

DYNAMICS AND PHENOMENA IN STATEFUL MULTI-AGENT SYSTEMS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Jason Gaitonde

August 2023

© 2023 Jason Gaitonde
ALL RIGHTS RESERVED

DYNAMICS AND PHENOMENA IN STATEFUL MULTI-AGENT SYSTEMS

Jason Gaitonde, Ph.D.

Cornell University 2023

An important goal at the intersection of theoretical computer science and economics is to understand the long-run outcomes of complex processes in multi-agent systems. In many cases, such social systems are inherently algorithmic and strategic; in others, agents may act in simpler behavioral ways in response to some underlying randomness in their environment. In real-world settings, of course, both of these features will likely arise. But regardless of the precise model of agent behavior, these local interactions give rise to important macroscopic outcomes that are vital to theoretically understand and quantify.

A subtle factor in many such real-world systems is that their evolution strongly depends on past outcomes. For instance, strategic agents may use learning algorithms that adapt to the actions taken by the other agents in previous rounds, which in turn evolve based on past outcomes and actions. The underlying environment, strategic or stochastic, may itself also change in response to the sequence of past outcomes, as in ridesharing or information diffusion on networks. The complexities of statefulness induce substantial technical challenges that preclude existing analyses from applying. In this thesis, we develop new techniques towards understanding these important and dynamic social systems in two key settings: price of anarchy bounds in repeated games with state, and polarization and discord in models of opinion formation.

In Part I, we consider the interplay between learning algorithms and welfare in *repeated games with state* (also known as stochastic games). In particular, we

prove optimal *price of anarchy* guarantees for natural learning algorithms in two prototypical settings. First, we consider the long-run stability of queuing systems with strategic agents. Our main result shows that queuing systems with no-regret learners strategically competing will be stable so long as the queuing system has just *twice* the necessary resources under complete coordination. Nonetheless, we also demonstrate the *myopia* of no-regret learning in games with state by comparing these guarantees with those that can be obtained at long-run patient equilibrium.

We then consider the long-run outcomes of learning dynamics in repeated auctions with budgets, an abstraction of the important setting of Internet ad auctions. We show that when all strategic agents use a form of budget pacing, a popular algorithmic approach to strategic bidding in both theory and practice, the resulting outcome of the auctions achieves at least *half* the optimal welfare attainable by any allocation. Collectively, these results highlight the surprising effectiveness of learning and welfare in stateful environments, as well as several challenges and subtleties.

In Part II, we aim to capture a different class of macroscopic phenomena: namely, the long-run emergence of discord and polarization in models of opinion dynamics in networks. In these models, agents belong to a social network and maintain numerical opinions about one or more issues. They then update their beliefs according to simple behavioral rules in response to stochastic events and the beliefs of their neighbors in the network. While classical opinions of opinion formation, like the well-known DeGroot model, suggest approximate consensus ought be reached asymptotically, these theoretical predictions are contradicted by recent empirical observations.

We make progress towards theoretically understanding these important

global phenomena. We first demonstrate how polarization and discord can be induced by adversarial perturbations of Friedkin-Johnsen dynamics, focusing on the interplay between the underlying network structure and the algorithmic power of the adversary. We then demonstrate how polarization and issue alignment can arise, endogenously and robustly, in geometric models of opinion formation.

A crucial theme of our analyses in all settings will be determining how *local, round-by-round* properties of social dynamics, which are easier to reason about, can coalesce to produce *global, long-run* properties. Our results suggest new avenues and techniques towards analyzing these kinds of social processes more broadly.

BIOGRAPHICAL SKETCH

Jason Gaitonde was raised in Centerville, OH. He received a B.S. in Mathematics and a B.A. in Economics from Yale University in 2018. After receiving his Ph.D. in Computer Science under the supervision of Éva Tardos, he will be joining the Massachusetts Institute of Technology as a Postdoctoral Associate in the Department of Mathematics.

To my family: Mom, Dad, Kevin, and Rocky.

ACKNOWLEDGEMENTS

While there are simply too many people to thank, I must start with my advisor, Éva Tardos. Her technical excellence is of course deeply impressive, but her equal, if not greater, commitment to the overall well-being of her students is truly inspiring. Despite her innumerable responsibilities for the department and the computer science community at large, she has always been consistently and unreasonably generous with her time, advice, and feedback on ideas, writing, presentations, or anything else I could possibly need help with. I am extremely grateful for her support and patience in helping me grow as a researcher while pursuing various directions throughout theoretical computer science. When I arrived at Cornell, I had only a vague inkling of what serious research entailed; five years later, thanks in no small part to her excellent mentorship and guidance, I feel confident and ready to keep going. In short, thank you, Éva, for being everything I could have hoped for in an advisor and then some.

Jon Kleinberg and Brendan Lucier have both been incredible role models for me, the latter starting with an internship at Microsoft Research New England. I am continually amazed not just by their preternatural research intuition (impressive as it is), but also by their compassion and humility. Both of them have influenced me in more ways, research and otherwise, than I can recount here.

I owe a debt of gratitude to Ahmed El Alaoui—his enthusiasm, knowledge, and generosity with time and ideas have been a godsend, especially at times when I was not feeling particularly optimistic about research in general. One of my bigger regrets in grad school was not meeting him earlier, but hopefully there are many future projects to come. Thanks also to Eshan Chattopadhyay for being both a friend and mentor since I first started at Cornell, not to mention for serving on my committee; I feel quite fortunate that he started at the same

time I did.

Many thanks also to Bobby Kleinberg and Karthik Sridharan for serving on my thesis committee and being excellent sources of advice and technical guidance. I learned something (or more accurately, dozens of somethings) new every time I entered Bobby's office, starting with the Applied Mathematics Visit Days five years ago. Much of what I know about the theoretical foundations of learning comes from Karthik's teaching and various conversations with him over the years. Thanks also to Alex Slivkins, who I also worked with at Microsoft Research, for influencing me to keep asking interesting questions and exploring what research directions might be possible.

It has been a true pleasure being part of the Cornell Theory group, even as familiar faces have come and gone. To put it briefly, the culture is quite special: thanks to Abhishek, Ayush, Emily, Giannis, Jesse, Jyun-Jie, Kate, Makis, Manish, Marios, Mohit, Oli, Princewill, Raunak, Shijin, Sloan, Shawn, Spencer, and Tegan. Thanks also to Noah Stephens-Davidowitz for his help during the postdoc process, as well as for his infectious enthusiasm for theory.

Special thanks go out to Erika Fowler-Decatur, Melody Padgett, Becky Stewart, and Stacey Stone; I'm extremely grateful for all the assistance they provided over the years. Their expertise and helpfulness on how to navigate various Cornell rules, guidelines, and processes made focusing on research so much easier.

Over the course of my grad school career, I've been fortunate to have worked with several other excellent coauthors: thanks much to Max Hopkins, Chin Ho Lee, Yingkai Li, Bar Light, Shachar Lovett, Tali Kaufman, and Ruizhe Zhang. I've thankfully been able to meet a few of them in person since pandemic restrictions have been relaxed; I hope to meet the rest in person soon!

Academia has not always been a goal of mine—if anything, thinking about

difficult mathematical concepts never came as naturally to me as other subjects growing up. Thanks to Kenneth Barrett for first showing me how fascinating mathematics can be, and thanks to Asher Auel, Miki Havlickova, and Eduardo Faingold for helping me believe that I could actually succeed at it. I would never have arrived at theoretical computer science had it not been for incredible courses by Daniel Spielman and Dana Angluin that first exposed me to the area.

The results in this thesis were supported by a Cornell University Fellowship, National Science Foundation (NSF) Grant CCF-1408673, Air Force Office of Scientific Research (AFOSR) Grant FA9550-19-1-0183, and an internship at Microsoft Research New England.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	viii
List of Figures	xi
1 Introduction	1
1.1 Overview of Results	5
1.1.1 Price of Anarchy in Stateful Systems	5
1.1.2 Opinion Dynamics in Networks: Polarization and Discord	16
1.2 Themes	21
2 Preliminaries	28
2.1 Notation and Terminology	28
2.2 Basic Inequalities	29
2.3 Probabilistic Tools	31
2.3.1 Useful Tail Bounds	31
2.3.2 Convergence of Stochastic Processes	33
2.4 (Spectral) Graph Theory	35
I Learning and Welfare in Dynamic Environments	37
3 Games, Learning, and Price of Anarchy	38
3.1 Games and Equilibria	40
3.2 Learning in Repeated Games: The No-Regret Property	44
3.3 Price of Anarchy	49
3.4 Challenges: Price of Anarchy in Dynamic Systems	53
3.5 Chapter Notes	55
4 Strategic Queuing Systems I: Stability of No-Regret Learning	56
4.1 Model	58
4.1.1 Queuing Model	59
4.1.2 Benchmark: Centralized Feasibility	61
4.1.3 Discussion	64
4.2 Main Results	67
4.3 Overview of Proof	70
4.4 Geometric Queuing Systems	75
4.5 Proof of Main Result	78
4.6 Optimality of Main Result	87
4.7 Chapter Notes	91
4.7.1 Deferred Proofs	91
4.7.2 On the No-Regret Assumption	95

4.7.3	Related and Subsequent Work	95
5	Strategic Queuing Systems II: Stability of Patient Equilibria	99
5.1	Patient Queuing Games	103
5.2	Main Results	105
5.3	Overview of Proofs	106
5.4	Algorithmic Formulation of Costs	112
5.5	Properties of Rate Functions	115
5.5.1	Analytic Properties	116
5.5.2	Game-Theoretic Properties	119
5.6	Proof of Main Result	121
5.7	Proof of Convergence	130
5.7.1	Proof of Rate Upper Bound	138
5.7.2	Proof of Rate Lower Bound	149
5.8	Table of Notation	152
5.9	Chapter Notes	154
5.9.1	Deferred Proofs	155
5.9.2	The Need for Packet Priority in Patient Queuing	157
5.9.3	Related Work	159
6	Repeated Auctions with Budgets: Pacing Dynamics and Welfare	161
6.1	Preliminaries: Auctions, Welfare, and Pacing	166
6.1.1	Repeated Auctions with Budgets	166
6.1.2	Gradient-Based Pacing Dynamics	173
6.1.3	Liquid Welfare	176
6.2	Main Result: Optimal Liquid Welfare for GBP	178
6.3	Overview of Proof and Challenges	180
6.4	Proof of the Main Result	187
6.4.1	Proof Preliminaries	187
6.5	Chapter Notes	195
6.5.1	Deferred Proofs	196
6.5.2	Related Work	197
II	Opinion Dynamics: Beyond Consensus	200
7	Opinion Dynamics	201
7.1	Opinion Evolution in Networks	202
7.2	A Classical Model of Opinion Formation: DeGroot Dynamics	204
7.3	Challenges: Consensus, Discord, and Polarization	206
8	Adversarial Perturbations of Friedkin-Johnsen Dynamics	209
8.1	Overview of Results and Techniques	211
8.1.1	Spectral Properties of Discord	211

8.1.2	Network Defense	213
8.1.3	Mixed-Graph Objectives	214
8.2	Preliminaries and Friedkin-Johnsen Dynamics	217
8.3	Adversarial Optimization in FJ Dynamics	218
8.3.1	Disagreement	220
8.3.2	Repeated Disagreement	223
8.3.3	Polarization and Disagreement	224
8.3.4	Absolute Displacement	227
8.3.5	Other Adversary Constraints	228
8.4	Defending the Network	236
8.5	Mixed-Graph Objectives	240
8.5.1	Spectral Similarity	246
8.5.2	Spectral Dissimilarity	251
8.6	Chapter Notes	254
8.6.1	Related Work	254
9	Geometric Opinion Models: Dynamics of Polarization	257
9.1	Overview of Results and Techniques	261
9.2	Preliminaries	263
9.3	Models with Strong Polarization	267
9.3.1	Strong Polarization in Signed HJMR Model	269
9.3.2	Strong Polarization in the Party Model	273
9.4	Markov Chains and Polarization	281
9.4.1	Polarization and Invariant Distributions	283
9.4.2	HJMR Model and Weak Polarization	288
9.5	Chapter Notes	293
9.5.1	Related Work	293
	Bibliography	295

LIST OF FIGURES

- 5.1 Four queues compete for three servers, with maximally tight sets marked. The two outer queues exclusively send to the outer servers, and the inner queues send to two servers each, as indicated by the figure. One can verify that this profile is Nash and the subsets marked are tight. Consider whether one of the inner queues would want to deviate from the current profile. If she shifts probability to the inner server, the smallest tight set will rise in aging rate. However, if she instead attempts to deviate to one of the outer servers, the rate of all four queues will rise. 123

CHAPTER 1

INTRODUCTION

An important goal at the intersection of computer science and economics is to understand the long-run outcomes of complex social systems. Many such systems are inherently *algorithmic and strategic*, governed by the interactions of individual, selfish agents that learn from their environment. Given the widespread prevalence of data and computing power, as well as the rapid development of machine learning algorithms, agents can adapt to their environment and each other's behaviors in extraordinarily sophisticated ways. Indeed, algorithmic decision-making and learning are increasingly employed in myriad domains like online marketplaces, routing, and multi-robot systems.

As a result, understanding what kinds of long-run outcomes might arise in such systems is of large practical importance, and requires simultaneously exploring the limits of algorithmic learning and the interplay *between* these algorithms. After all, the strategic incentives and behaviors of each agent may not be aligned with the goal of promoting the overall social welfare of the system. For instance, in routing problems, each agent aims to choose a path to their destination that minimizes their own delay without regard for the delays of the other agents; a socially optimal outcome would instead centrally coordinate the routes of each agent to minimize the total system congestion.

In other important real-world multi-agent systems, however, agents may not necessarily behave strategically *per se*. Rather, their behavior may instead follow simple probabilistic or behavioral rules that nonetheless yield intricate macroscopic and long-run effects. Important social processes like the evolution of opinions or the spread of disease between individuals in a network may instead

follow some localized random process along the edges of the graph, rather than from some explicit optimization of a strategic actor. Investigating the relationship between these underlying structures and these stochastic features provides insights into important social phenomena like polarization and contagion.

Both of these kinds of qualitative properties of real-world agents, strategic or stochastic, serve as key challenges in obtaining provable guarantees about the *long-run outcomes* of these systems. Strategic interactions can be incredibly adaptive and sophisticated, or even chaotic, making their local evolution difficult to reason about. However, we at least have assurances that strategic agents at least weakly attempt to optimize some objective function, so we have hope that their actions could also weakly promote desirable aggregate properties. On the other hand, behavioral or probabilistic updates may be simpler or more explicit to describe, but then understanding how these microcanonical update rules can combine to produce long-run, global effects may become more difficult.

Of course, in typical real-world settings, both kinds of complicating factors will be concurrent. For instance, in a model of ridesharing, one may imagine that riders arrive in the market according to some underlying stochastic process and are willing to pay for rides so long as prices are below some threshold. Drivers might strategically choose locations to maximize their overall earnings over time, and once having accepted a new ride, must strategically choose a path to the desired location to minimize their travel time so that they may accept new rides. The ridesharing platform strategically chooses prices so as to maximize revenues by balancing increased ridership with lower pricing. In this case, the various strategic and random components of such a model interact in important, but technically challenging and delicate ways.

Regardless of the precise mechanisms driving a given multi-agent system of interest, a further complication is **statefulness**. That is, a fundamental property of many real-world processes is that their evolution strongly depends on past outcomes. Strategic agents may use learning algorithms that adapt to the actions taken by the other strategic agents in previous rounds. The underlying environment that the agents interact in may itself also change in response to the sequence of past outcomes, as in ridesharing or information diffusion on networks. These changes will in turn affect the future behavior of the agents, altering the evolution of the environment, and so on.

From a technical level, understanding long-run properties and outcomes of these systems becomes significantly more complex with state, even beyond the complications imposed by the factors above. As we will see below, the strong correlations and interdependencies across time render the existing theory for the multi-agent systems we consider inapplicable.

In this thesis, our main task is therefore to provide new theoretical tools towards understanding the *global phenomena* that arise from strategic and probabilistic dynamics in important multi-agent systems with state. Abstractly, every setting we consider will take place as a discrete-time process for $t = 1, \dots, T$ where $T \in \mathbb{N} \cup \{\infty\}$. At each timestep, agents may observe some new, independent random information that may change the underlying state of the system. Each agent then chooses an action based on past events and the current observable state of the system: these actions may be strategic or instead follow some simple behavioral rule. These actions will then be partially observable by agents, and the underlying state of the world will change in response to these actions.

Our first instantiation of this paradigm will be our investigations of learning in stateful repeated games. Here, the global property of these dynamics that will be the object of our study is the *aggregate welfare* of the agents. While a deep theory of learning in repeated games has been developed in the last few decades (see, for instance [35, 121, 88]), as well as powerful extensions of the classical price of anarchy theory to capture learning outcomes [25, 114], this theory cannot give useful provable guarantees when the underlying strategic environment evolves over time. Our work in Part I will develop new techniques for analyzing the welfare of no-regret learning algorithms in *strategic queuing systems*; in particular, we show that global, long-run stability of these queuing systems can be ensured even with algorithmic and strategic queues with just a mild increase in system resources. We then complement this result by highlighting the *myopia* of these algorithms and precisely capture the quantitative inefficiency that arises.

We also give new welfare guarantees for a popular class of learning algorithms in the equally important setting of repeated auctions with budgets. These guarantees are also optimal, showing that the underlying state of these repeated games induced by budgets can be mitigated by simple learning algorithms. In all, our results and techniques in Part I shed significant light on the interplay between learning and welfare in strategic environments that can evolve over time.

In Part II, we then instantiate this framework in the setting of opinion dynamics in networks to understand other important long-run properties of these systems. These models view agents as holding numerical opinions on one or more issues, which are updated over time in response to the opinions held by

the agent’s neighbors in the graph or possibly to other random stimuli. Our goal is again to give provable guarantees about the macroscopic, asymptotic features of this opinion formation.

One such feature that has been the subject of immense popular and empirical interest is *opinion polarization*. However, classical tractable models of opinion formation suffer from the problem of (approximate) consensus: many such models inherently cause opinions to converge, rather than meaningfully polarize. We thus provide new analyses of mechanisms that provably induce polarization and discord in variations of existing models of opinion dynamics. Our results show that polarization is a *robust* phenomenon in certain models, providing new insight into how these kinds outcomes can arise in empirical opinion dynamics.

Organization. We give a high-level overview of the main results presented in this thesis in Section 1.1. In Section 1.2, we then touch on several of the technical and conceptual themes that underlie much of our analysis. We then begin in earnest with our study of price of anarchy bounds in stateful environments in Part I, before turning to theoretical mechanisms of polarization and discord in Part II.

1.1 Overview of Results

1.1.1 Price of Anarchy in Stateful Systems

In Part I, we study the long-run outcomes of stateful, multi-agent systems when the dynamics are governed by the actions of strategic and competitive agents.

In particular, we aim to provide strong quantitative bounds on the *overall quality* of systems with strategic agents compared to centrally coordinated agents. Such results are known as *price of anarchy* bounds, a notion first introduced by Koutsoupias and Papadimitriou [99] in the computer science literature.

Formally, given some quantitative notion of social welfare, the price of anarchy is classically defined as the ratio between the optimal social welfare attainable by any choice of actions by the agents and the *worst-case* welfare that is achieved among selfish behaviors belonging to some mathematical class of strategies, like Nash equilibrium or no-regret learning (see Chapter 3 for precise definitions). In our investigation in Part I, we more liberally use the phrase “price of anarchy” to refer to any quantitative comparison between strategic outcomes and optimal coordinated outcomes. For instance, suppose a given system or game is parametrized by some notion of resources. Suppose further that one can prove that the social welfare obtained by strategic agents when the resources are sufficiently augmented is at least the optimal social welfare without these extra resources. In this case, we will abuse terminology and label such a result a price of anarchy bound all the same.

For different games, the objectives of selfish agents may be only weakly aligned, if at all, with a particular notion of social welfare. In such cases, the outcomes reached by strategic agents may have low social welfare compared to the best outcome if agents could instead be centrally coordinated. But for games or strategic environments admitting a price of anarchy bound, one obtains a provable guarantee that the system design is *robust* to decentralized, strategic behavior: while agents are themselves not explicitly aiming to promote social welfare, their actions nonetheless will produce welfare that is comparable to the

optimum. Such systems can be safely decentralized without much loss, possibly by augmenting resources in the system.

To see these concepts in action, consider the setting of routing in networks by Roughgarden and Tardos [117]. In this game, each agent chooses a path in the network from their source vertex to their sink vertex with the objective of minimizing their experienced delay along this path. The delay of an edge is some monotonic function of the number of agents traversing this edge in their chosen path. At a *Nash equilibrium*, selfish agents each select the path with minimum delay *given the equilibrium behavior of the others*; in other words, each agent selfishly optimizes her own delay function given the actions of the others. These equilibrium outcomes can be very different than the globally optimal choice of routes that minimizes the *total delay* of all agents. However, Roughgarden and Tardos show that in the case of atomless selfish routing with affine delay functions, the ratio between the total delay at any Nash equilibrium and that of the global optimum is bounded by $4/3$. In fact, they further show that in such routing games, the cost of any Nash equilibrium is no more than that of the centralized optimum when routing twice as much flow non-selfishly. Such an analysis gives clearly actionable prescriptions: to attain good performance, one can simply augment the amount of resources in the system.

A large body of work has established price of anarchy bounds for various well-studied games, like routing [117, 113, 12], scheduling [6, 91], auctions [102, 126, 116, 95], and network design [7, 8, 36, 51], among many others. Establishing POA results is important for several reasons. First, POA analyses help us understand the performance of real-world systems that already exist “in the wild” [115]. Equally as important, the quantitative understanding given

by POA-style analyses yield crucial insights towards the design of more robust systems with respect to selfish behaviors, in addition to shedding light on real-world outcomes like Braess' paradox [29, 41].

For our purposes, the true power of POA bounds arises from the fact they often extend to the outcomes obtained by agents employing efficient learning algorithms [25, 114]. This paradigm is commonly known as *learning in repeated games*, and a classical assumption from economics and computer science is that strategic agents use algorithms satisfying the *no-regret property*. The *no-regret* guarantee is satisfied by a number of natural and efficient algorithms [35, 11, 97, 10]. Informally, the no-regret property simply ensures that the learning algorithm can compete with the performance of any *fixed* action, a reasonable behavioral assumption even in practice [109]. Such extensions of price of anarchy guarantees seemingly capture the aggregate welfare of outcomes that arise from algorithmic decision-making, thereby increasing the applicability of this theory to real-world and modern systems.

However, a critical assumption of this theory is that agents play an “independent” copy of the *same* game in each round. In other words, the past sequence of actions taken by agents does not fundamentally change the nature of the game that is repeated in each round. In some settings, this approximation might be valid: when considering the repeated routing in traffic networks, it will indeed be true that traffic one day will not directly affect the congestion the following day as it will have dissipated. But when considering routing in computer networks, which operate on much faster timescales, the outcomes of previous events can have much stronger influences on the underlying game dynamics themselves. For instance, in such systems, packets may regularly drop

if on a congested link, forcing them to be re-sent in later periods. The aggregate effect is that future rounds may experience higher total congestion, thereby fundamentally changing the parameters of the routing game.

When this independence assumption on the repeated games is not satisfied, the existing theory of learning and games alluded to above fails to provide provable guarantees. Therefore, developing a deeper theory of the efficiency of strategic agents in repeated games that retain state is of large importance. This motivates the first main question that we will study in this thesis:

Question 1.1. *Do natural repeated games with state admit strong price of anarchy guarantees? If so, under what kinds of strategic behaviors?*

In Part I, we resolve Question 1.1 affirmatively in two important settings: strategic queuing systems and repeated auctions with budgets. In both cases, we obtain optimal guarantees on the long-run welfare of these settings under a variety of strategic behaviors. Our analysis provides new approaches to understanding the global welfare properties of complex strategic environments with state.

Strategic Queuing Systems

A quintessential setting of repeated games with state that will occupy our focus in Chapter 4 and Chapter 5 is discrete-time strategic queuing systems. The results described in this section are based on joint work with Éva Tardos [71, 72].

In the model we consider, n queues receive packets at heterogeneous rates parametrized by a vector $\lambda = (\lambda_1, \dots, \lambda_n) \in (0, 1)^n$ in decreasing order. Queue

i receives a new packet with probability λ_i at the beginning of each round, independently of all past events. Each queue with an uncleared packet may choose one of exactly m servers, which have (possibly unknown) service rates $\boldsymbol{\mu} = (\mu_1, \dots, \mu_m) \in [0, 1]^m$ in decreasing order, to attempt to clear their oldest uncleared packet thus far. Each server j then attempts to clear only the *oldest* packet among those it receives if any, and succeeds with probability μ_j . Under some fixed set of policies by queues to choose servers, we will say that the queuing system is *stable* if the expected number of uncleared packets is $O(1)$ (uniform in time).

In this setting, our focus on the price of anarchy in these queuing systems in Question 1.1 specializes to the following: *under what conditions on the arrival rates $\boldsymbol{\lambda}$ and service rates $\boldsymbol{\mu}$ will these queuing systems remain stable under strategic behavior by the queues?*

To cast this question more precisely in the price of anarchy framework, we first must determine what this relationship should be for even completely coordinated dynamics to be stable. We establish the following baseline result on necessary and sufficient conditions:

Theorem 1.1. *A queuing system with parameters $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$ is stable under some centralized (coordinated) scheduling policy if and only if for all $1 \leq k \leq n$,*

$$\sum_{j=1}^{\min\{k,m\}} \mu_j > \sum_{i=1}^k \lambda_i. \quad (1.1)$$

However, gaining insight into the dynamics of these queuing systems for completely decentralized, strategic policies will require overcoming several technical challenges. First, even ignoring strategic considerations, the underlying evolution of the system is itself an inherently stochastic system due to

the randomness in arrivals and clearances. More importantly, these systems exhibit extremely strong interdependencies due to the assumed priority structure over queues based on past performance.¹ These dependencies manifest even from round-to-round: the priority structure may fluctuate greatly every period. Third, as in the classical learning in games abstraction, we will not impose any specific choice of algorithm by the queues. Our results will be more general, but cannot exploit any specific features of their strategic behaviors beyond their defining guarantees, like equilibrium or no-regret.

In Chapter 4, we redeem the paradigm of no-regret learning in games by showing that despite the statefulness of the queuing system, queuing systems with no-regret learners will remain stable so long as one augments the resources of the system by at most a factor of 2:

Theorem 1.2. *If each queue in the strategic queuing system satisfies the no-regret property in choosing servers at each round, and if for all $k \leq n$,*

$$\sum_{j=1}^{\min\{k,m\}} \mu_j > 2 \sum_{i=1}^k \lambda_i,$$

then the queuing system under these dynamics is stable.

Moreover, the constant 2 cannot be lowered; for any $c < 2$, there exists strategic queuing systems and no-regret policies that will not be stable with the constant 2 replaced by c .

The positive result above is agnostic to the precise algorithm used by each queue; that they satisfy the no-regret guarantee is enough. We remark that in

¹As we will discuss, this priority structure is in some sense necessary towards guaranteeing strong price of anarchy guarantees. Simpler choices by servers to select packets provably fail to ensure stability even with superconstant excess system capacity.

these queuing systems, the agents (queues) may not know *any* underlying parameters of the system, nor the identities or even number of other agents. As a result, agents cannot learn sufficient information to understand the precise strategic environment they are competing in. However, the no-regret guarantee still gives a meaningful performance guarantee that is pivotal in establishing the above price of anarchy result, justifying its use even in stateful environments.

Nonetheless, in Chapter 5, we initiate a closer investigation of this price of anarchy guarantee. What we find is that no-regret learning can exhibit considerable *myopia* due to the fact it cannot reason about the long-run evolution of the repeated queuing game. We thus turn to quantifying this inefficiency by formulating a *patient queuing games* where queues are restricted to time-invariant policies over queues and seek to minimize their *long-run growth rate* given the policies of the others. Such a game models strategic outcomes that can arise in stable *equilibrium*.

Even formulating such a static game proves to be a severe probabilistic challenge. A key technical component in our study provides the following purely probabilistic result:

Theorem 1.3. *There exists an explicit, continuous function $r : (\Delta^{m-1})^n \rightarrow \mathbb{R}_{\geq 0}^n$ such that, if queues independently randomize over servers according to $\mathbf{p} \in (\Delta^{m-1})^n$, then the (random) long-run growth rate of each queue i is $r_i(\mathbf{p})$ almost surely.*

Here, Δ^{m-1} denotes the space of distributions over servers; this forms the strategy space of each queue when required to choose time-invariant policies. With this result, the long-run, equilibrium behavior of patient queuing dynamics reduces to the analysis of a static game with cost functions given above.

After developing several analytical properties of this game, we can then exactly quantify the efficiency loss of no-regret dynamics compared to patient equilibria. When queues can strategically reason about long-run outcomes, as with Nash equilibria of the patient queuing game, the price of anarchy drops to $\frac{e}{e-1}$. Formally, we prove the following:

Theorem 1.4. *In the patient queuing game, Nash equilibria exist. If it holds that for all $k \leq n$,*

$$\sum_{j=1}^{\min\{k,m\}} \mu_j > \frac{e}{e-1} \sum_{i=1}^k \lambda_i,$$

then the queuing system under these dynamics is weakly stable² at every Nash equilibrium of the game.

Taken in tandem, our positive resolution to Question 1.1 in discrete-time queuing systems produces a curious compromise: no-regret learning can sometimes exploit the underlying structure of important stateful games, but may not do so optimally. These results lie in contrast to the standard theory of repeated games and learning where learning and equilibrium outcomes often achieve the same welfare guarantees, illustrating the myriad new subtleties of learning, games, and welfare in stateful environments.

Repeated Auctions with Budgets

The results of the previous section, on the whole, were rather promising towards a satisfactory answer for Question 1.1. While simple learning dynamics did not

²For technical reasons, this conclusion only assures that the almost sure growth of each queue is sublinear almost surely. This mismatch arises because queue incentives in the patient queuing game are only specified up to *linear* asymptotic growth—a more technical, but not particularly illuminating, version of the game where queue incentives include sensitivity to lower-order growth as well would recover the original stability notion.

necessarily achieve optimal efficiency guarantees, they did achieve highly non-trivial performance that is only slightly suboptimal. To what extent is this finding robust to other important games and natural learning algorithms? We therefore investigate similar questions in a model of repeated auctions with budget constraints. This setting forms a realistic abstraction of Internet ad auctions, a ubiquitous and predominant form of advertising that has proven a lodestar for algorithmic game theory broadly. The results described below are based on joint work with Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins [70].

In this model, agents repeatedly compete in auctions for T rounds, but are constrained by limited budgets. Their goal is to maximize the utility they receive over the course of these auctions, which for now can be viewed as their value for the items they obtain via auction minus the price they spend for these items. These budgets form *global* constraints that govern their strategic behavior; these agents cannot simply view each auction in isolation. Thus, the remaining budgets form the underlying state of these repeated games that dynamically evolves in direct response to past outcomes.

While the price of anarchy of repeated auctions without budget constraints has been thoroughly investigated (see, for instance [102, 126, 116, 95]), the strategic and computational complexities of auctions with budgets imposes significantly new challenges, both on a practical and theoretical level. In practice (like in ad auctions), agents must *learn* to effectively bid in an online fashion while respecting their long-run budget, already a challenging special case of the learning problem of “bandits with knapsacks” [15, 93].

To help mitigate the difficulties of strategic bidding, a popular learning approach both in theory and practice [16, 42, 43] is done via *budget pacing algo-*

rithms that can be efficiently tuned. In ad auctions, these autobidders are even provided by the platform itself. At a high-level, these algorithms aim for their long-run expenditure to match the budget. They do so by dynamically maintaining a “pacing multiplier” with which to shade valuations in each round to produce bids. A large class of *gradient-based pacing algorithms* do so in a very natural way: the algorithm simply increases/decreases the multiplier depending on the amount the previous round expenditure exceeds/falls short of the target spend rate. Prior work of Balseiro and Gur [17] shows how this algorithm can indeed be interpreted as gradient descent, justifying the terminology. But due to the popularity of this approach, our formulation of Question 1.1 in this setting becomes quite pressing: *what aggregate outcomes arise when all agents employ gradient-based pacing?*

Our main finding is that with respect to liquid welfare [54], a standard welfare notion in budgeted settings introduced by Dobzinski and Paes Leme, gradient-based pacing achieves optimal price of anarchy guarantees. Briefly, the liquid welfare of an allocation to an agent measures her willingness to pay for the allocation she receives while respecting budget constraints, so that the liquid welfare of a collective allocation is then the sum of individual liquid welfares. Our result can be stated informally as follows:

Theorem 1.5. *Suppose that all agents use gradient-based pacing in repeated core auctions³ with budgets. Then the expected liquid welfare obtained by the outcome of repeated auctions with budgets is at least half the expected liquid welfare obtained by any allocation rule. Moreover, this factor cannot be improved even for the class of offline equilibrium of these repeated games.*

³Core auctions are a large class capturing many settings of importance, like first-price, second-price, and generalized second-price auctions.

A key novelty of our approach in establishing this theorem is that it will hold without requiring any form of convergence of the learning dynamics.⁴ This restriction is actually crucial; recent complexity-theoretic results of Chen, Kroer, and Kumar imply that it is unlikely for *any* efficient learning to converge [38] in the case of repeated second-price auctions with budgets, which are captured by our result. Instead, we must argue directly using the explicit form of the local updates in gradient-based pacing to infer these aggregate price of anarchy guarantees. This result gives the first welfare guarantee for any natural learning dynamics in these general settings.

As with strategic queuing systems, we thus find that “simple” learning algorithms can achieve strong long-run aggregate welfare guarantees despite the complexities of evolving state. Unlike the queuing setting, though, pacing dynamics do not suffer any excess loss beyond what can be reasonably attained in any sufficiently strategic outcome. Our results and techniques in Part I, taken collectively, provide first steps towards a more general understanding of the aggregate guarantees of simple learning dynamics in stateful environments, as hoped for in Question 1.1.

1.1.2 Opinion Dynamics in Networks: Polarization and Discord

In Part II, we turn from the analysis of *strategic* multi-agent system to systems where agents instead follow explicit behavioral assumptions: in particular, we

⁴In standard repeated games, no-regret dynamics will typically not converge either in the usual sense. However, the dynamics converge to the set of coarse correlated equilibrium of the underlying game, and this fact is essential in proving price of anarchy guarantees in this theory.

will turn our attention to understanding *polarization and discord* in models of opinion formation in networks. While many of the same analytical tools will reappear for analyzing these dynamic systems, we can no longer connect strategic behaviors by the agents to these long-run phenomena. Instead, we must find new approaches to understanding these long-run properties.

To motivate these models, observe that people’s opinions naturally evolve in response to a variety of external factors, like interactions with other individuals, campaign messaging, media reports, political figures, and random events. Understanding how opinions evolve and the emergent qualitative features of these dynamics remains a subject of immense interest in the computer science, economics, and social science communities.

In recent years, a crucial phenomenon of interest is that of *polarization*, where agents roughly partition into groups holding diametric views. That is, rather than agents holding a rich spectrum of beliefs, individuals instead typically belong to opposite clusters even in cases where beliefs on separate topics ostensibly ought not be correlated. For instance, in the United States beliefs on gun laws and climate change are highly correlated despite the fact that these issues may not have an *a priori* connection. A related, but similar, phenomenon is *discord*, where agents maintain fairly different beliefs, but do not necessarily cluster or partition.

A prototypical way to model the evolution of the beliefs of agents, and thereby address phenomena like polarization, is to assume that the opinion of each agent is scalar- or vector-valued and evolves according to some prescribed update rule. Many classical models, like the DeGroot and Friedkin-Johnsen models [47, 63], study opinion evolution as a form of social learning where

agents update beliefs based on the graph-weighted opinions of their neighbors within an ambient network. However, the original forms of these models do not appear to provide a satisfactory mechanism to explain polarization or discord, as the dynamics provably lead to *less* polarization or disagreement than in the initial configuration.

Our main goal in Part II is to provide new insights into these phenomena. The specific question we study in Part II can thus be formulated as follows:

Question 1.2. *How do long-run phenomena like polarization and discord provably and naturally arise in models of opinion formation?*

The results stated in this section are based on joint work with Jon Kleinberg and Éva Tardos [67, 68].

Adversarial Perturbations of Friedkin-Johnsen Dynamics

In Chapter 8, we consider an approach to Question 1.2 motivated by prominent recent events wherein an external actor explicitly seeks to sow discord in a social network. Concretely, we consider the Friedkin-Johnsen dynamics mentioned above. In this model, agents lie on a weighted, connected graph $G = (V, E)$ and hold an immutable, internal scalar opinion. At each time, each agent updates their opinion to the graph-weighted average of their neighbor's opinions *as well as their own internal opinion*. The incorporation of these fixed opinions provides an anchor that prevents complete consensus, as in the related DeGroot model.

We model this problem as an external actor that can *seed* initial opinions subject to a budget constraint; in real-world settings, this might correspond to

targeted advertising or by interactions on social media sites. The adversary chooses these opinions to optimize their objective, knowing that these seeded opinions will diffuse in the graph according to the Friedkin-Johnsen dynamics.

Our main finding, informally stated, is that *the spectral properties of the underlying network govern the adversary's power to induce polarization and discord*. Moreover, the adversary's power to sow discord is often bounded *independent of graph topology*.⁵ This result holds for a variety of adversarial optimization problems, as well as a variety of budget constraints. Therefore, we obtain precise control of the adversary's power and strategies from the eigenstructure of the underlying graph; conversely, from this spectral characterization, we gain insights into the *combinatorial* structure of networks that are well-insulated from these kinds of adversarial attacks.

We also consider the algorithmic problem of *defending the network* from such perturbations. This setting forms a natural min-max game where the network defender insulates certain nodes from perturbations subject to a budget constraint (for instance, via online literacy campaigns), and then the adversary perturbs the opinions to optimize their objective. Here, we show that the network defense problem is efficiently solvable in a variety of settings via a reduction to convex programming.

To conclude, we also propose a novel mechanism for polarization by *decoupling* the network that measures discord and the network that diffuses these opinions. For instance, opinion formation may be driven by online interactions while disagreement may take place on the axis of one's physical social network. In this case, we show that the power of the adversary can greatly increase if

⁵That is, the adversary's power is exactly captured by spectral parameters, but these remain universally bounded for any graph.

these graphs have different combinatorial or spectral structure. These results provide a new explanation of real-world polarization via a misalignment of the modes of information diffusion and disagreement.

Polarization in Geometric Opinion Dynamics

To conclude, we analyze a fascinating proposal from a recent model of opinion formation by Hazła, Jin, Mossel, and Ramnarayan (HJMR) [89]. Their model views opinions as vector-valued on a high-dimensional sphere, where each component reflects the magnitude and sign of one's opinion of each issue. At each time step, all agents observe a random vector drawn uniformly from this sphere representing a random issue, political figure, or influencer. Each agent then evaluates the inner product between their current beliefs and this new issue; she then moves either toward or away from this new direction proportional to the magnitude of this inner product depending on the sign.

