seed in V_2^0 . Since the events determining this transfer of seeds are independent for different $v \in \text{GOOD}$, by a simple multiplicative Chernoff bound, we see that whp, we have at least $\frac{1}{2}q|\text{GOOD}| \geq n^{\delta/8}$ seeds into V_2^0 by time $\tau + 3$. $\square$

E.4 Proof of Main Theorems

E.4.1 No Interventions

Given the results from the previous section, the proof of Theorem 6.3.2 is a straightforward summary of the propositions. Part (a) is a consequence of Propositions E.3.2 and E.3.4; Part (b) follows from Proposition E.3.2; Part (c) follows from a few more: the upper bounds on the stopping times are given in Propositions E.3.4 and E.3.6, the lower bounds on the stopping times follow from Proposition E.2.4, the upper bound on $R_k(\infty)$ follows from Proposition E.3.3, and finally the bound on the termination time follows from combining the lower bounds on $\tau_k(r_\infty - \delta)$ with Proposition E.2.5, where we choose δ sufficiently small.

E.4.2 Interventions

In this section, we prove Theorems 6.3.6 and 6.3.7 on travel bans and social distancing.

Proof of Theorem 6.3.6. The core idea is that the travel ban has little effect on the spread within the static part of both populations, since they don't travel anyway; the first cross-community transmission usually happens much sooner than the travel ban can be implemented (or even when an outbreak is identified).

Suppose that at time $t_{\text{tb}} = \tau_1(\epsilon) < \infty$, we implement a travel ban (in particular, we condition on community 1 seeing a large outbreak). As in the proof of Proposition E.3.6,

notice that there are at least $\ln n$ individuals in V_2^0 infected by time $(\alpha + \delta)\frac{\ln n}{\lambda}$. Lemma E.1.2 then gives a lower bound on the final size of an outbreak and an upper bound on the time of an outbreak in community 2 that matches the lower bound on the final size from the two community model with no interventions.

For an upper bound on the growth of the epidemic, we use similar arguments involving a branching process approximation.

A nearly identical branching process analysis as in Section E.2 applies to the case with travel ban too, except that once the travel ban is implemented, the two communities decouple and can be independently bounded with a single-type branching process; this makes no difference, as the growth rate in either process remains $\lambda = c\beta - \beta - \gamma$ as in the original case. The eventual size of the outbreak in both communities also follows similarly as before using the traveler and non-traveler decomposition: It is straightforward to check that Proposition E.3.3 is valid in the context of travel ban, and finally, once a travel ban is implemented, the spread in both communities can last up to $O(\ln n)$ time, a natural extension of results in the SIR processes of a single community. $\square$

Proof of Theorem 6.3.7. Suppose that at time $t_{sd} = \tau_1(\epsilon)$, we implement social distancing in community 2 (in particular, conditioned on a large outbreak in community 1). Note that the lower bound for community 1 remains unchanged: ignoring the travelers still give a valid lower bound on the size of outbreak. Likewise, the upper bounds for community 1 also hold (again up to time $T = \ln^2 n$): the branching process (up to time $\tau_1(\epsilon)$) still gives a valid upper bound on the time of outbreak, and the bound on the $R_k(\infty)$ in Proposition E.3.3 remains unaffected. It suffices now to give a sublinear upper bound on the final size of outbreak in community 2. We prove such an upper bound in two steps: first, we bound the fractional size of outbreak in community 2 up to time T to be as small as we desire, and then use a subcritical branching process coupling to argue that the infection has died out by T whp. Here we provide a high level proof sketch.

For the first step, we recall that, under our condition of a large outbreak, $\tau_1(\epsilon) \leq \tau^+ = (1 + \delta)\frac{\ln n}{\lambda}$ whp by Proposition E.3.4, at which time $Y_2(\tau^+) \leq \tilde{O}(n^{1-\alpha+\delta}) \ll n$. With social distancing, every seed infection within V_2^0 will cause a constant number of infections in expectation. The total number of such seeds is at most $Y_2(\tau^+) + |\mathcal{T}_2| \ll n$. As a result, the number of infections in community 2 is $o(n)$ whp up to T .

For the second step, we use a subcritical branching process to dominate the growth after the social distancing is implemented, similar to our analysis in Section E.2.1. Like in Proposition E.2.5, we wait till time τ^+ , at which point the proportion of susceptible individuals in community 1 has fall below s_0 whp. We can construct a branching process with cloning as we did in Section E.2.1, this time with lowered transmission rate to β' in community 2 (including clones of it). Our assumption of $\lambda' = c\beta' - \beta' - \gamma < 0$ ensures that this process is again subcritical and thus dies out within $O(\ln n)$ time. The full details of the computation is nevertheless tedious and shall be omitted. $\square$

Appendix F

Appendix to Chapter 7

F.1 Proofs from Section 7.2.1

We start with the proof of the following proposition regarding the case $d = 1$, and then discuss about the transition dynamics.

Proposition F.1.1. *The token system is not stable for any $n \geq 2, P > 0$ and $Q > 0$ when $d = 1$.*

Proof. Fix an agent $i \in \mathcal{A}$. Then s_i^t is a lazy random walk in one dimension. In particular, we have

$$s_i^{t+1} = \begin{cases} s_i^t + 1 & \text{with probability } \sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i \\ s_i^t & \text{with probability } 1 - \sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i - \sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j \\ s_i^t - 1 & \text{with probability } \sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j \end{cases} \quad (\text{F.1})$$

for all $t \geq 0$. Note that if $\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i \neq \sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j$, then $\mathbb{E}[s_i^t]$ diverges as $t \rightarrow \infty$ (i.e., the random walk is transient). Now assume that

$$\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i = \sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j \text{ for all } i \in \mathcal{A}. \quad (\text{F.2})$$

Then s_i^t is a lazy symmetric random walk. It is a well-known fact that s_i^t will take all the values in $\mathbb{Z}$ with probability 1 (i.e., the symmetric random walk is recurrent). Moreover, even though s_i^t will be 0 infinitely often with probability 1, the expected return time to 0 is infinity (i.e., the symmetric random walk is null recurrent). Hence, (C1) and (C2) are not satisfied for any $n \geq 2, P > 0$ and $Q > 0$ when $d = 1$. $\square$

As an immediate corollary to Proposition F.1.1, consider the random tie-breaking selection rule, under which the policy selects a service provider among $d \geq 1$ available providers uniformly at random. The behavior of the random walk s_i^t follows a similar structure as in (F.1), where the transition probabilities are state-independent¹; unlike (F.3), where the tran-

¹The transition probabilities of both random walks are identical if the agents are symmetric.

sition probabilities are state-dependent and there is a possibility of stabilizing the random walk. Therefore, s_i^t is again either transient or null-recurrent random walk.