The key insight of HJMR is that in this geometric model of opinion formation, *agents intrinsically polarize with probability one*. In other words, polarization naturally and necessarily occurs purely due to the geometry and probabilistic features of this model. However, their analysis heavily relies on the spherical symmetry of the random issue distribution: this condition is required to show that the angle between two distinct opinions forms a martingale, leaving open how general this polarization phenomenon is. Moreover, this geometric model does not incorporate any network structure. The updates of the agents are only correlated through their joint movement towards/away from the random issue that arises, but the entire dynamics are otherwise in isolation. Therefore, while this model can provide satisfactory explanations of crucial long-run prop-

erties like polarization, it cannot simultaneously capture the natural intuition that opinion formation depends in part on social interactions.

We make significant progress towards addressing these gaps. To do so, we consider two new variations of the HJMR dynamics which we call the *signed HJMR* model and *party* model. This latter model notably incorporates network effects much like in DeGroot or Friedkin-Johnsen dynamics. Our main result shows that in both settings, polarization provably occurs under robust assumptions on the random issue vectors and in any dimension. In particular, the polarization observed by HJMR seems to be a significantly more *universal* phenomenon. To analyze these models, we provide novel and quite general reductions that reduce proving polarization to giving much simpler lower bounds on the probabilities of simpler geometric configurations.

1.2 Themes

Collectively, the results described above make significant progress towards understanding the macroscopic phenomena and dynamics of complex multi-agent systems. Underlying them are several unifying techniques and ideas, each of which suggest promising general approaches towards a more comprehensive study of important social systems that share similar challenging qualities, like state, stochasticity, and strategicness.

Local Dynamics to Long-Run Phenomena. A persistent theme of all of the main results presented in this thesis is the relation between *local, round-by-round* guarantees and *global, macroscopic* outcomes. Indeed, each of our results described in the previous section aims to provide strong quantitative and qualitative char-

acterizations of the long-run outcomes of natural processes in multi-agent systems, but these processes themselves are only specified in terms of their local evolution. A key contribution of this thesis is developing a variety of techniques to connect local guarantees to the long-run properties that we seek to understand, like stability in queuing systems or polarization in networks.

The simplest manifestation of this approach arises when the aggregate effects of local updates admit simple, deterministic closed-form representations. Such will be the case in our analysis of how discord and polarization can arise via adversarial perturbations of Friedkin-Johnsen dynamics in Chapter 8. The repeated averaging of this class of dynamics is well-known to admit nice linear-algebraic closed forms, making our formulation of the corresponding adversarial optimization problem quite clean. In particular, understanding this problem becomes tantamount to studying *quadratic forms* in the Laplacian of the underlying network. These representations naturally lead to connections to the spectral structure of the graph and will be the workhorse of our analysis of the algorithmic and analytical aspects of this setting.

A substantially more challenging variation of this kind of reasoning will arise in our study of the patient queuing game in Chapter 5. In that setting, any choice of strategies by the queuing agents induces a high-dimensional Markov chain whose transitions are quite difficult to control on the level of round-by-round sample paths. The primary bulk of our analysis will be to reduce the analysis of the dynamics of this family of high-dimensional Markov chains to the analysis of an explicit, though still high-dimensional, continuous function that gives the long-run costs in the patient queuing game. We do so by proving that while the sample paths of these Markov chains may be quite volatile on

smaller timescales, there is nonetheless sufficient structure in the evolution to infer almost surely deterministic long-run behavior that even admits an algorithmic description. We then complete our proof of the price of anarchy in such queuing systems by studying the various analytic properties of this explicit representation. While the resulting technical task is significantly more complex, this high-level approach bears conceptual similarities to the plan in Chapter 8.

In general, we will not be able to reduce the analysis of global features of these complex, evolving multi-agent systems to a completely static problem. No such representation will even exist, rendering this approach completely ineffective. In Chapter 4, we consider the stability of no-regret learners in these same queuing systems. We will not assume any specific structure on the choices made by the queuing agents other than they satisfy the no-regret property on sufficiently large, but finite, timescales. The generality of this assumption inherently prevents us from attempting to emulate the previous approaches simply because the round-by-round dynamics are not even well-specified. In this case, we leverage only these local no-regret guarantees to prove that a suitable potential function necessarily decreases in expectation whenever it is sufficiently large. By then employing deep supermartingale-type results, we show that these local guarantees obtained purely from the local no-regret property imply global stability of the dynamics.

We employ a combination of the previous two approaches to leveraging local structure in our analysis of polarization in Chapter 9. The underlying opinion dynamics form a Markov chain, albeit a somewhat complicated one that also lives in high-dimensions. However, our argument shows that locally approaching polarization forms a self-fulfilling prophecy; by making progress to polar-

ization in a step (which may hold even with low probability), the probability of making similar progress in the next step increases by a large enough factor. These local improvements can be glued together to argue that the probability of polarizing is at least a nonzero constant. Using general techniques from the theory of Markov chains and stochastic properties, we show that this is enough to establish almost sure polarization.

In contrast to the above approaches, our analysis of the long-run welfare of pacing dynamics in repeated auctions with budgets in Chapter 6 requires concrete arguments with the explicit form of the iterative updates taken by the agents. Our approach partitions time, separately for each agent, into periods where we can directly relate the local movement of the pacing multipliers to the value obtained by the agents via the explicit updating rule. This observation linking local dynamics to individual performance on longer scales forms the basis for our proof that the expected liquid welfare of these learning dynamics is high.

In all cases, the central crux of our analysis is developing an appropriate analytical approach to leverage the local properties towards our aims. These approaches may vary from problem to problem, and understanding the right foothold is essential to producing optimal bounds on the long-run outcomes of these stateful systems.

Stateful Systems as Structured Stochastic Processes. Several of the results in this thesis fundamentally exploit the interpretation of these models as structured stochastic processes. Rather than endeavoring to reason about all of the precise microcanonical or strategic aspects of the model, our approach instead relies heavily on the probabilistic properties. Thus, randomness is a powerful

unifying tool in our analysis.

Our primary work in Chapter 4 and Chapter 5 studying the stability of strategic queuing systems will be proving sub/supermartingale-type properties of appropriate *potential functions* by leveraging the local guarantees described above. Chapter 4 considers a potential function of the *squares* of the ages of the queues; a powerful result of Pemantle and Rosenthal then shows it is sufficient to argue that this potential decreases whenever it exceeds some threshold. This deep result thus reduces the problem of stability to proving the requisite supermartingale-type condition. Chapter 5, on the other hand, considers as potential functions the maximum queue age and an appropriate weighted average of queue ages to develop the explicit representation of the long-run outcomes of stationary queuing dynamics. In these cases, more elementary probabilistic techniques are devised to prove the supermartingale and submartingale properties, respectively.

In both chapters, an extremely useful resource is utilizing the principle of deferred decisions: doing so imposes independence in the local analyses, thereby enabling powerful tools like concentration of measure. Designing and analyzing these potential functions is, for now, more of an art than science; there does not appear to be a generic way to construct or study them, but we expect that these kinds of analyses could be of broader applicability in a broader theory of repeated games with state.

Chapter 9 employs the Markov structure of geometric opinion dynamics to great effect. The Markov structure, as well as powerful zero-one laws, enable us to reduce from proving polarization occurs almost surely to just with constant probability. These tools help show the robustness and universality of polar-

ization in these kinds of dynamics. Similar zero-one laws in tandem with the Markov property will also help us illustrate more general notions of polarization in these models.

The analysis of pacing dynamics in repeated auctions with budgets of Chapter 6, though also interpretable as a high-dimensional Markov chain, will leverage stochasticity in fundamentally different, but equally important, ways. The independence of valuations across time will ensure that these dynamics cannot concentrate on events where the valuations are adversarially correlated with the evolution of the bids of the agents. In particular, the inherent randomness of these systems in the valuations prevents pacing dynamics from conspiring to yield low welfare.

The Role of Combinatorial/Geometric Structure. Finally, an implicit or explicit, but extremely important, role is played by the underlying combinatorial or geometric structure of several of the models we consider. In both Chapter 4 and Chapter 5, we identify (approximate) *majorization* of the queue arrival rates and server rates as the canonical structural property that leads to stability. Indeed, the geometry of majorization is integral in establishing the limits of both centralized feasibility of our queuing systems and stability of strategic dynamics, albeit in different ways. The former is due to a classical connection obtained via the Birkoff-von Neumann theorem that connects majorization to matchings; however, the latter arises because of a convenient consequence of majorization to inner products in our proof of stability.

Our study of the patient queuing game further reveals a fascinating combinatorial structure on the set of queues under any fixed choice of randomizations over servers. Roughly speaking, given strategies for each queue, the set

of queues partitions into nontrivial subsets that each asymptotically age at the same rate; however, each subset in this partition will further refine in nontrivial ways to reflect subtle internal structure. In fact, within each subset of the partition, a striking lattice structure emerges due to the submodularity properties of the explicit, algorithmic representation we develop. This substructure proves instrumental in our price of anarchy argument.

In Chapter 8, the relevant combinatorial structure becomes the underlying network topology. We show that the spectral properties of the network govern the objective value of various adversarial optimization problems in the Friedkin-Johnsen dynamics. Indeed, this underlying geometry plays two important, but competing roles: the edges of the network play host to the diffusion process of opinions in the graph, but also the measurement of polarization and discord. These factors lie in fundamental tension, but lead to precise bounds on an adversary's power to induce discord. We actually generalize classical spectral analyses to account for multiple graphs simultaneously when these processes are decoupled, and show again how these competing geometries can affect an adversary's power.

The geometry of opinions on a high-dimensional sphere forms a pivotal aspect of our analysis in Chapter 9. Indeed, our approach to proving polarization in multiple models will rely on identifying favorable and well-behaved configurations on the sphere. This geometric analysis directly furnishes the requisite probabilistic estimates for our more robust approach to showing polarization provably occurs.

CHAPTER 2
PRELIMINARIES

In this section, we collect several useful, but fairly standard facts that we will need throughout this thesis. We also document some fairly common notational conventions. The reader should feel free to skip this section and use it as a reference if needed.

2.1 Notation and Terminology

In this thesis, random variables will typically be denoted by capital letters (i.e. X, Y, Z, \dots), while vectors will be bolded (i.e. $\mathbf{x}, \mathbf{y}, \mathbf{z}, \boldsymbol{\mu}, \boldsymbol{\lambda}$, etc). If a random variable X has some distribution \mathcal{D} , we write $X \sim \mathcal{D}$. We use the notation $\text{Geom}(p)$ to denote a geometric distribution with parameter p , $\text{Bern}(p)$ for a Bernoulli distribution that is 1 with probability p and 0 otherwise, and $\text{Bin}(n, p)$ for a binomial distribution with parameters n and p .

An event occurs *almost surely* if it has probability 1. We use standard $O(\cdot)$, $o(\cdot)$, and $\Theta(\cdot)$ notation. We will sometimes write $f(n) \asymp g(n)$ if $f(n) = \Theta(g(n))$. We will also consider the following norms: for a strictly positive vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_n)$, with $\lambda_1 \geq \dots \geq \lambda_n > 0$, we define the following two weighted ℓ_p norms on \mathbb{R}^n : $\|\mathbf{x}\|_{\boldsymbol{\lambda},1} \triangleq \sum_{i=1}^n \lambda_i |x_i|$ and $\|\mathbf{x}\|_{\boldsymbol{\lambda},2} \triangleq \sqrt{\sum_{i=1}^n \lambda_i x_i^2}$. It is easily seen that for any \mathbf{x} , $\|\mathbf{x}\|_{\boldsymbol{\lambda},1} \asymp \|\mathbf{x}\|_{\boldsymbol{\lambda},2}$ (where the constants depend on $\boldsymbol{\lambda}$) via Cauchy-Schwarz, see Lemma 2.5. We also denote the sparsity “norm” by $\|\mathbf{x}\|_0 = |\{i \in [n] : x_i \neq 0\}|$.

We will later use the following **fractional sum** operation $\oplus : \mathbb{R}_{\geq 0}^2 \times \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$:

$$\frac{a}{b} \oplus \frac{c}{d} \triangleq \frac{a+c}{b+d}.$$

We will later repeatedly use the following simple fact:

Fact 2.1. For all $a_1, \dots, a_n \geq 0$ and $b_1, \dots, b_n > 0$,

$$\min_{i \in [n]} \frac{a_i}{b_i} \leq \frac{a_1}{b_1} \oplus \dots \oplus \frac{a_n}{b_n} = \frac{\sum_{i=1}^n a_i}{\sum_{i=1}^n b_i} \leq \max_{i \in [n]} \frac{a_i}{b_i}.$$

Moreover, equality holds in either of the inequalities if and only if both inequalities are tight.

Given a $n \cdot m$ -dimensional vector $\mathbf{p} = (p_1, \dots, p_n)$, where $p_i \in \mathbb{R}^m$, we will write p_{ij} for the j th element of p_i . Given a vector $\mathbf{x} \in \mathbb{R}^n$ and a subset $I \subseteq [n]$, we write \mathbf{x}_I to denote the vector restricted to the components in I . Given a set S , we will write $\mathcal{P}(S)$ to denote the power set.

2.2 Basic Inequalities

We will need the notion of *majorization* of nonnegative vectors:

Definition 2.1. Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^n$, and assume that $x_1 \geq x_2 \geq \dots \geq x_n \geq 0$ and analogously for \mathbf{y} . Then \mathbf{x} **majorizes** \mathbf{y} (written $\mathbf{x} \geq \mathbf{y}$) if for each $1 \leq k \leq n$

$$\sum_{i=1}^k x_i \geq \sum_{i=1}^k y_i.$$

If the above inequalities are strict for each $1 \leq k \leq n$, then \mathbf{x} **strictly majorizes** \mathbf{y} (written $\mathbf{x} > \mathbf{y}$).¹

Fact 2.2. Suppose \mathbf{x} majorizes \mathbf{y} . Then for any nonnegative, monotone decreasing sequence $z_1 \geq \dots \geq z_n \geq 0$,

$$\sum_{i=1}^n z_i x_i \geq \sum_{i=1}^n z_i y_i$$

¹Majorization is often used for probabilities, and hence defined so that the total sums are equal; we omit this condition in our work.

Proof. Multiply the equations by appropriate scalars in the definition and sum to obtain the inequality. \square

Fact 2.3. Suppose $a, b, c \geq 0$ and that $a - b \leq c$. Then

$$\sqrt{a} - \sqrt{b} \leq \min \left\{ \frac{c}{2\sqrt{b}}, \sqrt{c} \right\}.$$

Proof. The first inequality arises from rearranging and concavity of the square-root function:

$$\sqrt{a} \leq \sqrt{b} \sqrt{1 + c/b} \leq \sqrt{b}(1 + c/2b).$$

The second follows from assuming without loss of generality that $a \geq b$ and observing the claim is implied by $\sqrt{a} - \sqrt{b} \leq \sqrt{a-b}$, which holds by squaring and simple algebra. \square

Fact 2.4. Suppose $a, b, c \geq 0$. Then $a - b \geq c$ implies

$$\sqrt{a} - \sqrt{b} \geq \frac{c}{2\sqrt{a}}.$$

Proof.

$$a - b = (\sqrt{a} - \sqrt{b})(\sqrt{a} + \sqrt{b}) \geq c \implies \sqrt{a} - \sqrt{b} \geq \frac{c}{\sqrt{a} + \sqrt{b}} \geq \frac{c}{2\sqrt{a}}. \quad \square$$

Recall that we defined the following two weighted ℓ_p norms on \mathbb{R}^n : $\|\mathbf{x}\|_{\lambda,1} \triangleq \sum_{i=1}^n \lambda_i |x_i|$ and $\|\mathbf{x}\|_{\lambda,2} \triangleq \sqrt{\sum_{i=1}^n \lambda_i x_i^2}$.

Lemma 2.5. For all $x \in \mathbb{R}^n$,

$$\sqrt{\lambda_n} \|x\|_{\lambda,2} \leq \|x\|_{\lambda,1} \leq \sqrt{\sum_{i=1}^n \lambda_i} \|x\|_{\lambda,2}.$$

Proof. For the first inequality,

$$\|x\|_{\lambda,1}^2 = \sum_{i,j=1}^n \lambda_i \lambda_j |x_i| |x_j| \geq \sum_{i=1}^n \lambda_i^2 x_i^2 \geq \lambda_n \sum_{i=1}^n \lambda_i x_i^2 = \lambda_n \|x\|_{\lambda,2}^2.$$

The second is a routine application of Cauchy-Schwarz:

$$\sum_{i=1}^n \lambda_i |x_i| = \sum_{i=1}^n \sqrt{\lambda_i} (\sqrt{\lambda_i} |x_i|) \leq \sqrt{\sum_{i=1}^n \lambda_i} \sqrt{\sum_{i=1}^n \lambda_i x_i^2} = \sqrt{\sum_{i=1}^n \lambda_i} \|x\|_{\lambda,2}. \quad \square$$

2.3 Probabilistic Tools

2.3.1 Useful Tail Bounds

We will use the following standard concentration results throughout this thesis. In all cases, these bounds are either well-known or can be easily derived from well-known concentration inequalities.

Lemma 2.6 (First Borel-Cantelli Lemma, Theorem 2.3.1 of [55]). *Let A_1, A_2, \dots be a sequence of events with $\sum_{i=1}^{\infty} \Pr(A_i) < \infty$. Then with probability one at most finitely many of the A_i occur.*

Lemma 2.7 (Azuma-Hoeffding). *Let $\{\mathcal{F}_k\}_{k \leq n}$ be any filtration and let A_k, B_k, Δ_k satisfy the following conditions:*

1. Δ_k is \mathcal{F}_k -measurable and $\mathbb{E}[\Delta_k | \mathcal{F}_{k-1}] = 0$. That is, the Δ_k form a martingale difference sequence.
2. A_k, B_k are \mathcal{F}_{k-1} -measurable and satisfy $A_k \leq \Delta_k \leq B_k$ almost surely.

Then

$$\Pr\left(\sum_{k=1}^n \Delta_k \geq t\right) \leq \exp\left(\frac{-2t^2}{\sum_{k=1}^n \|B_k - A_k\|_{\infty}}\right).$$

In the first part of this thesis, we will frequently need to appeal to concentration bounds for Bernoulli and geometric random variables. This first bound is obtained from Witt [131]:

Lemma 2.8 (Theorem 1 in [131]). *Let X_1, \dots, X_n be i.i.d. $\text{Geom}(\lambda)$ random variables, so that $\mathbb{E}[X_i] = \frac{1}{\lambda}$. Let $s = \frac{n}{\lambda^2}$ and $Z_n = \sum_{i=1}^n X_i$. Then for all $\delta > 0$,*

$$\Pr\left(Z_n - \frac{n}{\lambda} < -\delta\right) \leq \exp\left(\frac{-\delta^2}{2s}\right) \quad \text{and} \quad \Pr\left(Z_n - \frac{n}{\lambda} > \delta\right) \leq \exp\left(\frac{-\delta}{4} \min\{\delta/s, \lambda\}\right).$$

From Lemma 2.8, it is straightforward to derive the following using standard maximal inequalities (i.e. Etemadi's inequality, 22.5 in [22])

Corollary 2.9. *Under the assumptions and notation of Lemma 2.8, for any $\epsilon \in [0, 1]$,*

$$\Pr\left(\max_{1 \leq j \leq n} \left|Z_j - \frac{j}{\lambda}\right| > \frac{\epsilon n}{\lambda}\right) \leq 6 \exp\left(\frac{-\epsilon^2 n}{36}\right).$$

An immediate consequence of the previous corollary and a union bound is the following tail bound on a family of geometric random variables:

Corollary 2.10. *Let $\{G_k^i\}_{i \in [n], k \in [w]}$ be a family of independent geometric random variables such that for all i, k , $G_k^i \sim \text{Geom}(\lambda_i)$. Let $Z_q^i = \sum_{k=1}^q G_k^i$. Then for any $\epsilon \in [0, 1]$,*

$$\Pr\left(\exists i \in [n], q \in [w] : \left|Z_q^i - \frac{q}{\lambda_i}\right| \geq \frac{\epsilon w}{\lambda_i}\right) \leq 6n \exp\left(\frac{-\epsilon^2 w}{36}\right). \quad (2.1)$$

Similarly, a standard application of the multiplicative Chernoff bound and union bounds imply the following tail bounds on Bernoulli ensembles:

Lemma 2.11. *Let $\{I_k^j\}_{j \in [m], k \in [w]}$ be an independent Bernoulli ensemble such that for all j, k $I_k^j \sim \text{Bern}(\mu_j)$ with $\mu_1 \geq \mu_2 \geq \dots \geq \mu_m$. Then for all $\delta \in [0, 1]$,*

$$\Pr\left(\exists q \in [m] : \sum_{j=1}^q \sum_{k=1}^w I_k^j \leq (1 - \delta)w \left(\sum_{j=1}^q \mu_j\right)\right) \leq m \exp\left(\frac{-\delta^2 w \mu_1}{2}\right). \quad (2.2)$$

Finally, we will need the following crude bounds on the moments of geometric and binomial random variables.

Lemma 2.12. *Let $X \sim \text{Geom}(\lambda)$. Then for all $k \geq 1$, $\mathbb{E}[X^k] \leq \frac{c_k}{\lambda^k}$, where c_k is a constant depending on k but not on λ .*

Lemma 2.13. *Let $X \sim \text{Bin}(n, p)$, where $p \in (0, 1]$ is considered fixed. Then, for any fixed integer $k \geq 0$, $\mathbb{E}[X^k] \asymp n^k$, where the implicit constants depend on p and k , but not n .*

Proof. By definition, $X = \sum_{i=1}^n X_i$, where $X_i \sim \text{Bern}(p)$ are i.i.d. We clearly have

$$X^k = \sum_{1 \leq i_1, \dots, i_k \leq n} \prod_{j=1}^k X_{i_j}.$$

Note that products of these indicator variables remain indicator random variables, and it is easy to see that for any indices $1 \leq i_1, \dots, i_k \leq n$, $p^k \leq \mathbb{E}[\prod_{j=1}^k X_{i_j}] \leq p$. Taking expectations and summing, we obtain $p^k n^k \leq \mathbb{E}[X^k] \leq p n^k$. \square

2.3.2 Convergence of Stochastic Processes

In the first part of this thesis, we will repeatedly require the following powerful theorem of Pemantle and Rosenthal which asserts that a sufficiently nice random process has uniformly bounded moments so long as it has expected negative drift conditioned on being large:

Theorem 2.14 (Theorem 1 in [111]). *Let X_1, X_2, \dots be a sequence of nonnegative random variables with the property that*

1. *There exists constants $\alpha, \beta > 0$ such that if $x_t > \beta$, then*

$$\mathbb{E}[X_{t+1} - X_t | \mathcal{F}_t] < -\alpha,$$

where the σ -algebra $\mathcal{F}_t = \sigma(X_0, \dots, X_t)$ is the history until period t and $X_t = x_t$.

2. There exists $p > 2$ and $\theta > 0$ a constant such that for any history,

$$\mathbb{E}[|X_{t+1} - X_t|^p | \mathcal{F}_t] \leq \theta.$$

Then, for any $0 < r < p - 1$, there exists an absolute constant $M = M(\alpha, \beta, \theta, p, r)$ not depending on t such that $\mathbb{E}[X_t^r] \leq M$ for all t .

A somewhat more interpretable consequence of establishing uniformly bounded r th-moments for every $r \geq 0$ for a stochastic process, for instance using Theorem 2.14, is a bound on the growth rate of stochastic process that holds with probability 1:

Lemma 2.15. *Suppose a nonnegative sequence of random variables X_1, X_2, \dots satisfies $X_t \leq X_{t-1} + L$ surely for some fixed $L \geq 0$ and any t , as well as the moment condition $\sup_t \mathbb{E}[X_t^p] \leq C_p$ for some constant $C_p \geq 0$ for each $p \geq 1$. Then, for any $c > 0$, almost surely, $X_t = o(t^c)$.*

Proof. Fix $\epsilon > 0$. It suffices to prove the lemma for $0 < c < 1$, so take $0 < d < c$ and set $p = d^{-1}$. We do this by proving the desired asymptotics on a conveniently chosen subsequence, then interpolate to intermediate values. Indeed, by Markov's inequality, for each $k \geq 1$

$$\Pr(X_{k^{1+\epsilon}} > k^{(1+\epsilon)d}) = \Pr(X_{k^{1+\epsilon}}^p > k^{1+\epsilon}) \leq \frac{C_p}{k^{1+\epsilon}}.$$

Summing over k and observing the right side is summable, we deduce from the first Borel-Cantelli Lemma that almost surely, for all sufficiently large k , $X_{k^{1+\epsilon}} \leq k^{(1+\epsilon)d}$. To extend this to all large enough t , suppose that t is such that $k^{1+\epsilon} \leq t <$

$(k + 1)^{1+\epsilon}$. By the one-sided boundedness, we know that almost surely, for such t and all large enough k ,

$$X_t \leq L \cdot (t - k^{1+\epsilon}) + X_{k^{1+\epsilon}} \leq L \cdot (1 + \epsilon)(k + 1)^\epsilon + k^{(1+\epsilon)d} \leq L \cdot (1 + \epsilon)(t^{1/(1+\epsilon)} + 1)^\epsilon + t^d,$$

where the bound on $t - k^{1+\epsilon}$ arises from the Mean Value Theorem. Clearly this last expression is $O(t^{\epsilon/(1+\epsilon)} + t^d)$. As this holds for arbitrary $\epsilon > 0$, we may take ϵ small enough so that this expression is $o(t^\epsilon)$, as claimed. \square

2.4 (Spectral) Graph Theory

In Part II, we will consider simple, undirected, weighted graphs $G = (V, E, w)$, where $|V| = n$ and $w : E \rightarrow \mathbb{R}_{>0}$. We will usually write $m = \sum_{(i,j) \in E} w(i, j)$ for the sum of the weights of all edges in G ; in the unweighted case, this is just the total number of edges. One can equivalently think of G as being a complete graph, with $w(i, j) = 0$ if and only if $(i, j) \notin E$. We will usually require G to be connected.

The adjacency matrix $A \in \mathbb{R}^{n \times n}$ is defined by $A_{i,j} = A_{j,i} = w(i, j)$. Let D be the diagonal degree matrix given by $D_{i,i} = \sum_{j:(i,j) \in E} w(i, j)$ and 0 off the diagonal. Then the Laplacian matrix of G is given by $L := D - A$. It is well-known that $L = \sum_{(i,j) \in E} w(i, j)(\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^T$, where $\mathbf{e}_i \in \mathbb{R}^n$ is the i th standard basis vector, and that the quadratic form induced by L is given by

$$\mathbf{x}^T L \mathbf{x} = \sum_{(i,j) \in E} w(i, j)(\mathbf{x}(i) - \mathbf{x}(j))^2.$$

It is immediate that L is symmetric and positive semidefinite. In general, I will denote the identity matrix of appropriate dimension.

We use standard notation for the Loewner (positive semidefinite) order, i.e. $M_1 \geq M_2$ if and only if $M_1 - M_2 \geq 0$ if and only if $M_1 - M_2$ is positive semidefinite.

For any connected graph G as above, it is well-known that $L \geq 0$ and that

$$L = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^T,$$

where $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$, the \mathbf{v}_i are orthonormal eigenvectors, and $\mathbf{v}_1 = \mathbf{1} / \sqrt{n}$, where $\mathbf{1}$ is the all-ones vector in \mathbb{R}^n . We will write V_i for the set of vectors of unit length in the λ_i -eigenspace of L . Observe that if $\lambda_i \neq \lambda_j$ for all $i \neq j$, then $V_i = \{\pm \mathbf{v}_i\}$. For more on the spectral theory of graphs, see for instance [80, 76].

Given a symmetric matrix $X \in \mathbb{R}^{n \times n}$ with eigendecomposition as above (though with not necessarily nonnegative eigenvalues) and function $f : \mathbb{R} \rightarrow \mathbb{R}$, we define

$$f(X) = \sum_{i=1}^n f(\lambda_i) \mathbf{v}_i \mathbf{v}_i^T.$$

Note that if we stipulate that $f(y) \geq 0$ for $y \geq 0$, then if $X \geq 0$, $f(X) \geq 0$. For $X \geq 0$, we write $\|X\|$ for the operator norm, or equivalently, the largest eigenvalue.

Part I

Learning and Welfare in Dynamic Environments

CHAPTER 3

GAMES, LEARNING, AND PRICE OF ANARCHY

Here, we provide a very brief introduction to games, learning, and the price of anarchy. This chapter is meant to explain introductory themes and concepts in this area rather than serve as a comprehensive reference. The results in the remainder of Part I will also be self-contained, so the definitions and results provided in this chapter are not strictly necessary for understanding the remainder of this thesis.

Informally, game theory provides a principled and convenient mathematical formalism with which to reason about the objectives and behaviors of strategic agents. Under this formulation, a game is parametrized by the set of agents, each of whom has an action set, and utility functions that specify the mapping from a tuple of actions by each agent to a numerical value for each agent. The aim of each agent is to choose actions that maximize their own utility; the complicating factor is that their utility function will of course depend on the actions of the other agents. The objectives of the agents can be well-aligned or completely adversarial, leading to an extremely rich spectrum of outcomes that can arise.

A classical prediction from game theory is that when strategic agents play a game, they reach a *Nash equilibrium*, a solution concept that requires that no agent can benefit from deviating from their current strategy given the equilibrium behavior of the others. A seminal result in the theory of games is that Nash equilibria always exist in finite games, at least when allowing randomization over actions. These existential results are typically obtained as applications of powerful fixed point theorems.

However, as we will see, the concept of equilibrium suffers from a variety of shortcomings that sharply undermine its predictive power. A more plausible model of agent behavior instead considers *learning in repeated games*. This paradigm asserts that rather than viewing a game in isolation, we instead imagine the same set of agents repeatedly playing the *same game*. Over the course of this play, agents can at least partially observe their utilities and actions of the others, informing how they ought play in the future rounds.

How ought agents decide how to leverage this information of past play? Instead of equilibrium, an attractive model of agent behavior asserts that agents use efficient learning algorithms to choose actions. These learning algorithms take as input the past observations of the agent in these repeated games and output a distribution over actions to be played in the next round. It turns out that efficient and simple learning algorithms can satisfy powerful theoretical guarantees, called the *no-regret property*, that assert the actions taken by the learner are competitive against the *single best action* the agent could have taken, in hindsight. In particular, no-regret dynamics actually converge, in a suitable sense, to a generalization of Nash equilibrium, somewhat redeeming the predictive power of the concept.

Regardless of the precise strategic behavior under consideration, an equally important question is, how *good* are these strategic outcomes? To make sense of the question, one must provide a quantitative notion of “good” by specifying a notion of *welfare*. Quantifying the welfare loss of strategic outcomes compared to socially optimal outcomes is known as the *price of anarchy* [99], and has been well-studied in a long line of work in many settings of importance (like auctions, routing, scheduling, and many others). A deep theory has emerged of bounding

the price of anarchy of both equilibrium and no-regret dynamics in a variety of important settings.

However, despite this successful theory, a crucial feature of this model is that the agents repeatedly play the very *same game*. As we will discuss, a key challenge is determining to what extent these kinds of welfare guarantees can generalize to the setting where the repeated game can *itself evolve* in response to the past sequence of play. Investigating this question for natural learning algorithms and important strategic environments with state will occupy the remainder of Part I.

3.1 Games and Equilibria

We now formally define the notion of games. We stick to the simplest possible formulation to avoid technical details.

Definition 3.1. A (normal-form) game $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ is specified by:

1. a finite set I of players,
2. a set of actions \mathcal{A}_i for each agent $i \in I$, and
3. utility functions $u_i : \prod_{i \in I} \mathcal{A}_i \rightarrow \mathbb{R}$ mapping tuples of actions to real-valued utilities for each agent i .

We write $\mathcal{A} = \prod_{i \in I} \mathcal{A}_i$ to denote the set of tuples of actions.

Remark 3.1. When considering a given agent $i \in I$, we will often abuse notation and write $u_i(a_i, a_{-i})$ for the utility she obtains when she chooses actions a_i and the remaining agents choose their corresponding action in the vector a_{-i} . This will be clear from con-

text. We also write $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$ to denote the set of tuples of actions of agents other than i , so that $a_{-i} \in \mathcal{A}_{-i}$.

The interpretation is that each player $i \in I$ selects an action a_i , and then each agent $i \in I$ receives utility $u_i(\mathbf{a})$ depending on the joint tuple of actions played by each player. One can consider an analogous version defined with costs c_i instead of utilities in the natural way.

Example 3.2 (Rock-Paper-Scissors). *For a simple example of a game, consider the standard formulation of rock-paper-scissors. In this case, we can specify the game as follows:*

1. $I = [2]$, as it is a two-player game,
2. $\mathcal{A}_1 = \mathcal{A}_2 = \{\text{rock}, \text{paper}, \text{scissors}\}$,
3. the following utility functions:

$$(a) \ u_1(\text{rock}, \text{rock}) = -u_2(\text{rock}, \text{rock}) = 0,$$

$$(b) \ u_1(\text{rock}, \text{paper}) = -u_2(\text{rock}, \text{paper}) = -1,$$

$$(c) \ u_1(\text{rock}, \text{scissors}) = -u_2(\text{rock}, \text{scissors}) = 1,$$

$$(d) \ u_1(\text{paper}, \text{paper}) = -u_2(\text{paper}, \text{paper}) = 0,$$

$$(e) \ u_1(\text{paper}, \text{scissors}) = -u_2(\text{paper}, \text{scissors}) = -1,$$

$$(f) \ u_1(\text{scissors}, \text{scissors}) = -u_2(\text{scissors}, \text{scissors}) = 0,$$

The remaining payoffs can be obtained by setting $u_1(x, y) = u_2(y, x)$ by symmetry.

With this notion in hand, how ought agents choose actions? We will assume that agents seek to *maximize* their own utility (respectively, *minimize* their costs). The idea behind Nash equilibrium is based on the following idea:

a strategic agent should *best respond* to the actions chosen by the other agents. In particular, if *all* agents are strategic, the all agents *must* be best responding to the actions of the other agents. In particular, the agents must be at equilibrium, because none of them have an incentive to change their action if they are best responding. We thus arrive at the following definition:

Definition 3.2. A **pure Nash equilibrium** of a normal-form game $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ is an action profile $\mathbf{a} \in \mathcal{A}$ such that for all $i \in I$, it holds that

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \quad \forall a'_i \in \mathcal{A}_i.$$

Here, we write $\Delta(X)$ to denote the set of probability distributions over a set X ; in most applications, we take X to be a finite set to avoid measure-theoretic difficulties.

At present, there is one immediate issue with this definition: pure Nash equilibria need not exist! Indeed, in the example of rock-paper-scissors given above, this is quite easy to see: at any action profile \mathbf{a} , at least one of the agents must have an incentive to deviate to the action that dominates the action of the other in this configuration. However, if we allow *randomization* over actions, it is not difficult to see that *randomized strategy* that uniformly mixes over all actions is actually stable in this sense. This motivates the following generalization of pure Nash equilibria:

Definition 3.3. A **mixed Nash equilibrium** of a normal-form game $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ is a sequence of *distributions* $\sigma_i \in \Delta(\mathcal{A}_i)$ for each $i \in I$, such that for all $i \in I$, it holds that

$$\mathbb{E}_{a_i \sim \sigma_i, a_{-i} \sim \sigma_{-i}}[u_i(a_i, a_{-i})] \geq \mathbb{E}_{a_{-i} \sim \sigma_{-i}}[u_i(a'_i, a_{-i})] \quad \forall a'_i \in \mathcal{A}_i.$$

Here, σ_{-i} denotes the product distribution $\prod_{j \neq i} \sigma_j \in \prod_{j \neq i} \Delta(\mathcal{A}_j)$.

Note that a mixed Nash equilibrium is a *product* distribution over actions; each agent samples an action from their marginal distribution independently of the other agents.

It turns out that under quite general conditions, normal-form games will admit mixed equilibria. The classical approach to obtaining such results is to appeal to fixed point theorems, like the Brouwer fixed point theorem or the Kakutani fixed point theorem. We record here a slightly more powerful generalization of these results that establishes the existence of equilibria in a large class of normal-form games:

Theorem 3.3 (Debreu-Glicksberg-Fan, Theorem 1.2 of [66]). *Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ be a normal-form game with a finite set of agents, with convex and compact action spaces $\mathcal{A}_1, \dots, \mathcal{A}_n$. Suppose that the utility functions u_1, \dots, u_n are such that $u_i(\cdot, a_{-i})$ is continuous and quasiconcave on \mathcal{A}_i for each fixed $a_{-i} \in \prod_{j \neq i} \mathcal{A}_j$ and $i \in I$. Then, there exists a pure Nash equilibrium \mathbf{a}^* of \mathcal{G} .*

Corollary 3.4. *Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ be a normal-form game with finite action spaces. Then there exists a mixed Nash equilibrium of \mathcal{G} .*

Proof. Consider the extended game \mathcal{G}^* where the finite action spaces are replaced by the space of distributions over actions for each agent, with the utilities are extended by taking expectations over the sampling of actions from distributions. These utilities are affine in the randomizations by linearity of expectation, hence quasiconcave and continuous. Then Theorem 3.3 implies \mathcal{G}^* has a pure Nash equilibrium in the space of distributions, which exactly corresponds to a mixed Nash equilibrium of \mathcal{G} . □

3.2 Learning in Repeated Games: The No-Regret Property

In the previous section, we introduced the notion of game-theoretic equilibrium and saw, as a result of deep fixed point theorems, that such an equilibrium always exists. However, as we now argue, there are severe issues with using Nash equilibrium as a normative model of strategic behavior:

1. **(Multiplicity)** First, Nash equilibria need not be unique.¹ In the case that there are multiple equilibria, which one do agents arrive at? Which equilibrium is more likely to be selected by real-world agents, if any? Even if there is a satisfactory answer to this latter question, it would seemingly require understanding the analytical properties of *off-equilibrium* strategies.
2. **(Rationality and Information)** The formalism we provided above for game-theoretic behavior imposes an extreme amount of rationality on the agents; indeed, this was arguably the point of introducing the framework. But is it reasonable to expect that real-world agents are quite this analytical? Moreover, real-world games are significantly more dynamic and uncertain than provided in the above formulation, making it unrealistic to expect agents to converge to stable behavior. Agents may not fully know the various parameters of the underlying strategic environment, like the identities of the other agents or their payoff structure. While one can generalize the above model to account for information constraints, doing so only increases the stringent rationality demands on the agents.
3. **(Computational Resources)** Related to the above, real-world agents are subject to real-world computational constraints. In particular, their com-

¹An amusing result of Wilson [130] shows that almost every finite game actually has an odd number of equilibria.

putations towards choosing actions in a game ought be efficiently implementable to be a satisfactory prediction—if even computing an equilibrium is infeasible before the heat death of the universe, how can this be a plausible model of strategic behavior? A striking result of Daskalakis, Goldberg, and Papadimitriou [45] shows that indeed, Nash equilibria are computationally hard to even approximate: more formally, the problem of computing a constant approximation to Nash is PPAD-complete, which rules out efficient algorithms under widely-believed complexity-theoretic assumptions. These results have since been extended to even hold in smoothed, average-case settings [26].

If even computing Nash equilibria seems like a challenge in general games, how ought we model strategic behavior? To rectify these issues, an appealing alternative is the setting of *learning in repeated games* introduced in the economics and computer science communities. In the model of repeated games, agents *repeatedly* play the same, *fixed* game \mathcal{G} for T rounds. While agents may initially not know how to play the game, agents can use the information obtained over the sequence of repeated play to *learn and adapt* to the actions of the other agents. The basic version of this model is as follows:

Algorithm 1: Model of Repeated Games

Input: Normal-form game $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$

- 1 **for** round $t = 1, \dots, T$ **do**
 - 2 Each agent i chooses action $a_i^t \in \mathcal{A}_i$ based on past observations.
 - 3 Observe information about $u_i(\mathbf{a}^t)$ and \mathbf{a}^t .
-

In the above, we have left undetermined two crucial features, one in each step:

1. How do agents choose actions? What properties should these choices satisfy?
2. What information do agents receive? Can they observe the entire action vector \mathbf{a}^t or just their own utility, or some intermediate amount of information?

For the latter point, we will only make the minimal assumption that agents can observe their own utility $u_i(\mathbf{a}^t)$. Without information of this form, the agent cannot even evaluate the quality of their chosen actions, making strategic adaptation somewhat difficult. Of course, should the agent also observe the stronger information \mathbf{a}^t , she could potentially implement more sophisticated strategies.

For the former consideration, a natural proposal is that each agent uses an efficient learning algorithm that takes as input the past observations of the agent and outputs distributions over actions at each step. The agent then chooses her action a_i^t by sampling from the distribution she receives from her learning algorithm, observes information about her outcome, and updates her learning algorithm for the next round using these observations. A popular proposal is that each agent uses a *no-regret learning algorithm*. For our purposes, the following definition is sufficient²:

Definition 3.4. An online learning algorithm ALG satisfies the **no-regret** property if for any $\varepsilon > 0$, there exists a $T(\varepsilon)$ such that for any (possibly adaptive and adversarial) sequence of actions³ $a_{-i}^1, \dots, a_{-i}^T$ where $T \geq T(\varepsilon)$, the distributions

²It is possible to achieve algorithms that achieve low regret with high probability, not just in expectation (see for instance [11]). While we will need these kinds of guarantees later on, we defer a precise statement until necessary.

³By this, we mean that the sequence can be chosen arbitrarily and can depend on the previous (random) actions of the agent randomizing using the outputs of ALG. This assumption is required to apply learning algorithms to the game setting, where the other agents also adapt to the choices made by other agents.

$\sigma_1^i, \dots, \sigma_T^i$ output by ALG satisfy

$$\max_{a_i \in \mathcal{A}_i} \frac{1}{T} \sum_{t=1}^T u_i(a_i, a_{-i}^t) - \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{a_{-i}^t \sim \sigma_{-i}^t} [u_i(a_i, a_{-i}^t)] \leq \varepsilon.$$

The left-hand side of the above equation is known as the **regret** of the learning algorithm ALG.

Equivalently, it holds that⁴

$$\max_{a_i \in \mathcal{A}_i} \frac{1}{T} \sum_{t=1}^T u_i(a_i, a_{-i}^t) - \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{a_{-i}^t \sim \sigma_{-i}^t} [u_i(a_i, a_{-i}^t)] = o(1).$$

In words, the guarantee of a no-regret learning algorithm is that the time-averaged expected reward obtained by sampling actions according to the distributions output by the algorithm is comparable to the *single best fixed action*, in hindsight.

From a practical standpoint, no-regret algorithms are simple (particularly in the case of full information over opponent actions, though this is not necessary), computationally efficient, and encode natural and minimal behavioral assumptions of strategic agents. Indeed, the promise of a no-regret algorithm is simply that the algorithm should notice if a single fixed action would have been good to play in hindsight, an observation which has some empirical support [109]. This guarantee can be ensured by running any of a large set of learning algorithms (see [121, 11, 97, 10], among many others). The study of learning in games has a long history, dating back to the early work of Brown [31] and Robinson [112], see also [65]. Moreover, no-regret learning also has an intrinsic game-theoretic interpretation: if all players employ a no-regret learning strategy, then the play converges to a form of **correlated equilibrium** of the game [84]:

⁴Formally, this equivalence does not quite hold if the ALG takes $\varepsilon > 0$ as an input; however, it turns out that providing it is not necessary to get the same guarantees.

Definition 3.5. Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$. Then a joint probability distribution $\sigma \in \Delta(\mathcal{A})$ is a **coarse correlated equilibrium** if for each $i \in I$,

$$\mathbb{E}_{\mathbf{a} \sim \sigma}[u_i(\mathbf{a})] \geq \mathbb{E}_{\mathbf{a}_{-i} \sim \sigma_{-i}}[u_i(a_i, \mathbf{a}_{-i})] \quad \forall a_i \in \mathcal{A}_i.$$

The crucial feature of a coarse correlated equilibrium that begets tractability is that the distribution σ is not necessarily required to be a product distribution, unlike a mixed Nash equilibrium. With this definition, the convergence of no-regret learning can be stated as follows:

Theorem 3.5 ([83]). *Suppose that all agents employ no-regret learning algorithms, each with regret at most ε after T iterations. Let $\sigma_t = \prod_{i \in I} \sigma_t^i$ denote the product of the distributions output by each agent's algorithm at time t , and define $\sigma = \frac{1}{T} \sum_{t=1}^T \sigma_t$ denote the time average of the product distributions. Then σ forms a ε -approximate coarse correlated equilibrium.*

In other words, *any* no-regret dynamics necessarily and naturally converge (in an appropriate sense) to a version of equilibrium, circumventing the equilibrium selection problem of Nash.

To summarize, while we saw there are a variety of difficulties with using Nash equilibrium as a prediction of agent behavior, viewing strategic agents as computationally efficient learners gives a significantly more tractable and flexible model that serves as a reasonable approximation to real-world, algorithmic agents. Moreover, strong individual guarantees, via the no-regret property, are simple to attain under minimal information and rationality assumptions, and no-regret dynamics themselves have favorable game-theoretic properties via the convergence to a meaningful generalization of Nash equilibrium.

3.3 Price of Anarchy

The abstractions of the previous sections have now provided a plethora of strategic behaviors that capture many aspects of how selfish and rational agents ought behave. In particular, the model of learning in repeated games naturally aligns with the prevalence of algorithmic decision-making.

However, the main question we will be interested in answering in Part I is not how do strategic agents act. Rather, we will take these game-theoretic models as primitives in our analysis. Our main question will be determining how *good* are these strategic outcomes compared to optimal, coordinated outcomes?

To even make sense of this question, one must first specify a metric of social welfare:

Definition 3.6. Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ be a game. A **social welfare function** is a map $W : \mathcal{A} \rightarrow \mathbb{R}_+$ from actions to overall welfare.

A canonical example of a social welfare function is simply the *aggregate utility* of the agents:

$$W(\mathbf{a}) := \sum_{i \in I} u_i(\mathbf{a}).$$

In this case, the utilities of the agents directly contribute to the social welfare. However, in the remainder of Part I, we will consider substantially more general notions of welfare.

The relationship between optimal actions and strategic actions with respect to welfare is classically called the *price of anarchy*, a term introduced by Koutsoupias and Papadimitriou [99]. In Part I, we will actually use “price of anarchy” to describe more general quantitative relationships between strategic behavior

and optimal outcomes. However, the classical definition introduced by Koutsoupias and Papadimitriou [99] more specifically considers the *worst-case ratio* between the optimal social welfare and Nash welfare:

Definition 3.7 (Price of Anarchy, Nash Version). Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ be a finite game and let $W : \prod_{i \in I} \mathcal{A}_i \rightarrow \mathbb{R}_+$ be a social welfare function. Let $\mathcal{N}_x(\mathcal{G})$ denote the set of Nash equilibria of \mathcal{G} , where $x \in \{\text{PNE}, \text{MNE}, \text{CCE}\}$ denotes pure Nash, mixed Nash, or coarse correlated equilibria.⁵ Then the **price of anarchy** of \mathcal{G} with respect to Nash equilibria is

$$\max_{\mathbf{a}^* \in \mathcal{A}, \sigma \in \mathcal{N}_x(\mathcal{G})} \frac{W(\mathbf{a}^*)}{\mathbb{E}_{\mathbf{a} \sim \sigma}[W(\mathbf{a})]} = \frac{\max_{\mathbf{a}^* \in \mathcal{A}} W(\mathbf{a}^*)}{\min_{\sigma \in \mathcal{N}_x(\mathcal{G})} \mathbb{E}_{\mathbf{a} \sim \sigma}[W(\mathbf{a})]}$$

Here, the numerator is the optimal social welfare that can be achieved by any action profile, while the denominator gives the worst-case welfare that can be attained a Nash equilibrium (within the class of equilibria under consideration).⁶ This ratio is necessarily at least 1, and thus a bound on the price of anarchy of a game gives a *lower bound* on the welfare of *every* Nash equilibrium. In particular, a price of anarchy bound inherently avoids issues of equilibrium selection: it is agnostic to which exact equilibrium agents reach.

Because $\mathcal{N}_{\text{PNE}}(\mathcal{G}) \subseteq \mathcal{N}_{\text{MNE}}(\mathcal{G}) \subseteq \mathcal{N}_{\text{CCE}}(\mathcal{G})$, note that proving price of anarchy bounds for coarse correlated equilibria yields a stronger guarantee. In fact, by the convergence of no-regret dynamics, it can be shown that a price of anarchy bound that holds for coarse correlated equilibrium also yields the same bound for no-regret dynamics up to an additive $o(1)$ term that tends to zero with the time horizon T . As we will see shortly, a very general technique, called

⁵Recall that the latter two sets are nonempty.

⁶Note the minimum is actually achieved by compactness of the sets under consideration, at least in finite games.

“smoothness”, for proving price of anarchy bounds is *robust*: the framework applies to coarse correlated equilibrium just as well as it does to even pure Nash equilibrium. In particular, for most games, price of anarchy bounds obtained using arguments for pure Nash equilibria (which may not even exist) typically extend to even no-regret learning dynamics.

The philosophy of price of anarchy bounds is that in many cases, even though agents behave strategically and selfishly, their objectives are sufficiently well-aligned with the social welfare function to obtain nontrivial performance guarantees. By proving a price of anarchy bound for a game, one obtains a theoretical guarantee that the efficiency of the system (at least with respect to the welfare metric) will not substantially degrade even in the presence of selfish actors. Thus, proving price of anarchy results give actionable insights into which systems can be safely decentralized and remain robust to strategic behaviors, and conversely, which systems may require additional resources, centralization, or regulation.

We now turn to a brief discussion on a general framework for proving price of anarchy bounds. A by now standard technique for price of anarchy bounds, which unifies several existing analyses in the literature, is the “smoothness” framework of Roughgarden [114]. To set up the result, we require the following definition:

Definition 3.8. Let $\mathcal{G} = (I, \{\mathcal{A}_i\}_{i \in I}, \{u_i\}_{i \in I})$ be a normal-form game. Let $W : \prod_{i \in I} \mathcal{A}_i \rightarrow \mathbb{R}_+$ be a welfare function satisfying $W(\mathbf{a}) \geq \sum_{i \in I} u_i(\mathbf{a})$ for all $\mathbf{a} \in \prod_{i \in I} \mathcal{A}_i$. Then \mathcal{G} is (λ, μ) -smooth with respect to W if it holds that

$$\sum_{i \in I} u_i(a_i^*, a_{-i}) \geq \lambda W(\mathbf{a}^*) - \mu W(\mathbf{a}) \quad (3.1)$$

for some $\mathbf{a}^* \in \arg \max_{\mathbf{z} \in \prod_{i \in I} \mathcal{A}_i} W(\mathbf{z})$ and all $\mathbf{a} \in \prod_{i \in I} \mathcal{A}_i$.

Note that the definition of smoothness depends only on the underlying game \mathcal{G} , and moreover, only considers unilateral deviations by the agents from any strategy profile to a different, socially optimal profile. However, with this definition in hand, Roughgarden's smoothness framework yields the following bound of the price of anarchy with respect to even coarse correlated equilibria of any smooth game, which can be shown to extend to no-regret dynamics.

Theorem 3.6 (Theorem 3.3 of [114]). *Let \mathcal{G} be (λ, μ) -smooth with respect to a welfare function W . Then the price of anarchy with respect to $\mathcal{N}_{\text{CCE}}(\mathcal{G})$ is at most $\frac{\lambda}{1+\mu}$.*

As a corollary, the (time-averaged) price of anarchy of no-regret learning dynamics on \mathcal{G} with respect to W is at most $\frac{\lambda}{1+\mu} + o(1)$.

Proof. The proof is immediate: for σ a coarse correlated equilibrium, we have

$$\mathbb{E}_{\mathbf{a} \sim \sigma}[W(\mathbf{a})] \geq \sum_{i \in I} \mathbb{E}_{\mathbf{a} \sim \sigma}[u_i(\mathbf{a})] \geq \sum_{i \in I} \mathbb{E}_{\mathbf{a}_{-i} \sim \sigma_{-i}}[u_i(a_i^*, \mathbf{a}_{-i})] \geq \lambda W(\mathbf{a}^*) - \mu \mathbb{E}_{\mathbf{a} \sim \sigma}[W(\mathbf{a})],$$

where we use the fact $W \geq \sum u_i$, the equilibrium condition, and then smoothness, as well as linearity of expectation. Rearranging gives the claim, and Theorem 3.5 gives the no-regret guarantee. \square

The smoothness framework thus provides an extremely convenient recipe for proving price of anarchy bounds for learning: instead of needing to carefully reason about the precise evolution of the dynamics, one instead need only verify a *static* inequality as given by Equation (3.1). In particular, analyzing dynamics reduces to bounding the effect of the discrete deviations considered in Equation (3.1), which in turn relies only on the structural properties of the game rather than the precise learning mechanisms of the agents. Learning dynamics can thus be viewed as a black-box guarantee in the analysis and otherwise abstracted away.

3.4 Challenges: Price of Anarchy in Dynamic Systems

The results described above yield a rather deep and fruitful theory of games, learning, and welfare that has enjoyed spectacular successes over the last few decades. However, we now discuss two vital inadequacies that limit the predictive power of this theory:

Broader Notions of Welfare. First, standard notions of social welfare that are most amenable to this theory typically hold in aggregate over the agents. For instance, the canonical example of social welfare described above simply took the sum of the utilities. This notion of welfare, however, cannot immediately say anything concrete about *any given agent's welfare*—while one could use a Markov argument to assert that a typical agent must have high welfare if a price of anarchy bound holds in this setting, these guarantees fundamentally cannot address these stronger conditions. For instance, in many real-world settings where fairness of outcomes is a vital concern, a more appropriate notion of social welfare may instead consider the *minimum* of the agent's individual welfares.

Relatedly, important notions of welfare may not additively decompose across time, as is often necessary to employ the smoothness framework or related variants. For instance, Chapter 4, our notion of welfare in queuing systems will be posed as *stability* of the system. In this case, social welfare will be a binary variable; the underlying system is either stable or it is not, and understanding which is the case requires finer-grained analysis than simply viewing each round of the system in isolation. Our notion of stability will also require strong performance guarantees for *each* player in the system—it will not be enough that most players perform well. Neither of these features is well-suited to existing

frameworks in the price of anarchy theory as developed above.

Evolving Games. In the above model of learning in repeated games, a pivotal feature is that agents must repeat the *same fixed game*. That is, while the agents may themselves learn and adapt based on the past sequence of play, the underlying game \mathcal{G} must remain fixed. However, this assumption often will not hold in practice: in many systems of important, the underlying strategic environment itself may evolve in direct response to the past history of play. In such cases, this theory of learning in repeated games is rendered incapable of giving any answers on the quality of strategic outcomes.

To make this more concrete, consider the setting of routing in networks (as considered in the work of Roughgarden and Tardos [117]). Formulating this as a game, each agent chooses a path in the network from their source vertex to their sink vertex with the objective of minimizing their experienced delay along this path. The delay of an edge is some monotonic function of the number of agents traversing this edge in their chosen path, and the overall cost to the agent for choosing a given path is simply the sum of the delays of the edges in this path (which, we stress, depends on the number of other agents using these same edges). In the classical theory of repeated games described above, at each time, agents choose a (possibly new) path incorporating new knowledge about the actions of the competing players, but the underlying mapping from choices from paths to costs remains the same in each round.

In some settings, this approximation might be valid. On the scale of the morning rush-hour traffic, it indeed may be the case that each morning, the underlying routing game is essentially the same: it almost always holds that any traffic on Monday morning will have cleared by Tuesday morning, at which

point agents again choose paths in a fresh version of the game. But on the other hand, consider modeling packet routing in computer networks, a system that operates on incredibly fast scales. If a packet gets dropped in a round, then this packet will be re-sent in future rounds and thereby increases the total congestion going forward. In particular, the assumption that the game itself has not changed is violated.

Both of these challenges limit the applicability of the price of anarchy theory developed above. To circumvent these difficulties, we will require more fine-grained reasoning about the underlying dynamics of these repeated, strategic interactions. In the remainder of Part I, we will provide new techniques and analyses to specifically overcome these deficiencies in important settings. As we will see, even in simple games with state, these new complications lead to fundamentally new analyses that cannot be obtained via the existing theory of games, learning, and price of anarchy.

3.5 Chapter Notes

The definitions in this chapter are standard and can be found in any textbook on game theory or learning. For the reader interested in the interplay between games and learning, we recommend the excellent textbook of Cesa-Bianchi and Lugosi [35]. For a more recent monograph of a wide range of topics in algorithmic game theory, we also recommend the text of Roughgarden [115].

CHAPTER 4
STRATEGIC QUEUING SYSTEMS I: STABILITY OF NO-REGRET
LEARNING

In the next two chapters, we consider the dynamics of *strategic queuing systems*, which will serve as an illuminating case study in understanding the subtleties of learning and welfare in stateful environments. In our model of strategic queuing, queues act as selfish agents that compete for servers. The goal of each queue is to efficiently clear arriving packets by sending an uncleared packet to a server in each round. To model the inherent stochasticity of the queuing system, we will assume that a new packet may randomly arrive at each queue in each round with heterogeneous rates. Servers will have varying (possibly unknown) quality, and even more importantly, can only clear a single arriving packet in each round. We will make the assumption that each server attempts to clear only the *oldest* arriving packet it receives, breaking ties arbitrarily.¹ Every uncleared packet in a round is returned to its original queue, to be cleared in the future rounds, and each queue only receives feedback whether or not their packet succeeded in clearing at their chosen server.

A critical feature of such systems is that, no matter how queues decide which servers to choose, the underlying state of the system is highly dependent on the outcomes of previous rounds. In standard repeated games, the underlying game-theoretic environment is independent of past actions; the outcomes of previous rounds can only influence the future decisions of the agents. In the above model, however, the outcomes of previous rounds will also influence the implicit priority structure that arises due to the server preferences in attempt-

¹See below for more discussion on this assumption.

ing to clear the oldest packet it receives. This priority structure may fluctuate rapidly from round-to-round necessitating a delicate understanding of the long-term behaviors of the dynamics. Such queuing systems highlight precisely the kinds of strategic, dynamic environments that existing price of anarchy theory cannot address, and their study will require the development of several new techniques.

Our goal in the next two chapters will be to establish precise conditions, depending only on the underlying parameters of the system, under which two natural classes of strategic behaviors by the queues exhibit *stability* (formally defined below): no-regret learning and game-theoretic equilibrium. Intuitively, our notion of stability captures the notion that the number of uncleared packets does not diverge over time, at least in expectation. As discussed further below, due to the underlying randomness of these systems, this guarantee in expectation is essentially the best possible because with probability one, queue sizes can be arbitrarily large infinitely often. We will nonetheless prove a stronger form of stability that will imply sub-polynomial bounds on the growth of this quantity along each sample path (see below for precise statements).

We further remark that our notion of stability in such systems, while quite natural in this context, is qualitatively quite different from standard welfare notions in the price of anarchy literature in several ways. First, as our queuing systems have an infinite time horizon, stability requires uniform control of the state of the system at each time, while standard notions of welfare are typically measured in terms of total aggregate utility of all agents up to some large, but fixed time horizon T . Moreover, stability trivially implies that *every* queue remains bounded in expectation uniformly over time, which requires extremely

fine-grained control on the performance of *each* queue in the system. By contrast, standard notions of welfare in the price of anarchy literature are usually aggregate measures like above that typically cannot provide precise guarantees on the performance of any given agent. Beyond the difficulties induced by the randomness of these systems, these stark differences in the welfare metric will require the introduction of several new techniques to analyze the quality of these systems.

In this chapter, we focus on the first of the above assumptions on strategic behavior: no-regret dynamics. Our main result, informally stated, is as follows:

Consider any queuing system with the property that it would remain stable under optimal, centrally coordinated scheduling even if the server qualities were halved. Then, the original queuing system will remain stable even if queues strategically compete so long as their actions (i.e. choices of servers) satisfy a suitable quantitative version of the no-regret property.

Moreover, we will further show this multiplicative factor of 2 is optimal in the sense that no-regret dynamics cannot guarantee stability under any smaller number. In all, our result ensures that the price of decentralizing such queuing systems is fairly small; to safeguard against this kind of strategic behavior, one only needs to augment the resources in the system by a factor of 2.

4.1 Model

We now more formally define our model of strategic queuing and stability.

4.1.1 Queuing Model

We consider the following, discrete-time queuing system that is a decentralized version of a queuing system introduced by Krishnasamy, et al [100].

Definition 4.1 (Bernoulli Queuing System). A **Bernoulli queuing system** with n queues and m servers is defined by parameters $\lambda \in (0, 1)^n$ and $\mu \in [0, 1]^m$ satisfying $\lambda_1 \geq \dots \geq \lambda_n$ and $\mu_1 \geq \dots \geq \mu_m$. During each discrete time step $t = 0, 1, \dots$, the following occurs:

1. Each queue i receives a new packet with a fixed, time-independent probability λ_i . We model this via an independent random variable $B_t^{(i)} \sim \text{Bern}(\lambda_i)$. This packet has a timestamp that indicates that it was generated at time t .
2. Any queue with an uncleared packet selects one server and sends their oldest unprocessed packet (in terms of timestamp) to that server.
3. Each server j attempts to process the *oldest* packet it receives (breaking ties arbitrarily), succeeding in clearing this packet with a fixed, time-independent probability μ_j independent of all previous events.
4. All unprocessed packets are then sent back to their respective queues still uncompleted.
5. Each queue i only receives feedback $X_t^{(i)} \in \{0, 1\}$, where 1 (0) denotes that their packet did (not) clear in round t .

We let $\mathcal{F}_t^{(i)} = \sigma(\{(B_s^{(i)}, X_s^{(i)})_{s=0}^t\})$ to be the σ -algebra generated by the observable history of queue i up to and including time t , along with any internal randomness they use. We require that each queue i 's choice of server at time $t + 1$

be $\mathcal{F}_t^{(i)}$ -measurable (again, along with any internal randomness). Finally, let $\mathcal{F}_t = \sigma(\mathcal{F}_t^{(1)}, \dots, \mathcal{F}_t^{(n)})$ denote the observed history of the queuing system up to time t .

Note that in Definition 4.1, we have left the choice of how queues select servers unspecified. All of our definitions and results will implicitly be with respect to either some fixed dynamics or apply to a particular class of dynamics satisfying certain properties, which we will make clear in context.

Let $Q_t^{(i)}$ denote the number of unprocessed packets of queue i at the *beginning* of time t (before sampling new packets) and $\mathbf{Q}_t = (Q_t^{(1)}, \dots, Q_t^{(n)})$ denote the vector of queue sizes at time t . Define $Q_t = \sum_{i=1}^n Q_t^{(i)}$ as the total number of unprocessed packets in the system at time t . Using the notation in Definition 4.1, the number of uncleared packets at queue i satisfies the recurrence

$$Q_{t+1}^{(i)} = Q_t^{(i)} + B_t^{(i)} - X_t^{(i)}, \quad (4.1)$$

where we note that $X_t^{(i)}$ is necessarily 0 if $Q_t^{(i)} + B_t^{(i)} = 0$. By construction, note that each $Q_t^{(i)}$ is integral and nonnegative.

As mentioned earlier, our main objective will be to understand the *stability* properties of such systems under natural strategic assumptions.

Definition 4.2 (Stability). A Bernoulli queuing system under some given dynamics is **r -stable** if

$$\mathbb{E}[(Q_t)^r] \leq C(r, \boldsymbol{\mu}, \boldsymbol{\lambda}),$$

where $C(r, \boldsymbol{\mu}, \boldsymbol{\lambda})$ is an absolute constant depending only on $r, \boldsymbol{\mu}$ and $\boldsymbol{\lambda}$.² A Bernoulli queuing system is **strongly stable** if it is r -stable for every $r > 0$.

²Note that any r -stable process is trivially r' -stable for any $r' \leq r$ by Jensen's inequality.

A Bernoulli queuing system is **weakly stable** if it holds that

$$\mathbb{E}[Q_t] = o(t).$$

4.1.2 Benchmark: Centralized Feasibility

To get a baseline measure for the outcomes that could arise under strategic behavior, we must first understand when a queuing system is stable under centralized coordination. It turns out that the obvious necessary condition on μ and λ is also sufficient. For intuition, consider a one-queue, one-server system parametrized by λ and μ . Clearly no strategic behavior or learning is necessary in such a system, and the asymptotics are given by the following cases:

1. $0 < \lambda < \mu \leq 1$: in this case, the $Q_t^{(1)}$ follows a random walk on the nonnegative integers biased towards 0. This can easily be shown to imply strong stability.
2. $0 < \lambda = \mu < 1$: in this case, the essentially unbiased random walk $Q_t^{(1)}$ satisfies $\mathbb{E}[Q_t^{(1)}] = \Theta(\sqrt{t})$ so that the system is weakly stable, but not 1-stable.
3. $0 < \mu < \lambda < 1$: in this case, the random walk satisfies $Q_t^{(1)} \approx (\lambda - \mu) \cdot t$ up to sublinear terms almost surely, essentially by the law of large numbers. Simply put, packets arrive faster than they can be cleared.

For centralized queuing systems, a natural generalization of the above gives the necessary and sufficient conditions for stability:

Lemma 4.1. *A Bernoulli queuing system with parameters μ and λ is strongly stable³ under some centralized (coordinated) scheduling policy if and only if for all $1 \leq k \leq n$,*

$$\sum_{j=1}^{\min\{k,m\}} \mu_j > \sum_{i=1}^k \lambda_i. \quad (4.2)$$

We say μ **strictly majorizes** λ if (4.2) holds, and **weakly majorizes** if the inequalities in (4.2) only weakly hold.

Proof. Sufficiency: By padding with zeros, we assume $m = n$; this does not change (4.2) nor the space of centralized scheduling policies so long as we allow queues to abstain from sending a packet in a round. Suppose μ strictly majorizes λ . It follows there (see, e.g. [105]) that exists some doubly stochastic P such that $P\mu > \lambda$ element-wise. By the Birkhoff-von Neumann Theorem, there exists a distribution π over the set \mathcal{P} of permutation matrices in $\mathbb{R}^{n \times n}$ such that $P\mu > \lambda$ componentwise, where $P = \mathbb{E}_{\Pi \sim \pi}[\Pi]$.

Consider the following oblivious, centralized scheduling algorithm: at each time t , independently sample a permutation matrix Π_t from π , and schedule queues via the associated matching on the bipartite graph of queues and servers. For each queue i , the associated marginal distribution on servers it sends to in each round is given by the i th row of P . Given that queue i has a packet to send at time t , the probability of successfully clearing a packet is exactly $(P\mu)_i > \lambda_i$ because queue i is the only queue sending to this server due to the matching. It follows that the coordinate process $Q_t^{(i)}$ follows a homogeneous random walk on the half-line biased towards 0 exactly like in the one-queue, one-server case with parameters λ and $(P\mu)_i$. This process is ergodic with a stationary distribution with geometric tails. It is not difficult to show that any distribution on the

³An analogous statement for weak stability holds if the inequalities in Equation (4.2) only weakly hold by the same proof.

natural numbers with geometric tails has bounded r th moments for any $r \geq 0$.⁴ Because $(Q_t)^r \leq n^{r-1} \sum_{i=1}^n (Q_t^{(i)})^r$, this implies strong stability.

Necessity: Suppose that there is some minimal $k \leq n$ such that $\sum_{i=1}^k \lambda_i \geq \sum_{i=1}^k \mu_i$; note that $\lambda_k > 0$. Let $Q_t^{(\leq k)} = \sum_{i=1}^k Q_t^{(i)}$. It suffices to show $\mathbb{E}[Q_t^{(\leq k)}] \rightarrow \infty$. To that end, observe that

$$\mathbb{E}[Q_{t+1}^{(\leq k)} - Q_t^{(\leq k)}] = \mathbb{E}\left[\sum_{i=1}^k B_t^{(i)}\right] - \mathbb{E}\left[\sum_{i=1}^k X_t^{(i)}\right] \geq \sum_{i=1}^k \lambda_i - \sum_{i=1}^k \mu_i,$$

because the expected number of arriving packets is the sum over μ_i , while for any measurable policy $\sigma_t : [k] \rightarrow [m]$

$$\mathbb{E}\left[\sum_{i=1}^k X_t^{(i)}\right] \leq \sum_{j \in [m]: |\sigma_t^{-1}(j)| \geq 1} \mu_j \leq \sum_{j=1}^k \mu_j.$$

In particular, $Q_t^{(\leq k)}$ is a nonnegative submartingale for any measurable scheduling policy. If $\lim_{t \rightarrow \infty} \mathbb{E}[Q_t^{(\leq k)}] = \sup_t \mathbb{E}[Q_t^{(\leq k)}] < \infty$, then the Martingale Convergence Theorem (Theorem 4.2.11 of [55]) implies that there exists an almost surely finite random variable $Q_\infty^{(\leq k)}$ such that $\lim_{t \rightarrow \infty} Q_t^{(\leq k)} \rightarrow Q_\infty^{(\leq k)}$ almost surely. But $Q_{t+1}^{(\leq k)} - Q_t^{(\leq k)}$ is integer-valued and at least one with constant probability as $\lambda_i > 0$ for all $i \leq k$. This implies the pointwise limit cannot exist unless it is infinite, violating the almost sure finiteness of $Q_\infty^{(\leq k)}$. \square

This simple result justifies the following definition:

Definition 4.3. A Bernoulli queuing system is **centrally feasible** if (4.2) holds.

⁴This can also easily be seen directly using Theorem 2.14. Negative drift when exceeding $Q_t = 0$ is obvious, and as queue sizes can change by at most n in total between steps, increments are clearly bounded in L^p for any $p \geq 0$.

4.1.3 Discussion

We now briefly comment on the various features of this queuing model.

On the Need for Packet Priorities

For central stability, as established in Lemma 4.1, the precise mechanism by which servers select among arriving packets is immaterial: the scheduling policy constructed there ensures that there are never any collisions at any server. On the other hand, this mechanism clearly will affect the stability of decentralized dynamics where collisions are inevitable. In strategic environments where queues selfishly aim to clear their own packets, queues may naturally collide at the relatively effective servers. Our model above that explicitly biases older packets (that is, queues that have received insufficient service relative to their arrival rate) is intuitively more conducive towards being robust to strategic behavior; while queues will still compete for the most effective resources, the priority afforded to queues that have been harmed by this strategic behavior can facilitate their recovery.

A simpler model, however, would be to assume that servers select among arriving packets uniformly at random. It is easy to see that if a queuing system is feasible even if μ scaled down by a factor of n , then it will remain a stable queuing system with reasonably strategic queues. By central feasibility, $\mu_1 > n \cdot \lambda_1$, so that $\mu_1 > \sum_{i=1}^n \lambda_i$, so any queue that always sends to the top server will clear with probability at least $\mu_1/n > \lambda_i$. Analogous arguments to Lemma 4.1 then imply that any such queue will remain stable, regardless of the strategic behaviors of the others.

It is natural to ask if a better factor is attainable in this alternate model, perhaps even a constant. It is possible to show that in general, a polynomial in n is required:

Theorem 4.2. *Suppose servers select among arriving packets uniformly at random. Then for large enough n , there exists a centrally feasible queuing system with n queues and servers with the following properties:*

1. *the system remains feasible even if μ is scaled down by $\Omega(n^{1/3})$,*
2. *there exists a sequence of choices for every queue and times t such that their choice maximizes the probability of clearing a packet given the choices of the other queues at every time t , but the system is not weakly stable.⁵*

Proof. Let $\lambda_1 = 2/n^{1/3}$, while $\lambda_2 = \dots = \lambda_n = 1/n^{2/3}$; let $\mu_1 = 1/2$ and $\mu_2 = \dots = \mu_n = c/n^{1/3}$, where $c = c(n) = \Theta(1)$ is such that

$$\frac{1}{n^{1/3} + 2} < \frac{c}{n^{1/3}} < \frac{1}{n^{1/3}}.$$

We call the first queue and server “high-rate” and the remaining queues and servers low-rate.

Consider the following adversarial scheduling policy: in each round, the scheduler chooses $n^{1/3}/2 - 1$ of the low-rate agents arbitrarily to send to the unique high rate server, if that many low-rate agents have uncleared packets, as well as the high-rate queue. All other low-rate agents send to distinct low-rate servers. If fewer than $n^{1/3}/2 - 1$ low rate servers are active, then the scheduler

⁵This condition combined with the Azuma-Hoeffding inequality (Lemma 2.7) further implies that each queue is using a strategy satisfying the no-regret condition given by Assumption 4.1. Our main result shows that in the Bernoulli queuing system with timestamp priorities, any strategy satisfying Assumption 4.1 is enough so long as there is just a factor of 2 slack.

schedules all active queues to the high rate server. Then the following can be shown to hold:

1. For any round t , with exceedingly high probability at least $n^{1/3}$ low-rate queues receive a new packet by standard Chernoff bounds, in which case the first option in the above scheduling policy is taken.
2. On the event that this occurs, the above inequalities ensure that all queues choose servers that maximize the probability of clearing their packet at time t conditioned on the choices of the others because each server uniformly chooses among arriving packets. This remains true even on the event that this does not occur by construction.
3. By the law of total probability, at any time t such that the high-rate queue has an uncleared packet, her probability of clearing a packet in the round is at most $3/2n^{1/3}$ because the first scheduling option is taken with overwhelming probability. In particular, $\mathbb{E}[Q_{t+1}^{(1)} - Q_t^{(1)}] \geq 1/2n^{1/3}$, so $\mathbb{E}[Q_t^{(1)}] = \Theta(t)$. This means the system is not weakly stable.

Finally, note that the system would remain feasible even if μ were scaled down by a factor of $n^{1/3}/4$, as needed. □

In words, the basic reason why this instability can occur is that low arrival rate queues can saturate the high success rate servers, making it impossible for high arrival rate queues to clear fast enough to offset their higher arrival. Our priority model, while more difficult to analyze, results in older queues gaining an advantage on young queues causing the young queues to prefer lower quality servers.

On Stability

As discussed above, our notion of 1-stability exactly captures that the queuing system does not diverge in expectation with time. As explained above, it is impossible to remove the “in expectation,” as the underlying randomness in the queuing system will almost surely cause large fluctuations in the number of un-cleared packets. Nonetheless, in our main results, we will actually prove strong stability of strategic queuing dynamics. As a straightforward consequence of Lemma 2.15, this will imply that $Q_t = o_\varepsilon(t^\varepsilon)$ for every $\varepsilon > 0$, where the implicit constants are random and depend on ε . That is, our main results on strong stability will imply almost sure sub-polynomial growth of the number of un-cleared packets as a function of t . It would be quite interesting to determine the best possible asymptotic upper bound that holds almost surely (for instance, like the Law of the Iterated Logarithm for Brownian motion).

4.2 Main Results

Our first main result shows that, if the queuing system has enough slack and all queues satisfy an appropriate high-probability no-regret guarantee, then the queuing system is strongly stable. To this end, we make the following feasibility assumption:

Assumption 4.1 (Feasibility). *There exists $\eta > 0$ such that for all $k \in [n]$,*

$$\frac{1}{2}(1 - \eta) \sum_{j=1}^{\min\{k,m\}} \mu_j \geq \sum_{i=1}^k \lambda_i.$$

We will generically use η to denote the maximum such value that this inequality holds. The parameter η controls the quality of learning required for

our results.

Next, we will assume that each queue will satisfy a no-regret learning guarantee on the number of packets served during each window of w time length. To formalize this assumption, we need a few definitions.

Definition 4.4. Fix an interval $I = [\ell, u]$. Let $X_t^{(i),j}$ be the indicator variable that queue i would have succeeded in clearing a packet at server j at time t had she sent there⁶, and let $\sigma_i(t)$ be the identity of the server that queue i chooses at time t .⁷ Then the **regret** of queue i on this window, denoted $\text{Reg}_i(I)$, is defined as

$$\text{Reg}_i(I) \triangleq \max_{j \in [m]} \sum_{t=\ell}^u X_t^{(i),j} - \sum_{t=\ell}^u X_t^{(i),\sigma_i(t)} = \max_{j \in [m]} \sum_{t=\ell}^u X_t^{(i),j} - \sum_{t=\ell}^u X_t^{(i)}. \quad (4.3)$$

Note that all these random variables are with respect to the same sample path; the $X_t^{i,j}$ will depend on all previous randomizations and choices by the queues, as these implicitly yield the priorities of the queues. In words, $\text{Reg}_i([\ell, u])$ of queue i is defined to be the (random) difference between the number of packets queue i cleared on this interval compared to the backward-looking number of packets she would have cleared had she simply always sent to the best single server for her on this interval; that is, for each time t and server j we have $X_t^{i,j} = 1$ if at time t server j was successful (regardless of if a packet was sent there then), and the packet that queue i sent had priority over any packet sent there at that time.⁸

An important definition that will be crucial is the following:

⁶Note that this is well-defined as a function of whether server j would be successful, queue i 's packet timestamp, and any given tiebreaking rule of server j .

⁷Note that if queue i has no packets at time t , then $X_t^{i,j} = 0$ for all j .

⁸Note that this notion of regret does *not* take into account that had a queue i cleared a packet at at server j instead of another queue i' , at a later time i' would have had an older packet and therefore higher priority.

Definition 4.5. For any Bernoulli queuing system, the **age** of queue i at time t is the difference between t and the timestamp of queue i 's oldest uncleared packet at the beginning of time t . By convention, this is zero if the queue has no uncleared packet.

The age of a queue at time t simply measures how long their oldest uncleared packet has remained in the system. Under our queuing dynamics, the ages of the queues at each time are precisely what determines the priority structure in the given round.