Corollary F.1.2. *Consider a token system (n, P, Q, d) under the random tie-breaking selection rule. Then the token system is not stable for any $n \geq 2, P > 0$ and $Q > 0$ when $d \geq 1$.*

Transition dynamics. For future calculations, let us denote the following outcome by $(i, (j, k))$: $i \in \mathcal{A}$ is the service requester, and $\{j, k\} \subset \mathcal{A}$ is the subset of agents who are available to provide service. At each time period t , the outcome $(i, (j, k))$ occurs with probability $p_i 2q_j q_k$ if $j \neq k$ and with probability $p_i q_j^2$ if $j = k$. Hence, $i \in \mathcal{A}$ is the service requester at time t , $j \in \mathcal{A}$ is the service provider if $s_k^t > s_j^t$, $k \in \mathcal{A}$ is the service provider if $s_j^t > s_k^t$, and one of $j, k \in \mathcal{A}$ is selected as the service provider uniformly at random if $s_j^t = s_k^t$. Let r_{jk}^t , $j, k \in \mathcal{A}$, $j \neq k$, be the probability that given that only agents j and k are available to provide service at time t , agent j is the service provider. Then, we have

$$r_{jk}^t = \begin{cases} 1 & \text{if } s_k^t > s_j^t \\ \frac{1}{2} & \text{if } s_k^t = s_j^t \\ 0 & \text{if } s_k^t < s_j^t \end{cases}.$$

Fix $i \in \mathcal{A}$. We have $s_i^{t+1} = s_i^t + 1$ if one of the following outcomes occurs at time t :

- $(j, (i, i))$, where $j \neq i$, which happens with probability $\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i^2$,
- $(j, (i, k))$, where $j, k \neq i$, $s_k^t \geq s_i^t$ and agent i wins the tiebreak if any, which happens with probability $\sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i\}} p_j 2q_i q_k r_{ik}^t$.

Similarly, $s_i^{t+1} = s_i^t - 1$ if one of the following outcomes occurs at time t :

- $(i, (j, j))$, where $j \neq i$, which happens with probability

$$\sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j^2,$$

- $(i, (j, k))$, where $j, k \neq i$, $j \neq k$, which happens with probability

$$\sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i, j\}} p_i q_j q_k,$$

- $(i, (i, j))$, where $j \neq i$, $s_i^t \geq s_j^t$ and agent i loses the tiebreak if any, which happens with probability

$$\sum_{j \in \mathcal{A} \setminus \{i\}} p_i 2q_i q_j r_{ji}^t.$$

Therefore, we have

$$s_i^{t+1} = \begin{cases} s_i^t - 1 & \text{with probability } \frac{\sum_{j \in \mathcal{A} \setminus \{i\}} p_i q_j^2 + \sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i, j\}} p_i q_j q_k + \sum_{j \in \mathcal{A} \setminus \{i\}} p_i 2q_i q_j r_{ji}^t}{\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i^2 + \sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i\}} p_j 2q_i q_k r_{ik}^t} \\ s_i^t + 1 & \text{with probability } \frac{\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i^2 + \sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i\}} p_j 2q_i q_k r_{ik}^t}{\sum_{j \in \mathcal{A} \setminus \{i\}} p_j q_i^2 + \sum_{j \in \mathcal{A} \setminus \{i\}} \sum_{k \in \mathcal{A} \setminus \{i\}} p_j 2q_i q_k r_{ik}^t} \\ s_i^t & \text{otherwise} \end{cases} \quad (\text{F.3})$$

Note that (F.3) behaves similar to (F.1); the main difference is that the transition probabilities in (F.3) change as t changes because of the r_{ij}^t terms, which are time-dependent. Processes such as (F.3) referred as heterogeneous random walks, where the transition probabilities are state-dependent.

F.2 Proofs from Section 7.3

We start with the proof of Proposition 7.3.1. First, we describe the model here for convenience. Assume that there are 2 agents ($n = 2$). We analyze the following discrete time birth-death process with the state space $\mathcal{S} = \{(a, b) : a \geq 1, a \in \mathbb{N}, b = 1, 2\} \cup \{(0, 0)\}$, which captures the system: the state (a, b) represents the case in which agent b has $a > 0$ tokens, and $(0, 0)$ is the initial state. Denote this birth-death process by $(Z_t : t \geq 0)$. Let $X^t = \max_{i \in \mathcal{A}} s_i^t$ and $Y^t = \min_{i \in \mathcal{A}} s_i^t$.

Proof of Proposition 7.3.1. The transition probabilities for the process $(Z_t : t \geq 0)$ are as follows:

$$\begin{aligned} \mathbb{P}(Z_{t+1} = (1, 1) \mid Z_t = (0, 0)) &= p_2 \left(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_1^i q_2^{d-i} \right), \\ \mathbb{P}(Z_{t+1} = (0, 0) \mid Z_t = (1, 1)) &= p_1 (1 - q_1^d), \\ \mathbb{P}(Z_{t+1} = (1, 2) \mid Z_t = (0, 0)) &= p_1 \left(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_2^i q_1^{d-i} \right), \\ \mathbb{P}(Z_{t+1} = (0, 0) \mid Z_t = (1, 2)) &= p_2 (1 - q_2^d), \\ \mathbb{P}(Z_{t+1} = (a+1, 1) \mid Z_t = (a, 1)) &= p_2 q_1^d \text{ for all } a \geq 1, \\ \mathbb{P}(Z_{t+1} = (a+1, 2) \mid Z_t = (a, 2)) &= p_1 q_2^d \text{ for all } a \geq 1, \\ \mathbb{P}(Z_{t+1} = (a-1, 1) \mid Z_t = (a, 1)) &= p_1 (1 - q_1^d) \text{ for all } a \geq 2, \\ \mathbb{P}(Z_{t+1} = (a-1, 2) \mid Z_t = (a, 2)) &= p_2 (1 - q_2^d) \text{ for all } a \geq 2. \end{aligned}$$