We now make the following assumption on the regret of queuing strategies:

Assumption 4.2 (High-Probability No-Regret of Queues). *All queues select servers using a strategy satisfying the following **no-regret** guarantee: given fixed $\delta \in (0, 1)$ and a window length w , for any interval $I = [\ell, u]$ with $u - \ell = w$, it holds that*

$$\text{Reg}_i(I) \leq \varphi_\delta(w)$$

with probability at least $1 - \delta$ conditioned on all events before time ℓ . Here, the probability is taken only over their own randomization during this window and $\varphi_\delta(w) = o(w)$ is an explicit function where the constant factors depend on δ and m , but not w .

Moreover, we require that the choices of the queue depend only on their feedback and their past history of ages, but not on their history of queue sizes.

For instance, this assumption holds with EXP3.P.1. In this case, the regret scales like $\sqrt{wm \ln(mw/\delta)} = o(w)$ [11]. Note that, as with all standard applications of learning in multi-player games this high-probability guarantee is possible in our setting even in the priority model where the random variables of success at each server from the perspective of each queue at each time step depend on

all previous actions (via timestamps and priorities), as well as the actions of the other queues in the current time period; see for instance the discussion in Section 9 of Auer, et al [11]. Using EXP3.P.1 ensures that such a guarantee holds simultaneously for each window of this length, and not only a fixed window, so the players would not have to be aware which window of size w is relevant for our analysis. This is true as EXP3.P.1 mixes in uniform exploration to guarantee that the probabilities remain high enough throughout the algorithm, allowing us to adapt the classical no-regret analysis starting at any time step for the window of the next w time steps. In our analysis below, we actually only require that Assumption 4.2 holds for a fixed sequence of *consecutive* windows of length w , and this guarantee is achieved if each queue simply restarts their algorithm at the beginning of each window.

With these definitions in order, we may formally state our main result on stability of no-regret learners:

Theorem 4.3. *Suppose that Assumption 4.1 holds with parameter η and suppose all queues use strategies satisfying Assumption 4.2 with a regret upper bound of $\phi_\delta(w)$ given δ and w .*

Then, there exists $w = w(\eta, \mu, \lambda, n)$ large enough such that if each queue satisfies the guarantee of Assumption 4.2 with parameter w and failure probability $\gamma = \frac{\eta}{256n}$, then the queuing system is strongly stable under these dynamics.

4.3 Overview of Proof

Before we proceed with the formal proof, we provide a high-level overview of the argument and techniques. The main probabilistic tool that will be used to

establish strong stability is the powerful result of Pemantle and Rosenthal, as stated in Theorem 2.14. For the reader's convenience, we restate their result below:

Theorem 4.4 (Theorem 1 in [111]). *Let X_1, X_2, \dots be a sequence of nonnegative random variables with the property that*

1. *There exists constants $\alpha, \beta > 0$ such that if $x_t > \beta$, then*

$$\mathbb{E}[X_{t+1} - X_t | \mathcal{F}_t] < -\alpha,$$

where the σ -algebra $\mathcal{F}_t = \sigma(X_0, \dots, X_t)$ is the history until period t and $X_t = x_t$.

2. *There exists $p > 2$ and $\theta > 0$ a constant such that for any history,*

$$\mathbb{E}[|X_{t+1} - X_t|^p | \mathcal{F}_t] \leq \theta.$$

Then, for any $0 < r < p - 1$, there exists an absolute constant $M = M(\alpha, \beta, \theta, p, r)$ not depending on t such that $\mathbb{E}[X_t^r] \leq M$ for all t .⁹

Framed in our terminology from above, their result shows that any nonnegative stochastic process, subject to a somewhat technical condition on the p -th moment of the conditional increments, will be r -stable for any $r < p - 1$ so long as it is expected to decrease by a constant whenever it exceeds some threshold. Strong stability will follow if Theorem 2.14 can be applied for every even $p \geq 2$.

Therefore, our main goal in applying this result is defining a suitable stochastic process whose stability will imply the stability of the random sequence we care about, namely Q_1, Q_2, \dots . A naïve attempt at this would be to try to directly apply this result to the sequence $Q_1, Q_2, \dots, Q_t, \dots$. However, we already

⁹Note that it suffices to verify these conditions conditionally on a sequence of σ -algebra \mathcal{G}_t such that $\mathcal{F}_t \subseteq \mathcal{G}_t$ by the tower law for conditional expectations.

encounter challenges with this before accounting for how the no-regret assumption factors into the evolution of this process. It is not entirely clear how to relate this process, which only measures the number of uncleared packets, to the underlying priority structure that drives the dynamics, which depend on the ages of the queues. Verifying the first condition itself also becomes rather challenging, as we must condition on arbitrary histories and reason about the timestamps of all packets that have already arrived in the system. While we do have control on the distribution of packet arrivals and server successes, we would have to also reason about atypical, low-probability events.

Our first idea is to take a dual perspective on the underlying queuing dynamics using the *principle of deferred decisions*. More precisely, we will consider the same queuing system where we only keep track of the *age* of each queue at each time t , rather than the round-to-round evolution of $Q_t^{(i)}$. This shift in perspective will lead to several advantages:

1. First, the ages of the queues at each time t give sufficient and direct information to argue about the priority structure among queues.
2. The evolution of ages will also follow a much lighter-weight recursion that arises from the following observation: whenever queue i clears their oldest packet, the age of their next oldest packet can be obtained by subtracting an independent $\text{Geom}(\lambda_i)$ random variable. This holds because the gaps between successful Bernoulli random trials is governed by the Geometric distribution.
3. Moreover, under Assumption 4.2, the stability of both systems is identical because the conditional distribution of the number of uncleared packets given the age of a queue at any time t is straightforward to determine.

4. Finally, because we only keep track of the ages of each queue at each time t , we no longer have to reason about the packets that have arrived up to time t until they “are needed” – indeed, with this Geometric formulation, these packets have not been sampled yet. In particular, if we take the window length w in Assumption 4.2 large enough, standard concentration bounds will ensure that every queue has no-regret on the interval and that the realizations of the Geometric random variables for age decreases and the Bernoulli random variables for server successes do not deviate much from their expectations. This will allow us to reason essentially deterministically about the changes in the ages over the course of the window using the no-regret condition.

By virtue of these properties, it will be more convenient but equivalent to reason about this dual system, which we will call **Geometric queuing systems**.

To carry out the argument, the key difficulty now becomes: what potential function or stochastic process related to the queue ages should we choose that provably has negative drift when it is large? The simplest possible choice is the maximum queue age, which can be viewed as an ℓ_∞ potential. We will be able to show that many old packets will be cleared throughout the interval, which suggests that the ages of the oldest queues decrease *in aggregate* using the no-regret condition. This part of the argument will be quite natural and will follow by establishing a simple dichotomy. Any queue that remains old throughout a window would be chosen at any high-rate server at a round where no old packet is sent there. By the no-regret condition, this queue’s performance must also be at least their performance by deviating to any high-rate server up to sublinear additive terms. These two facts will imply that for each high-rate server, either

an old packet was often sent there or each continually old queue could have often been chosen at the server and thus would have received good service in expectation. Because w is chosen large, concentration of the number of successes at each server around their expectations and summing over a suitable choice of high-rate servers will thus imply a lower bound on the number of old packets cleared by the system throughout this period with high probability.

The only issue with this approach is that the dependencies from the priority scheme and the no-regret dynamics make analyzing how this decrease is spread *amongst* old queues rather tricky: this is necessary to control the growth of the oldest queue age. Moreover, this potential does not alone sufficiently keep track of the full state of the system, in that it does not capture the behavior of any other queue. To try to benefit from the performance of all queues using no-regret algorithms, one could instead try an ℓ_1 -style potential function. However, this potential has a different problem; the gains by older queues could be washed out by the aging of young queues that do not have priority for this choice of potential.

It will thus be most convenient to consider an ℓ_2^2 potential function of queue ages. This choice of potential function is natural in several ways: first, using the squared queue ages ensures that the oldest queues contribute more to the potential, modeling the advantages of an ℓ_∞ potential analysis described above. On the other hand, this potential will also suitably account for the aggregate decrease in queue ages at *all scales* like in an ℓ_1 potential. Indeed, the ℓ_2^2 potential can be obtained by ℓ_1 potential over just those queues above each age threshold. A careful application of these insights will lead to the desired drift condition. To satisfy the second condition of the Pemantle and Rosenthal theorem, we will

only need to translate back to a suitable ℓ_2 -style norm of queue ages.

4.4 Geometric Queuing Systems

In this section, we carry out the first part of the above plan by showing that the Bernoulli queuing system can be equivalently described as a Geometric queuing system using the principle of deferred decisions.

Definition 4.6. Consider a Bernoulli queuing system with parameters μ and λ . Furthermore, suppose that each queue selects servers at each round t only using their past feedback of successes and past history of ages. Then the **Geometric queuing system** is described as follows. For each discrete time step $t = 0, \dots$, the following occurs:

1. At each time t , $\tilde{T}_t^{(i)}$ is the **timestamp** of the oldest unprocessed packet of queue i at time t . $T_t^{(i)} = \max\{0, t - \tilde{T}_t^{(i)}\}$ is the **age** of the current oldest packet of queue i at time step t . That is, $T_t^{(i)}$ measures how long the current oldest unprocessed packet for queue i has been in the system at time t .¹⁰
2. Queue i can send a packet to any server j in this time step if $t - \tilde{T}_t^{(i)} \geq 0$. Each server j attempts to serve only the oldest packet it receives, and succeeds with probability μ_j . If queue i 's packet is successfully served, set $\tilde{T}_{t+1}^{(i)} = \tilde{T}_t^{(i)} + G^{(i)}$, where $G^{(i)} \sim \text{Geom}(\lambda_i)$ is independent of all past events¹¹, and otherwise $\tilde{T}_{t+1}^{(i)} = \tilde{T}_t^{(i)}$.

¹⁰Note that while $T_t^{(i)} \geq 0$ by definition, it is possible that $\tilde{T}_t^{(i)} > t$. The interpretation is that the queue has cleared all of her packets at time t and will receive her next one at time $t = \tilde{T}_t^{(i)}$, or equivalently, in $\tilde{T}_t^{(i)} - t$ steps in the future from the perspective at time t .

¹¹This independence is where we use the assumption that each queue selects servers that is measurable with respect to just their observed feedback of successes and past history of ages.

We let $\mathcal{G}_t^{(i)} = \sigma(\{(T_s^{(i)}, X_s^{(i)})_{s=0}^t, T_{t+1}^{(i)}\})$ be the σ -algebra generated by the observable history of queue i in the Geometric system up to and including time t , along with any internal randomness they use. We require that each queue i 's choice of server at time $t + 1$ be $\mathcal{G}_t^{(i)}$ -measurable (again, along with any internal randomness). Finally, let $\mathcal{G}_t = \sigma(\mathcal{G}_t^{(1)}, \dots, \mathcal{G}_t^{(n)})$ denote the history of the Geometric queuing system up to time t .

Let $\mathbf{T}_t = (T_t^{(1)}, \dots, T_t^{(n)}) \in \mathbb{N}^n$ denote the vector of current ages of oldest packets, and let $T_t = \sum_{i=1}^n T_t^{(i)}$. To see the equivalence with the Bernoulli queuing system, consider any Bernoulli queuing system with Bernoulli random variables $\{B_t^{(i)}\}_{i \in [n], t \geq 0}$ for packet generation. To get a coupled Geometric system for the same system, use an independent sequence $\{G_s^{(i)}\}_{i \in [n], s \geq 0}$ with the interpretation that $G_s^{(i)} \sim \text{Geom}(\lambda_i)$ is the size of the s th gap between successes in the B_t^i . When queue i succeeds in clearing her s -th packet, her new oldest timestamp increases by $G^{(i)} = G_s^{(i)}$ as described above. As such gaps between timestamps in the Bernoulli model have $\text{Geom}(\lambda_i)$ distributions, the Geometric system gives the ages of each queue in the Bernoulli system at all times and gives an explicit coupling.

The key feature is that, under the assumption that each queue i chooses servers at time t using a measurable policy with respect to $\mathcal{G}_t^{(i)}$, all choices by queues are the same conditioned on just the current timestamp and past feedback as it is conditioned on all the past information in the Bernoulli queuing model (which also includes arrivals received after the current oldest packet). In particular, we have $\mathcal{F}_t^{(i)} \supseteq \mathcal{G}_t^{(i)}$ but all choices by the queues at time $t + 1$ are measurable with respect to either history (possibly using any internal randomness). This ensures that $G_s^{(i)}$ will be independent of \mathcal{G}_t until the stopping time of clear-

ing her s -th packet; in other words, the timestamp of queue i 's $s + 1$ -th arriving packet is not known until the time queue i clears her s -th packet.

In the Geometric system, we can define stability in the same way as before:

Definition 4.7 (Geometric Stability). A Geometric queuing system under some given dynamics is **r -stable** if

$$\mathbb{E}[(T_i)^r] \leq C(r, \boldsymbol{\mu}, \boldsymbol{\lambda}),$$

where $C(r, \boldsymbol{\mu}, \boldsymbol{\lambda})$ is an absolute constant depending only on $r, \boldsymbol{\mu}$ and $\boldsymbol{\lambda}$. A Geometric queuing system is **strongly stable** if it is r -stable for every $r > 0$.

A Geometric queuing system is **weakly stable** if it holds that

$$\mathbb{E}[T_i] = o(t).$$

Because heuristically $Q_i^{(i)} \approx \lambda_i T_i^{(i)}$, it is not difficult to show that our notions of strong stability are equivalent whenever both systems correspond to the same random process.

Lemma 4.5. *Suppose that each queue i selects servers using strategies at time $t + 1$ that are measurable with respect to the filtration $\mathcal{G}_t^{(i)}$. Then strong (weak) stability in the Bernoulli system holds iff and only if strong (weak) stability holds in the Geometric system.*

Proof. Under the measurability assumption, the distribution of $Q_i^{(i)}$ conditioned on $\mathcal{G}_t^{(i)}$ is $\text{Bin}(T_i^{(i)}, \lambda_i)$ because the subsequent Bernoulli arrivals since the oldest packet current packet arrived are independent of $\mathcal{G}_t^{(i)}$. By the Law of Iterated Expectations,

$$\mathbb{E}[(Q_i^{(i)})^p] = \mathbb{E}[\mathbb{E}[(Q_i^{(i)})^p | \mathcal{G}_t^{(i)}]] \asymp \mathbb{E}[(T_i^{(i)})^p],$$

where we use Lemma 2.13 in the last step. Therefore, the two systems have equivalent stability properties. \square

4.5 Proof of Main Result

In this section, we can finally prove our main result as stated in Theorem 4.3. First, by Assumption 4.2 and Lemma 4.5, it suffices to analyze the stability of the Geometric queuing system. To apply Theorem 2.14, we must define an appropriate potential function of queue ages that satisfies the negative drift and bounded moments condition.

For each $\tau \in \mathbb{N}$, define the following potential functions:

$$\Phi_\tau(\mathbf{T}_t) \triangleq \sum_{i \in [n]: T_t^{(i)} \geq \tau} \lambda_i (T_t^{(i)} - \tau), \quad (4.4)$$

$$\Phi(\mathbf{T}_t) \triangleq \sum_{\tau=1}^{\infty} \Phi_\tau(\mathbf{T}_t) = \sum_{\tau=1}^{\infty} \sum_{i \in [n]: T_t^{(i)} \geq \tau} \lambda_i (T_t^{(i)} - \tau) = \frac{1}{2} \sum_{i=1}^n \lambda_i T_t^{(i)} (T_t^{(i)} - 1). \quad (4.5)$$

In words, $\Phi_\tau(\cdot)$ denotes the expected number of total packets in the system aged above τ , conditioned just on the ages \mathbf{T}_t . We also set the following parameters:

$$\delta = \frac{\eta}{4}, \quad (4.6)$$

$$\epsilon = \frac{\delta \mu_1}{4n}, \quad (4.7)$$

$$\epsilon_i = \frac{\epsilon}{\lambda_i} \quad \forall i \leq n. \quad (4.8)$$

To apply Theorem 2.14, we define the stochastic process Z_0, Z_1, \dots by $Z_\ell = \sqrt{\Phi(\mathbf{T}_{w\ell})}$. That is, Z_ℓ is the ‘‘snapshot’’ of the potential function $\sqrt{\Phi}$ when evaluated on $\mathbf{T}_{w\ell}$ that occurs every w steps. The filtration is given by $\mathcal{H}_\ell = \mathcal{G}_{\ell, w}$, where

\mathcal{G}_t is the corresponding information of the Geometric system at time t available to the queues.

Summary of the Main Ideas: Before we go through the detailed proof, we offer an outline of the main ideas. To establish the negative drift, we will focus on the w -long interval between two Z_ℓ and $Z_{\ell+1}$. In this w -long window, we use the no-regret condition, as well as concentration bounds on the behavior of queues and servers. The main idea of the proof is to consider all queues that have retained an old packet throughout the period. A server either clears many such old packets, or many times during this period no old packet is sent to it. In the second case, we can use the no-regret condition for any queue that still has very old packets, as they would have priority at the server, so these bounds in tandem will imply that many old packets must have cleared. To aid the analysis, we also lower bound the total decrease in ages from clearing packets on this window *before* accounting for the w extra steps of aging, only accounting for aging at the end; this allows us to consider the clearing process and aging from time passage separately. Finally, when concentration or the no-regret condition fails, we can trivially upper bound what this contributes to the expected drift and this will be subsumed by the low probability that this occurs in the overall expectation.

Organizing Randomness. Let us first set up how we model the actual queuing process on each consecutive window of w steps between Z_ℓ and $Z_{\ell+1}$ for the probabilistic analysis. In the spirit of “organizing randomness,” at step ℓ of this process (step $\ell \cdot w$ of the actual queuing process), sample up front an independent geometric ensemble $\{G_k^{(i)}\}_{i \in [n], k \in [w]}$ with

$$G_k^{(i)} \sim \text{Geom}(\lambda_i), k = 1, \dots, w$$

as well as an independent Bernoulli ensemble $\{I_k^{(j)}\}_{j \in [m], k \in [w]}$ with

$$I_k^{(j)} \sim \text{Bern}(\mu_j), k = 1, \dots, w.$$

The interpretation is that the $I_k^{(j)}$ are random indicators if the j th server is able to clear a packet, *regardless of whether a packet is sent there*, at the k th step of this block of w steps. The $G_k^{(i)}$ have the interpretation that, when queue i clears her k th packet on this window, her age decreases by $G_k^{(i)}$ (without accounting for the aging from passage of time). Crucially, as queue i clears packets on this window of w steps, her age decreases *by a sum of a prefix of $G_1^{(i)}, \dots, G_k^{(i)}$* (before accounting for aging as time passes). Observe that this independence arises precisely because of the independence of the geometric ensemble of timestamp differences from the filtration $\mathcal{H}_\ell = \mathcal{G}_{\ell-w}$ of the Geometric system that only conditions on past feedback and the realized past sequence of ages.

Good Event. We now define the “good” event that asserts that the realized sequences of random variables on this window are close to their expected values. First, we claim that given our choice of parameters in Equation (4.6), there exists $w = w(\eta, \mu, \lambda, n)$ large enough such that with probability at least $1 - \eta/128$, all of the following events hold on this window:

$$n\varphi_{\frac{\eta}{256n}}(w) + n \leq \frac{w\delta\mu_1}{4}, \quad (4.9)$$

and

$$\left| \sum_{k=1}^q G_k^{(i)} - \frac{q}{\lambda_i} \right| < \epsilon_i w \quad \forall i \in [n], q \in [w], \quad (4.10)$$

$$\sum_{j=1}^q \sum_{k=1}^w I_k^{(j)} > w(1 - \delta) \sum_{j=1}^q \mu_j \quad \forall q \in [m], \quad (4.11)$$

$$\text{Reg}_i(w) \leq \varphi_{\eta/(128n)}(w) \quad \forall i \in [n] \implies n + \sum_{i=1}^n \text{Reg}_i(w) \leq n + n\varphi_{\eta/(128n)}(w) \leq \frac{w\delta\mu_1}{4}. \quad (4.12)$$

This is a consequence of our strong concentration and high-probability no-regret bounds. Equation (4.9) holds for large enough w by the sublinearity in the no-regret condition. We may also take w large enough so that the Equation (4.10) and Equation (4.11) hold simultaneously with probability at least $1 - \eta/256$ via Corollary 2.10 and Lemma 2.11, and also large enough so that Equation (4.12) holds with probability at least $1 - \eta/(256n)$ for each queue and applying a union bound. Notice that Equation (4.10) asserts that *every prefix of each of the geometric ensembles is additively not too far from the expectation, relative to w .*

Threshold Value for Z_ℓ . We will show that under the threshold assumption that

$$Z_\ell > \frac{w}{\sqrt{2\lambda_n}} \max \left(\frac{16}{\eta} \left(\sum_{i=1}^n \lambda_i \right), 16n^2 \right), \quad (4.13)$$

then the drift condition in the Pemantle-Rosenthal result holds. The significance of this value is that it will imply there exists at least one rather old queue, and that the expected total number of packets in the system is large relative to the window size. We will later use these facts as demonstrated by the following simple claim, whose proof we defer to Section 4.7.1:

Claim 4.6. *Under Equation (4.13), both of the following statements hold:*

1. *There exists some $i \in [n]$ such that $\lambda_i T_{\ell, w}^i > 16nw$.*
2. $\sum_{i=1}^n \lambda_i T_{\ell, w}^i \geq \frac{16}{\eta} w \sum_{i=1}^n \lambda_i$.

No-Regret Implies Many Old Packets Clear. Continuing with the proof, we first analyze what happens on the “good” event of Equations (4.10) to (4.12). Let $\tau^{(i)}$ be the age of the oldest unprocessed packet of queue i at the end of this window of w steps, *measured with respect to the beginning of the window without*

accounting for the w steps of aging. If queue i cleared all her packets that were received before the beginning of this window, then we say $\tau^{(i)} = 0$. Let J_τ be the set of queues that at the end of the w steps still have packets that are at least τ -old with respect to the beginning of the window. Let $Y_{j,k}^\tau$ be an indicator variable that some packet that was at least τ -old with respect to the beginning of the considered interval was sent to server j at the k th step in this window. As such queues in J_τ evidently have packets that are at least τ -old throughout this interval, priority and the regret bound Equation (4.12) implies that the number of packets cleared by any such queue is at least, for any server $j \in [m]$

$$\sum_{k=1}^w I_k^{(j)} (1 - Y_{j,k}^\tau) - \varphi_{\eta/(128n)}(w). \quad (4.14)$$

This is because a queue that is always at least τ -old throughout the interval would succeed on any server j that is successful on a time step k (indicated by $I_k^{(j)}$) when no τ -old packets were sent there.

Let N_τ be the number of packets that were at least τ -old with respect to the beginning of the interval that were cleared in the interval and $N_\tau^{(i)}$ the number of such packets cleared by queue i . Then we clearly have

$$N_\tau = \sum_{j \in [m]} \sum_{k \in [w]} I_k^{(j)} Y_{j,k}^\tau \geq \sum_{j=1}^{\min\{m, |J_\tau|\}} \sum_{k \in [w]} I_k^{(j)} Y_{j,k}^\tau. \quad (4.15)$$

As every packet processed by queues in J_τ contribute to N_τ , by instantiating Equation (4.14) for each queue in J_τ with each of the top $\min\{m, |J_\tau|\}$ servers and summing, we also obtain

$$\min\{m, |J_\tau|\} \cdot N_\tau \geq \min\{m, |J_\tau|\} \sum_{i \in J_\tau} N_\tau^{(i)} \geq |J_\tau| \sum_{j=1}^{\min\{m, |J_\tau|\}} \left(\sum_{k=1}^w I_k^{(j)} (1 - Y_{j,k}^\tau) - \varphi_{\eta/(128n)}(w) \right). \quad (4.16)$$

Multiplying Equation (4.15) by $|J_\tau|$ and summing with the previous equation,

we obtain

$$N_\tau \geq \left(\frac{|J_\tau|}{|J_\tau| + \min\{m, |J_\tau|\}} \right) \sum_{j=1}^{\min\{m, |J_\tau|\}} \left(\sum_{k=1}^w I_k^{(j)} - \varphi_{\eta/(128n)}(w) \right) \quad (4.17)$$

$$\geq \frac{1}{2} \sum_{j=1}^{\min\{m, |J_\tau|\}} \sum_{k \in [w]} I_k^{(j)} - n\varphi_{\eta/(128n)}(w) \quad (4.18)$$

$$\geq \frac{1}{2} w(1 - \delta) \sum_{j=1}^{\min\{m, |J_\tau|\}} \mu_j - n\varphi_{\eta/(128n)}(w), \quad (4.19)$$

where the last inequality uses Equation (4.11).

Large Φ -Drift on Good Event. Observe that from the construction of Φ_τ , when queue i manages to process a packet that is at least τ -old, Φ_τ decreases either by $\lambda_i G_k^i$ for some k if the new age remains above τ , or the term vanishes in which case Φ_τ may decrease by less. Crucially, this latter possibility can only happen at most once. Again, write $N_\tau^{(i)}$ for the number of packets that queue i clears during this interval that are at least τ -old. Then as $\sum_{i=1}^n N_\tau^{(i)} = N_\tau$, the decrease in Φ from Φ_τ , denoted Δ_τ , satisfies

$$\begin{aligned} \Delta_\tau &\geq \sum_{i=1}^n \lambda_i \left(\sum_{k=1}^{N_\tau^{(i)}-1} G_k^{(i)} \right) \\ &\geq \sum_{i=1}^n \lambda_i \left(\frac{N_\tau^{(i)} - 1}{\lambda_i} - \epsilon_i w \right) \\ &= \sum_{i=1}^n (N_\tau^{(i)} - 1 - \lambda_i \epsilon_i w) \\ &= N_\tau - n - n\epsilon w, \end{aligned}$$

where we apply Equation (4.10) in the second inequality. By applying the lower bound derived in Equation (4.19), and then Equation (4.9) with the definition of

ϵ , we further obtain

$$\begin{aligned} N_\tau - n - n\epsilon w &\geq \frac{1}{2}w(1 - \delta) \sum_{j=1}^{\min\{m, J_\tau\}} \mu_j - n\varphi_{\eta/(128n)}(w) - n - n\epsilon w \\ &\geq \frac{1}{2}w(1 - \delta) \sum_{j=1}^{\min\{m, J_\tau\}} \mu_j - \frac{\delta w \mu_1}{2}. \end{aligned}$$

By absorbing the last term into the sum, we thus conclude that

$$\Delta_\tau \geq \frac{1}{2}w(1 - 2\delta) \sum_{j=1}^{\min\{m, J_\tau\}} \mu_j. \quad (4.20)$$

Summing this over all τ , the decrease in Φ before considering aging satisfies

$$\Delta\Phi \triangleq \sum_{\tau=1}^{\infty} \Delta_\tau \geq \frac{1}{2}w(1 - 2\delta) \sum_{\tau=1}^{\infty} \sum_{j=1}^{\min\{m, J_\tau\}} \mu_j = \frac{1}{2}w(1 - 2\delta) \sum_{j=1}^{\min\{m, n\}} \tau_j \mu_j, \quad (4.21)$$

where τ_j is the j 'th largest of the $\tau^{(i)}$.

Effect of Aging. We now account for the increase due to aging by w over the course of this interval. The increase in Φ from this is upper bounded by

$$\begin{aligned} \frac{1}{2} \sum_{i=1}^n \lambda_i (\tau^{(i)} + w)(\tau^{(i)} + w - 1) - \frac{1}{2} \sum_{i=1}^n \lambda_i \tau^{(i)} (\tau^{(i)} - 1) &= w \sum_{i=1}^n \lambda_i \tau^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i w(w - 1) \\ &\leq w \sum_{i=1}^n \lambda_i \tau^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i w^2. \end{aligned}$$

Note that this is only exact for those τ_i that are nonzero, while is an upper bound for those that are zero. Combining this with Equation (4.21), we see that the potential decrease is at least

$$\frac{1}{2}w(1 - 2\delta) \sum_{j=1}^{\min\{m, n\}} \tau_j \mu_j - \left(w \sum_{i=1}^n \lambda_i \tau^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i w^2 \right) \geq \frac{\eta}{4}w \sum_{i=1}^n \lambda_i \tau^{(i)} - \frac{1}{2} \sum_{i=1}^n \lambda_i w^2, \quad (4.22)$$

as $2\delta = \eta/2$ and using Assumption 4.1 with the fact that the sum of a product of two nonnegative sequences is maximal when both are in the same sorted order

(see Fact 2.2).

Relating $\tau^{(i)}$ and $T_{\ell-w}^{(i)}$. We now need to transfer this decrease back to the $T_{\ell-w}^{(i)}$ random variables. For this, we have the following claim that relates the $\tau^{(i)}$ and $T_{\ell-w}^{(i)}$:

Claim 4.7. *If there exists an $i \in [n]$ such that $\lambda_i T_{\ell-w}^{(i)} > 16nw$ under the good event assumptions, then*

$$\sum_{i=1}^n \lambda_i \tau^{(i)} \geq \frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}.$$

By Claim 4.6, the precondition of Claim 4.7 holds for our threshold value, so Claim 4.6, Claim 4.7 and Equation (4.22) imply the decrease in Φ on this good event is at least

$$\frac{\eta}{8} w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} - \frac{1}{2} \sum_{i=1}^n \lambda_i w^2 \geq \frac{\eta}{16} w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}.$$

Translating this into the decrease in $\sqrt{\Phi}$, Fact 2.4 implies that the contribution towards the expected decrease on this event, which occurs with probability at least $1 - \eta/128 \geq 1/2$, is at least

$$\left(\frac{1}{2}\right) \frac{\frac{\eta}{16} w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}}{2 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} (T_{\ell-w}^{(i)} - 1)}} = \frac{\eta w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}}{64 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} (T_{\ell-w}^{(i)} - 1)}} \quad (4.23)$$

Considering “bad” events. We now analyze the bad event where any of these assumptions fails: the worst case is that all queues clear no packets, and so each $T_{\ell-w}^{(i)}$ increases by w on the next w steps. The increase in Φ is thus at most

$$\frac{1}{2} \sum_{i=1}^n \lambda_i (T_{\ell-w}^{(i)} + w)(T_{\ell-w}^{(i)} + w - 1) - \frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} (T_{\ell-w}^{(i)} - 1) \leq w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i w^2 \leq 2w \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)},$$

where the last inequality is again Claim 4.6. Translating to squareroots again, on this bad event which occurs with probability at most $\eta/128$, the contribution

of increase to the expected change in $\sqrt{\Phi}$ is at most

$$\left(\frac{\eta}{128}\right) \frac{2w \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)}}{2 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}} = \frac{\eta w \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)}}{128 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}}, \quad (4.24)$$

by Fact 2.3. Summing Equation (4.23) and Equation (4.24), it follows $\sqrt{\Phi}$ decreases in expectation by at least

$$\frac{\eta w \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)}}{128 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}} \geq \frac{\eta w \sqrt{\lambda_n}}{128},$$

where the last inequality is Lemma 2.5. This proves that the drift condition holds for this stochastic process with the threshold given above.

Bounded p th Moments: The last thing to check to apply Theorem 2.14 is that the increments $Z_{\ell+1} - Z_\ell$ have conditionally bounded p th moments for each even integer $p \geq 2$ to obtain boundedness of our sequence in L^r for all $r \geq 0$. This is relatively straightforward, but tedious: therefore, we prove the following claim in Section 4.7.1:

Claim 4.8. *For each even integer $p \geq 2$, there exists a constant $C_{p,n,w} > 0$ which depends only on n, w, p and the parameters of the system such that for all $\ell \geq 0$,*

$$\mathbb{E}[|Z_{\ell+1} - Z_\ell|^p | \mathcal{H}_\ell] \leq C_{p,n,w}.$$

Putting It Together: Therefore, Theorem 2.14 applies to the random process Z_ℓ , and we conclude that for each $r \geq 0$, there exists some absolute constant C_r such that for all $\ell = 0, 1, \dots$,

$$\mathbb{E}[Z_\ell^r] \leq C_r.$$

In particular, this means that for each $\ell \geq 0$,

$$\mathbb{E} \left[\left[\sqrt{\sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)} \right]^r \right] \leq C_r.$$

To extend this to all $t \geq 0$ not necessarily a multiple of w , it is clear that deterministically, $\sqrt{\Phi(\mathbf{T}_t)} \asymp \|\mathbf{T}_t\|_{\lambda,2} \asymp \|\mathbf{T}_t\|_{\lambda,1}$ up to additive and multiplicative constants as in Lemma 2.5, from which we can conclude

$$\mathbb{E} \left[\left(\sum_{i=1}^n \lambda_i T_t^{(i)} \right)^r \right] \leq C'_r$$

for some other constant C'_r for each $t = \ell w$. Note that for $\ell w \leq t < (\ell + 1)w$, each term can increase by at most w compared to the value at ℓw , and therefore we can conclude that for all $t \geq 0$,

$$\mathbb{E} \left[\left(\sum_{i=1}^n T_t^{(i)} \right)^r \right] \leq C''_r$$

for some constant C''_r independent of t . This concludes the proof of strong stability.

4.6 Optimality of Main Result

We now provide a simple construction showing that a partial converse to Theorem 4.3: $\frac{1}{2}$ is the best constant that can appear in Assumption 4.1 for a similar no-regret condition to be sufficient for stability as in Theorem 4.3. To set it up, let $w_k = k^2$ for each $k \geq 1$. Then we have the following:

Theorem 4.9. *Partition time $t = 0, 1, \dots$ into consecutive windows, where the k th window has length $w_k = k^2$. Then there exists a family of queuing systems with n queues and servers for each $n \geq 1$ satisfying Assumption 4.1 with $\frac{1}{2} + o_n(1)$ in place of $\frac{1}{2}$ with the following properties: almost surely, each queue has zero regret on all but at most finitely many of the windows, but the system is not weakly stable.*

The formal details are slightly technical, but the high-level idea is quite natural: for each $n \geq 1$, consider the following system on n queues and n servers

where we set $\lambda = (\frac{n+1}{n^2}, \dots, \frac{n+1}{n^2})$ and $\mu = (1, \frac{n-1}{n^2}, \dots, \frac{n-1}{n^2})$. Consider the strategy where every queue always sends to the rate 1 server. It is easy to see purely from expectations that the queue lengths are unbounded in expectation, as the sum of arrival rates strictly exceeds 1. On the other hand, it is intuitive that this strategy will “usually” be zero-regret; if all the queues are similarly aged at the start of some window, then they should expect to clear roughly $1/n$ fraction of the time on this window using this strategy, which strictly exceeds what they would get at any other server.

Proof of Theorem 4.9. Let $W_k = \sum_{i=1}^{k-1} w_i$. Note $W_k = \Theta(k^3) = \Theta(w_k^{3/2})$. W_k is the actual time step at the end of $k - 1$ of the consecutive windows of length w_i for $i = 1, \dots, k - 1$. Note also that $W_{k+1} - W_k = w_k$.

For each $n \geq 1$, consider the following system on n queues and n servers: set $\lambda = (\frac{n+1}{n^2}, \dots, \frac{n+1}{n^2})$ and $\mu = (1, \frac{n-1}{n^2}, \dots, \frac{n-1}{n^2})$. This system satisfies Assumption 4.1 with factor $\frac{1}{2} - o_n(1)$. We consider the simple strategy where every queue always sends to the rate 1 server. Under these oblivious dynamics, in expectation the total number of packets grows by $\frac{1}{n}$ with every step, and therefore this system is not even weakly stable. What we must show is that almost surely, this fixed strategy is zero regret for every queue for all but finitely many of the windows.

Concentration of Packet Arrivals. We first show almost sure concentration of the arrivals of new packets. Let $\{B_t^i\}_{i \in [n], t \geq 1}$ be the independent random variables for arrivals as usual. Now, for each queue $i \in [n]$ and $\ell \geq 0$, we have

$$\Pr \left(\left| \sum_{t=1}^{\ell} B_t^i - \lambda_i \ell \right| \geq \sqrt{\ell \ln(\ell)} \right) \leq \frac{2}{\ell^2}, \quad (4.25)$$

where we use the additive form of the Chernoff bound. As the same holds for all queues, the probability this event happens for any of the n queues is at most

$2n/\ell^2$. As this is summable in ℓ , we may sum over all $\ell \geq 1$ to deduce from the Borel-Cantelli lemma that almost surely, for all sufficiently large ℓ , all $i \in [n]$ satisfy

$$\sum_{t=1}^{\ell} B_t^i = \lambda_i \ell \pm O(\sqrt{\ell \ln(\ell)}). \quad (4.26)$$

Note that this also implies that almost surely, for all large ℓ , $\sum_{t=1}^{\ell} \sum_{i=1}^n B_t^i \geq (1 + \frac{1}{2n}) \cdot \ell$ by the choice of λ_i . Moreover, under this fixed strategy where everyone always sends to the rate 1 server, at most ℓ packets can be cleared by time ℓ .

Proportional Instability. Next, we show that almost surely, there is a large backup proportional to the current time period. Let t_k be the last timestamp the rate 1 server clears up to time W_k . As all queues send there under this fixed strategy, at this point, all queues only have packets that were received after t_k by priority. On the one hand, it is not difficult to see that deterministically $t_k \geq W_k/n$ (equality happens in the worst case where every queue received a packet in every step up to W_k). On the other hand, in light of our results above, almost surely, for all but finitely many of the k ,

$$t_k < \frac{1}{1 + \frac{1}{2n}} W_k = (1 - \Omega(1)) W_k. \quad (4.27)$$

This is because at least W_k packets have been received up to time $\frac{1}{1 + \frac{1}{2n}} W_k$, and because the server can only have cleared at most W_k packets up to time W_k , the oldest timestamp the server could have cleared by time W_k can be at most this quantity.

Concentration of Server Successes. Next, we show almost sure concentration of the nontrivial server success rates. Let I_t^j be the indicator that server j would succeed at clearing a packet at time t (regardless of if one is sent there; under this strategy no queue ever sends to $j \neq 1$). A similar application of the Chernoff bound and union bound with the Borel-Cantelli lemma implies that almost

surely, for all but finitely many of the k , and for each server $j \in [n]$, we have

$$\sum_{t=W_k+1}^{W_{k+1}} I_t^j = w_k \mu_i \pm O(\sqrt{w_k \ln w_k}). \quad (4.28)$$

Note that the increasing nature of the w_k is needed here for this to be valid (and in fact, this statement will be false with probability one if interval sizes are kept fixed by independence and the second Borel-Cantelli lemma). Thus, almost surely, for all large enough ℓ and k , all of these events happen simultaneously.

Zero Regret All but Finitely Many Times. As $t_k \geq W_k/n$, almost surely for large enough k , t_k eventually exceeds the random time ℓ at which Equation (4.26) holds. Consider any subsequent window of length w_k . Our goal is to use these facts to show that on these windows, all queues have zero regret. First, we show that each queue clears $(\frac{1}{n} - o(1))w_k$ packets on each such window. Let $c = \frac{1}{n\lambda_i} < 1$ (note this is independent of i). We know from Equation (4.27) that $t_k + w_k < (1 - \Omega(1))W_k + w_k < W_k$; moreover, using Equation (4.26) and the fact $t_k \geq \ell$,

$$\begin{aligned} \sum_{t=t_k+1}^{t_k+c \cdot w_k} B_t^i &= \sum_{t=1}^{t_k+c \cdot w_k} B_t^i - \sum_{t=1}^{t_k} B_t^i \\ &= \frac{1}{n} \cdot w_k \pm O(\sqrt{(t_k + c \cdot w_k) \ln(t_k + c \cdot w_k)}) \\ &= \frac{1}{n} \cdot w_k \pm O(\sqrt{W_k \ln W_k}) \\ &= \frac{1}{n} \cdot w_k \pm \tilde{O}(w_k^{3/4}), \end{aligned}$$

where the last line uses the relationship between W_k and w_k . As $t_k + w_k < W_k$, all of these packets were evidently received before the start of the given window, and therefore, every queue is backed up throughout the period, and by virtue of the previous equation, each queue has $\frac{1}{n} - o(1)$ fraction of the next w_k packets

that will be cleared by this top server on this window. Therefore, each queue clears at least $(\frac{1}{n} - o(1)) \cdot w_k$ packets on such windows under this fixed strategy.

Finally, had any queue deviated on such a window to a single fixed low rate server, in light of Equation (4.28), she would have cleared $(\frac{n-1}{n^2} + o(1)) \cdot w_k$ packets, which is linearly smaller than the amount she actually cleared. Therefore, almost surely, on all but finitely many of the windows, every queue actually has zero regret. \square

4.7 Chapter Notes

The content in this section originally appears in [71], joint with Éva Tardos. Compared to the original version, the measurability assumptions underlying the equivalence between Bernoulli and Geometric systems and stability definitions are made more explicit. Some of the proofs have been somewhat modified for readability.

4.7.1 Deferred Proofs

In this subsection, we return to proving some auxiliary claims from the proof of Theorem 4.3.

Claim 4.10 (Claim 4.6, restated). *Under Equation (4.13), both of the following statements hold:*

1. *There exists some $i \in [n]$ such that $\lambda_i T_{\ell-w}^{(i)} > 16nw$.*

$$2. \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} \geq \frac{16}{\eta} w \sum_{i=1}^n \lambda_i.$$

Proof. Equation (4.13) immediately implies by definition of Z_ℓ and Φ that

$$\sqrt{\sum_{i=1}^n \lambda_i (T_{\ell-w}^{(i)})^2} \geq \frac{w}{\sqrt{\lambda_n}} \max\left(\frac{16}{\eta} \left(\sum_{i=1}^n \lambda_i\right), 16n^2\right).$$

From Lemma 2.5, this implies that

$$\sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)} \geq w \max\left(\frac{16}{\eta} \left(\sum_{i=1}^n \lambda_i\right), 16n^2\right),$$

from which both parts follow, the first from averaging. \square

Claim 4.11 (Claim 4.7, restated). *If there exists an $i \in [n]$ such that $\lambda_i T_{\ell-w}^{(i)} > 16nw$ under the good event assumptions, then*

$$\sum_{i=1}^n \lambda^{(i)} \tau_i \geq \frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}.$$

Proof. By Equation (4.10), we must have

$$\lambda_i \tau^{(i)} \geq \lambda_i T_{\ell-w}^{(i)} - w - \lambda_i \epsilon_i w,$$

even if queue i clears a packet every step in the window. Observe that if $\lambda_i T_{\ell-w}^{(i)} \geq 8w$, then

$$w + \lambda_i \epsilon_i w \leq 2w \leq \frac{1}{4} \lambda_i T_{\ell-w}^{(i)},$$

and so $\lambda_i \tau^{(i)} \geq \frac{3}{4} \lambda_i T_{\ell-w}^{(i)}$. We also have

$$\frac{1}{2} \sum_{i: \lambda_i T_{\ell-w}^{(i)} < 8w} \lambda_i T_{\ell-w}^{(i)} < 4nw.$$

In particular, if there exists some i such that $\lambda_i T_{\ell-w}^{(i)} > 16nw$, then

$$\lambda_i \tau^{(i)} \geq \frac{3}{4} \lambda_i T_{\ell-w}^{(i)} > \frac{1}{2} \lambda_i T_{\ell-w}^{(i)} + 4nw \geq \frac{1}{2} \lambda_i T_{\ell-w}^{(i)} + \frac{1}{2} \sum_{i: \lambda_i T_{\ell-w}^{(i)} < 8w} \lambda_i T_{\ell-w}^{(i)}.$$

It follows that if this holds, then indeed,

$$\sum_{i=1}^n \lambda_i \tau^{(i)} = \sum_{i: \lambda_i T_{\ell-w}^{(i)} \geq 8w} \lambda_i \tau^{(i)} + \sum_{i: \lambda_i T_{\ell-w}^{(i)} < 8w} \lambda_i \tau^{(i)} \geq \frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell-w}^{(i)}.$$

□

Claim 4.12 (Claim 4.8, restated). *For each even integer $p \geq 2$, there exists a constant $C_{p,n,w} > 0$ which depends only on n, w, p and the parameters of the system such that for all $\ell \geq 0$,*

$$\mathbb{E}[|Z_{\ell+1} - Z_\ell|^p | \mathcal{H}_\ell] \leq C_{p,n,w}.$$

Proof. By the Triangle Inequality, it is easy to see that as random variables, the change in $T_{\ell-w}^{(i)}$ is at most $G^{(i)} \triangleq \sum_{k=1}^w G_k^{(i)}$. Then the change in Φ is again at most

$$\frac{1}{2} \sum_{i=1}^n \lambda_i (T_{\ell-w}^{(i)} + G^{(i)}) (T_{\ell-w}^{(i)} + G^{(i)} - 1) - \frac{1}{2} \sum_{i=1}^n \lambda_i (T_{\ell-w}^{(i)}) (T_{\ell-w}^{(i)} - 1) \leq \sum_{i=1}^n \lambda_i G^{(i)} T_{\ell-w}^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i (G^{(i)})^2,$$

as random variables. We treat two different cases separately:

1. Suppose there does not exist $i \in [n]$ such that $\lambda_i T_{\ell-w}^{(i)} > 1$. Then the change in Φ is at most

$$\sum_{i=1}^n G^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i (G^{(i)})^2.$$

From Fact 2.3, this means the change in $\sqrt{\Phi}$ is upper bounded as random variables by

$$\sqrt{\sum_{i=1}^n G^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i (G^{(i)})^2}.$$

Raising this to the p th power, expanding, and taking expectations, this term is at most $C_{p,n,w} / \lambda_n^{2p}$ for some constant $C_{p,n,w}$ depending only on n, w , and p by Lemma 2.12.

2. Suppose there does exist $i \in [n]$ such that $\lambda_i T_{\ell,w}^{(i)} > 1$. We claim this implies that for all $j \in [n]$

$$\frac{\lambda_j T_{\ell,w}^{(j)}}{2 \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}} \leq \sqrt{\lambda_j}$$

First, note that for any $i \in [n]$, $T_{\ell,w}^{(i)} \geq 2$ implies

$$\frac{1}{2} \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1) \geq \frac{1}{4} \lambda_i (T_{\ell,w}^{(i)})^2, \quad (4.29)$$

as can be confirmed from basic algebra. As $\lambda_i \leq 1/2$ by feasibility (as $\mu_1 \leq 1$), our assumption implies $T_{\ell,w}^{(i)} > 2$, and so

$$2 \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)} > 1.$$

To prove the claim, we split into more cases: if $T_{\ell,w}^{(j)} \leq 1/\sqrt{\lambda_j}$, the claim holds using the last inequality in the denominator. Otherwise, we must have $T_{\ell,w}^{(j)} \geq 2$ by integrality, in which case by Equation (4.29),

$$\frac{\lambda_j T_{\ell,w}^{(j)}}{2 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}} \leq \frac{\lambda_j T_{\ell,w}^{(j)}}{\sqrt{\lambda_j (T_{\ell,w}^{(j)})^2}} = \sqrt{\lambda_j}.$$

Thus, in this case, we have

$$\frac{\sum_{i=1}^n \lambda_i G^{(i)} T_{\ell,w}^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i (G^{(i)})^2}{2 \cdot \sqrt{\frac{1}{2} \sum_{i=1}^n \lambda_i T_{\ell,w}^{(i)} (T_{\ell,w}^{(i)} - 1)}} \leq \sum_{i=1}^n \sqrt{\lambda_i} G^{(i)} + \frac{1}{2} \sum_{i=1}^n \lambda_i (G^{(i)})^2.$$

By Fact 2.3, this is an upper bound as random variables of the change in $\sqrt{\Phi}$, so taking p th powers, expanding, and taking expectations, we get an upper bound of $C_{p,n,w}/\lambda_n^{2p}$ by Lemma 2.12 for some constant $C_{p,n,w}$ depending only on n, w, p, λ .

□

4.7.2 On the No-Regret Assumption

We now provide a brief discussion of our assumption of no-regret in these queuing systems. As mentioned in Chapter 3, no-regret learning is a classical modeling assumption in the context of repeated, independent games that often attains equivalent price-of-anarchy-style guarantees to that of Nash equilibria. Moreover, the convergence of no-regret learning to a correlated equilibrium gives an intrinsic game-theoretic justification for using it as a behavioral model of the agents. In our setting, because queues only receive bandit feedback on their actions, but otherwise may not know the service rates μ nor any other information on the number or identities of the other queues in the systems or their choices of servers, it is natural to assume that they may use an adaptive learning algorithm to select servers in each round. In large systems, reasoning about explicit game-theoretic actions on a round-by-round basis may prove difficult or impossible due to these information constraints, but no-regret learning is nonetheless achievable and gives useful performance guarantees.

4.7.3 Related and Subsequent Work

Queuing theory is the subject of immense study due to the enormous practical applications; while we cannot provide a complete survey of this area, we refer the reader to [123]). More specifically, though, there has been more recent work focusing on efficiency loss due to selfishness in different classical queuing systems, as well as the role of learning in such systems. For work on price of anarchy in queuing systems, see the book of Hassin [85] and survey of Hassin and Haviv [86], and for a very recent tutorial on the role of learning and information

in queuing systems see Walton and Xu [129]. Closest to our model from this literature is the work of Krishnasamy, et al [100]. Our model of strategic queuing is based off their model; however, they consider *centralized* learning algorithms for learning server parameters instead of strategic dynamics. Their main results the “queue-regret” of centralized learning— In contrast, we assume that each queue separately learns to selfishly make sure its own packets are served at highest possible rate, offering a strategic model of scheduling packets in a queuing system.

While our results fall in a long tradition of establishing price of anarchy bounds for various games [99], our results here appear to be one of the only provable guarantees of learning in games with carryover effects between rounds. As discussed in Chapter 3, the crucial novelty of our results is that we no longer make the assumption that games at different rounds are independent, in stark contrast to classical results in learning in games [25, 114, 126]. Studying this model requires us to combine ideas from the price of anarchy analysis of games with the theory of stochastic processes. On a technical level, perhaps the closest work to ours is work of Borodin, et al [28] on adversarial queuing systems, who also use the Pemantle and Rosenthal [111] theorem to establish bounded queue sizes in expectation.

Our focus on infinite-time processes and welfare objectives substantially differ from the price of anarchy literature, which are typically framed in terms of terms of aggregate total welfare achieved by all agents over the course of the (finite) dynamics. By contrast, our stability notion is substantially more complex for two reasons: first, proving stability inherently requires performance guarantees for *all agents separately*, not just in aggregate. The latter is far easier

to achieve in our setting, and our emphasis on stability necessitates the priority structure we consider. Second, stability is a binary, probabilistic notion that depends on the entire course of the (infinite horizon) dynamics, necessitating our probabilistic analyses that are not present in the price of anarchy literature. However, our results qualitatively resemble the bicriteria result of Roughgarden and Tardos [117], which shows that in nonatomic routing, the cost incurred at any Nash flow is at most the optimal cost when twice the flow is routed.

The repeated queuing game we consider is a special case of a *stochastic game* with partial information, a generalization of Markov decision processes where multiple players competitively and jointly control the actions and transitions (see for instance [59, 110]). Very recent work has sought to determine the algorithmic and complexity-theoretic limits of such strategic settings (see for instance [46, 94, 124]). However, our work differs from this line in multiple ways, making the formalism of stochastic games inappropriate for our analysis. In our model, queues are fundamentally unaware of any system parameters—this restriction precludes hope for our stability results to be obtained via general-purpose learning algorithms. Even leaving aside the current gaps in our algorithmic understanding of learning in these settings, our aim is to understand the overall welfare of natural learning algorithms, rather than the limits of individual guarantees of learning. Our results justify the no-regret condition as a desirable, achievable learning in special stateful environments.

Subsequent Work: Following the initial appearance of these results, several works have built on the model and ideas we presented here. On the strategic side, Fu, Hu, and Lin [64] extend the probabilistic techniques we introduced to more general queuing networks once the feasibility conditions under coor-

minated scheduling are suitably adjusted, which naturally arise via duality arguments. Their work also requires variations on the regret and priority notions considered here to enforce stability. Baudin, Scarsini, and Venel [20] propose an alternative, episodic queuing system where agents have incentives to hold jobs in an episode before sending to a central server, but suffer penalties should their jobs not be completed before the end of the episode. Their main conclusion, similar to ours, is that both equilibrium and versions of no-regret dynamics ensure a related form of stability so long as these costs are sufficiently large. However, their model has no underlying stochasticity or heterogeneity among queues, and there is a substantial gap in their positive and negative stability results. On the learning-theoretic side, Sentenac, Boursier, and Perchet [120] as well as Freund, Lykouris, and Weng [62] consider the problem of decentralized learning dynamics in bipartite queuing systems that attain near-optimal performance, extending the original work of centralized learning by Krishnasamy, et al [100]. These algorithmic advances are incomparable with our results, which are inherently non-cooperative and strategic in nature. The former paper also shows that the more refined guarantee of *policy regret* [48] cannot improve the factor of 2 obtained by no-regret learning.

CHAPTER 5
STRATEGIC QUEUING SYSTEMS II: STABILITY OF PATIENT
EQUILIBRIA

In this chapter, we continue our investigation of the stability of strategic queuing systems. In Chapter 4, we showed that any queuing system that would remain centrally feasible even if server rates were halved will be stable under a suitable form of no-regret dynamics; moreover, this factor of $1/2$ was essentially optimal. To our knowledge, this result is one of the only known price of anarchy bounds for natural learning dynamics in stateful environments where the underlying game itself evolves depending on the actions of the agents.

As we will extensively discuss in this chapter, this result is only the beginning of the story of learning and welfare in dynamic environments. While no-regret algorithms clearly do provide useful performance guarantees, we show that the no-regret condition alone can be highly *myopic* in evaluating the performance of their actions in these kinds of queuing systems.

To understand this phenomenon, consider the following queuing system and policies: in this example, there is a unique no-regret policy for each agent, given the behavior of the other agent, *but both agents would have been better off long-term had one even slightly deviated to an inferior server and the other stayed the same*. In the classical setting with “independent” repeated games, this cannot occur.

Example 5.1. *Suppose there are two queues with arrival rates $\lambda_1 = \lambda_2 = 1/2 + \varepsilon$ and two servers with $\mu_1 = 1$ and $\mu_2 = 1/2 - \varepsilon$ for small $\varepsilon > 0$. In this case, each queue receives a new packet slightly more than once every two periods on average. It is possible to*

prove that if both servers deterministically send packets to the first server, the sequence of play will satisfy the no-regret property (as in Assumption 4.2), as they roughly split the top server equally. Each server will clear packets at a rate of $1/2$, strictly better than deviating to the lower server. But this system will not be stable because packets arrive at a total rate of $1 + 2\varepsilon$ while are cleared at a rate of 1 in expectation, leading to linear growth in packets over time almost surely. However, if one queue slightly deviates towards the inferior server with probability $\Theta(\varepsilon)$, both queues will actually be stable.¹

The primary reason for this short-sightedness is because of the long-run dependencies that govern the evolution of these systems via the priority structure. In particular, strategic behaviors that are justified with respect to satisfying the no-regret condition cannot reason about how these actions affect the underlying strategic environment itself. The queue that deviates to the second server in the example above will have regret in the classical sense on shorter timescales, even though this behaviour would have yielded long-run benefits.

The problem of designing algorithms that take into account the effect that actions have on the future strategic environment is known as *multi-agent reinforcement learning* (MARL), a topic of intense study in the past few years. Unlike the no-regret condition, provable guarantees with completely decentralized dynamics are somewhat nascent. Moreover, we stress that in our queuing model, queues only receive very limited feedback: they may have no a priori information about any of the system parameters, and may not even know the identities or number of other queues in the system. From a learning perspective, our queuing model is an extremely challenging form of *partially observable multi-agent reinforcement learning*, making it unlikely that existing general-purpose al-

¹The results obtained in the next two chapters will formally justify this last claim.

gorithms can give improved welfare guarantees.

In this chapter, we precisely quantify the inefficiency that arises due to the myopia of no-regret algorithms. Our approach will be to formulate a *patient queuing game* that isolates this inefficiency due to the inability to reason about long-run outcomes. In our model, queues know all parameters of the system, but must choose a single distribution of choices over servers to be repeated in every round. Their objective will be to minimize their asymptotic long-run growth rate, which depends on the policies chosen by all agents. In the Geometric queuing model, the system will therefore evolve according to a time-homogeneous Markov chain determined by these policies. Using our terminology from the previous chapter, a queue will experience zero cost precisely if their age are weakly stable in this Markov chain.

The intuition behind this formulation of the queuing game is as follows: suppose that the distributions over servers have stabilized for each queue. Then, it must be the case that no queue can asymptotically benefit by altering this distribution

The main result of this chapter, informally stated, is the following analogue of Theorem 4.3 for this patient queuing game:

Consider any queuing system with the property that it would remain stable under optimal, centrally coordinated scheduling even if server qualities are scaled by $\frac{e-1}{e} \approx 0.632$. Then, the original queuing system will remain stable at every Nash equilibrium of the patient queuing game.

The key takeaway of this result is that strictly *fewer* extra resources are required for stability of patient equilibria than for no-regret dynamics. One interpretation

of this gap is as the price of myopia of no-regret learning in queuing systems.

At this stage, the reader may notice several peculiar features of the patient queuing game as currently described. First, it is not immediately clear that the cost functions, which are defined to be asymptotic long-run rates of Markov chains, are even well-defined given the policies of the agents. Even if they are well-defined, this will not be enough to do any serious game-theoretic analysis of these system. We thus provide an equivalent, *algorithmic* characterization of the cost functions. We defer the proof to the end of this chapter: the probabilistic analysis is the most technically involved component of this thesis and readers interested primarily in the game-theoretic analysis may freely skip that section.

With this characterization in hand, we can begin understanding the analytic properties of these cost functions. But now we face a second immediate question: why do there even exist Nash equilibria? We will prove that this is so by appealing to standard fixed-point theorems after proving that the algorithmic cost functions are quasiconvex. By proving the existence of equilibria, our price of anarchy bounds will become non-vacuous. Equally as important, our analysis will reveal fascinating structural properties of these queuing systems that will prove essential in obtaining the improved price of anarchy bound.

The final step is to show the desired bound of $\frac{e}{e-1}$ on the excess server capacity beyond the centrally feasible case is sufficient for stability of equilibria. Our approach to showing this will be to provide a novel *deformation argument*. Specifically, we will show that it is always possible to continuously interpolate any Nash equilibrium to a different configuration of policies such that the long-run growth rate of the fastest growing queues are nondecreasing along this path. This step will require an extremely careful analysis of both the Nash condition

and the underlying combinatorial structure of equilibria, in a sense to be made precise later. We will be able to directly show that this terminal configuration is stable in an appropriate sense, completing the proof of our main result.

5.1 Patient Queuing Games

In this section, we introduce and establish structural results about a patient version of our queuing game. Our notation will be consistent with that of the previous chapter. In this game, each queue chooses a *fixed* distribution over servers that will be used in all rounds to optimize the long-run age of the packets in the queue. To begin, we formulate the game in a manner that is well-defined *a priori*:

Definition 5.1. The **patient queuing game** $\mathcal{G} = ([n], (c_i)_{i=1}^n, \mu, \lambda)$ is defined as follows: the strategy space for each queue is Δ^{m-1} . Let $\mathbf{p} = (p_1, \dots, p_n) \in (\Delta^{m-1})^n$ denote the vector of fixed distributions for all queues over servers. The cost function c_i for queue i is given by

$$c_i(p_i, p_{-i}) \triangleq \limsup_{t \rightarrow \infty} \mathbb{E} \left[\frac{T_t^{(i)}}{t} \right],$$

where $T_t^{(i)}$ is the age of queue i at time t in the random queuing process induced by running a Bernoulli queuing system² with parameters μ and λ and when each queue chooses a server by independently randomizing according to \mathbf{p} at each time t that she has an uncleared packet.

A tuple $\mathbf{p} = (p_1, \dots, p_n)$ is a **Nash equilibrium** of \mathcal{G} if for all $i \in [n]$, $p_i \in \arg \min_{p'_i \in \Delta^{m-1}} c_i(p'_i, p_{-i})$, i.e. each queue chooses p_i to minimize their cost function

²Note that because queue policies are time-independent in this formulation, Lemma 4.5 implies that we may study either the Geometric or Bernoulli system.

conditioned on the strategies p_{-i} of the other queues.

Importantly, this cost function explicitly characterizes the long-run incentives of queues to efficiently clear packets up to first order in time.³ Note that the cost function defined above is clearly well-defined as the limsup of expected values. However, we will actually show that *the limit of the random quantity $T_i^{(i)}/t$ (without expectations) is almost surely equal to a deterministic constant depending on \mathbf{p}, λ, μ* (see Theorem 5.4). By deriving an alternate, explicit characterization of these values, we show that Nash equilibria exist in Theorem 5.11. Because the cost functions are explicit functions of the randomizations and the parameters μ and λ , we omit the notation $(c_i)_{i=1}^n$ when instantiating a game \mathcal{G} .

Our main focus in Section 5.6 will be to give guarantees on the quality of all Nash equilibria in this game. In a slight abuse of the price of anarchy terminology, we make the following definition:

Definition 5.2. Suppose μ majorizes λ . For $\alpha \geq 1$, let $\mathcal{G}(\alpha) = ([n], \mu, \alpha^{-1}\lambda)$, and let $\mathcal{N}(\alpha)$ denote the set of Nash equilibria of $\mathcal{G}(\alpha)$. The **price of anarchy** of $\mathcal{G} = \mathcal{G}(1)$ is defined as the supremum of α values such that there exists a Nash equilibrium $\mathbf{p}^* \in \mathcal{N}(\alpha)$ and some $i \in [n]$ such that $c_i(\mathbf{p}^*) > 0$ in $\mathcal{G}(\alpha)$.⁴ The price of anarchy of the patient queuing game is the supremum over all instances of \mathcal{G} .

³In this formulation, queues have strict preferences over the linear growth rate of their age/number of uncleared packets, but are otherwise indifferent between differences in sub-linear terms.

⁴In words, the price of anarchy of a centrally feasible system is the supremum of values of α such that, when all queue arrival rates are scaled down by α , there nonetheless exists a Nash equilibrium and some queue that suffers nonzero linear aging.

5.2 Main Results

With these preliminaries in order, we may now formally state our main result.

As stated before, our main result is the following:

Theorem 5.2 (Main). *In any patient queuing game \mathcal{G} , there exists a Nash equilibrium.*

Moreover, for any patient queuing game $\mathcal{G} = (\boldsymbol{\mu}, \boldsymbol{\lambda})$, if for all $k \leq n$ it holds that

$$\sum_{j=1}^{\min\{m,k\}} \mu_j \geq \frac{e}{e-1} \sum_{i=1}^k \lambda_i,$$

then for every Nash equilibrium \mathbf{p}^ of \mathcal{G} , the corresponding queuing system is weakly stable.*

In particular, the price of anarchy of the patient queuing game is $\frac{e}{e-1}$.

While the existence of Nash equilibria is not strictly necessary towards proving our price of anarchy bound (which would be vacuously true in this case), this result shows that strategically stable behavior exists in patient queuing games, and moreover, such behavior will also be probabilistically stable in the sense we have defined so long as the system resources slightly exceed what is required for central feasibility. Taken in tandem with our learning-theoretic results in the previous chapter, these results showcase the power of patience: when queues are able to evaluate the long-term effects of their chosen policies, the underlying queuing system requires fewer resources to ensure stability in the presence of strategic behavior.

To establish Theorem 5.2, we require several probabilistic and game-theoretic properties of these queuing systems. A major focus of our work will be in establishing an alternative, algorithmic characterization of the cost functions, which are currently defined abstractly as the long-run aging (or equivalently,

growth by Lemma 4.5) rates. Our main technical result is the following, stated informally:

Theorem 5.3 (Algorithmic Form of Cost Functions). *There exists an explicit, continuous function $r : (\Delta^{m-1})^n \rightarrow \mathbb{R}_{\geq 0}^n$ such that, if queues independently randomize over servers according to $\mathbf{p} \in (\Delta^{m-1})^n$, then the (random) long-run aging rate of each queue i is $r_i(\mathbf{p})$ almost surely. Moreover, this function r admits an algorithmic characterization.*

Proving this result will occupy a large part of this chapter. Let us briefly comment on the implications of this result: Theorem 5.3 states that not only does the long-run aging rate exist almost surely (so that there are no linear oscillations preventing convergence), but that this limiting value is almost surely equal to some deterministic constant that we will be able to compute algorithmically. This step will be absolutely vital towards establishing our main price of anarchy result: it gives us sufficiently strong control over the long-run behaviors of these patient queuing systems to understand their properties at strategic equilibrium. In particular, working with this algorithmic description of the costs, we will be able to establish several *analytic* properties of these games that will all be needed in the proof of Theorem 5.2.

5.3 Overview of Proofs

In this section, we provide a somewhat detailed overview of the arguments used to establish Theorem 5.2.

As stated in the previous section, the first key step in understanding this result is deriving an alternative representation of the cost functions of the pa-

tient queuing games as in Theorem 5.3. We do so by giving the following algorithmic description of the evolution of the queuing system given any choice $\mathbf{p} = (p_1, \dots, p_n)$ of the queues that arises via the following heuristic. Under this choice of policies, it is natural to imagine that some subset S_1 of queues must asymptotically age at the same, fastest rate, and all other queues age at a strictly lower rate.

Under what conditions could such a configuration persist? Because every queue in S_1 has priority over the rest of the queues in the long run, we can now imagine that these queues in S_1 are alone in the system—the priority structure ensures that when any queue in this group collides with any other at a server, they will necessarily have priority due to their asymptotically larger age. In this case, what would their aging rate be?

An simple heuristic is to suppose that $S_1 = \{i^*\}$ is a singleton, and even more conveniently, suppose $p_{i^*} = \delta_{j^*}$ for some $j^* \in [m]$ where δ_j is the j th unit basis vector. In this case, it is not difficult to see that her aging rate should be

$$r_{i^*}(\mathbf{p}) = \min \left\{ 1 - \frac{\mu_{j^*}}{\lambda_{i^*}}, 0 \right\}.$$

The reason is as follows: if i^* is not weakly stable in this configuration, then she must have uncleared packets essentially every round. In a given round t where she has priority over all other queues, her age increases by 1 deterministically due to the passage of time, and then decreases by a $\text{Geom}(\lambda_{i^*})$ random variable on the event that server j^* processed her packet. This event occurs with probability μ_{j^*} and because her $\text{Geom}(\lambda_{i^*})$ age decrement is independent, her expected decrease in age becomes μ_{j^*}/λ_{i^*} . Putting this together with the fact that aging rates are nonnegative, we arrive at the claimed formula.

For a general subset S_1 that would all equally age at a common, fastest rate,

one generalization of this formula suggests that if all queues in S_1 have a packet at time t , their common growth rate ought to be

$$\min \left\{ 1 - \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i \in S_1} (1 - p_{ij}))}{\sum_{i \in S_1} \lambda_i}, 1 \right\}.$$

The numerator is a simple calculation via linearity of expectation of the expected number of packets that the queues in S_1 collectively clear (assuming they all have a packet in a round, as well as priority over all other queues), while the denominator becomes the expected number of packets they receive.

At present, it may not be quite clear why this is the right generalization of the simple calculation for a single queue. Another way to interpret this formula is that if we *enforce* that all queues in S_1 age at the same rate by fiat, this is the quantity that would arise if we could somehow contract S_1 into a single queue with arrival rate $\sum_{i \in S_1} \lambda_i$ that can clear packets at a rate given by the numerator.

Leaving this discussion aside for why this is the correct generalization of the formula, we now discuss what further properties we must impose of S_1 to ensure that this really is the fastest aging subset. For each queue in S_1 to age at this rate, we require that no strict subset of S_1 break off and age at a strictly faster rate if they were to now have priority over all other queues. Moreover, if S_1 has priority over all other queues in the long run, all other queues face reduced service but still must not ever “catch up” to S_1 . In all, we arrive at the proposal that the fastest growing subset should satisfy

$$S_1 = \arg \min_{S \subseteq [n]} \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i \in S} (1 - p_{ij}))}{\sum_{i \in S} \lambda_i}.$$

While the minimizer will be unique for a generic choice of \mathbf{p} , this may not be so at special values like Nash equilibrium. In general, we will show that the set of minimizers forms a *Boolean lattice*: they will be closed under (non-disjoint)

union and intersection using the structural form of the objective function. It will be convenient to define S_1 to be the subset with largest cardinality, which will also be the union of all minimizers. We will call any subset attaining the minimum a **tight** subset.

To extend this analysis, we reason inductively: under the ansatz that the queues in S_1 will be linearly older than every queue not in S_1 , the remaining queues will effectively face a queuing system where server capacities are discounted by the probability any queue in S_1 sends a packet there. We may therefore do precisely the same reasoning in this modified queuing system to determine the next fastest aging subset, and so on. Eventually, we arrive at a partitioning of the queues $S_1 \sqcup S_2 \sqcup \dots \sqcup S_k$ such that the aging rates are equal within each S_ℓ and are strictly decreasing in ℓ (if we use the convention of taking S_ℓ to be the maximal fastest aging subset at stage ℓ of this recursion).

Formally establishing these relations in Theorem 5.3 will require rather intense probabilistic reasoning, deferred to the end of this chapter. The intuition is as follows: a simple first moment argument will show that an appropriate *weighted average* of the ages of queues in S_1 must grow at least at the rate given above. A much more difficult task is showing that no queue in $[n]$ grows at a strictly larger rate. This argument will require a very careful induction on the number of very old queues. If a single queue is much older than the rest, this analysis will be straightforward using the minimality of S_1 in the above algorithmic description and the priority structure to argue this queue must clear sufficiently many packets with high probability. A very careful induction on the size of the old subset will be required to extend this analysis to arbitrary subsets where the internal priority structure can evolve rapidly. Once we do so, we

will have shown that a weighted average of the queue ages and the maximum queue age among S_1 grows at the same rate. We can then immediately conclude that *all* queues in S_1 age at this common rate. Conditioning on this event and reasoning recursively on the rest of the system concludes the proof.

Once we establish this characterization, we may begin to analyze the properties of the cost functions. The same considerations that lead to the lattice property above will enable us to show continuity of the cost functions. To establish the existence of Nash equilibria, we will rely on somewhat standard fixed point theorems from the game theory literature. However, it turns out that fixing the strategies p_{-i} of all queues but i , it will *not* be true that $r_i(\cdot, p_{-i})$ is convex. Instead, it will be true that there are no *local maxima* on the interior of any line segment.⁵ This allows us to apply the Debreu-Glicksberg-Fan Fixed-Point Theorem to conclude the existence of pure equilibria.

Finally, we will turn to the bulk of the argument in Theorem 5.2 of showing the price of anarchy bound. The basic idea is, of course, to use the Nash condition to argue about the quality of the solution. Our approach will be to take an arbitrary Nash equilibrium \mathbf{p}^* and assume that the aging rate of S_1 is strictly positive and that the $\frac{e}{e-1}$ -majorization condition holds. We will show that in this case, we can continuously deform the Nash profile \mathbf{p}^* to a carefully constructed profile \mathbf{p}^{**} in a way such that the calculated rate for S_1 can only increase along this path. This final configuration will be chosen so that the *sum* of probabilities $\sum_{i \in S_1} p_{ij}^{**}$ are balanced in an appropriate sense across $j \in [m]$ and then applying standard inequalities to lower bound the quality.

The key challenge is then to argue monotonicity along this path. Several

⁵Formally, this condition is known as *quasiconvexity*.

difficulties arise in the course of proving this result. The first is that even restricting to S_1 , queues may belong to several tight subsets which each constrain their costs in different ways. In other words, the local optimality of Nash for a given queue i may arise from several different constraints, making it difficult to connect their incentives at \mathbf{p}^* to the rate of S_1 rather than just some tight proper subset containing i —we need to do so to argue about the effect of our interpolation. This problem is exacerbated when we try to deform the policies of multiple queues, as the intermediate profiles will not even be Nash.

Our idea in resolving this is to leverage the lattice structure of the tight subsets to find a **level decomposition** of S_1 . While all queues in S_1 age at the fastest rate, one can essentially view this level decomposition as a graded version of the Hasse diagram of this lattice under inclusion, where the gradation is given by the distance to \emptyset in an appropriate sense. The “level-1” subsets are *minimal* tight subsets, while the “level-2” subsets are minimal supersets that also age at this rate because given that the “level-1” subsets age at this maximal rate, the added queues in these subsets must also age at this rate, and so on. The level of a queue will be defined as the minimal level of a tight subset it is contained in. This refinement of S_1 enable us to schedule the deformations in a precise way by descending level. These queues belong to the fewest tight subsets, so their incentives will be easier to reason about to show monotonicity. At each subsequent level, we will be able to argue that if the deformation actually lowered the rate achieved by S_1 , then the same deformation will also (at least locally) lower the rate for every subset containing the queue at the original Nash \mathbf{p}^* , a violation of the Nash condition. We conclude that these deformations can only increase the aging rate of S_1 , as was needed.

5.4 Algorithmic Formulation of Costs

In this section, we begin our study of the cost functions $c(\mathbf{p})$ as described in the previous section, which are currently written as the lim sup of the expected value of the random linear aging rate of each queue. These cost functions are well-defined but quite unwieldy at present. Our first task is thus to provide an alternative, algorithmic description of $c(\mathbf{p})$, which we initially denote $r(\mathbf{p})$ (for “rates”) in Section 5.4. We show that r has significant analytic structure that will help establish various game-theoretic properties of this system. In particular, we show that the level subsets (in $[n]$) of $r(\mathbf{p})$ enjoy convenient closure properties, which will be enough to establish continuity and other properties, which we use to prove the existence of equilibria. We will return to proving that this function is equal to c in Section 5.7.

As stated, we now construct a function $r : (\Delta^{m-1})^n \rightarrow [0, 1]$ that we will show is equivalent to c . We will show that for any fixed \mathbf{p} , the set $[n]$ of queues partitions into subsets S_1, S_2, \dots , where each queue in S_i group has the same aging rate and S_1 ages the fastest, then S_2 , etc, according to r (and so for c as well). To get a sense of the quantities that will arise before considering the general case, consider the simplest setting of a single queue and a single server (where there are no nontrivial strategies nor competition), with rates $\lambda > \mu$. In any round where the queue has an uncleared packet, the age will first increase by 1 deterministically. With probability μ , the queue will succeed in clearing this packet, and the age will go down in expectation by $\mathbb{E}[G] = 1/\lambda$, where $G \sim \text{Geom}(\lambda)$ is independent of whether or not the server succeeds. Therefore, the expected change in this queue’s age will be $1 - \mu/\lambda > 0$, and we expect that the queue will asymptotically age at this rate.

In general, with multiple queues and servers, the actual values of c_i are best described via a recursive algorithm (Algorithm 2) that computes the rates. See Section 5.3 for intuition on where this construction comes from.

Algorithm 2: Algorithmic Rates of Patient Queuing

Input: Queue arrival rates λ , server rates μ , strategies $\mathbf{p}_1, \dots, \mathbf{p}_n$ over servers

```

1 Initialize  $k = 1$  and  $I = [n]$ ;
2 while  $I \neq \emptyset$  do
3   Compute
      
$$\min_{S \subseteq I} \left\{ \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i \in S} (1 - p_{i,j}))}{\sum_{i \in S} \lambda_i} \right\}. \quad (5.1)$$

      ; // Expected jobs cleared by  $S$  over total arrival rate.
4   if (5.1)  $\geq 1$  then
5     Set  $S_k \leftarrow I$ ,  $r_i(\mathbf{p}) \leftarrow 0$  for all  $i \in S_k$ , and  $I \leftarrow \emptyset$ ;
6   else
7     Let  $S_k$  be minimizer of Equation (5.1) of largest cardinality;
8     for  $i \in S_k$  do
9       Set
          
$$r_i(\mathbf{p}) \leftarrow 1 - \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i \in S_k} (1 - p_{i,j}))}{\sum_{i \in S_k} \lambda_i}.$$

          For  $k = 1$ , we refer to any subset with the minimum ratio as a
          tight, or minimizing, subset;
10    Update  $\mu_j \leftarrow \mu_j \prod_{i \in S_k} (1 - p_{i,j})$  for  $j \in [m]$ ; //  $\mu_j$  discounted by
          probability queue from  $S_k$  sends to  $j$ 
11    Update  $I \leftarrow I \setminus S_k$ ,  $k \leftarrow k + 1$ , and recurse on  $I$  with  $\mu$  and  $\mathbf{p}_I$ .

```

As many of these quantities will appear often, we make the following con-

ventions: for any subsets S, S' such that $S \subseteq [n] \setminus S'$, define $\lambda(S)$ as the sum of arrival rates of packets to a set of queues S , and $\alpha(S|\mathbf{p}, \boldsymbol{\mu}, S')$ as the expected number of packets cleared from queues in S with service rates $\boldsymbol{\mu}$, if the queues in S' have priority, S has priority over all other queues, and all queues in $S \cup S'$ send packets in the round. That is, define

$$\alpha(S|\mathbf{p}, \boldsymbol{\mu}, S') \triangleq \sum_{j=1}^m \mu_j \prod_{i \in S'} (1 - p_{i,j}) (1 - \prod_{i \in S} (1 - p_{i,j})), \quad \lambda(S) \triangleq \sum_{i \in S} \lambda_i,$$

and then let

$$f(S|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda}, S') \triangleq \frac{\sum_{j=1}^m \mu_j \prod_{i \in S'} (1 - p_{i,j}) (1 - \prod_{i \in S} (1 - p_{i,j}))}{\sum_{i \in S} \lambda_i} = \frac{\alpha(S|\mathbf{p}, \boldsymbol{\mu}, S')}{\lambda(S)},$$

denote the ratio of expected number of packets cleared by S when having priority over all members but S' , normalized by the expected number of new packets received in each round by S .