Assume that the steady-state exists, and denote the steady-state vector by π . The detailed balance equations are:

$$\pi_{(0,0)} p_2 \left(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_1^i q_2^{d-i} \right) = \pi_{(1,1)} p_1 (1 - q_1^d), \quad (\text{F.4})$$

$$\pi_{(0,0)}p_1 \left(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_2^i q_1^{d-i} \right) = \pi_{(1,2)}p_2(1 - q_2^d), \quad (\text{F.5})$$

$$\pi_{(a,1)}p_2q_1^d = \pi_{(a+1,1)}p_1(1 - q_1^d) \text{ for all } a \geq 1, \quad (\text{F.6})$$

$$\pi_{(a,2)}p_1q_2^d = \pi_{(a+1,2)}p_2(1 - q_2^d) \text{ for all } a \geq 1, \quad (\text{F.7})$$

$$\pi_{(0,0)} + \sum_{b=1}^2 \sum_{a=1}^{\infty} \pi_{(a,b)} = 1. \quad (\text{F.8})$$

It follows by (F.6) and (F.7) that

$$\pi_{(a,1)} = \pi_{(1,1)} \left(\frac{p_2q_1^d}{p_1(1 - q_1^d)} \right)^{a-1} \text{ for all } a \geq 1, \quad (\text{F.9})$$

$$\pi_{(a,2)} = \pi_{(1,2)} \left(\frac{p_1q_2^d}{p_2(1 - q_2^d)} \right)^{a-1} \text{ for all } a \geq 1. \quad (\text{F.10})$$

Using (F.4), (F.5), (F.8), (F.9) and (F.10), first note that since the infinite geometric series in (F.8) must converge, the following are necessary and sufficient conditions for $(Z_t : t \geq 0)$ to have a steady-state:

$$\frac{p_2q_1^d}{p_1(1 - q_1^d)} < 1, \text{ or, equivalently } q_1^d < p_1, \quad (\text{F.11})$$

$$\frac{p_1q_2^d}{p_2(1 - q_2^d)} < 1, \text{ or, equivalently } q_2^d < p_2, \quad (\text{F.12})$$

where the equivalences in (F.11) and (F.12) follow from the fact that $p_1 + p_2 = 1$. Using (F.8), (F.11), and (F.12), we have

$$\pi_{(0,0)} + \pi_{(1,1)} \left(\sum_{a=1}^{\infty} \left(\frac{p_2q_1^d}{p_1(1 - q_1^d)} \right)^{a-1} \right) + \pi_{(1,2)} \left(\sum_{a=1}^{\infty} \left(\frac{p_1q_2^d}{p_2(1 - q_2^d)} \right)^{a-1} \right) = 1. \quad (\text{F.13})$$

Using (F.4), (F.5) and (F.13), we have

$$\pi_{(0,0)} = \left(1 + \frac{p_2(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_1^i q_2^{d-i})}{p_1 - q_1^d} + \frac{p_1(\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_2^i q_1^{d-i})}{p_2 - q_2^d} \right)^{-1}. \quad (\text{F.14})$$

Note that $\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_1^i q_2^{d-i} = q_1 \sum_{i=1}^d \binom{d-1}{i-1} q_1^{i-1} q_2^{d-i} = q_1(q_1 + q_2)^{d-1} = q_1$. Similarly,

$$\sum_{i=1}^d \frac{i}{d} \binom{d}{i} q_2^i q_1^{d-i} = q_2.$$

Hence, (F.14) simplifies to

$$\pi_{(0,0)} = \left(1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}\right)^{-1}. \quad (\text{F.15})$$

Once we have $\pi_{(0,0)}$ as a function of p_i 's and q_i 's, we can write all the steady-state probabilities as a function of p_i 's and q_i 's. Note that in the steady-state, X^t is bounded above by $M \in \mathbb{Z}_+$ with probability $1 - \sum_{a=M+1}^{\infty} \pi_{(a,1)} - \sum_{a=M+1}^{\infty} \pi_{(a,2)}$, which is equal to

$$1 - \frac{\left(\frac{p_2 q_1}{p_1 - q_1^d} \left(\frac{p_2 q_1^d}{p_1(1 - q_1^d)}\right)^M + \frac{p_1 q_2}{p_2 - q_2^d} \left(\frac{p_1 q_2^d}{p_2(1 - q_2^d)}\right)^M\right)}{\left(1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}\right)}. \quad (\text{F.16})$$

Since $X^t + Y^t = 0$ for all $t \geq 0$, for all $M \in \mathbb{Z}_+$, $\mathbb{P}_\pi(|s_i^t| \leq M)$ is given by (F.16) for $i = 1, 2$.

It is a well-known fact that starting from a state (a, b) , the expected time of the first occurrence of state (a, b) is $\frac{1}{\pi_{(a,b)}}$. Hence, starting from state $(0, 0)$, the expected time of the first occurrence of state $(0, 0)$ is

$$\frac{1}{\pi_{(0,0)}} = \left(1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}\right). \quad (\text{F.17})$$

□

Remark F.2.1. $\mathbb{E}(T_0)$ in (C2) follows from (F.17). The constant a in (C1) can be found using (F.16) as follows. Define

$$x = \frac{p_2 q_1^d}{p_1(1 - q_1^d)}, y = \frac{p_1 q_2^d}{p_2(1 - q_2^d)}, c_1 = \frac{\frac{p_2 q_1}{p_1 - q_1^d}}{1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}}, \text{ and } c_2 = \frac{\frac{p_1 q_2}{p_2 - q_2^d}}{1 + \frac{p_2 q_1}{p_1 - q_1^d} + \frac{p_1 q_2}{p_2 - q_2^d}}.$$

Then (F.16) becomes $1 - c_1 x^M - c_2 y^M$. We want to find $0 < a < 1$ such that $1 - c_1 x^M - c_2 y^M \geq 1 - a^M$, or $a^M \geq c_1 x^M + c_2 y^M$ for all $M \in \mathbb{Z}_+$. Note that $c_1, c_2 < 1$. Hence, a can be chosen to be $x + y$ if $x + y < 1$. Consider the case when $x + y > 1$. Let $z = \max\{x, y\}$. Then $c_1 x^M + c_2 y^M < 2z^M$. Clearly, there exists $M' \in \mathbb{Z}_+$ such that $2z^M < 1$ for all $M \geq M'$. Pick $0 < a_1 < 1$ such that $a_1 > z$ and $a_1^{M'} > 2z^{M'}$. Consider the inequalities $a^M \geq c_1 x^M + c_2 y^M$ for all $M < M'$. Since there are finitely many inequalities, we can pick $0 < a_2 < 1$ such that these inequalities are satisfied. Thus in this case, a can be chosen to be $\max\{a_1, a_2\}$.