Let $S_k(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})$ be the k th set output by the above algorithm. When $\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda}$ are clear from context, we will suppress them. We write $U_k = \cup_{\ell=1}^k S_\ell$ as the set of queues in the top k groups outputted by the algorithm, with $U_0 = \emptyset$. We will write $f_k = f(S_k|U_{k-1})$, and we use $g_k = \max\{0, 1 - f_k\}$ for the rate of the k th outputted set, which is equal to $r_i(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})$ for any $i \in S_k(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})$. From the recursive construction,

$$S_{k+1}(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda}) = S_1(\mathbf{p}_{[n] \setminus U_k}, \boldsymbol{\mu}', \boldsymbol{\lambda}_{[n] \setminus U_k}) \quad (5.2)$$

where $\mu'_j = \mu_j \prod_{i \in U_k(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})} (1 - p_{ij})$ for all $j \in [m]$. In words, having found U_k , S_{k+1} is the largest minimal set among the remaining elements, but where the $\boldsymbol{\mu}$ rates have been reweighed by the probability no element of U_k sends to each server. These quantities are compiled in a table in Section 5.8 for easy reference.

Our main probabilistic result about the function r is that this is indeed equivalent to the cost function c of the patient queuing games. In fact, we prove the following, stronger result:

Theorem 5.4 (Almost Sure Asymptotic Convergence). *Let $\mathcal{G} = ([n], \mu, \lambda)$ be a patient queuing game. For each fixed \mathbf{p} and all $i \in [n]$, almost surely the long-run aging rate of queue i satisfies*

$$c_i(\mathbf{p}) = \lim_{t \rightarrow \infty} \frac{T_t^{(i)}}{t} = r_i(\mathbf{p}).$$

We will prove our game-theoretic results assuming this theorem; however, as the proof is quite nontrivial and rather involved, we defer the proof to Section 5.7.

5.5 Properties of Rate Functions

In this section, we prove various properties about the rate functions given in the previous section that will be crucial in the analyses to come.

We repeatedly use the following fact, which can be seen simply by expanding the definition of f (see Section 2.1 for notation):

Fact 5.5. *Suppose S, S', T are such that $S, S' \subseteq [n] \setminus T$ and are disjoint. Writing f in the form of the quotient α/λ , then*

$$f(S \cup S'|T) = f(S|T) \oplus f(S'|S \cup T).$$

Throughout this section, we will view f as the quotient α/λ when invoking Fact 5.5.

5.5.1 Analytic Properties

First, we characterize some structure in the minimizing subsets at each step of the algorithm. This structure will allow us to choose the S_k canonically as the largest cardinality minimizer. To do this, we first show that the function $\alpha(\cdot)$ is *submodular*:

Lemma 5.6 (Submodularity). *For fixed S' , \mathbf{p} , $\boldsymbol{\mu}$, $\boldsymbol{\lambda}$ the function $\alpha(S|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda}, S')$ is submodular in S , i.e. for any $S, T \subseteq [n] \setminus S'$, $\alpha(S \cap T|S') + \alpha(S \cup T|S') \leq \alpha(S|S') + \alpha(T|S')$.*

Proof. Fix S' , \mathbf{p} , $\boldsymbol{\mu}$, $\boldsymbol{\lambda}$, and we suppress the dependence in α . To prove submodularity, recall that an equivalent definition is that for any S, T satisfying $S \subseteq T$, and $i \notin T$, then $\alpha(S \cup \{i\}) - \alpha(S) \geq \alpha(T \cup \{i\}) - \alpha(T)$ [118]. Let $S \subseteq T$ and $i \notin T$. A simple computation shows that for any $V \subseteq [n] \setminus S'$,

$$\alpha(V \cup \{i\}) - \alpha(V) = \sum_{j=1}^m \mu'_j p_{ij} \prod_{k \in V} (1 - p_{kj}),$$

where μ'_j is $\mu_j \cdot \prod_{i \in S'} (1 - p_{ij})$ from priority. As each factor in the product is at most 1, this is clearly decreasing in V as a set function, establishing submodularity. \square

Now, recall that the relevant functions in the construction of the above algorithm is the set function $f = \alpha/\lambda$. As a consequence of the fact that this function is the ratio of a submodular function with a modular function, we will be able to gain significant closure properties of the tight subsets (as defined above), which will end up being critical in establishing both game-theoretic and probabilistic properties of our systems.

Lemma 5.7 (Closure). *For each fixed \mathbf{p} and $k \geq 1$, the set of minimizers of $f(\cdot|U_{k-1})$ in $\mathcal{P}([n] \setminus U_{k-1})$ are closed under union and non-disjoint intersection; that is, if $S, S' \subseteq$*

$[n] \setminus U_{k-1}$ are minimizers, then so is $S \cup S'$, as well as $S \cap S'$ if nonempty. Moreover, if $S \cap S'$ is empty, then the queues in S and S' must send to disjoint subsets of servers.

In particular, the minimizing set with largest cardinality is unique, and is the union of all minimizing sets at step k . If S is considered at step k of the algorithm, but S is not a subset of S_k , then $f(S|U_{k-1}) > f(S_k|U_{k-1})$.

Proof. The last statement is an immediate consequence of the first, so we focus on the first part. We show this just for $k = 1$; the general case for $k > 1$ follows from the recurrence in Equation (5.2).

By Lemma 5.6, α is submodular, and it is immediate that $\lambda(\cdot)$ is a *modular* function, i.e. $\lambda(S \cup S') + \lambda(S \cap S') = \lambda(S) + \lambda(S')$. We claim that if a function f is a ratio of a nonnegative submodular function and a nonnegative supermodular function, then the set of minimizers of f is closed under union and non-disjoint intersection. To see this, suppose S, S' are minimizers. Then we see the following inequalities, writing f always in the form of the quotient α/λ , and using Fact 5.5 and Fact 2.1,

$$\max\{f(S), f(S')\} \geq f(S) \oplus f(S') \geq f(S \cap S') \oplus f(S \cup S') \geq \min\{f(S \cap S'), f(S \cup S')\}$$

(where we omit $S \cap S'$ if it is empty), as the inequalities in the numerator and denominator go the correct way by Lemma 5.6, and then using Fact 2.1. But as S, S' were minimizers, these must be equalities, which occurs iff $S \cap S'$ (if nonempty) and $S \cup S'$ are both minimizers. As $f = \alpha/\lambda$ here, this applies for our functions.

If S, S' are both minimizers and are disjoint, then from Fact 5.5, it follows that $f(S \cup S') = f(S) \oplus f(S'|S)$. As $S, S', S \cup S'$ are evidently minimal, this equation

can only hold if $f(S'|S) = f(S)$, which occurs if and only if S and S' disjointly mix among servers. \square

From Lemma 5.7, it will nearly immediately follow that the outputted rates are strictly monotonic decreasing in the groups: as mentioned, $[n] = S_1 \sqcup S_2 \sqcup \dots$ is meant to give a partition into groups that age together, where S_1 is the fastest aging group, S_2 the next fastest, etc. As such, the disjoint subsets iteratively output by the algorithm satisfy the intuition that motivates the construction.

Lemma 5.8 (Monotonicity). *Let S_1, S_2, \dots be the outputs of the algorithm in order. Then $g_k > g_{k+1}$ for each $k \geq 1$.*

Proof. Consider $S_k \cup S_{k+1}$: this set is considered at stage k of the algorithm, and evidently was not selected. From the previous lemma, we must have

$$f(S_k|U_{k-1}) < f(S_k \cup S_{k+1}|U_{k-1}) < f(S_{k+1}|U_{k-1} \cup S_k) = f(S_{k+1}|U_k).$$

The first inequality follows from the selection criteria of the algorithm (and maximality of S_k), while the second follows from writing (via Fact 5.5)

$$f(S_k \cup S_{k+1}|U_{k-1}) = f(S_k|U_{k-1}) \oplus f(S_{k+1}|U_{k-1} \cup S_k).$$

As we have already proven the first inequality, the second must follow from Fact 2.1. This yields the claim. \square

With these basic properties, we can show the intuitive, but rather technical, fact that the algorithm given above is continuous in \mathbf{p} . We defer the proof to Section 5.9.1:

Proposition 5.9 (Continuity). *The function $r : (\Delta^{m-1})^n \rightarrow [0, 1]^n$ given by $r(\mathbf{p}) = (r_1(\mathbf{p}), \dots, r_n(\mathbf{p}))$ is continuous.*

5.5.2 Game-Theoretic Properties

With these structural results, we can turn to showing our first game-theoretic property of this game, for now assuming that the costs are given by r , the output of the algorithm of Section 5.4: namely, that equilibria exist. While the cost functions are not quite convex, by restricting each component to a line that varies only a single queue's strategy, one can deduce enough structure that allows for an application of Kakutani's Theorem.

Fix any queue i , as well as any fixed probability choices $p_{-i} \in (\Delta^{m-1})^{n-1}$ by the other players, and any two $p, p' \in \Delta^{m-1}$. Define for $t \in [0, 1]$

$$h(t) = r_i(tp + (1 - t)p', p_{-i}).$$

Then we have the following lemma, whose proof is deferred to Section 5.9.1:

Lemma 5.10. *For any fixed i , $p_{-i} \in (\Delta^{m-1})^{n-1}$, and $p_i, p'_i \in \Delta^{m-1}$, the function $h(t)$ is piecewise linear and has no local maxima on the interior. In particular, $r_i(\cdot, p_{-i})$ is quasiconvex.*

Proof. Let $\mathbf{p}(t) = (tp_i + (1 - t)p'_i, p_{-i})$. By Proposition 5.9, h is continuous as the restriction of a continuous function, and is easy to see it must be piecewise linear in t by inspection. Indeed, as the algorithmic description of c takes minimums and maximums of finitely many linear functions, this yields a piecewise linear function with no jump discontinuities.

We now prove the last claim. It is sufficient to show that if h is increasing at t' , then it is increasing for all $t'' > t'$. Suppose that this is violated for some $t' < t''$; by piecewise linearity, there must exist some t^* such that as $t' < t^* < t''$ where

two lines intersect in the graph, and so that as $t \rightarrow t^{*-}$, the slope is increasing while it is non-increasing for $t \rightarrow t^{*+}$.

Suppose that for all t that are sufficiently close to t^* from the left, i is outputted at step k of the algorithm. The only reason the slope can go from positive to nonpositive at t^* is there is a change in which sets are outputted in the algorithm at some step $\ell \leq k$, which can happen only if some new set S including i gets selected for $t \geq t^*$. However, as the rates of all sets not including i fixing any other disjoint set having priority are all constants with respect to t , this can only occur because at t^* , some linear function $f(S|\mathbf{p}(t), S')$ went below the $f(S''|\mathbf{p}(t), S')$ that was previously selected at step ℓ , where S' is the union of all sets outputted prior at t and S'' is the set that was outputted next for all t close enough to the left of t^* . If S'' included i , this can only occur if $r(S|\mathbf{p}(t), S') := \max\{0, 1 - f(S|\mathbf{p}(t), S')\} = 1 - f(S|\mathbf{p}(t), S')$ has larger positive slope than $r(S''|\mathbf{p}(t), S')$, so the slope of h would be strictly larger (and in particular, increasing) for all t sufficiently close to t^* on the right, contradicting our assumption that it is non-increasing. If S'' does not include i , $r(S''|\mathbf{p}(t), S')$ is a constant with respect to t , so for $r(S|\mathbf{p}(t), S')$ to exceed it for t larger than t^* but be lower for t less than t^* , the slope of $r(S|\mathbf{p}(t), S')$ must also be positive, another contradiction. Both cases lead to a contradiction, proving the claim.

The definition of quasiconvexity is equivalent to establishing no local maxima along the restriction to a segment can exist in the interior, as quasiconvexity is a two-dimensional notion. This is precisely what was achieved above. \square

The existence of equilibria is now an immediate consequence of our previously attained results:

Theorem 5.11 (Existence of Equilibria). *There exists an equilibrium of the game with*

costs given by $r : (\Delta^{m-1})^n \rightarrow [0, 1]^n$.

Proof. Recall the following statement of the Debreu-Glicksberg-Fan fixed point theorem (Theorem 3.3) for any strategic-form game with n agents, convex and compact action spaces I_1, \dots, I_n , and cost functions c_1, \dots, c_n such that $c_\ell(\cdot, p_{-\ell})$ is continuous and quasiconvex on I_ℓ for each fixed $p_{-\ell}$ and $\ell \in [n]$, there exists a pure Nash equilibrium \mathbf{p}^* . The theorem is now an immediate consequence of Theorem 5.4, Proposition 5.9 and Lemma 5.10, noting that the action spaces in the patient queuing game are clearly convex and compact. \square

5.6 Proof of Main Result

In this section, we turn to the game-theoretic problem of understanding what condition ensures the stability at any equilibrium profile assuming Theorem 5.4—we return to proving that this is valid in the next Chapter.

Concretely, we prove the following instance-dependent bound from which the claimed factor immediately follows:

Theorem 5.12 (Main). *Let \mathbf{p} be any Nash equilibrium of \mathcal{G} , and let S_1 be as defined before. Then⁶*

$$f(S_1|\mathbf{p}) \geq \min \left\{ 1, \min_{k \leq n} \max_{\mathbf{x} \in \mathbb{R}_{\geq 0}^m : \sum_{j=1}^m x_j = k} \frac{\sum_{j=1}^m \mu_j (1 - (1 - x_j/k)^k)}{\sum_{i=1}^k \lambda_i} \right\}.$$

⁶For any fixed k , it is not difficult to determine the optimal value of \mathbf{x} to give the tightest lower bound. It suffices to maximize the numerator, which is concave. By standard KKT conditions at optimality, for all $j, j' \in [m]$ such that $x_j > 0$, we must have $\mu_j (1 - x_j/k)^{k-1} = \mu_{j'} (1 - x_{j'}/k)^{k-1}$, and $x_\ell = 0$ for all lower indices.

Corollary 5.13 (Theorem 5.2, restated). *Let \mathbf{p} be a Nash equilibrium of \mathcal{G} , and suppose that for each $1 \leq k \leq n$,*

$$\sum_{j=1}^{\min\{k,m\}} \mu_j \geq \left(\frac{e}{e-1}\right) \sum_{i=1}^k \lambda_i.$$

Then every queue is weakly stable at \mathbf{p} . In particular, the price of anarchy of the patient queuing game is exactly $\frac{e}{e-1}$.

Proof. In Theorem 5.12, for any $k \leq n$, we may set $x_j = 1$ for $1 \leq j \leq k$. Note that $(1 - x_j/k)^k < e^{-1}$, so if $\boldsymbol{\mu}$ majorizes $\boldsymbol{\lambda}$ by a factor of at least $\frac{e}{e-1}$, Theorem 5.12 implies that $f(S_1|\mathbf{p}) \geq 1$. From Theorem 5.4 and Lemma 5.8, we conclude that all queues are weakly stable. \square

The following simple example shows that this is the best possible constant factor: fix $\epsilon > 0$ small and suppose there are n queues and n servers, with $\boldsymbol{\lambda} = (1 - 1/e + \epsilon, \dots, 1 - 1/e + \epsilon)$ and $\boldsymbol{\mu} = (1, \dots, 1)$, and \mathbf{p} has every queue uniformly mixing among the servers. It is easy to see by symmetry that this system is Nash with $S_1 = [n]$, for if a queue deviates from this uniform distribution, this does not change the worst ratio in the algorithm. Moreover, for any fixed $\epsilon > 0$, as $n \rightarrow \infty$, this system becomes unstable. One can check that

$$f(S_1|\mathbf{p}) = f([n]|\mathbf{p}) = \frac{\sum_{j=1}^n (1 - \prod_{i=1}^n (1 - 1/n))}{n(1 - 1/e + \epsilon)} \rightarrow \frac{1 - 1/e}{1 - 1/e + \epsilon} < 1,$$

so that $r(S_1) = \max_i c_i(\mathbf{p}) > 0$.

We now prepare for the proof of Theorem 5.12. The idea will be to continuously deform the Nash profile towards a highly symmetrized strategy vector while only weakly decreasing $f(S_1)$. At the end of this process, we obtain a lower bound on this value at Nash. To do this deformation and ensure monotonicity of the growth rate, we must at some point use the Nash property. The




Figure 5.1: Four queues compete for three servers, with maximally tight sets marked. The two outer queues exclusively send to the outer servers, and the inner queues send to two servers each, as indicated by the figure. One can verify that this profile is Nash and the subsets marked are tight. Consider whether one of the inner queues would want to deviate from the current profile. If she shifts probability to the inner server, the smallest tight set will rise in aging rate. However, if she instead attempts to deviate to one of the outer servers, the rate of all four queues will rise.

difficulty lies in the form of the f functions; recall that as S_1 is the set of all queues growing at the fastest rate as the union of all tight subsets, it can have many proper tight subsets, and each queue $i \in S_1$ thus has to locally optimize all of the functions $f(S|\mathbf{p})$ with $S \ni i$ simultaneously at Nash (see Figure 5.1 for an interesting example). In particular, if queue $i \in S_1$ at Nash, one possible deviation may weakly decrease $f(S)$ for some tight subset $S \ni i$, while another deviation may weakly decrease $f(S')$ for some *different* tight subset $S' \ni i$. In other words, each queue may be constrained by multiple different objective functions at Nash, making it difficult to generically argue about *why* any given deviation decreases performance. We overcome this barrier via Proposition 5.14 by connecting the incentives for each queue in S_1 with the structure guaranteed by Lemma 5.7:

Proposition 5.14. *Let \mathbf{p} be any arbitrary strategy vector by the queues, and without loss of generality, let $[k]$ be the maximal tight subset after relabeling. Then, for some $s > 0$, there exists a **level partition** of $[k]$ into s levels with the following property: if a queue $i \in [k]$ belongs to a level- ℓ subset, then for any deviation by i that shifts probability mass from one server to another and does not increase $f(S)$ for some tight subset $S \ni i$, there exists a tight subset S' containing all queues at all levels $j \leq \ell$ such that $f(S')$ must not increase.*

Proof. We construct the desired decomposition of $[k]$ by instead recursively considering *minimal* tight subsets. In particular, let the level-1 subsets be the set of all minimal tight subsets (notice that by Lemma 5.7 and minimality, each subset of queues must be disjoint and disjointly mixing). Note that to clear packets at the rate of these subsets, they all must have priority over any other queues. Then given the level- j subsets for all $j \leq \ell$ for some $\ell \geq 1$, we recursively define the level- $(\ell + 1)$ subsets as the set of minimal tight subsets of the remaining queues conditioned on all subsets at lower levels having priority. The same argument as used by Lemma 5.7 implies that these subsets at each level must all be disjoint and disjointly mixing.⁷ It is easy to verify by iteratively using Fact 5.5 that for any level- ℓ , the union of all subsets at levels $j < \ell$ with any subset of the level- ℓ subsets must be tight. As $[k]$ is tight and any tight subset must be contained in $[k]$ by the maximality guaranteed in Lemma 5.7, it is immediate that this decomposition exhausts $[k]$.

Let $i \in [k]$ belong to some level- ℓ subset, for some $\ell \geq 1$. Suppose that i has some deviation from server j to another server j' that causes $f(S)$ to not increase

⁷In the example given before of a system with multiple tight subsets, the level-1 subset is the subset of two queues that split between the inner and outer servers. The level-2 subsets are the two singleton sets of the outer queues sending each to their own outer server. Notice that these two subsets indeed disjointly mix.

for some tight subset $S \subseteq [k]$. As $f(S)$ is a linear function in the randomizations by queue i holding each other queue fixed, this must hold for any sufficiently small deviation as well. By expanding the partial derivatives $\frac{\partial f(S)}{\partial p_{ij}}$ and $\frac{\partial f(S)}{\partial p_{i'j}}$ and clearing denominators, this holds if and only if

$$\mu_j \prod_{r \in S \setminus \{i\}} (1 - p_{rj}) \geq \mu_{j'} \prod_{r \in S \setminus \{i\}} (1 - p_{rj'}) \quad (5.3)$$

Define $V_{\leq \ell}^i$ to be the set of all queues at level at most ℓ that do not belong to the level- ℓ subset i belongs to. Then $V_{\leq \ell}^i$ is tight, and by construction $S \setminus V_{\leq \ell}^i$ is nonempty as i belongs to this set. As tight subsets are preserved under unions and intersections, we have $S \cup V_{\leq \ell}^i$ and $S \cap V_{\leq \ell}^i$ are tight. On the one hand, we have by Fact 5.5 that

$$f(S) = f(S \cap V_{\leq \ell}^i) \oplus f(S \setminus V_{\leq \ell}^i | S \cap V_{\leq \ell}^i).$$

On the other, we have

$$f(S \cup V_{\leq \ell}^i) = f(V_{\leq \ell}^i) \oplus f(S \setminus V_{\leq \ell}^i | V_{\leq \ell}^i).$$

Comparing these two expressions, the left-hand sides are minimal, as well as the first terms on the right, which by Fact 2.1 implies the latter terms are equal. But we know from definition that

$$f(S \setminus V_{\leq \ell}^i | S \cap V_{\leq \ell}^i) \geq f(S \setminus V_{\leq \ell}^i | V_{\leq \ell}^i),$$

with equality if and only if $V_{\leq \ell}^i \setminus S$ disjointly mixes from $S \setminus V_{\leq \ell}^i$ in \mathbf{p} . As i belongs to this set, we can combine this with Equation (5.3) to see

$$\mu_j \prod_{r \in (V_{\leq \ell}^i \cup S) \setminus \{i\}} (1 - p_{rj}) = \mu_j \prod_{r \in S \setminus \{i\}} (1 - p_{rj}) \geq \mu_{j'} \prod_{r \in S \setminus \{i\}} (1 - p_{rj'}) \geq \mu_{j'} \prod_{r \in (V_{\leq \ell}^i \cup S) \setminus \{i\}} (1 - p_{rj'}).$$

This implies that $f(S \cup V_{\leq \ell}^i)$ must also not increase from this deviation. The last thing to check is that the level- ℓ subset i belongs to is contained in S , so that

$S \cup V_{\leq \ell}^i$ contains all queues up to level ℓ . But this follows because $S \setminus V_{\leq \ell}^i$ consists of queues at level at least ℓ and intersects this level- ℓ subset nontrivially at i , and therefore must contain it by minimality in our construction. Therefore, for this choice of deviation by queue i , the subset $S \cup V_{\leq \ell}^i$ certifies the desired claim. \square

With this result, we may finally return to the proof of Theorem 5.12.

Proof of Theorem 5.12. Let \mathbf{p} be any Nash equilibrium, and suppose that S_1 is the maximal tight subset. If $f(S_1) \geq 1$, then we are done, so suppose that $f(S_1) < 1$. The Nash assumption then implies that any deviation by a queue in S_1 cannot decrease the value of each tight subset it is part of; note that we need the $f(S_1) < 1$ assumption, as incentives are about the rates, and when $f(S_1) \geq 1$, the rate may remain 0 even if f decreases.

For convenience, reindex and relabel so that $|S_1| = k$ and $S_1 = [k]$. Fix any $\mathbf{x} \in \mathbb{R}_{\geq 0}^m$ such that $\sum_{i=1}^m x_i = k$. It suffices to show that

$$f([k]|\mathbf{p}) \geq \frac{\sum_{j=1}^m \mu_j (1 - (1 - x_j/k)^k)}{\sum_{i=1}^k \lambda_i}.$$

From now on, we omit the dependence on \mathbf{p} in f unless explicitly needed.

Consider the level partition of $[k]$ guaranteed by Proposition 5.14 and suppose that there are s levels. We continuously deform the Nash solution while monotonically decreasing $f([k])$, so that at the end of this process, we have a lower bound on the value of Nash. Given any strategy profile \mathbf{p} , we say a server j is *oversaturated* if $\sum_{i \in [k]} p_{ij} > x_j$ and *undersaturated* if $\sum_{i \in [k]} p_{ij} < x_j$. We will continuously move probability mass from the queues from oversaturated to undersaturated servers. If no server is oversaturated, we will be done; notice that if a

server is oversaturated, an easy averaging argument implies some server must be undersaturated.

Suppose that there exists some oversaturated server. Let $i \in [k]$ be a queue at level- s , the top level. If i nontrivially sends to an oversaturated server j so that $p_{ij} > 0$, we continuously decrease p_{ij} and increase $p_{ij'}$ for some undersaturated server j' , until either j stops being oversaturated, j' stops being undersaturated, or p_{ij} hits zero. We claim that this deformation cannot increase $f([k])$. To see this, observe that because \mathbf{p} is Nash, we know that any deviation by i from one server it is nontrivially mixing at to another cannot increase $f(S)$ for all tight subsets S containing i , hence there must be some $S \ni i$ such that $f(S)$ does not increase. But then Proposition 5.14 implies that some subset S' containing all queues up to level- s must have $f(S')$ not increase either. As $[k]$ is the only such subset, this deformation could not have actually increased $f([k])$.

Moreover, we claim that we can do this for all level- s queues one-by-one without increasing $f([k])$. While the intermediate profiles are not Nash, because we only move probability mass from oversaturated to undersaturated servers, each oversaturated queue only has at most the same probability mass as it did at Nash while each undersaturated queue only has additional probability mass compared to what it had at Nash. As we have shown any such deviation by a level- s server from an oversaturated queue to an undersaturated queue at Nash cannot increase $f([k])$, and now deviations are only worse at this intermediate stage while deforming the level- k queue strategies, each such deformation still cannot increase $f([k])$. Therefore, we can continuously shift all probability mass from level- s queues at oversaturated servers to undersaturated servers while never increasing $f([k])$.

Suppose we have now done this for all levels at least $\ell + 1$ for some $\ell < s$ while not increasing $f([k])$, and we want to continue this process at level- ℓ . Let \mathbf{p}' be this intermediate strategy vector, where we note that for any queue i below level $\ell + 1$, $\mathbf{p}'_i = \mathbf{p}_i$. Again, if no server is oversaturated, we are done. Otherwise, suppose some queue i at level- ℓ still sends to an oversaturated server j , and we again try to decrease p'_{ij} and increase $p'_{ij'}$ for some undersaturated server j' as before until the same stopping criterion. We must show that this too cannot increase $f([k])$.

Suppose otherwise that it did indeed increase $f([k])$ with respect to \mathbf{p}' . For a contradiction, it suffices to show that this implies that this same deviation, with respect to the original Nash solution \mathbf{p} , must have increased $f(S)$ for every subset S containing all queues up to level- ℓ . This is sufficient to obtain a contradiction as then Proposition 5.14 implies that *every* subset containing i improves at Nash with respect to this deviation, which violates the Nash property.

To prove this claim, let $S \subseteq [k]$ be any arbitrary tight subset at Nash containing all queues up to level- ℓ . Because we assume that this deviation improves $f([k])$ with respect to \mathbf{p}' , it follows from taking partial derivatives that

$$\mu_j \prod_{r \in [k] \setminus \{i\}} (1 - p'_{rj}) < \mu_{j'} \prod_{r \in [k] \setminus \{i\}} (1 - p'_{rj'}).$$

However, note that at \mathbf{p}' , as j is still oversaturated, $p'_{rj} = 0$ for all queues r that are at strictly higher levels. As all queues at level- ℓ and below have $\mathbf{p}'_i = \mathbf{p}_i$ and S contains all such queues, this inequality implies

$$\mu_j \prod_{r \in S \setminus \{i\}} (1 - p_{rj}) < \mu_{j'} \prod_{r \in [k] \setminus \{i\}} (1 - p'_{rj'}).$$

Moreover, as j' is undersaturated, we must have $p'_{rj'} \geq p_{rj'}$ for all $r \in [k]$ from the construction of this process, and removing terms in the product only increases

the right side. Therefore, we deduce that

$$\mu_j \prod_{r \in S \setminus \{i\}} (1 - p_{rj}) < \mu_{j'} \prod_{r \in S \setminus \{i\}} (1 - p_{rj'}).$$

This implies that this deviation also increases $f(S)$ at Nash. As S was an arbitrary tight subset containing all queues up to level- ℓ , the claim is proved. By the reduction described above, this is a contradiction, and therefore $f([k])$ must further decrease with respect to \mathbf{p}' . The argument extends analogously at all intermediate points of this process at level- ℓ by the same reasoning as before.

Therefore, by induction, it follows that we may continuously deform probability mass from oversaturated servers to undersaturated servers while only decreasing $f([k])$. At the end of this process, there cannot be any oversaturated servers, otherwise the process could have continued. In particular, if \mathbf{p}'' is the final probability vector at the end of this process, we have shown that $\sum_{i \in [k]} p''_{ij} = x_j$ for all servers j and that $f([k]|\mathbf{p}'') \leq f([k]|\mathbf{p})$. We have

$$\begin{aligned} f([k]|\mathbf{p}) &= \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i=1}^k (1 - p_{ij}))}{\sum_{i=1}^k \lambda_i} && \text{(by definition)} \\ &\geq \frac{\sum_{j=1}^m \mu_j (1 - \prod_{i=1}^k (1 - p''_{ij}))}{\sum_{i=1}^k \lambda_i} && \text{(by construction)} \\ &\geq \frac{\sum_{j=1}^m \mu_j (1 - (1 - x_j/k)^k)}{\sum_{i=1}^k \lambda_i} && \text{(as symmetric profile maximizes product).} \end{aligned}$$

The second inequality holds because given $\sum_{i \in [k]} p''_{ij}$, the maximizer of $\prod_{i \in [k]} (1 - p''_{ij})$ is attained when each term is equal. As \mathbf{x} was arbitrary, we may take the maximum of the right side over all \mathbf{x} satisfying the constraints, and then the minimum over k . Because \mathbf{p} was an arbitrary Nash profile, this concludes the proof. \square

5.7 Proof of Convergence

In this section, we formally prove the key probabilistic result given by Theorem 5.4. The reader primarily interested in the game-theoretic results obtained above may safely skip this section.

We restate the main probabilistic result here for convenience:

Theorem 5.15 (Theorem 5.4, restated). *Let $\mathcal{G} = ([n], \boldsymbol{\mu}, \boldsymbol{\lambda})$ be a patient queuing game. For each fixed \mathbf{p} and all $i \in [n]$, almost surely the long-run aging rate of queue i satisfies*

$$c_i(\mathbf{p}) = \lim_{t \rightarrow \infty} \frac{T_t^{(i)}}{t} = r_i(\mathbf{p}),$$

where r_i is the output of Alg.

The high-level idea is to show that this identity holds for all $i \in S_1$, then S_2 , and so on. We will do so in two steps: the first step is showing that the maximum queue age grows by at most the desired rate on each long-enough window with high probability.

Proposition 5.16. *Fix $\epsilon > 0$. For any integer $a \in \mathbb{N}$, let $w = a \cdot \lceil \frac{a}{\epsilon} \rceil^{n-1}$. Suppose it holds at time t that $\max_{i \in [n]} T_t^{(i)} \geq w \cdot f_1$. Then*

$$\max_{i \in [n]} T_{t+w}^{(i)} - \max_{i \in [n]} T_t^{(i)} \leq (1 - (1 - \epsilon) \cdot f_1) \cdot w$$

with probability at least $1 - C_1 \exp(-C_2 a)$, where $C_1, C_2 > 0$ are absolute constants depending only on $n, \epsilon, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{p}$, but not on a . More generally, for each $s \geq 1$, if $\max_{i \notin U_{s-1}} T_t^{(i)} \geq w \cdot f_s$, then

$$\max_{i \notin U_{s-1}} T_{t+w}^{(i)} - \max_{i \notin U_{s-1}} T_t^{(i)} \leq (1 - (1 - \epsilon) \cdot f_s) \cdot w$$

with probability at least $1 - C_1 \exp(-C_2 a)$, where $C_1, C_2 > 0$ are absolute constants depending only on $n, \epsilon, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{p}$, but not on a .

We prove Proposition 5.16 in Section 5.7.1 using a delicate argument relying on a variety of concentration bounds. The key insight is that if a subset of much older queues S is likely to have priority on a long window of length w , the quantity $w \cdot f(S_1) \cdot \lambda(S)$ is a *lower bound* on the expected number of packets cleared *collectively* by S on this window by definition of S_1 . The analysis gets complicated when there are multiple old queues, as while we know these queues *collectively* have priority over all young queues, we must argue about priorities *within* this subset to bound the growth of the maximum queue age. We deal with this by induction by carefully chaining together large windows to obtain a win-win analysis.

For Proposition 5.16 to yield anything useful, we will need a corresponding lower bound that asserts roughly that if groups have separated according to what the algorithm asserts, then the aging rate of the *average* queue in a group grows at the conjectured rate. To that end, we prove the following result in Section 5.7.2 which shows that, *if* we have the conjectured aging separation between groups U_{k-1} and S_k , then some weighted combination of the queue ages in S_k (whose significance will prove apparent momentarily) must rise quickly:

Proposition 5.17. *For any $s \geq 1$ and any fixed $\epsilon > 0$, the following holds: suppose that at time t , it holds that*

$$\min_{i \in U_s} T_t^{(i)} - \max_{i \in S_{s+1}} T_t^{(i)} \geq 2 \cdot \frac{w}{\lambda_n}.$$

Then with probability $1 - A \exp(-Bw)$ where $A, B > 0$ are absolute constants not depending on w , we have

$$\sum_{i \in S_{s+1}} \lambda_i T_{t+w}^{(i)} - \sum_{i \in S_{s+1}} \lambda_i T_t^{(i)} \geq (1 - (1 + \epsilon)f_{s+1}) \cdot w \cdot \left(\sum_{i \in S_{s+1}} \lambda_i \right).$$

Moreover, for any fixed $\epsilon > 0$, with probability at least $1 - A \exp(-Bw)$ it holds that

$$\sum_{i \in S_1} \lambda_i T_w^{(i)} \geq (1 - (1 + \epsilon)f_1) \cdot w \cdot \left(\sum_{i \in S_1} \lambda_i \right).$$

Combined with Proposition 5.16, this will allow us to conclude that, because the average queue and oldest queue in S_1 ages at the desired rate almost surely, *all* queues in S_1 must age at this rate almost surely. To extend this analysis to lower groups S_2 , etc, we will use a similar analysis to show that the maximum age over every queue not in S_1 grows at most like $r(S_2)$. Then, because we know that every queue in S_1 grows by a $r(S_1) > r(S_2)$ rate, almost surely, eventually every queue in S_1 will be much older than every queue not in S_1 , giving priority. We leverage this fact to show again that the average queue in S_2 must grow by at least $r(S_2)$, and therefore every queue in S_2 grows at this rate almost surely. The proof for the lower groups S_3, \dots is completely analogous. We now proceed to make this argument formal to prove Theorem 5.4.

Proof of Theorem 5.4. By the Dominated Convergence Theorem, it suffices to show the second equality. We will show the desired statement holds for each $i \in S_1$, then S_2 , and so on. We first treat the case that the last outputted group S_k satisfies $g_k = 0$, or equivalently that $f_k \geq 1$. Fix $\epsilon > 0$ and partition time into consecutive windows of size $w_\ell = \ell \cdot \lceil \frac{\ell}{\epsilon} \rceil^{n-1}$. Let $W_\ell = \sum_{q=1}^{\ell-1} w_q$ be the time period at the beginning of the ℓ th window, and note that $w_\ell = \Theta(W_\ell^{1/2})$.

Consider the following events for $\ell = 1, 2, \dots$

$$A_\ell = \left\{ \max_{i \in [n] \setminus U_{k-1}} T_{W_{\ell+1}}^{(i)} - \max_{i \in [n] \setminus U_{k-1}} T_{W_\ell}^{(i)} > (1 - (1 - \epsilon) \cdot f_k) \cdot w_\ell \right\}$$

$$B_\ell = \left\{ \max_{i \in [n] \setminus U_{k-1}} T_{W_\ell}^{(i)} \geq w_\ell \cdot f_k \right\}$$

$$C_\ell = A_\ell \cap B_\ell.$$

Clearly, $\Pr(C_\ell) \leq \Pr(A_\ell|B_\ell)$. But by Proposition 5.16, we know that for some constants $C_1, C_2 > 0$ independent of ℓ , that

$$\Pr(A_\ell|B_\ell) \leq C_1 \exp(-C_2 \cdot \ell).$$

The sum over ℓ is thus finite and so the first Borel-Cantelli lemma (Lemma 2.6) implies that almost surely at most finitely many of the C_ℓ occur. Equivalently, almost surely, for all but finitely many of the ℓ , either

$$\max_{i \in [n] \setminus U_{k-1}} T_{W_{\ell+1}}^{(i)} - \max_{i \in [n] \setminus U_{k-1}} T_{W_\ell}^{(i)} \leq (1 - (1 - \epsilon) \cdot f_k) \cdot w_\ell \quad \text{or} \quad \max_{i \in [n] \setminus U_{k-1}} T_{W_\ell}^{(i)} < w_\ell \cdot f_k.$$

Observe that for each of the intervals where the latter holds, the value during the interval is at most $w_\ell \cdot f_k + w_{\ell+1} = O(W_\ell^{1/2})$. In particular, it is not difficult to see that almost surely $\max_{i \in [n] \setminus U_{k-1}} T_{W_\ell}^{(i)}$ is either $o(W_\ell)$, in which case we are done, or grows by at most a rate of $(1 - (1 - \epsilon) \cdot f_k) \cdot w_\ell$. Either way, as $\epsilon > 0$ was arbitrary, we may take $\epsilon \rightarrow 0$ to deduce the desired result that almost surely⁸

$$\limsup_{t \rightarrow \infty} \frac{\max_{i \in [n] \setminus U_{k-1}} T_t^{(i)}}{t} = 0 = g_k,$$

using $f_k \geq 1$. As ages of queues are nonnegative, the lower bound of 0 is trivial.

Now we turn to the rest of the groups, and we now assume that $g_k > 0$. We do this inductively. For S_1 , fix $\epsilon > 0$ and define

$$A_t = \left\{ \sum_{i \in S_1} \lambda_i T_t^{(i)} < (1 - (1 + \epsilon)f_1) \cdot t \cdot \left(\sum_{i \in S_1} \lambda_i \right) \right\}.$$

By Proposition 5.17, we know $\Pr(A_t) \leq A \exp(-Bt)$ for some constants $A, B > 0$ independent of t . Therefore, $\sum_{t=1}^{\infty} \Pr(A_t) < \infty$, from which the Borel-Cantelli

⁸For any $\epsilon > 0$, we have directly shown that the statement holds for t of the form $t = W_\ell$ for $\ell \geq 1$. For any t such that $W_\ell \leq t < W_{\ell+1}$, $T_t^{(i)}$ cannot be more than w_ℓ from the value at W_ℓ as ages can increase by at most 1 in each period. This implies that on any such intermediate time, the difference in the numerator from the value at $t = W_\ell$ is $O(w_\ell) = O(t^{1/2}) = o(t)$ and thus vanishes in the limit when divide by t , so the lim sup may be taken over all t , not just the sparsified sequence.

lemma implies that almost surely, for all but finitely many t ,

$$\sum_{i \in \mathcal{S}_1} \lambda_i T_t^{(i)} \geq (1 - (1 + \epsilon)f_1) \cdot t \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right).$$

Taking $\epsilon \rightarrow 0$, we obtain almost surely

$$\liminf_{t \rightarrow \infty} \frac{\sum_{i \in \mathcal{S}_1} \lambda_i T_t^{(i)}}{t} \geq g_1 \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right). \quad (5.4)$$

Next, note that deterministically, we have from Fact 2.1

$$\min_{i \in \mathcal{S}_1} T_t^{(i)} \leq \frac{\sum_{i \in \mathcal{S}_1} \lambda_i T_t^{(i)}}{\sum_{i \in \mathcal{S}_1} \lambda_i} \leq \max_{i \in \mathcal{S}_1} T_t^{(i)}. \quad (5.5)$$

In particular, we deduce that almost surely,

$$\liminf_{t \rightarrow \infty} \frac{\max_{i \in \mathcal{S}_1} T_t^{(i)}}{t} \geq g_1 > 0. \quad (5.6)$$

For the upper bound, let the w_ℓ and W_ℓ be as before, and now define

$$A_\ell = \left\{ \max_{i \in [n]} T_{W_{\ell+1}}^{(i)} - \max_{i \in [n]} T_{W_\ell}^{(i)} > (1 - (1 - \epsilon) \cdot f_1) \cdot w_\ell \right\} \quad B_\ell = \left\{ \max_{i \in [n]} T_{W_\ell}^{(i)} \geq w_\ell \cdot f_1 \right\}$$

$$C_\ell = A_\ell \cap B_\ell.$$

Again, $\Pr(C_\ell) \leq \Pr(A_\ell | B_\ell)$. By Proposition 5.16, a now routine application of the Borel-Cantelli lemma implies that almost surely, for all but finitely many ℓ , either

$$\max_{i \in [n]} T_{W_{\ell+1}}^{(i)} - \max_{i \in [n]} T_{W_\ell}^{(i)} < (1 - (1 - \epsilon) \cdot f_1) \cdot w_\ell \quad \text{or} \quad \max_{i \in [n]} T_{W_\ell}^{(i)} < w_\ell \cdot f_1.$$

But the latter event cannot happen infinitely often with positive probability, as this would imply $\max_{i \in [n]} T_{W_\ell}^{(i)} = o(W_\ell)$ infinitely often with nonzero probability, which violates (5.6). Therefore, it must be the case that almost surely, for all but finitely many ℓ , the former event holds. This implies that almost surely

$$\limsup_{t \rightarrow \infty} \frac{\max_{i \in [n]} T_t^{(i)}}{t} \leq (1 - (1 - \epsilon) \cdot f_1);$$

taking $\epsilon \rightarrow 0$ implies that

$$\limsup_{t \rightarrow \infty} \frac{\max_{i \in [n]} T_t^{(i)}}{t} \leq g_1.$$

As clearly the left side is an upper bound for the lim sup of only those queues in S_1 , almost surely

$$\limsup_{t \rightarrow \infty} \frac{\max_{i \in S_1} T_t^{(i)}}{t} \leq g_1.$$

Combining this with (5.6), we finally deduce that almost surely

$$\lim_{t \rightarrow \infty} \frac{\max_{i \in S_1} T_t^{(i)}}{t} = g_1.$$

By Equation (5.4) and Equation (5.5), we can also conclude that almost surely the same holds for the minimum $i \in S_1$. Thus $r_i(\mathbf{p}) = g_1$ for all $i \in S_1(\mathbf{p})$ by definition of g_1 , proving the theorem for all queues in S_1 .

We now show how to extend this inductively to higher values of k with $g_k > 0$. Suppose that we have shown for all $i \in U_{k-1}$ that the desired almost sure limit holds, and now consider S_k . A completely analogous argument using the windows w_ℓ as above with Proposition 5.16 via the Borel-Cantelli lemma implies that almost surely

$$\limsup_{t \rightarrow \infty} \frac{\max_{i \in [n] \setminus U_{k-1}} T_t^{(i)}}{t} \leq g_k. \quad (5.7)$$

Now, with these same windows, fix $\epsilon > 0$ and let

$$A_\ell = \left\{ \sum_{i \in S_k} \lambda_i T_{W_{\ell+1}}^{(i)} - \sum_{i \in S_k} \lambda_i T_{W_\ell}^{(i)} < (1 - (1 + \epsilon)f_k) \cdot w_\ell \cdot \left(\sum_{i \in S_k} \lambda_i \right) \right\}$$

$$B_\ell = \left\{ \min_{i \in U_{k-1}} T_{W_\ell}^{(i)} - \max_{i \in S_k} T_t^{(i)} \geq 2 \cdot \frac{w_\ell}{\lambda_n} \right\}$$

$$C_\ell = A_\ell \cap B_\ell.$$

Another completely analogous application of Proposition 5.17 and the Borel-Cantelli lemma implies that almost surely, at most finitely many of the C_ℓ occur.

That is, almost surely, for all but finitely many ℓ , either

$$\sum_{i \in S_k} \lambda_i T_{W_{\ell+1}}^{(i)} - \sum_{i \in S_k} \lambda_i T_{W_\ell}^{(i)} \geq (1 - (1 + \epsilon)f_k) \cdot w_\ell \cdot \left(\sum_{i \in S_k} \lambda_i \right) \quad \text{or} \quad \min_{i \in U_{k-1}} T_{W_\ell}^{(i)} - \max_{i \in S_k} T_t^{(i)} < 2 \cdot \frac{w_\ell}{\lambda_n}.$$

But the latter event cannot happen infinitely often with any nonzero probability by virtue of the inductive hypothesis and (5.7), as $g_{k-1} > g_k$ by Lemma 5.8, which implies that these timestamps cannot be so close infinitely often. Therefore, it must be the case that for all but finitely many of the ℓ , the former event holds. As usual, this immediately implies that

$$\liminf_{t \rightarrow \infty} \frac{\sum_{i \in S_k} \lambda_i T_t^{(i)}}{t} \geq (1 - (1 + \epsilon)f_k) \left(\sum_{i \in S_k} \lambda_i \right).$$

Again taking $\epsilon \rightarrow 0$ thus implies that almost surely

$$\liminf_{t \rightarrow \infty} \frac{\sum_{i \in S_k} \lambda_i T_t^{(i)}}{t} \geq g_k \left(\sum_{i \in S_k} \lambda_i \right),$$

which again coupled with (5.7) and Fact 2.1 yields that almost surely

$$\lim_{t \rightarrow \infty} \frac{\min_{i \in S_k} T_t^{(i)}}{t} = \lim_{t \rightarrow \infty} \frac{\max_{i \in S_k} T_t^{(i)}}{t} = g_k.$$

The extension to all $i \in S_k$ follows in the same manner as before by comparing with the average, completing the proof. \square

Observe that Theorem 5.4 rather strongly and explicitly characterizes the *linear* almost sure asymptotic growth rates of each queue for any choices of randomizations. Our main result in Theorem 5.12 showed that with a small slack in the system capacity, each queue will be guaranteed *sublinear* asymptotic growth almost surely in any equilibrium. While our objective function emphasizes the physical interpretation for each queue as asymptotic linear growth rate, these incentives impose that queues are indifferent between sublinear growth rates. One could instead define the game just using the f_k quantities directly, rather

than taking the max with 0 as is needed to argue about the asymptotic growth rates via r . If queues started out equally backed up, the f_k quantities measure the linear speed at which their ages descend to zero. In this setting, we provide the following stronger conclusion:

Corollary 5.18. *Fix \mathbf{p} and suppose that for some group S_k output by the algorithm, $f_k > 1$, so that $1 - f_k < 0$. Then, for each $i \in S_k$, $T_t^{(i)}$ is strongly stable.*

Proof Sketch. It suffices to show this for the random variable $\max_{i \notin U_{k-1}} T_t^{(i)}$. Let $\epsilon > 0$ be small enough such that $(1 - (1 - \epsilon) \cdot f_k) < \eta < 0$. Then let $w = a \cdot \lceil \frac{6}{\epsilon} \rceil^{n-1}$ be large enough so, on the event that $\max_{i \notin U_{k-1}} T_{\ell \cdot w}^{(i)} \geq f_k \cdot w$, then $\mathbb{E} \left[\max_{i \notin U_{k-1}} T_{(\ell+1) \cdot w}^{(i)} - \max_{i \notin U_{k-1}} T_{\ell \cdot w}^{(i)} \middle| \mathcal{F}_{\ell \cdot w} \right] < \beta < 0$ for some $\beta < 0$, where $\mathcal{F}_{\ell \cdot w}$ is the filtration of the Geometric system; this can be done by Proposition 5.16, noting that on the event where the proposition fails, the queue age can increase at most by w , and this can be drowned out in the expectation by the exponential decay of the probability bound by taking w large enough. This yields negative drift for the random process $Y_\ell := \max_{i \notin U_{k-1}} T_{\ell \cdot w}^{(i)}$ with threshold value $f_k \cdot w$.

Then, for any even $p \geq 0$, $\mathbb{E} \left[\left| \max_{i \notin U_{k-1}} T_{(\ell+1) \cdot w}^{(i)} - \max_{i \notin T_{j-1}} T_{\ell \cdot w}^{(i)} \right|^p \middle| \mathcal{F}_{\ell \cdot w} \right]$ is bounded by some constant $C_p > 0$ for each p depending only on n, w, λ . This is because the difference between these is crudely upper bounded as random variables by a sum of at most $n \cdot w$ geometric random variables in the case that queues somehow clear a packet every round, which are easily seen to have bounded moments. By Theorem 2.14, this implies stochastic stability as the p th moment condition holds for arbitrarily large p . \square

5.7.1 Proof of Rate Upper Bound

In this section, we turn to proving the first main claim above, Proposition 5.16. This result asserts that with high probability, the maximum queue age increases at a rate of at most $(1 - (1 - \epsilon) \cdot f_1)$ on the next w steps for a large enough w . In fact, more generally, the following holds:

Proposition 5.19 (Proposition 5.16, restated). *Fix $\epsilon > 0$. For any integer $a \in \mathbb{N}$, let $w = a \cdot \lceil \frac{6}{\epsilon} \rceil^{n-1}$. Suppose it holds at time t that $\max_{i \in [n]} T_t^{(i)} \geq w \cdot f_1$. Then*

$$\max_{i \in [n]} T_{t+w}^{(i)} - \max_{i \in [n]} T_t^{(i)} \leq (1 - (1 - \epsilon) \cdot f_1) \cdot w$$

with probability at least $1 - C_1 \exp(-C_2 a)$, where $C_1, C_2 > 0$ are absolute constants depending only on $n, \epsilon, \lambda, \mu, \mathbf{p}$, but not on a . More generally, for each $s \geq 1$, if $\max_{i \notin U_{s-1}} T_t^{(i)} \geq w \cdot f_s$, then

$$\max_{i \notin U_{s-1}} T_{t+w}^{(i)} - \max_{i \notin U_{s-1}} T_t^{(i)} \leq (1 - (1 - \epsilon) \cdot f_s) \cdot w$$

with probability at least $1 - C_1 \exp(-C_2 a)$, where $C_1, C_2 > 0$ are absolute constants depending only on $n, \epsilon, \lambda, \mu, \mathbf{p}$, but not on a .

Note that the first part is simply the $j = 1$ case of the more general statement. This proposition will follow from the following lemma that we will prove inductively:

Lemma 5.20. *Fix $s \geq 1$ and $\epsilon > 0$. Then, for each $1 \leq \tau \leq n$ and all $a \in \mathbb{N}$, the following holds: let $w = a \cdot \lceil \frac{6}{\epsilon} \rceil^{\tau-1}$, and suppose that at time t , $M^* := \max_{i \notin U_{s-1}} T_t^{(i)} \geq w \cdot f_s$, and that the set*

$$J = \{i \notin U_{s-1} : T_t^{(i)} \geq M^* - w \cdot f_s\}$$

has $|J| \leq \tau$. Then, with probability at least $1 - C_1 \exp(-C_2 a)$, we have

$$\max_{i \notin U_{s-1}} T_{t+w}^{(i)} - \max_{i \notin U_{s-1}} T_t^{(i)} \leq (1 - (1 - \epsilon) \cdot f_s) \cdot w,$$

where $C_1, C_2 > 0$ are absolute constants depending only on $n, \epsilon, \lambda, \mu, \mathbf{p}$, but not a .

Proposition 5.16 follows as for any $s \geq 1$ and $a \in \mathbb{N}, |J| \leq n$. We now turn to proving Lemma 5.20 inductively. The case for $\tau = 1$ will turn out to be relatively easy; this case just says that there is a single very old queue among those not in U_{s-1} , so we will be able to lower bound the number of packets she clears by simply assuming every queue in U_{s-1} is older than her. To extend this to higher τ will be more difficult. To do this, we will chunk together many windows that we know have this property for smaller values of τ and then leverage two facts to get a win-win situation. We will be able to easily show that at least one queue in J is always decreasing at the correct rate. If all queues in J are “close,” they also are at the correct rate. If not, then on the next chunk, they will clear at the correct rate on the next chunk inductively with high probability.

We now carry out this high-level plan. For reference, we will use the following similar notation to that used in the main text, but extended to more general windows:

1. $w := B \cdot L$ will denote a given *window* length composed of B consecutive *blocks* of L steps. As we will be considering the behavior of the process on some fixed window, we may as well reindex $t = 1$ for convenience, so that each window we consider will go from $t = 1$ to w . We will reserve the superscript $t = 0$ to denote the value of the ages at the very beginning of the window we consider.
2. Recall the shorthand $f_s := f(S_s | U_{s-1})$ and $g_s := \max\{0, 1 - f_s\}$.
3. With this convention, we will often define (and will make clear from context) at the beginning of some considered window of fixed length w , fixed

$s \geq 1$, and fixed $\epsilon > 0$:

$$M^* := \max_{i \in [n] \setminus U_{s-1}} T_0^{(i)} \quad T^* := M^* - (1 - \epsilon) \cdot w \cdot f_s.$$

We will often refer to T^* as the *target* value for this window, which does not change over the course of the window (notice it is measured at the beginning of the window). Then, define

$$J = \{i \notin U_{s-1} : T_0^{(i)} \geq M^* - w \cdot f_s\}.$$

That is, J is the set of queues whose age is within $w \cdot f_s$ of the oldest age, measured at the beginning of the window. Our goal will be to eventually show if w is sufficiently large, then with high probability, every queue in J has age below T^* at the end of the next w steps before accounting for w steps of aging, and of course all queues not in J are already strictly below T^* by definition. This will imply that the maximum age grows by at most $(1 - (1 - \epsilon) \cdot f_s) \cdot w$ once we account for the w steps of aging over this window.

4. Given a window of length $w = B \cdot L$, $\mathcal{F}_\ell^{(b)}$ is the filtration of σ -algebras generated up to step ℓ in the b th block, for $b = 1, \dots, B$. In particular, $\mathcal{F}_0^{(1)} \subseteq \mathcal{F}_1^{(1)} \subseteq \dots \subseteq \mathcal{F}_L^{(1)} \subseteq \mathcal{F}_0^{(2)} \subseteq \dots \subseteq \mathcal{F}_\ell^{(b)} \subseteq \dots \subseteq \mathcal{F}_L^{(B)}$.
5. $X_{b,\ell}^i$ will be the indicator that queue i cleared a packet in timestep ℓ in the b th block. $X_{b,\ell}^i$ is $\mathcal{F}_\ell^{(b)}$ -measurable.
6. $Y_{b,\ell}^i$ will be a sequence of random variables, with same interpretation of the indices, defined as follows: for b, ℓ such that every queue in J is still above $M^* - w \cdot f_s$ at the start of the ℓ th step of the b th block, set $Y_{b,\ell}^i = X_{b,\ell}^i$. If this does not hold for some b, ℓ , then let the $Y_{b,\ell}^i$ be arbitrary indicator random variables satisfying

$$\mathbb{E} \left[\sum_{i \in I} Y_{b,\ell}^i \middle| \mathcal{F}_{\ell-1}^{(b)} \right] \geq f_s \cdot \left(\sum_{i \in I} \lambda_i \right).$$

Note that $Y_{b,\ell}^i$ is $\mathcal{F}_\ell^{(b)}$ -measurable. These random variables are purely for technical convenience because they have an *a priori* lower bound on the conditional expectation, which is not always true of X (if for instance, queues have already cleared a lot and thus some queues in J have lost priority over those not in J).

7. $G_{b,\ell}^i$ will be i.i.d. $\text{Geom}(\lambda_i)$ random variables in the same way for $\ell \in [L], b \in [B]$. We define the partial sums on each window $Z_{b,k}^i := \sum_{\ell=1}^k G_{b,\ell}^i$. For each $b = 1, \dots, B$, we make the convention that the $G_{b,\ell}^i$ are sampled for all $i \in [n]$ and $\ell \in [L]$ at the beginning of the b th block, so that they are all $\mathcal{F}_0^{(b)}$ -measurable; there is no corresponding queuing step. When queue i clears her k th packet in the b th block, her timestamp decreases by $G_{b,k}^i$; equivalently, when this happens, her timestamp will have decreased on the b th block by $Z_{b,k}^i$ so far.

Proof of Lemma 5.20. Fix any group $s \geq 1$ and $\epsilon > 0$. Note that if $f_s = 0$ (i.e. queues never clear packets), then the lemma holds trivially with probability 1 as every queue in this subset increases deterministically by w in age on any window of length w , so we suppose otherwise. The proof is by induction on $\tau = |J|$.

Base Case: $\tau = 1$ (One Old Queue). We first consider $\tau = 1$. Let $a \in \mathbb{N}$ be arbitrary, and then set $w = a$. If $i^* \in J$ is the unique queue in the subset, then $i^* = \arg \max_{i \in [n] \setminus U_{s-1}} T_0^i$. We must show that at the end of w steps, this queue has decreased age (before accounting for aging) by at least $(1 - \epsilon) \cdot w \cdot f_s$ with high probability; the desired conclusion then follows from adding back in the w steps of aging. For simplicity, write $\lambda = \lambda_{i^*}$ as the arrival rate of this unique queue. We prove the claim in the following two steps:

First, we show that with high probability, the number of packets this queue must clear on this window to get below the target is not too large. Recall that to model this random process on the next w steps, let G_1, \dots, G_w be i.i.d. $\text{Geom}(\lambda)$ random variables (write $w = 1 \cdot w$ to indicate our window is just one block of w steps, so we omit the superscripts for blocks). As usual, write $Z_k = \sum_{\ell=1}^k G_\ell$. Recall that when this queue clears her k th packet on this window, her timestamp decreases by G_k ; equivalently, clearing k packets decreases her timestamp by Z_k collectively. Sample all of these random variables before hand, and let \mathcal{F}_0 be the σ -algebra generated by the previous history, as well as these random variables. \mathcal{F}_ℓ will be the filtration generated by all prior events up to the ℓ th step in this window.

Next, define the random variable $K^* = \min\{k \in [w] : Z_k \geq (1 - \epsilon) \cdot w \cdot f_s\}$, with the convention that if this set is empty, then $K^* = w + 1$. Observe that K^* is exactly the number of packets that this queue must clear to clear the target. We claim that with high probability, $K^* \leq K := \lambda \cdot (1 - \epsilon/2) \cdot w \cdot f_s$. To see this, apply Lemma 2.8 with $K = (1 - \epsilon/2) \cdot w \cdot f_s$ and $\delta = \frac{wf_s\epsilon}{2}$ to see that for this choice of K ,

$$\Pr(K^* > K) = \Pr(Z_K < (1 - \epsilon) \cdot w \cdot f_s) \quad (5.8)$$

$$= \Pr\left(Z_K < \frac{\lambda \cdot (1 - \epsilon/2) \cdot w \cdot f_s}{\lambda} - \frac{\epsilon \cdot w \cdot f_s}{2}\right) \quad (5.9)$$

$$\leq \exp\left(\frac{-w^2 f_s^2 \lambda^2 \epsilon^2}{8 \lambda w f_s (1 - \epsilon/2)}\right) \quad (5.10)$$

$$\leq \exp\left(\frac{-\lambda w f_s \epsilon^2}{8}\right). \quad (5.11)$$

As $f_s > 0$, with high probability, the queue will only need to clear at most K packets to get below the desired target.

Next, we show that with high probability, this queue will clear at least as many packets as required by the previous claim. Let X_1, \dots, X_w be the indicator variables for if the queue cleared a packet on each step of the window. Note that these random variables are path-dependent, but X_ℓ is \mathcal{F}_ℓ -measurable. Then the queue's timestamp decreases by at least $(1 - \epsilon) \cdot w \cdot f_s$ if and only if $\sum_{\ell=1}^w X_\ell \geq K^*$ by definition. Therefore, the probability the queue's timestamp decreases by at least $(1 - \epsilon) \cdot w \cdot f_s$ on the next w steps is

$$\Pr\left(\sum_{\ell=1}^w X_\ell \geq K^*\right) = \sum_{k=1}^w \Pr\left(\sum_{\ell=1}^w X_\ell \geq K^* \mid K^* = k\right) \Pr(K^* = k)$$

Consider the family Y_1, \dots, Y_w of indicator random variables that we couple with X_1, \dots, X_w as follows: while $\sum_{q=1}^{\ell-1} X_q < K^*$, set $Y_\ell = X_\ell$. While $\sum_{q=1}^{\ell-1} X_q \geq K^*$ on a sample path, let Y_ℓ be an arbitrary indicator random variable satisfying $\mathbb{E}[Y_\ell | \mathcal{F}_{\ell-1}] \geq \lambda \cdot f_s$. Notice that by construction, we always have $\mathbb{E}[Y_\ell | \mathcal{F}_{\ell-1}] \geq \lambda \cdot f_s$; this is because if $\sum_{q=1}^{\ell-1} X_q < K^*$, then the queue is still above the target, and therefore by assumption has priority. As this queue is the oldest not in U_{s-1} , she has priority over all other queues. If $V \subseteq U_{s-1}$ is some subset of queues with priority over her before she reaches her target, we know that in this case,

$$\mathbb{E}[Y_\ell | \mathcal{F}_{\ell-1}] \mathbb{E}[X_\ell | \mathcal{F}_{\ell-1}] = \lambda \cdot f(\{i^*\} | V) \geq \lambda \cdot f(\{i^*\} | U_{s-1}) \geq \lambda \cdot f_s,$$

where we use set monotonicity and the fact that f_s is the minimal value of $f(\cdot | U_{s-1})$ over all subsets contained in the complement.

Of course, if $\sum_{q=1}^{\ell-1} X_q \geq K^*$, $\mathbb{E}[Y_\ell | \mathcal{F}_{\ell-1}] \geq \lambda f_s$ simply by construction. However, because X_ℓ and Y_ℓ are equal while the queue is above the target, or equivalently before having cleared K^* packets, it follows that the events that $\sum_{\ell=1}^w X_\ell \geq K^*$ and $\sum_{\ell=1}^w Y_\ell \geq K^*$ have the same probability. Recall again that $K := \lambda \cdot (1 - \epsilon/2) \cdot w \cdot f_s$.

We obtain the probability the queue's timestamp gets below the target is at least

$$\begin{aligned}
\sum_{k=1}^K \Pr\left(\sum_{\ell=1}^w X_\ell \geq k \mid K^* = k\right) \Pr(K^* = k) &= \sum_{k=1}^K \Pr\left(\sum_{\ell=1}^w Y_\ell \geq k \mid K^* = k\right) \Pr(K^* = k) \\
&\geq \sum_{k=1}^K \Pr\left(\sum_{\ell=1}^w Y_\ell \geq \sum_{\ell=1}^w \mathbb{E}[Y_\ell \mid \mathcal{F}_{\ell-1}] - \frac{\lambda w f_s \epsilon}{2} \mid K^* = k\right) \Pr(K^* = k) \\
&\geq \left(1 - \exp\left(-\lambda^2 w f_s^2 \epsilon^2 / 4\right)\right) \Pr(K^* \leq K) \\
&\geq \left(1 - \exp\left(-\lambda^2 w f_s^2 \epsilon^2 / 4\right)\right) \left(1 - \exp\left(\frac{-\lambda w f_s \epsilon^2}{8}\right)\right) \\
&\geq 1 - 2 \exp\left(\frac{-\lambda^2 w f_s^2 \epsilon^2}{8}\right).
\end{aligned}$$

We use Azuma-Hoeffding in the fourth line applied to $\Delta_\ell := Y_\ell - \mathbb{E}[Y_\ell \mid \mathcal{F}_{\ell-1}]$, which is surely between -1 and 1 . As $w = a = a \cdot \lceil \frac{6}{\epsilon} \rceil^{1-1}$, it is clear that the probability this occurs is of the claimed form. To be safe, one should take $\lambda = \min_i \lambda_i$ to make the bound hold uniformly, independent of the identity of this queue.

Inductive Step for $\tau > 1$. Suppose that the proposition holds up to τ , and now we must show it holds for $\tau + 1$. Let $a \in \mathbb{N}$, and then $w = a \cdot \lceil \frac{6}{\epsilon} \rceil^\tau = \lceil \frac{6}{\epsilon} \rceil \cdot (a \cdot \lceil \frac{6}{\epsilon} \rceil^{\tau-1})$. In our notation, we have $w = B \cdot L$, where $B = \lceil \frac{6}{\epsilon} \rceil$ and $L = a \cdot \lceil \frac{6}{\epsilon} \rceil^{\tau-1}$.

Define as usual

$$\begin{aligned}
M^* &= \max_{i \in [n] \setminus U_{s-1}} T_0^{(i)} \\
T^* &:= M^* - (1 - \epsilon) \cdot w \cdot f_s \\
J &= \{i \in [n] \setminus U_{s-1} : T_t^{(i)} \geq M^* - w \cdot f_s\},
\end{aligned}$$

and suppose now that $|J| \leq \tau + 1$ and $M^* \geq w \cdot f_s$. For all $i \in J$, $\ell = 1, \dots, L$, and $b = 1, \dots, B$, define the random variables $X_{b,\ell}^i$ and $Y_{b,\ell}^i$ as described above, as well as the σ -algebras $\mathcal{F}_\ell^{(b)}$.

First, note that for any b, ℓ , if no queue in J has timestamp below $M^* - w \cdot f_s$ at the ℓ th step of the b th block (without accounting for aging), then every queue in J has priority over queues not in J . Then for the same reason as in the base case we have

$$\mathbb{E} \left[\sum_{i \in J} X_{b,\ell}^i \middle| \mathcal{F}_{\ell-1}^{(b)} \right] \geq f_s \cdot \left(\sum_{i \in J} \lambda_i \right).$$

Then by construction, we always have

$$\mathbb{E} \left[\sum_{i \in J} Y_{b,\ell}^i \middle| \mathcal{F}_{\ell-1}^{(b)} \right] \geq f_s \cdot \left(\sum_{i \in J} \lambda_i \right),$$

and this holds regardless of the conditioning. In particular, it follows that for every $b = 1, \dots, B$, we have

$$\sum_{\ell=1}^L \mathbb{E} \left[\sum_{i \in J} Y_{b,\ell}^i \middle| \mathcal{F}_{\ell-1}^{(b)} \right] \geq L \cdot f_s \cdot \left(\sum_{i \in J} \lambda_i \right).$$

Now, we define the following events for $b = 1, \dots, B$:

$$\begin{aligned} D_b &= \left\{ T_{b,L}^{(i)} \leq M^* - (b-3) \cdot L \cdot f_s (1 - \epsilon/2), \forall i \in J \right\} \\ E_b &= \left\{ Z_{b,\ell}^i \geq \frac{\ell}{\lambda} - \frac{\epsilon L f_s}{4}, \forall i \in J, \ell = 1, \dots, L \right\} \\ F_b &= \left\{ \sum_{\ell=1}^L \sum_{i \in J} Y_{b,\ell}^i \geq L \cdot f_s (1 - \epsilon/4) \left(\sum_{i \in J} \lambda_i \right) \right\}. \end{aligned}$$

We make the convention that D_0, E_0, F_0 are just the trivial event Ω that happens surely. Then define inductively:

$$A_0 = \Omega \quad A_{b+1} = A_b \cap D_{b+1} \cap E_{b+1} \cap F_{b+1}, \quad b < B.$$

We note that A_B implies D_B by construction so that if A_B holds, then at the end of this window of w steps (and again without accounting for w steps of aging), for all $i \in J$

$$T_w^{(i)} = T_{B,L}^{(i)} \leq M^* - (B-3) \cdot L \cdot f_s (1 - \epsilon/2).$$

Recall that our goal is that $T_w^{(i)} \leq M^* - B \cdot L \cdot f_s(1 - \epsilon)$ (before accounting for w steps of aging); note that our choice of $B = \lceil \frac{6}{\epsilon} \rceil$ satisfies

$$(B - 3) \cdot L \cdot f_s \cdot (1 - \epsilon/2) \geq B \cdot L \cdot f_s \cdot (1 - \epsilon).$$

Now, we lower bound the probability that A_b holds inductively. As $A_{b+1} \subseteq A_b$, we have

$$\Pr(A_{b+1}) = \Pr(D_{b+1} \cap E_{b+1} \cap F_{b+1} | A_b) \Pr(A_b) \geq (\Pr(D_{b+1} | A_b) - \Pr(E_{b+1}^c \cup F_{b+1}^c | A_b)) \Pr(A_b),$$

where we just use a union bound. By a union bound and Corollary 2.10,

$$\Pr(E_{b+1}^c | A_b) \leq 6(\tau + 1) \exp\left(\frac{-L \cdot \lambda^2 f_s^2 \epsilon^2}{144}\right),$$

where $\lambda = \min_{i \in [n]} \lambda_i$. A familiar argument by Azuma-Hoeffding gives

$$\Pr(F_{b+1}^c | A_b) \leq \exp\left(-\left(\sum_{i \in J} \lambda_i\right)^2 L f_s^2 \epsilon^2 / 16n\right) \leq \exp(-\lambda^2 L f_s^2 \epsilon^2 / 16n),$$

where we note that $\sum_{i \in J} Y_{b,\ell}^i - \mathbb{E}[\sum_{i \in J} Y_{b,\ell}^i]$ is surely between $-n$ and n for that extra factor. Therefore, a union bound implies

$$\Pr(A_{b+1}) \geq \left(\Pr(D_{b+1} | A_b) - 6(\tau + 2) \exp\left(\frac{-\lambda^2 L f_s^2 \epsilon^2}{144n}\right)\right) \Pr(A_b)$$

We now show that $\Pr(D_{b+1} | A_b)$ is large using a case analysis:

Case 1: No Gap. First suppose that after the b th block, there is no large gap between the maximum and minimum timestamp in J ; that is, (without accounting for aging, which affects queues equally)

$$\max_{i \in J} T_{b-L}^{(i)} - \min_{i \in J} T_{b-L}^{(i)} \leq L \cdot f_s(1 - \epsilon/2).$$

We show that this, along with the other assumptions in A_b , already imply the event D_{b+1} , so there is no need to analyze what happens on the $b + 1$ th block. Note that because there is no large gap, D_{b+1} will be implied by

$$\min_{i \in J} T_{b \cdot L}^{(i)} \leq M^* - b \cdot L \cdot f_s(1 - \epsilon/2), \quad (5.12)$$

because then (again, as we only account for aging at the end)

$$\max_{i \in J} T_{(b+1) \cdot L}^{(i)} \leq \max_{i \in J} T_{b \cdot L}^{(i)} \leq M^* - (b - 1) \cdot L \cdot f_s(1 - \epsilon/2).$$

We now show Equation (5.12) is indeed implied by A_b , which we recall implies E_b and F_b for all $q \leq b$. This means that for each $q \leq b, i \in J, \ell = 1, \dots, L$,

$$Z_{q, \ell}^i \geq \frac{\ell}{\lambda} - \frac{\epsilon L f_s}{4} \quad (5.13)$$

$$\sum_{q=1}^b \sum_{\ell=1}^L \sum_{i \in J} Y_{q, \ell}^i \geq b \cdot L \cdot f_s(1 - \epsilon/4) \left(\sum_{i \in J} \lambda_i \right). \quad (5.14)$$

If $\min_{i \in J} T_{b \cdot L}^i \leq M^* - B \cdot L \cdot f_s$, we are done by the above discussion, so suppose not; this means no queue in J has timestamp below $M^* - w f_s$ (before accounting for aging). By our coupling, we have $X_{q, \ell}^i = Y_{q, \ell}^i$ up until now as no queue in J has lost priority to those outside of J , so this last inequality is equivalent to

$$\sum_{q=1}^b \sum_{\ell=1}^L \sum_{i \in J} X_{q, \ell}^i \geq b \cdot L \cdot f_s(1 - \epsilon/4) \left(\sum_{i \in J} \lambda_i \right).$$

By averaging, this implies that there exists some $i \in J$ such that

$$\sum_{q=1}^b \sum_{\ell=1}^L X_{q, \ell}^i \geq b \cdot \lambda_i \cdot L \cdot f_s(1 - \epsilon/4).$$

By Equation (5.13), this implies this queue i has decreased in age by at least

$$b \cdot L \cdot f_s(1 - \epsilon/4) - \frac{b \cdot \epsilon \cdot L f_s}{4} = b \cdot L \cdot f_s(1 - \epsilon/2).$$

Because all queues in J have timestamp at most M^* at the beginning of this process, this implies that

$$\min_{i \in J} T_{b \cdot L}^{(i)} \leq M^* - b \cdot L \cdot f_s(1 - \epsilon/2),$$

(before accounting for aging), which combined with the no-gap assumption implies D_{b+1} . Thus, the no-gap condition implies the conditional probability of D_{b+1} is 1.

Case 2: Large Gap. Suppose after the b th block, there is a large gap between the maximum and minimum timestamp in J ; specifically, suppose that (again, without aging)

$$\max_{i \in J} T_{b-L}^i - \min_{i \in J} T_{b-L}^i > L \cdot f_s(1 - \epsilon/2).$$

The previous case showed that A_b itself implies

$$\min_{i \in J} T_{b-L}^{(i)} \leq M^* - b \cdot L \cdot f_s(1 - \epsilon/2)$$

As we have conditioned on A_b , $\max_{i \in J} T_{b-L}^i \leq M^* - (b-3) \cdot L f_s(1 - \epsilon/2)$; if this held instead with $b-2$, we would already be done. If not, as we have set $L = a \cdot \lceil \frac{6}{\epsilon} \rceil^{\tau-1}$ and evidently the set of queues with timestamp within $L \cdot f_s$ of the maximum in J has size at most τ , we may apply the inductive hypothesis to see that with probability at least $1 - C_1 \exp(-C_2 a)$, for some absolute constants C_1 and C_2 depending only on $n, \epsilon, \lambda, \mu, \mathbf{p}$ that every such queue decreases by at least $L \cdot f_s \cdot (1 - \epsilon/2)$ on this block by our choice of L ; therefore, conditioned on A_b , with probability at least $1 - C_1 \exp(-C_2 a)$, these queues decrease by enough to satisfy D_{b+1} .

Combining these cases, we conclude that

$$\Pr(A_{b+1}) \geq \left(1 - C_1 \exp(-C_2 a) - 6(\tau + 2) \exp\left(\frac{-\lambda^2 L f_s^2 \epsilon^2}{144n}\right)\right) \Pr(A_b)$$

Unravelling this recurrence and using $\Pr(A_0) = 1$, we conclude that

$$\Pr(A_B) \geq 1 - B(C_1 \exp(-C_2 a)) - 6B(\tau + 2) \exp\left(\frac{-\lambda^2 L f_s^2 \epsilon^2}{144n}\right).$$

As $L = a \cdot \lceil \frac{6}{\epsilon} \rceil^{\tau-1}$ and $B = \lceil \frac{6}{\epsilon} \rceil$, this evidently has the claimed form of $1 - C'_1 \exp(-C'_2 a)$ for absolute constants $C'_1, C'_2 > 0$ depending only on system parameters and not a , completing the inductive step. With this, the lemma is proved. \square

5.7.2 Proof of Rate Lower Bound

Next, we return to proving the corresponding lower bound of Proposition 5.17.

Proposition 5.21 (Proposition 5.17, restated). *For any $s \geq 1$ and any fixed $\epsilon > 0$, the following holds: suppose that at time t , it holds that*

$$\min_{i \in U_s} T_t^{(i)} - \max_{i \in S_{s+1}} T_t^{(i)} \geq 2 \cdot \frac{w}{\lambda_n}.$$

Then with probability $1 - A \exp(-Bw)$ where $A, B > 0$ are absolute constants not depending on w , we have

$$\sum_{i \in S_{s+1}} \lambda_i T_{t+w}^{(i)} - \sum_{i \in S_{s+1}} \lambda_i T_t^{(i)} \geq (1 - (1 + \epsilon)f_{s+1}) \cdot w \cdot \left(\sum_{i \in S_{s+1}} \lambda_i \right).$$

Moreover, for any fixed $\epsilon > 0$, with probability at least $1 - A \exp(-Bw)$ it holds that

$$\sum_{i \in S_1} \lambda_i T_w^{(i)} \geq (1 - (1 + \epsilon)f_1) \cdot w \cdot \left(\sum_{i \in S_1} \lambda_i \right).$$

Proof. We prove the second statement first, as it is slightly simpler and the main idea will reappear. As usual, let X_t^i denote the indicator variable that queue i cleared a packet at time t . Similar to the previous proof, we have for every $t \geq 1$,

$$\mathbb{E} \left[\sum_{i \in S_1} X_t^i \middle| \mathcal{F}_{t-1} \right] \leq \left(\sum_{i \in S_1} \lambda_i \right) \cdot f_1.$$

This is an upper bound as other queues may have priority at time t and some queues in S_1 may be empty at time t .