Intermediate availability. Suppose that at each time period, we have $d = 2$ with probability β and $d = 1$ with probability $1 - \beta$, independently, for some $\beta \in (0, 1)$. We can capture this system with the same birth-death process $(Z_t : t \geq 0)$ with updated transition probabilities. Following similar calculations as in the proof of Proposition 7.3.1, it follows that the following are necessary and sufficient conditions for the process to have a steady-state:

$$\frac{\beta p_2 q_1^d + (1 - \beta) p_2 q_1}{\beta p_1 (1 - q_1^d) + (1 - \beta) p_1 q_2} < 1, \quad (\text{F.18})$$

$$\frac{\beta p_1 q_2^d + (1 - \beta) p_1 q_2}{\beta p_2 (1 - q_2^d) + (1 - \beta) p_2 q_1} < 1. \quad (\text{F.19})$$

It also follows that

$$\pi_{(0,0)} = \left(1 + \frac{p_2 q_1}{\beta(p_1 - q_1^d) + (1 - \beta)(p_1 q_2 - p_2 q_1)} + \frac{p_1 q_2}{\beta(p_2 - q_2^d) + (1 - \beta)(p_2 q_1 - p_1 q_2)} \right)^{-1},$$

and in the steady-state, X^t is bounded above by $M \in \mathbb{Z}_+$ with probability

$$1 - \pi_{(0,0)}(A + B), \quad (\text{F.20})$$

where

$$A = \frac{p_2 q_1}{\beta(p_1 - q_1^d) + (1 - \beta)(p_1 q_2 - p_2 q_1)} \left(\frac{\beta p_2 q_1^d + (1 - \beta) p_2 q_1}{\beta p_1 (1 - q_1^d) + (1 - \beta) p_1 q_2} \right)^M,$$

$$B = \frac{p_1 q_2}{\beta(p_2 - q_2^d) + (1 - \beta)(p_2 q_1 - p_1 q_2)} \left(\frac{\beta p_1 q_2^d + (1 - \beta) p_1 q_2}{\beta p_2 (1 - q_2^d) + (1 - \beta) p_2 q_1} \right)^M.$$

Similarly, since $X^t + Y^t = 0$ for all $t \geq 0$, for all $M \in \mathbb{Z}_+$, $\mathbb{P}_\pi(|s_i^t| \leq M)$ is given by (F.20) for $i = 1, 2$. When agents are symmetric and $d = 2$, (F.20) becomes $1 - \frac{2}{2+\beta} \left(\frac{2-\beta}{2+\beta} \right)^M$. As Figure F.1 shows, even for small values of β , the token distribution is fairly balanced with high probability.

F.3 Proof of Theorem 7.4.1

Proof of Theorem 7.4.1. Denote by k^t the agent chosen to request service at time t , and denote by I^t the set of agents chosen to be available to provide service at time t . Let $|I^t|$ denote the size of I^t , and note that $|I^t|$ takes values in $\{1, 2, \dots, d\}$. Let $j^t \in I^t$ be the agent chosen to provide service, i.e., an agent chosen uniformly from the agents in $i \in I^t$ that minimizes s_i^t . Hence $s_{j^t}^t = \min_{i \in I^t} s_i^t$, so that in particular, the service provider j^t is an agent with the minimum number of tokens among the available agents in I^t .

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

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

Note that by symmetry, we have $v^t = n\mathbb{E}[(s_1^t)^2]$. Let E_t be the event $\{k^t = j^t\}$, i.e., the event that the service provider and requester are the same agent. Let E_t^c be the event $\{k^t \neq j^t\}$.




Figure F.1: (F.20) as a function of β for several values of M when the agents are symmetric and $d = 2$.

Since k^t is uniformly distributed and k^t and j^t are independent, the probability of E_t is $1/n$. Note that conditioned on E_t , it holds that $s^{t+1} = s^t$. Then we have

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

Now,

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

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

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

and so

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

Without loss of generality, assume that $I^t = \{1, 2, \dots, |I^t|\}$. Now with probability $n^{-(d-1)}$, we have that $|I^t| = 1$, in which case $I^t = \{j^t\} = \{1\}$. Hence $\mathbb{E}[s_{j^t}^t \mid |I^t| = 1] = 0$. With probability $1 - n^{-(d-1)}$, we have that $|I^t| \geq 2$, in which case $s_{j^t}^t = \min\{s_1^t, s_2^t, \dots, s_{|I^t|}^t\} \leq \min\{s_1^t, s_2^t\}$. Using the fact that $\min\{a, b\} = \frac{1}{2}(a + b - |a - b|)$ for all $a, b \in \mathbb{R}$, we get $2\mathbb{E}[\min\{s_1^t, s_2^t\}] = -\mathbb{E}[|s_1^t - s_2^t|]$, and we get

$$\mathbb{E}[s_{j^t}^t \mid |I^t| \geq 2] \leq -\frac{1}{2} \mathbb{E}[|s_1^t - s_2^t|].$$

Thus,

$$\mathbb{E}[s_{j^t}^t \mid |I^t| \geq 2] \leq -\frac{1}{2}(1 - n^{-(d-1)})\mathbb{E}[|s_1^t - s_2^t|],$$

and we get

$$v^{t+1} \leq v^t - \frac{n-1}{n}(1 - n^{-(d-1)})\mathbb{E}[|s_1^t - s_2^t|] + 2.$$

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

$$\mathbb{E}\left[\left|Y - \frac{1}{n} \sum_{i=1}^n Z_i\right|\right] \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[|Y - Z_i|], \quad (\text{F.21})$$

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

$$\mathbb{E}[|s_1^t - s_2^t|] = \frac{1}{n-1} \sum_{i=2}^n \mathbb{E}[|s_1^t - s_i^t|].$$

Therefore by (F.21) and using the fact that $\sum_{i=2}^n s_i^t = -s_1^t$ for all $t \geq 0$, we have

$$\mathbb{E}[|s_1^t - s_2^t|] \geq \mathbb{E}\left[\left|s_1^t - \frac{1}{n-1} \sum_{i=2}^n s_i^t\right|\right] = \mathbb{E}\left[\left|s_1^t + \frac{s_1^t}{n-1}\right|\right] = \frac{n}{n-1} \mathbb{E}[|s_1^t|],$$

Therefore, we have

$$v^{t+1} \leq v^t - (1 - n^{-(d-1)})\mathbb{E}[|s_1^t|] + 2,$$

and via recursion, we get

$$v^t \leq 2t - (1 - n^{-(d-1)}) \sum_{u=0}^{t-1} \mathbb{E}[|s_1^u|].$$

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

$$\frac{1}{t} \sum_{u=0}^{t-1} \mathbb{E}[|s_1^u|] \leq \frac{2}{1 - n^{-(d-1)}} \leq 4, \quad (\text{F.22})$$

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

Denote the distribution of s^t by $\mu^t \in \Delta(\mathbb{Z}^n)$, and let $\nu^t = \frac{1}{t} \sum_{u=0}^{t-1} \mu_u$. Let Y^t be a random variable with distribution ν^t for all $t \geq 1$, and denote by Y_i^t the i 'th index of Y^t for all $i \in \mathcal{A}$. Then per (F.22), we have

$$\mathbb{E}[|Y_i^t|] = \frac{1}{t} \sum_{u=0}^{t-1} \mathbb{E}[|s_i^u|] \leq 4, \quad (\text{F.23})$$

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

Suppose towards a contradiction that the Markov chain has no stationary probability measure. Then for any finite subset $E \subset \mathbb{Z}^n$, it holds that $\lim_{t \rightarrow \infty} \mathbb{P}(s_i^t \in E) = 0$ by standard arguments about Markov chains. In particular, $\lim_{u \rightarrow \infty} \mathbb{E}[|s_i^u|] = \infty$, which contradicts (F.23). Thus, we have stability.

Since the Markov chain has a stationary distribution, and since it is irreducible and aperiodic, the distribution of s_i^t will converge to the stationary distribution. It follows from (F.23) that $\lim_{t \rightarrow \infty} \mathbb{E}[|s_i^t|] \leq 4$, and so $\mathbb{E}[|s_i^t|] \leq 5$ for all t large enough. Hence, it follows from Markov's inequality that for all t large enough and for all $i \in \mathcal{A}$, we have

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

□

F.4 Proof of Theorem 7.4.2

Recall that for $\lambda \in \mathbb{R}$, $Z_\lambda^t = \sum_{i=1}^n e^{\lambda s_i^t}$ and $E_\lambda^t = \mathbb{E}[Z_\lambda^t]$. As a first step, we show that for λ sufficiently close to zero, E_λ^t has a negative drift and thus a finite limit.

Lemma F.4.1 (Negative drift). *For any $\lambda \in \mathbb{R}$,*

$$E_\lambda^{t+1} \leq C_\lambda + (1 - \gamma_\lambda) E_\lambda^t, \quad (\text{F.24})$$

where $C_\lambda := \frac{|1-e^\lambda|}{n}$ and $\gamma_\lambda := \frac{|1-e^\lambda|}{n^2} + \frac{2-2\cosh \lambda}{n}$. In particular, when $|\lambda| \lesssim 1/n$, we have $\gamma_\lambda > 0$ and thus, $\limsup_{t \rightarrow \infty} E_\lambda^t \leq C_\lambda / \gamma_\lambda < \infty$.

Proof. Throughout the analysis, we will assume without loss of generality that $s_1^t \leq s_2^t \leq \dots \leq s_n^t$. As in the proof of Theorem 7.4.1, we denote by k^t and j^t the agent chosen to request service and the one chosen to provide service at time t , respectively.

By the law of total probability, we have

$$E_\lambda^{t+1} = \mathbb{E}[\mathbb{E}[Z_\lambda^t | s^t]]. \quad (\text{F.25})$$

First, consider $\lambda > 0$. Note that, through some straightforward albeit tedious arithmetic manipulation, we have

$$\begin{aligned} \mathbb{E}[Z_\lambda^{t+1} | s^t] &= Z_\lambda^t + \mathbb{E}\left[\left(e^{\lambda(s_{j^t}^t+1)} - e^{\lambda s_{j^t}^t} + e^{\lambda(s_{k^t}^t-1)} - e^{\lambda s_{k^t}^t}\right) \cdot \mathbb{1}_{j^t \neq k^t} \middle| s^t\right] \\ &= Z_\lambda^t + (e^\lambda - 1)\mathbb{E}\left[\left(e^{\lambda s_{j^t}^t} - e^{\lambda(s_{k^t}^t-1)}\right) \cdot \mathbb{1}_{j^t \neq k^t} \middle| s^t\right] \end{aligned} \quad (\text{F.26})$$

$$\begin{aligned} &= Z_\lambda^t + (e^\lambda - 1)\mathbb{E}\left[\left(e^{\lambda s_{j^t}^t} - e^{\lambda(s_{k^t}^t-1)}\right) \middle| s^t\right] - \\ &\quad (e^\lambda - 1)\mathbb{E}\left[\left(e^{\lambda s_{j^t}^t} - e^{\lambda(s_{k^t}^t-1)}\right) \cdot \mathbb{1}_{j^t = k^t} \middle| s^t\right] \\ &= Z_\lambda^t + (e^\lambda - 1)\mathbb{E}[e^{\lambda s_{j^t}^t} | s^t] - (1 - e^{-\lambda})\mathbb{E}[e^{\lambda s_{k^t}^t} | s^t] - \\ &\quad (e^\lambda - 1)\mathbb{E}\left[e^{\lambda s_{j^t}^t} \cdot \mathbb{1}_{j^t = k^t} \middle| s^t\right] + (1 - e^{-\lambda})\mathbb{E}\left[e^{\lambda s_{j^t}^t} \cdot \mathbb{1}_{j^t = k^t} \middle| s^t\right] \\ &= Z_\lambda^t + (e^\lambda - 1)\mathbb{E}[e^{\lambda s_{j^t}^t} | s^t] - \frac{1 - e^{-\lambda}}{n} Z_\lambda^t + \\ &\quad \frac{2 - e^\lambda - e^{-\lambda}}{n} \mathbb{E}\left[e^{\lambda s_{j^t}^t} \middle| s^t\right] \\ &\leq \left(1 - \frac{1 - e^{-\lambda}}{n}\right) Z_\lambda^t + (e^\lambda - 1)\mathbb{E}[e^{\lambda s_{j^t}^t} | s^t], \end{aligned} \quad (\text{F.27})$$

where in the second last equality we use the fact that k^t is uniform and independent of j^t , and in the last step we drop the last term as it is non-positive.

Claim F.4.2. When $d \geq 2$, for $\lambda > 0$,

$$\mathbb{E}[e^{\lambda s_{j^t}^t} | s^t] - \mathbb{E}[e^{\lambda s_{k^t}^t} | s^t] \leq \frac{1}{n^2} \sum_{1 \leq i < j \leq n} \left(e^{\lambda s_i^t} - e^{\lambda s_j^t}\right) \leq \frac{1}{n} - \frac{1}{n^2} Z_\lambda^t \leq 0. \quad (\text{F.28})$$

Similarly, for $\lambda < 0$,

$$\mathbb{E}[e^{\lambda s_{j^t}^t} | s^t] - \mathbb{E}[e^{\lambda s_{k^t}^t} | s^t] \geq \frac{1}{n^2} \sum_{i=1}^n \left(e^{\lambda s_i^t} - e^{\lambda s_n^t}\right) \geq \frac{1}{n^2} Z_\lambda^t - \frac{1}{n} \geq 0. \quad (\text{F.29})$$

As a result of the claim, for $\lambda > 0$,

$$\begin{aligned} \mathbb{E}[Z_\lambda^{t+1} | s^t] &\leq \left(1 - \frac{1 - e^{-\lambda}}{n}\right) Z_\lambda^t + \\ &\quad (e^\lambda - 1) \left(\mathbb{E}[e^{\lambda s_{k^t}^t} | s^t] + \frac{1}{n} - \frac{1}{n^2} Z_\lambda^t\right) \\ &= \left(1 - \frac{e^\lambda - 1}{n} (e^{-\lambda} - 1 + \frac{1}{n})\right) Z_\lambda^t + \frac{e^\lambda - 1}{n}. \end{aligned}$$

By the law of total probability,

$$E_\lambda^{t+1} \leq \frac{e^\lambda - 1}{n} + \left(1 - \frac{e^\lambda - 1}{n} \left(e^{-\lambda} - 1 + \frac{1}{n}\right)\right) E_\lambda^t. \quad (\text{F.30})$$

For λ sufficiently small (namely, $\lambda \lesssim 1/n$), the coefficient $\gamma_\lambda := \frac{e^\lambda - 1}{n} (e^{-\lambda} - 1 + \frac{1}{n}) > 0$.

Similarly, for $\lambda < 0$,

$$\begin{aligned} \mathbb{E}[Z_\lambda^{t+1} | s^t] &\leq \left(1 - \frac{1 - e^{-\lambda}}{n}\right) Z_\lambda^t + \\ &\quad (e^\lambda - 1) \left(\mathbb{E}[e^{\lambda s_{k_t}^t} | s^t] + \frac{1}{n^2} Z_\lambda^t - \frac{1}{n}\right) \\ &= \left(1 - \frac{1 - e^\lambda}{n} \left(1 - e^{-\lambda} + \frac{1}{n}\right)\right) Z_\lambda^t + \frac{1 - e^\lambda}{n}. \end{aligned}$$

Hence,

$$E_\lambda^{t+1} \leq \frac{1 - e^\lambda}{n} + \left(1 - \frac{1 - e^\lambda}{n} \left(1 - e^{-\lambda} + \frac{1}{n}\right)\right) E_\lambda^t. \quad (\text{F.31})$$

Again, the recurrent coefficient is less than 1 when $|\lambda| \lesssim 1/n$, finishing our proof. $\square$

Proof of Claim F.4.2. Since

$$s_{j_t}^t = \min_{1 \leq \ell \leq d} s_{I_\ell^t}^t \quad \text{with } I_1^t, \dots, I_d^t \stackrel{i.i.d.}{\sim} \text{Uniform}(\{1, \dots, n\}), \quad (\text{F.32})$$

the distribution of $e^{\lambda s_{j_t}^t}$ is dominated by (resp. dominates) the case where $d = 2$ for $\lambda > 0$ (resp. $\lambda < 0$). It is sufficient to consider $d = 2$. Again, we assumed without loss of generality that the indices are ordered such that at time t we have $s_1^t \leq s_2^t \leq \dots \leq s_n^t$. The difference in the two expectations can be evaluated as

$$\begin{aligned} \mathbb{E}[e^{\lambda s_{j_t}^t} | s^t] - \mathbb{E}[e^{\lambda s_{k_t}^t} | s^t] &= \frac{1}{n^2} \sum_{1 \leq i < j \leq n} \left(e^{\lambda s_i^t} - e^{\lambda s_j^t}\right) \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \left(e^{\lambda s_1^t} - e^{\lambda s_i^t}\right) \\ &= \frac{1}{n} - \frac{1}{n^2} Z_\lambda^t. \end{aligned} \quad (\text{F.33})$$

The case for $\lambda < 0$ is analogous. $\square$

Lemma F.4.3. For $\lambda \in (-\ln 1.5, \ln 2)$, $\limsup_{t \rightarrow \infty} E_\lambda^t < \infty$.

Proof. Define $\bar{\lambda}$ as

$$\bar{\lambda} := \sup\{\lambda : \limsup_{t \rightarrow \infty} E_\lambda^t < \infty\}, \quad (\text{F.34})$$

and similarly define $\underline{\lambda}$ as

$$\underline{\lambda} := \inf\{\lambda : \limsup_{t \rightarrow \infty} E_\lambda^t < \infty\}. \quad (\text{F.35})$$

Lemma F.4.1 implies that $\underline{\lambda} < 0 < \bar{\lambda}$ for any $n \in \mathbb{Z}_+$. Note that for any $n, t \in \mathbb{Z}_+$, the function $\lambda \mapsto E_\lambda^t$ is convex. As a result, for any $\lambda \in (\underline{\lambda}(n), \bar{\lambda}(n))$, we have $\limsup_{t \rightarrow \infty} E_\lambda^t < \infty$. It now suffices to show that $\bar{\lambda} \geq \log 2$ and $\underline{\lambda} \leq -\log 1.5$.

For the moment, let us assume that n is even for convenience, and write $n = 2L$ with $L \geq 2$; generalizing this is straightforward.

Consider $\lambda_0 \in (\underline{\lambda}(n), 0)$. We will investigate for what values of $\lambda > 0$ we have a finite $\limsup_{t \rightarrow \infty} E_\lambda^t$.

Again, we consider the recursive relationship between E_λ^t and E_λ^{t+1} , and assume without loss of generality that $s_1^t \leq \dots \leq s_n^t$. Let $\epsilon \in (0, 1)$ be a constant to be specified, and denote by $\mathcal{E}_t$ the event that $e^{\lambda s_L^t} < \epsilon Z_\lambda^t/n$. We then decompose $E_\lambda(t+1)$ into the combined contribution from two terms as

$$\begin{aligned} E_\lambda(t+1) &= \mathbb{E}[Z_\lambda(t+1)] \\ &= \mathbb{E}[Z_\lambda(t+1)|\mathcal{E}_t] + \mathbb{E}[Z_\lambda(t+1)|\mathcal{E}_t^c]. \end{aligned} \quad (\text{F.36})$$

For the first term in (F.36), note that whenever $e^{\lambda s_L^t} < \epsilon Z_\lambda^t/n$, we have

$$\begin{aligned} \sum_{1 \leq i < j \leq n} e^{\lambda s_i^t} - e^{\lambda s_j^t} &\leq \sum_{i=1}^L \sum_{j=L+1}^n e^{\lambda s_i^t} - e^{\lambda s_j^t} \\ &\leq L(n-L) \cdot Z_\lambda^t \cdot \left(\frac{\epsilon}{n} - \frac{1}{n-L} \left(1 - \frac{\epsilon L}{n} \right) \right) \\ &= -(1-\epsilon)LZ_\lambda^t. \end{aligned} \quad (\text{F.37})$$

Thus, in this case, Claim F.4.2 can be refined into

$$\mathbb{E}[e^{\lambda s_{j^t}^t} - e^{\lambda s_{k^t}^t} | \mathcal{E}_t, s^t] \leq \frac{1}{n^2} \sum_{1 \leq i < j \leq n} e^{\lambda s_i^t} - e^{\lambda s_j^t} \leq -\frac{1-\epsilon}{n^2} LZ_\lambda^t. \quad (\text{F.38})$$

Similar to (F.27), we now have

$$\begin{aligned} \mathbb{E}[Z_\lambda^{t+1} | \mathcal{E}_t, s^t] &\leq \left(1 - \frac{1-e^{-\lambda}}{n} \right) Z_\lambda^t + (e^\lambda - 1) \mathbb{E}[e^{\lambda s_{j^t}^t} | \mathcal{E}_t, s^t] \\ &\leq \left(1 - \frac{1-e^{-\lambda}}{n} \right) Z_\lambda^t + (e^\lambda - 1) \left(\mathbb{E}[e^{\lambda s_{k^t}^t} | \mathcal{E}_t, s^t] - \frac{1-\epsilon}{n^2} LZ_\lambda^t \right) \\ &= \left(1 - \frac{e^\lambda - 1}{n} \left(e^{-\lambda} - 1 + \frac{(1-\epsilon)L}{n} \right) \right) Z_\lambda^t. \end{aligned}$$

Thus,

$$\mathbb{E}[Z_\lambda^{t+1} | \mathcal{E}_t] \leq \left(1 - \frac{e^\lambda - 1}{n} \left(e^{-\lambda} - 1 + \frac{(1-\epsilon)L}{n} \right) \right) E_\lambda^t. \quad (\text{F.39})$$

For the choice of $L = n/2$, the coefficient is strictly less than 1 for any $\lambda < \log 2$ granted that we choose ϵ sufficiently small.

Now we consider the second term in (F.36) under the event $\mathcal{E}_t^c$. By a natural extension of Lemma F.4.1, we have

$$\mathbb{E}[Z_\lambda^{t+1} | \mathcal{E}_t^c] \leq \mathbb{E}\left[(C_\lambda + (1 - \gamma_\lambda)Z_\lambda^t) \cdot \mathbb{1}_{e^{\lambda s_L^t} \geq \epsilon Z_\lambda^t/n}\right] \leq C_\lambda + (1 - \gamma_\lambda)\mathbb{E}\left[Z_\lambda^t \cdot \mathbb{1}_{e^{\lambda s_L^t} \geq \epsilon Z_\lambda^t/n}\right].$$

The expectation term on the right-hand side can be further expanded as

$$\mathbb{E}\left[Z_\lambda^t \cdot \mathbb{1}_{e^{\lambda s_L^t} \geq \epsilon Z_\lambda^t/n}\right] = \int_0^\infty \mathbb{P}(Z_\lambda^t \cdot \mathbb{1}_{e^{\lambda s_L^t} \geq \epsilon Z_\lambda^t/n} \geq x) dx \leq n + \int_n^\infty \mathbb{P}\left(Z_\lambda^t \geq x, e^{\lambda s_L^t} \geq \frac{\epsilon}{n} Z_\lambda^t\right) dx.$$

When $e^{\lambda s_L^t} \geq \epsilon Z_\lambda^t/n$, we must have

$$\prod_{i=1}^{L-1} e^{-\lambda s_i^t} = \prod_{i=L}^n e^{\lambda s_i^t} \geq \left(\frac{\epsilon}{n} Z_\lambda^t\right)^{n-L+1}, \quad (\text{F.40})$$

or equivalently for $\lambda_- \in (\underline{\lambda}, 0)$,

$$\prod_{i=1}^{L-1} e^{\lambda_- s_i^t} \geq \left(\frac{\epsilon}{n} Z_\lambda^t\right)^{\frac{-\lambda_-(n-L+1)}{\lambda}}.$$

By the AM-GM inequality,

$$Z_{\lambda_-}(t) \geq \sum_{i=1}^{L-1} e^{\lambda_- s_i^t} \geq (L-1) \left(\frac{\epsilon}{n} Z_\lambda^t\right)^{\frac{-\lambda_-(n-L+1)}{\lambda(L-1)}}.$$

Putting things together, we obtain

$$\begin{aligned} \mathbb{E}[Z_\lambda(t+1); \mathcal{E}_t^c] &\leq C_\lambda + (1 - \gamma_\lambda) \left(n + \int_n^\infty \mathbb{P}\left(Z_{\lambda_-}(t) \geq (L-1) \left(\frac{\epsilon}{n} x\right)^{\frac{-\lambda_-(n-L+1)}{\lambda(L-1)}}\right) dx \right) \\ &\leq C_\lambda + (1 - \gamma_\lambda) \left(n + \int_n^\infty \frac{E_{\lambda_-}(t)}{L-1} \cdot \left(\frac{\epsilon}{n}\right)^{\frac{\lambda_-(n-L+1)}{\lambda(L-1)}} \cdot x^{\frac{\lambda_-(n-L+1)}{\lambda(L-1)}} dx \right) \\ &= C_\lambda + (1 - \gamma_\lambda) \left(n + \frac{E_{\lambda_-}(t)}{L-1} \cdot \left(\frac{\epsilon}{n}\right)^{\frac{\lambda_-(n-L+1)}{\lambda(L-1)}} \int_n^\infty x^{\frac{\lambda_-(n-L+1)}{\lambda(L-1)}} dx \right), \quad (\text{F.41}) \end{aligned}$$

where the last inequality uses Markov's inequality. The entire coefficient in front of the integral is uniformly bounded in t since $\limsup_{\tau \rightarrow \infty} E_{\lambda_-}^\tau < \infty$; whenever $\lambda < -\lambda_- \frac{n-L+1}{L-1}$, the integral converges. Combining (F.39) with (F.41), we have

$$E_\lambda^{t+1} \leq (1 - \gamma') E_\lambda^t + C'$$

for some positive constants γ' and C' , whenever $\lambda < \min\{-\frac{n-L+1}{L-1}\lambda_-, \log 2\}$. Thus,

$$\limsup_{t \rightarrow \infty} E_\lambda^t < \infty.$$

Sending $\lambda_- \rightarrow \underline{\lambda}$ gives

$$\bar{\lambda} \geq \min \left\{ -\frac{n-L+1}{L-1}\underline{\lambda}, \log 2 \right\}. \quad (\text{F.42})$$

With an analogous argument by considering $\lambda_+ \in (0, \bar{\lambda}(n))$ and $\lambda < 0$, we obtain

$$E_\lambda(t+1) \leq (1 - \gamma'')E_t + C''$$

for some $\gamma'', C'' > 0$, granted that $\lambda > \max\{-\frac{n-L+1}{L-1}\lambda_+, \log \frac{2}{3}\}$. Sending $\lambda_+ \rightarrow \bar{\lambda}$ gives

$$\underline{\lambda} \leq \max \left\{ -\frac{n-L+1}{L-1}\bar{\lambda}, \log \frac{2}{3} \right\}. \quad (\text{F.43})$$

(F.42) and (F.43) together form a system of two piecewise linear inequalities. Solving it gives

$$\bar{\lambda} \geq \log 2 \quad \text{and} \quad \underline{\lambda} \leq \log \frac{2}{3}$$

as desired. $\square$

Proof of Theorem 7.4.2. The theorem follows directly from Lemmas 7.4.3 and F.4.3. $\square$

F.5 Proofs from Section 7.4

Proof of Lemma 7.4.8. We first prove that $\pi_0 > \frac{1}{2}$. Note that (7.7) can be written as

$$\sum_{M \geq 1} (\pi_M + \pi_{-M+1} - 1) = 0. \quad (\text{F.44})$$

We claim that if $\pi_1 + \pi_0 = \pi_0^2 + \pi_0 - 1 < 0$, which implies $\pi_0 \leq \frac{1}{2}$, then all the terms in the summation (F.44) is negative, which is a contradiction. We have $\pi_M + \pi_{-M+1} - 1 = \pi_0^{2^M} + \pi_0^{2^{-M+1}} - 1$ by (7.8). Let $f(M) = 1 - \pi_0^{2^M} - \pi_0^{2^{-M+1}}$. Assume to the contrary that $\pi_0 \leq \frac{1}{2}$. Then $f(1) > 0$, and clearly we have $\lim_{M \rightarrow \infty} f(M) = 0$. We will show that $f(M) \geq f(M+1)$ for all $M \in \mathbb{Z}_+$. The derivative of $f(M)$ with respect to M is

$$\frac{df(M)}{dM} = -2^M \cdot \log(2) \cdot \pi_0^{2^M} \cdot \log(\pi_0) + 2^{-M+1} \cdot \log(2) \cdot \pi_0^{2^{-M+1}} \cdot \log(\pi_0), \quad (\text{F.45})$$

where $\log$ is the natural logarithm. Since $\pi_0 \leq \frac{1}{2}$ by assumption, we have $\log(2) \cdot \log(\pi_0) < 0$. Thus, (F.45) has the same sign with

$$2^M \cdot \pi_0^{2^M} - 2^{-M+1} \cdot \pi_0^{2^{-M+1}} = \frac{2^{2M-1} \cdot \pi_0^{2^M} - \pi_0^{2^{-M+1}}}{2^{M-1}}. \quad (\text{F.46})$$

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

$$2^{2M-1} \cdot \pi_0^{2^M} - \pi_0^{2^{-M+1}} = (2\pi_0)^{2M-1} \cdot \pi_0^{2^M-2M+1} - \pi_0^{2^{-M+1}} \leq \pi_0^{2^M-2M+1} - \pi_0^{2^{-M+1}} < 0 \quad (\text{F.47})$$

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

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

$$\sum_{i \geq 1} \pi_i = \sum_{i \geq 1} \pi_0^{2^i} \geq 0.8. \quad (\text{F.48})$$

Now we want to find an upper bound for

$$\sum_{i \leq 0} (1 - \pi_i) = \sum_{i \geq 0} (1 - \pi_0^{2^{-i}}). \quad (\text{F.49})$$

Since $\frac{1 - \pi_0^{2^{-i}}}{1 - \pi_0^{2^{-i-1}}} = 1 + \pi_0^{2^{-i-1}}$ is an increasing function for $i \geq 0$, we can upper bound (F.49) by the following geometric series

$$(1 - \pi_0) + (1 - \pi_0^{2^{-1}}) + (1 - \pi_0^{2^{-2}})(1 + r + r^2 + \dots) \leq 0.6, \quad (\text{F.50})$$

where $r = \frac{1 - \pi_0^{2^{-2}}}{1 - \pi_0^{2^{-3}}}$. But, (F.48) and (F.50) contradict to (7.7). $\square$

Lemma F.5.1 (Lipschitz condition). *The finite model satisfies the Lipschitz condition in L_1 -distance.*

Proof. Let $x = (x_i)_{i \in \mathbb{Z}}$ and $y = (y_i)_{i \in \mathbb{Z}}$ be two states of the finite model. By (7.3), we have

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

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