Recall G_ℓ^i are i.i.d. geometric random variables with parameter λ_i for $\ell = 1, \dots, w$. Again, the interpretation is that when queue i clears her ℓ th packet, her age decreases by G_ℓ^i ; in particular, the cumulative decrease from clearing k packets is then $Z_k^i := \sum_{\ell=1}^k G_\ell^i$. Another familiar application of Corollary 2.10 implies that with probability at least $1 - A \exp(-Bw)$ (where A, B are absolute constants depending only on n, ϵ, λ , not w),

$$\left| Z_k^i - \frac{k}{\lambda_i} \right| \leq \frac{(f_1 \cdot (\sum_{i \in S_1} \lambda_i) \cdot \epsilon/2) \cdot w}{\lambda_i}.$$

Because the expected number of packets cleared by the queues in S_1 is at most $\lambda(S_1) \cdot f_1$ by definition, another familiar application of the Azuma-Hoeffding inequality also implies with probability at least $1 - A' \exp(-B'w)$ (where A', B' do not depend on w) that

$$\sum_{t=1}^w \sum_{i \in S_1} X_t^i \leq (1 + \epsilon/2) \cdot \left(\sum_{i \in S_1} \lambda_i \right) \cdot f_1 \cdot w.$$

Combining these two estimates, we recall that

$$T_w^{(i)} = \max \left\{ 0, w - \sum_{\ell=1}^{\Gamma_w^i} G_\ell^i \right\} \geq w - \sum_{\ell=1}^{\Gamma_w^i} G_\ell^i,$$

where $\Gamma_w^i := \sum_{t=1}^w X_T^{(i)}$. Thus, it follows under these good events that

$$\begin{aligned}
\sum_{i \in \mathcal{S}_1} \lambda_i T_w^{(i)} &\geq w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) - \sum_{i \in \mathcal{S}_1} \lambda_i \left(\sum_{\ell=1}^{\Gamma_w^i} G_{i,\ell} \right) \\
&\geq w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) - \sum_{i \in \mathcal{S}_1} \lambda_i \cdot \left(\frac{\Gamma_w^i}{\lambda_i} + \frac{(f_1 \cdot (\sum_{i \in \mathcal{S}_1} \lambda_i) \cdot \epsilon/2) \cdot w}{\lambda_i} \right) \\
&= w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) - \sum_{i \in \mathcal{S}_1} \left(\Gamma_w^i + \left(f_1 \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) \cdot \epsilon/2 \right) \cdot w \right) \\
&\geq w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) - w \cdot f_1 \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) \cdot \epsilon/2 - \sum_{t=1}^w \sum_{i \in \mathcal{S}_1} X_t^i \\
&\geq w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) - w \cdot f_1 \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) \epsilon/2 - w \cdot (1 + \epsilon/2) \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) \cdot f_1 \\
&= w \cdot \left(\sum_{i \in \mathcal{S}_1} \lambda_i \right) \cdot (1 - (1 + \epsilon) \cdot f_1).
\end{aligned}$$

By a union bound, we thus find this occurs with probability $1 - A \exp(-Bw)$ for some possibly different constants $A, B > 0$ that do not depend on w , concluding the proof of the second statement.

The proof for $s \geq 1$ is very similar, just with one extra condition. Suppose at time t , we have

$$\min_{i \in U_s} T_t^{(i)} - \max_{i \in \mathcal{S}_{s+1}} T_t^{(i)} \geq 2 \cdot \frac{w}{\lambda_n}.$$

Consider now the next w steps, and let G_ℓ^i be as above, with $\ell = 1$ to w being the same geometric random variables, with Z_k^i being the k th partial sum. Another application of Corollary 2.10 implies that with probability at least $1 - A_1 \exp(-A_2 w)$,

$$Z_w^i < 1.5 \cdot \frac{w}{\lambda_n}, \quad \forall i \in U_s.$$

Moreover, this event occurring implies that no queue in U_s can possibly become younger than the any queue in \mathcal{S}_{s+1} on the next w steps, even if they clear a

packet in every step by our assumption. Therefore, on this event, we can apply the same analysis with the variables X_t^i as above, just noting that conditioned on this occurring,

$$\mathbb{E} \left[\sum_{i \in S_{s+1}} X_{t+\ell}^i \middle| \mathcal{F}_{t+\ell-1} \right] \leq \left(\sum_{i \in S_{s+1}} \lambda_i \right) \cdot f_{s+1},$$

as every queue in U_s will have priority over every queue in S_{s+1} on this window. An extremely similar argument via Azuma-Hoeffding and Corollary 2.10 and taking a union bound so that the concentration of queues in U_s above implies that the desired result holds with probability at least $1 - A \exp(-Bw)$ for some constants $A, B > 0$ not depending on w (but again, on $n, \epsilon, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{p}$). \square

5.8 Table of Notation

For convenience, we repeat the definitions of various expressions considered in Section 5.1.

Symbol	Formula	Definition
$\boldsymbol{\lambda}$		Vector of length n of queue arrival rates in descending order.
Δ^{m-1}		Probability simplex over m element set.
$\boldsymbol{\mu}$		Vector of length m of server success rates in descending order.
\mathbf{p}		Vector of queue randomizations over servers in $(\Delta^{m-1})^n$.

$\widetilde{T}_t^{(i)}$		Timestamp of oldest packet at queue i at time step t .
$T_t^{(i)}$	$\max\{0, t - \widetilde{T}_t^{(i)}\}$	Age of queue i at time t .
$\alpha(S \mathbf{p}, \mu, S')$	$\sum_{j=1}^m \mu_j \prod_{i \in S'} (1 - p_{i,j}) (1 - \prod_{i \in S} (1 - p_{i,j}))$	Expected number of packets cleared by queues in S if all have packets in a round and have priority over all queues except for those in S' and each such queue also has packets in the round.
$\lambda(S)$	$\sum_{i \in S} \lambda_i$	Sum of arrival rates of queues in S .
$f(S \mathbf{p}, \mu, \lambda, S')$	$\alpha(S \mathbf{p}, \mu, S')/\lambda(S)$	Ratio of expected packets cleared by S with priority over all queues except S' to total arrival rate of S .
$S_k(\mathbf{p}, \mu, \lambda)$		k th subset output in the algorithm of Section 5.5.
$U_k(\mathbf{p}, \mu, \lambda)$	$\cup_{\ell=1}^k S_\ell(\mathbf{p}, \mu, \lambda)$	Union of first k outputted subsets in the algorithm of Section 5.5.
$r_i(\mathbf{p}, \mu, \lambda)$		Outputted aging rate of queue i in the algorithm of Section 5.5.
$f_k(\mathbf{p}, \mu, \lambda)$	$f(S_k(\mathbf{p}, \mu, \lambda) \mathbf{p}, \mu, \lambda, U_{k-1})$	Value of f for S_k when U_{k-1} has priority.

$g_k(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})$	$\max\{0, 1 - f_k(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda})\}$	Outputted rate for S_k ; equivalently, value of r_i for any $i \in S_k$.
---	--	---

When the values of, or dependencies on, $\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\lambda}$ are clear from context, we omit them for notational ease.

5.9 Chapter Notes

The results in this section were originally obtained in [72], joint with Éva Tardos. Compared to the original version, the proof of equilibrium existence (Theorem 5.11) is somewhat simplified by appealing to the more powerful Debreu-Glicksburg-Fan Theorem rather than more standard results like the Kakutani Fixed-Point Theorem.

The proof of Theorem 5.4 given in this section is arguably the most technically involved part of this thesis. We suspect that there should be a simple probabilistic proof that long-run rates almost surely converge to a deterministic constant, for instance using appropriate convergence theorems of Markov chains or zero-one laws. In this case, one could perhaps simplify the proof by then giving an independent argument that the algorithmic rates are the only plausible constant.

5.9.1 Deferred Proofs

In this section, we fill in the deferred proofs showing analytic properties of r , the output of the algorithm described in Section 5.4.

Proposition 5.22 (Proposition 5.9, restated). *The function $r : (\Delta^{m-1})^n \rightarrow [0, 1]^n$ given by $r(\mathbf{p}) = (r_1(\mathbf{p}), \dots, r_n(\mathbf{p}))$ is continuous.*

Proof. Fix \mathbf{p}^* , a point we wish to show continuity at, and let $\mathbf{p}^k \rightarrow \mathbf{p}^*$ be a convergent sequence in $(\Delta^{m-1})^n$. It is easy to see that because the function $\mathbf{p} \mapsto \min_{S \subseteq [n]} f(S|\mathbf{p})$ is clearly continuous as the minimum of finitely many continuous functions, the function $\mathbf{p} \mapsto \max\{1 - f(S_1(\mathbf{p})|\mathbf{p}), 0\} = g_1(S_1(\mathbf{p}))$ is continuous. Therefore, if $\max_{i \in [n]} r_i(\mathbf{p}^*) = g_1(\mathbf{p}^*) = 0$, then by Lemma 5.8, as $\max_i r_i(\mathbf{p}^k) \rightarrow 0$, monotonicity yields $r_i(\mathbf{p}^k) \rightarrow 0$ along the sequence for every $i \in [n]$, proving continuity. We will now assume the harder case $g_1(\mathbf{p}^*) > 0$.

Before proceeding, define $\delta > 0$ to be the minimal nonzero gap between $f(S|S', \mathbf{p}^*)$ and $f(T|S', \mathbf{p}^*)$ over all choices of S, T, S' such that $S, T \subseteq [n] \setminus S'$, i.e.

$$\delta \triangleq \min_{S, T, S': f(S|S', \mathbf{p}^*) \neq f(T|S', \mathbf{p}^*)} |f(S|S', \mathbf{p}^*) - f(T|S', \mathbf{p}^*)|. \quad (5.15)$$

Note that δ is strictly positive as there are only finitely many choices of S, T, S' .

Fix $0 < \varepsilon < \delta$. Clearly, for any fixed S, S' such that $S \subseteq [n] \setminus S'$, the function $f(S|\mathbf{p}, S')$ is continuous as a function of \mathbf{p} . For this choice of ε , we may restrict to a tail of the sequence $\{\mathbf{p}^k\}$ and reindex so that for all $k \geq 1$, and all S, S'

$$|f(S|\mathbf{p}^k, S') - f(S|\mathbf{p}^*, S')| < \varepsilon/3. \quad (5.16)$$

We now claim that for every $k \geq 1$, the following holds in the algorithm's outputted rates on \mathbf{p}^k : while there exists any element $i \in S_1(\mathbf{p}^*)$ that has not

been outputted, the union of outputted subsets to that point must itself be a minimizing subset with respect to $f(\cdot)$ evaluated with profile \mathbf{p}^* , and that each element outputted so far has f value at p^k at most $f(S_1|\mathbf{p}^*) + \varepsilon/3$.

To see this, we proceed inductively: at the beginning of the algorithm, for every tight subset $S \subseteq S_1(\mathbf{p}^*)$ (by Lemma 5.7), we have by Equation (5.16)

$$f(S|\mathbf{p}^k) < f(S|\mathbf{p}^*) + \varepsilon/3 = f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + \varepsilon/3,$$

while for every subset T that is not minimal, we have by Equations (5.15) and (5.16) that

$$f(T|\mathbf{p}^k) > f(T|\mathbf{p}^*) - \varepsilon/3 \geq f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + \delta - \varepsilon/3 > f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + 2\varepsilon/3.$$

In particular, the first outputted subset must be a tight subset for \mathbf{p}^* , and the rate of each element in that subset is at least the desired amount.

Suppose this holds inductively, and now let $S \subseteq S_1(\mathbf{p}^*)$ be the union of the initial outputted sets, which we know is tight. If $S = S_1(\mathbf{p}^*)$, we are done, so suppose there exists $i \in S_1(\mathbf{p}^*) \setminus S$. Suppose T is disjoint and such that $S \cup T \subseteq S_1(\mathbf{p}^*)$ is tight. Such sets exist, for instance $T = S_1(\mathbf{p}^*) \setminus S$. From Fact 5.5,

$$f(S \cup T|\mathbf{p}^*) = f(S|\mathbf{p}^*) \oplus f(T|S, \mathbf{p}^*),$$

and by Fact 2.1, the fact that both the left side and the first term on the right are minimal implies $f(T|S, \mathbf{p}^*) = f(S_1(\mathbf{p}^*)|\mathbf{p}^*)$. Then at the next step of the algorithm, we again have by Equation (5.16)

$$f(T|S, \mathbf{p}^k) < f(T|S, \mathbf{p}^*) + \varepsilon/3 = f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + \varepsilon/3.$$

For any disjoint subset T' such that $S \cup T'$ is not tight, note that by Fact 5.5,

$$f(S \cup T'|\mathbf{p}^*) = f(S|\mathbf{p}^*) \oplus f(T'|S, \mathbf{p}^*).$$

By Fact 2.1, minimality of the first term on the right, and the fact that the left term is *not* minimal, it follows that $f(T'|S, \mathbf{p}^*) \geq f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + \delta$ by definition. We then have by Equation (5.16) that

$$f(T'|S, \mathbf{p}^k) > f(T'|S, \mathbf{p}^*) - \varepsilon/3 \geq f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + \delta - \varepsilon/3 > f(S_1(\mathbf{p}^*)|\mathbf{p}^*) + 2\varepsilon/3.$$

In particular, the next outputted set is such that $S \cup T \subseteq S_1(\mathbf{p}^*)$ is tight, and the rate condition holds as well. We can iteratively apply this while $S_1(\mathbf{p}^*)$ is not exhausted, proving that every element of $S_1(\mathbf{p}^*)$ is outputted before any other element when running the algorithm on \mathbf{p}^k , with rate within $\varepsilon/3$ of $r_i(\mathbf{p}^*)$. As ε was arbitrary, this shows continuity on all components of $S_1(\mathbf{p}^*)$.

Because we have shown that every queue of $S_1(\mathbf{p}^*)$ is outputted before every queue not in S_1 , we can apply the recurrence as discussed in Equation (5.2) to show continuity for each queue in $S_2(\mathbf{p}^*)$, just discounting μ as usual. The same argument restricted to $[n] \setminus S_1(\mathbf{p}^*)$ nearly shows continuity; the only difference is the discounting of μ by the queues in $S_1(\mathbf{p}^*)$ depends on \mathbf{p}^k , not \mathbf{p}^* , but as each $f(S|S_1(\mathbf{p}^*), \mathbf{p}, \mu')$ is jointly continuous in \mathbf{p}, μ' , and the composition of continuous functions is continuous, the same argument holds with minimal modification. This proves continuity for the components of each subsequent group recursively, and thus of each component in $[n]$. \square

5.9.2 The Need for Packet Priority in Patient Queuing

Here, we argue that even with the restriction to stationary strategies in the patient queuing game, the priority scheme by servers to attempt to serve the oldest packet is necessary to obtain constant price of anarchy bounds.

We claim that if servers choose packets uniformly at random among those

it receives in each round, the price of anarchy can be polynomially large in n in the sense of Theorem 4.2 when the costs are defined to be the asymptotic aging rates as in Definition 5.1. To see this, consider a queuing system with one queue with arrival rate $C/n^{1/3}$ for some large constant $C > 0$ to be determined, and $n - 1$ queues with arrival rate $1/n^{2/3}$. Suppose that there is one server with success rate 1, and then n servers with success rate $1/n^{1/3}$. There exist bad Nash equilibria in stationary strategies of the following form: each small queue evenly mixes between a personal server with success rate $1/n^{2/3}$ and the top server with success rate 1. It is clear that each small queue will be stable for large enough n (provided no other queue shares her personal server) by this choice of constants, so she has no incentive to deviate because her cost in this game is her long-run aging rate of zero.

By standard Chernoff bounds, in any given round, there are at least $cn^{1/3}$ small queues that will send to the top server under these stationary strategies with overwhelming probability for some absolute constant $c > 0$; therefore, with very high probability in each round, the large queue can succeed in clearing a packet by sending to the large server with probability at most $1/cn^{1/3}$. This is strictly smaller than her rate if $C > 0$ is taken sufficiently large. At any other server, the queue can get rate at best $1/n^{1/3}$, which clearly does not ensure stability for large enough $C > 0$. Therefore, the system must remain unstable if she best responds to this behavior by the other queues (note further that this best response will never involve sending to a personal server of any of the other queues, as there are n such servers). No queue has any incentive to deviate according to the cost functions as defined in Definition 5.1 under this alternate server selection choice, so these strategies constitute a Nash equilibrium. This system would remain centrally feasible even if server rates were scaled down

by $\Omega(n^{1/3})$, and hence the price of anarchy of such a patient queuing game can be polynomially large in n .

In this example, small queues mix between the top server and their own server equally because their long-run growth rate is zero in any case. It is possible to modify this example so that all queues mix among servers that offer equal long-run probability of success. This can be done by having each small queue send to the top server with probability $p(n)$ and to a personal server with rate $1/n^{1/3}$ with probability $1 - p(n)$, while the top queue sends deterministically to the last server with rate $1/n^{1/3}$ (note that the effectiveness of the top server is at best $1/n^{1/3}$ if the top queue were to deviate to there, by construction, so this is a best response). The parameter $p(n)$ can be chosen so that the long-run average number of small queues sending to the top server is precisely $n^{1/3}$ almost surely using the strong law of large numbers for Markov chains, so that each queue indeed mixes among servers that offer long-run success probability $1/n^{1/3}$.

5.9.3 Related Work

Our main game-theoretic result of this chapter (Theorem 5.12) falls within the large body of work on price of anarchy bounds in games [99]. As mentioned in the previous chapter, to be more precise, our main result is more similar in spirit to the bicriteria result of Roughgarden and Tardos [117] for selfish routing. However, our proof of Theorem 5.12 differs substantially from most price of anarchy analyses, which often follow Roughgarden's *smoothness* framework [114]. In such arguments, one must establish an inequality based on *discrete, individual* deviations from Nash to argue about the overall quality. Our argu-

ment is instead an *equilibrium analysis* that is more similar to that of Johari and Tsitsiklis [96], who establish equilibrium conditions and modify their problem while maintaining the equilibrium condition to arrive at a version that is easy to analyze. In our argument, we also modify the equilibrium itself towards a more tractable solution, but the intermediate points in this deformation will *not* be Nash, requiring additional arguments.

While our main goal was to establish price-of-anarchy-style bounds in dependent systems, this necessarily required a careful understanding of random queuing dynamics as in Theorem 5.4. In the probability literature, the long-run aging rates that form the *incentives* are known as (linear) escape rates of random walks [104]. A rich theory has emerged to study this in special networks, but it is unclear how to apply these techniques in our setting. Our work relies on careful, self-contained estimates leveraging concentration, coupling, and supermartingale arguments.

Subsequent Work: Following the original appearance of this work, Fu, Hu, and Lin [64] extend our results on the patient queuing game to general bipartite queuing networks. They show that, under a suitable generalization of the feasibility condition to the non-complete case, the price of anarchy is at most 2. Their argument relies on discrete deviations as in more traditional smoothness approaches, in contrast to the continuous deformations considered here. It remains an open question whether our bound of $\frac{e}{e-1}$ extends to the non-complete setting.

CHAPTER 6
REPEATED AUCTIONS WITH BUDGETS: PACING DYNAMICS AND
WELFARE

In the previous two chapters, we have studied the outcomes of strategic behaviors in queuing systems as a prototypical example of repeated games with state. As we saw, while no-regret algorithms obtain very nontrivial aggregate performance guarantees, these bounds were suboptimal compared to even restricted classes of policies that more explicitly reasoned about long-run dynamics. But how general are these findings? Do they extend to other dynamic games with state?

We thus turn to another such system of great significance, both in theory and practice: repeated auctions with budgets. In the formulation we consider, n agents repeatedly interact for T rounds.¹ For our purposes in this discussion, at each round t , we assume an auctioneer puts a single item up for auction.² Each agent observes their own value $v_{i,t}$ for the item; we assume that the joint random vector $\mathbf{v}_t = (v_{1,t}, \dots, v_{n,t})$ is drawn from some time-invariant distribution F , but may be arbitrarily correlated across agents. Each agent i then submits a bid $b_{i,t}$, which can be chosen as any arbitrary function of their observed history and their current valuation. The auctioneer takes in the bid vector $\mathbf{b}_t = (b_{1,t}, \dots, b_{n,t})$ and then outputs an allocation $\mathbf{x}_t = (x_{1,t}, \dots, x_{n,t})$ and prices $\mathbf{p}_t = (p_{1,t}, \dots, p_{n,t})$; the interpretation of $x_{i,t}$ is the amount of the item agent i receives at time t , which $p_{i,t}$ is the price paid to the auctioneer for this allocation. In standard auctions, like first- or second- price, $x_{i,t} = 1$ if and only if $b_{i,t}$ is maximal, in which case

¹We will define this process more formally later on. Note also that unlike the results in the previous two chapters, we are now back in the finite horizon setting.

²Our analysis will be able to handle somewhat more general formats.

the corresponding price to agent i is either her own bid or the next highest bid, respectively. The quasilinear utility of the agent is then defined to be $x_{i,t}v_{i,t} - p_{i,t}$: the first term measures their utility for the amount of the item they receive, while the latter their disutility for spending money. The key feature of these repeated auctions is that each agent i is also constrained to spend at most B_i units over the course of the T rounds; we label B_i as agent i 's **budget**. Agents therefore cannot view each repeated auction in isolation, but rather must implicitly reason about how to balance expenditures in the present with future consumption.

This model of repeated auctions forms a convenient abstraction of Internet ad auctions, which has become the dominant mode of advertising. At a high level, when users search for a keyword in search engines like Google or Bing, advertisers may bid for the right to display their advertisement to this user. Based on various known features of the user provided by the Internet platform, advertisers can form estimates of the expected value they would obtain from the user clicking their ad. These estimates motivate the advertisers' bids in the auctions.³

In practice, the task of bidding strategically in Internet ad auctions with the goal of maximizing utility is an immensely difficult task even beyond reasoning about these consumer signals. First, these auctions occur at rapid scales; indeed, tens of thousands of queries are conducted on Google every second. Moreover, advertisers often participate in auctions corresponding to several different relevant keywords and in a variety of online platforms, each with different payment structures or interfaces. For these reasons, any adaptive bidding strategy of a

³In practice, Internet platforms charge advertisers per click, not per impression. Therefore, platforms must incorporate estimates of advertiser quality in the form of *clickthrough rates* in allocating the ad space; however, we abstract away these concerns and deal with the formalism of general auctions.

particular advertiser ought be extremely efficient to implement. These auctions are, again, all strategically linked by the global budget for the advertising campaign, further complicating their underlying optimization problem.

This kind of optimization problem is a special case of (*adversarial*) *bandits with knapsacks* [93, 34, 58], a rather challenging learning problem. In a general setting of bandits with knapsacks, the learner must sequentially choose actions (here, bids), which have associated rewards (here, quasilinear utilities) and resource consumption (here, money) until the end of the time horizon or until some resource has been exhausted. The goal of the learner is to remain competitive against the value obtained by every fixed distribution over actions, which is taken until the time some resource is exhausted or the end of the time horizon. Even though each agent in our model will only have a single resource constraint, such learning problems typically face an inherent “spend-or-save” conundrum. If rewards and resource consumptions can be chosen adversarially (non-stochastically), even obliviously, it is known that no algorithm can achieve better than a $O(\log T)$ ratio of the optimal attainable value [93] This stands in stark contrast to standard bandits problems, where one can typically obtain $o(T)$ *regret* without any approximation ratio.

To mitigate the difficulty of this task, most mature Internet platforms now provide automated bidding services. These autobidders take as input high-level information, like the overall budget constraints and maximum valuations for different kinds of queries, and then automatically submit bids in these auctions dynamically on behalf of the advertiser. Due to the computational challenges described above, a simple but popular strategy is “budget pacing”: the autobidder shades these maximum valuations in each round dynamically based on previ-

ous outcomes with the goal of adaptively bidding in a way so that the realized expenditure is close to the target budget. In this way, the autobidder adaptively learns and maintains a one-dimensional parameter, known as a “pacing multiplier,” which is used to determine the bid in the next round.

Budget pacing is proven to be a popular approach for repeated bidding under budget constraints, both in practice and in theory [16, 42, 43]. The advantages of this form of campaign management are manifold: first, while certain larger advertisers may be capable of employing more sophisticated algorithms, autobidding enables smaller advertisers to participate with low barrier to entry. As mentioned above, these automated strategies require only high-level parameters from the advertiser. These autobidders, which are provided by the platform, also can more flexibly adapt to market signals. From a theoretical point of view, in the case of second-price auctions, it is actually known that these kinds of linear policies are optimal in the following sense: for each advertiser, given the competing bids of all other agents, there exists a single, fixed pacing multiplier such that bidding the value multiplied by pacing multiplier would have been optimal in hindsight.

Due to the enormous success of these autobidding strategies, the main goal of this chapter is to prove aggregate welfare guarantees in systems where *all agents* uses some form of dynamic pacing algorithm as described above. From the theoretical point of view, the fundamental challenge of repeated auctions with budgets is that each agent cannot view each auction in isolation, must adaptively learn bid over time, and exert different influences in the market due to heterogeneities both in valuations and budgets. We will study a popular and general class of *gradient-based pacing algorithms*, which to our knowledge was

first introduced by Balseiro and Gur [17] In their work, these algorithms are derived via stochastic gradient descent for a particular Lagrangian dual of a bidder’s individual utility maximization problem in second-price auctions. As we shall see, these algorithms are quite natural, extend a much broader class of auctions, and are extremely simple to implement: each agent can easily update their bidding strategies using their own expenditures and allocations.

Our main result shows that pacing dynamics attain provably strong welfare guarantees with respect to **liquid welfare**, a natural notion of welfare that is adapted to budgeted settings introduced by Dobzinski and Paes Leme [54]. Informally, the liquid welfare of an agent, given an allocation of items, is the total amount the agent would have been willing to pay for this allocation. With quasilinear utilities and budgets, this is precisely the minimum of her true value for the items and her budget. Our main result (theorem 6.4) be stated informally as follows:

Suppose that all agents use gradient-based pacing algorithms when competing in repeated *core auctions*⁴ with budgets. Then, the expected liquid welfare (taken over the random valuations) obtained by the agents is at least half of the expected liquid welfare that is attainable via any allocation rule, up to sublinear terms in T .

We reiterate that this result is agnostic to the underlying correlation structure of the valuations across agents. This feature is quite desirable; in most applications, like ad auctions, all agents will receive correlated signals about the quality of the item for auction. Whereas many other results must assume independent valuations (like [17]), our analysis completely bypasses the technical challenges

⁴We define this class below, but core auctions includes all examples mentioned above.

that this imposes.

As we will discuss more formally later on after introducing various preliminaries, a key subtlety of the analysis is that we will analyze aggregate welfare *without requiring convergence of the dynamics to any notion of equilibrium*. This point is critical towards overcoming known complexity-theoretic obstacles [38], as will be discussed below. Moreover, our analysis does not rely directly on regret minimizing properties of the learning algorithms we study. In fact, the no-regret condition is *not sufficient* to proving any constant liquid welfare guarantee. These two challenges necessitate new approaches to analyzing the dynamics: we must instead carefully and directly link the welfare obtained by the agents to the evolution of their learning algorithms to establish our bounds.

6.1 Preliminaries: Auctions, Welfare, and Pacing

We now turn to formalizing our model of repeated auctions with budgets. We then define the welfare notion that will be the subject of our analysis, as well as the general class of learning algorithms that our results will apply to.

6.1.1 Repeated Auctions with Budgets

Our model of repeated auctions with budgets is quite similar to that introduced by Balseiro and Gur [17]. We there is one seller (the platform) that repeatedly sells items to n agents. For each round t , there exists a set of valid allocation profiles $X \subseteq [0, 1]^n$ such that X is compact, convex, downward-closed, and independent of t . A feasible allocation profile at time t is a vector $\mathbf{x}_t = (x_{1,t}, \dots, x_{n,t}) \in X$,

and a feasible allocation sequence is a sequence of allocations $(\mathbf{x}_1, \dots, \mathbf{x}_T)$.

An instance of repeated auctions with budgets with n bidders is thus parameterized by:

1. a time horizon $T \in \mathbb{N}$,
2. budgets B_i for each agent $i \in [n]$, with corresponding **target spend rate** $\rho_i \triangleq B_i/T$,
3. a joint distribution F over valuation profiles $\mathbf{v} = (v_1, \dots, v_n)$ of the agents⁵, where $v_i \leq \bar{v}$ surely for some parameter $\bar{v} \geq 1$ (by normalizing),
4. a set of valid allocation profiles $X \subseteq [0, 1]^n$, and
5. an allocation rule $\mathbf{x} : \mathbb{R}_+^n \rightarrow X_t$ and pricing rules $\mathbf{p} : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ that takes in nonnegative bids and outputs nonnegative allocations and payments for each agent. We assume that \mathbf{x}_t and \mathbf{p}_t are weakly monotone increasing in each component.

The model then works as follows: for each time $t = 1, \dots, T$,

1. A value profile $\mathbf{v}_t \sim F$ is drawn independently of all previous events. Each agent i observes only $v_{i,t}$.
2. Each agent i then chooses a bid $b_{i,t}$ as some measurable function of their observed history. We denote $\mathbf{b} = (b_{1,t}, \dots, b_{n,t})$ for the bidding profile at time t .
3. The auction then computes the allocation $\mathbf{x}_t \triangleq \mathbf{x}(\mathbf{b}_t)$ and payments $\mathbf{p}_t \triangleq \mathbf{p}(\mathbf{b}_t)$.

⁵We note that F need not be a product distribution; in particular, valuations can be arbitrarily correlated.

4. Each agent i then observes $x_{i,t}$ and $z_{i,t} \triangleq p_{i,t}$. We impose that if the total payment of an agent i exceeds B_i , then she may not bid in any future rounds.

Examples: Before proceeding, let us see how to instantiate the ad auction examples via the above framework.

1. (Single-Slot Ad Auctions) Each round corresponds to a single ad impression and single slot available for an advertisement. Each agent observes a signal $\theta_t \in \Theta$, where Θ is some type space parametrizing observable features of the impression like the keyword or user features. In each round t , θ_t is drawn i.i.d. from some distribution over Θ . Each agent i has a value function $v_i : \Theta \rightarrow \mathbb{R}_+$ from signals to valuations. The joint value profile at time t is then $\mathbf{v}_t = (v_1(\theta_t), \dots, v_n(\theta_t))$. In this case, because only a single advertisement can be shown, the set of allocation profiles X is the (downward closure of the) simplex $\{\mathbf{x} \in [0, 1]^n : \sum_{i \in [n]} x_i \leq 1\}$.
2. (Multiple-Slot Pay-Per-Click Ad Auctions) Using the same type and valuation terminology as above, multiple ad slots can be formulated as follows: there exist $m \geq 1$ slots available for each round labelled $\{1, \dots, m\}$. The quality, or click rate, of each slot is weakly decreasing in the slot index and are denoted $1 \geq \alpha_1 \geq \dots \geq \alpha_m \geq 0$. If agent i is allocated the k th slot, the value agent i obtains is $v_i(\theta_t) \cdot \alpha_k$; the first term in the product is the advertiser's expected value per impression, while the latter term is the slot-specific click rate per impression. In this case, the set of feasible allocation profiles X forms the polymatroid [77]

$$X = \left\{ \mathbf{x} \in [0, 1]^n : \sum_{k \in S} x_k \leq \sum_{i=1}^{|S|} \alpha_i \quad \forall S \subseteq [n] \right\},$$

where we assume $\alpha_k = 0$ if $k > m$.

In both examples, note that allowing valuations to be correlated across agents is vital; it is unreasonable to think that the valuations of advertisers are independent conditional on the signal, which we assume to be i.i.d.

In the above, we have left the precise auction rules in somewhat abstract form. We now make more concrete the class of auctions that we will consider in this chapter: our results will apply to any *core auction*.

Definition 6.1. A **core auction** is an allocation function $\mathbf{x} : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ and pricing rule $\mathbf{p} : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ that satisfies:

1. (Individual Rationality) For each $i \in [n]$ and any bidding profile \mathbf{b} , it holds that

$$p_i(\mathbf{b}) \leq b_i x_i(\mathbf{b}).$$

In words, the payment of each agent does not exceed her declared welfare for her allocation. With quasilinear utilities, this means that the utility of an agent can never be negative with truthful bidding.

2. (Coalition Robustness) The seller and any subset of agents $S \subseteq [n]$ cannot strictly benefit by abandoning the auction outcome and deviating to another outcome. Formally, for any bidding profile \mathbf{b} , any subset $S \subseteq [n]$, and any allocation profile $\mathbf{y} \in X$,

$$\sum_{k \notin S} p_k(\mathbf{b}) + \sum_{k \in S} b_k x_k(\mathbf{b}) \geq \sum_{k \in S} b_k y_k.$$

To parse this condition, observe that the left-hand side yields the total welfare obtained by the seller and agents in S , while the right-hand side denotes the total welfare of the agents of the seller and agents in S under the allocation \mathbf{y} . We use the fact that with quasilinear utilities, the utility loss

of payments by agents in S is exactly compensated by the revenue of the seller from agents in S and thus cancel on both sides.

To reiterate, core auctions are such that no coalition of players, possibly including the seller, could jointly obtain higher utility by renegotiating the auction outcome among themselves.

We now provide examples of core auctions:

1. **(First-Price Auctions)** The allocation rule is a welfare-maximizing allocation $\mathbf{x}(\mathbf{b}) \in \arg \max_{\mathbf{x} \in X} \{\sum_k b_k x_k\}$. In the case the set of valid allocations is the downward closed simplex and bids are unique, this amounts to allocating to the highest bidder. The payment of each agent k is then her bid for the allocation she receives; that is, $p_k(\mathbf{b}) = x_k b_k$.

First-price auctions clearly satisfy the core auction definition: by definition, the payment is exactly the declared welfare, and for any subset $S \subseteq [n]$, by definition we have

$$\sum_{k \notin S} p_k(\mathbf{b}) + \sum_{k \in S} b_k x_k(\mathbf{b}) = \sum_{k \in [n]} b_k x_k(\mathbf{b}) \geq \sum_{k \in [n]} b_k y_k \geq \sum_{k \in S} b_k y_k,$$

where we simply use the payment definition, then the definition of the allocation rule, and finally the nonnegativity of bids and allocation profiles in the last inequality.

2. **(Second-Price Auctions with A Single Item)** In this case, the set of valid allocation profiles is $X = \{\mathbf{x} \in [0, 1]^n : \sum_{i \in [n]} x_i \leq 1\}$. The allocation rule is the same as above, while the payment of each agent k is x_k multiplied by the second largest bid, which we denote $\text{smax}(\mathbf{b})$.

To verify the core auction properties, note that in this case, the item is completely allocated to (a subset) of bidders with highest bids: in this case,

her payment is at most her payment in the first-price case which is at most her declared valuation. To verify the second property, first observe that

$$\sum_{k \notin S} p_k(\mathbf{b}) + \sum_{k \in S} b_k x_k(\mathbf{b}) = \sum_{k \notin S} x_k(\mathbf{b}) \text{smax}(\mathbf{b}) + \sum_{k \in S} b_k x_k(\mathbf{b}). \quad (6.1)$$

If the maximizing bidder is unique, in which case we denote her by i^* we may proceed by cases: if $i^* \in S$, then the core auction property holds with equality. If not, then $\max_{k \in S} b_k \leq \text{smax}(\mathbf{b})$, in which case the inequality holds again. If the maximizing bidder is not unique, then whenever $x_k(\mathbf{b}) > 0$, it must hold that $x_k(\mathbf{b}) \text{smax}(\mathbf{b}) = x_k(\mathbf{b}) b_k$, as the second highest price is also the highest price, which must be the bid of any agent with nonzero allocation. In this case, Equation (6.1) is equivalent to

$$\sum_{k \in [n]} x_k(\mathbf{b}) b_k \geq \sum_{k \in [n]} b_k y_k \geq \sum_{k \in S} b_k y_k,$$

for the same reasons as above.

3. **(Generalized Second-Price (GSP) Auctions for Multiple Ad Slots)** Consider again the multi-slot environment described above. The GSP allocation goes as follows: slots are allocated greedily to the highest bidders in order, and if an agent is allocated the k th slot, her price per unit is equal to the $(k + 1)$ th (next) highest bid. Equivalently, letting $\pi : [n] \rightarrow [n]$ be a permutation such that $b_{\pi(1)} \geq \dots \geq b_{\pi(n)}$, agent $\pi(k)$ is allocated the k th slot (so that $x_{\pi(k)} = \alpha_k$) and pays $\alpha_k b_{\pi(k+1)}$.

In this case, verifying that the multi-slot GSP auction is somewhat trickier:

Proposition 6.1. *For any choice of $1 \geq \alpha_1 \geq \dots \geq \alpha_n \geq 0$ for slot click rates, the multi-slot GSP auction is a core auction.*

Proof. For convenience of notation, reorder bids so that $b_1 \geq \dots \geq b_n \geq 0$ and let $b_{n+1} = \alpha_{n+1} = 0$. Note that we implicitly assume that $m = n$ by

adding slots with zero click rate or deleting extra slots (that will never be allocated, as each agent can receive at most one slot).

We now show that the GSP auction is a core auction. First, since $b_{i+1} \leq b_i$ for all i , we have that $b_{i+1}\alpha_i \leq b_i\alpha_i$, and hence each bidder pays at most her declared welfare for the allocation received.

It remains to show the second property of a core auction. Choose any subset of bidders $S \subseteq [n]$. The allocation \mathbf{y} to agents in S that maximizes declared welfare is the one that allocates greedily in index order. More formally, for each $i \in S$, let $\sigma(i)$ be 1 plus the number of elements of S with index less than i . For example, if $S = \{2, 6, 7\}$, then $\sigma(2) = 1$, $\sigma(6) = 2$, and $\sigma(7) = 3$. Then the declared-welfare-maximizing allocation \mathbf{y} to agents in S is such that $y_i = \alpha_{\sigma(i)}$ for each i , for a total declared welfare of $\sum_{i \in S} b_i \alpha_{\sigma(i)}$. The core auction property on subset of bidders S therefore reduces to showing that

$$\sum_{i \notin S} b_{i+1} \alpha_i + \sum_{i \in S} b_i \alpha_i \geq \sum_{i \in S} b_i \alpha_{\sigma(i)}. \quad (6.2)$$

To establish inequality (6.2), we first note that

$$\begin{aligned} \sum_{i \in S} b_i \alpha_{\sigma(i)} - \sum_{i \in S} b_i \alpha_i &= \sum_{i \in S} b_i (\alpha_{\sigma(i)} - \alpha_i) \\ &= \sum_{i \in S} \sum_{j=\sigma(i)}^{i-1} b_i (\alpha_j - \alpha_{j+1}) \\ &\leq \sum_{i \in S} \sum_{j=\sigma(i)}^{i-1} b_{j+1} (\alpha_j - \alpha_{j+1}). \end{aligned}$$

This final double summation contains an instance of the term $b_{j+1}(\alpha_j - \alpha_{j+1})$ for each $i \in S$ such that $\sigma(i) \leq j < i$. But for each j and each i such that $\sigma(i) \leq j < i$ (i.e., such that the “new” allocation to agent i under \mathbf{y} is slot j or better, and the “old” allocation is worse than slot j), there must be

some $k \leq j$ such that $k \notin S$. The number of such i is therefore at most the number of agents at index j or less that are not in S . More precisely, observe that $j \geq |\{i \in S : i \leq j\}| + |\{i \in S : \sigma(i) \leq j < i\}|$. This holds as $\sigma(i) \leq i$ for all i and is injective, so the sets on the right hand side are evidently disjoint and σ maps each such element to a unique index in $[j]$. But $j = |\{i \in S : i \leq j\}| + |\{k \notin S : k \leq j\}|$, so cancelling terms shows that $|\{k \notin S : k \leq j\}| \geq |\{i \in S : \sigma(i) \leq j < i\}|$. Thus, by rearranging the order of summation, we have

$$\begin{aligned}
\sum_{i \in S} \sum_{j=\sigma(i)}^{i-1} b_{j+1}(\alpha_j - \alpha_{j+1}) &\leq \sum_{j=1}^m b_{j+1}(\alpha_j - \alpha_{j+1}) \times |\{k \leq j : k \notin S\}| \\
&= \sum_{k \notin S} \sum_{j \geq k} b_{j+1}(\alpha_j - \alpha_{j+1}) \\
&\leq \sum_{k \notin S} \sum_{j \geq k} b_{k+1}(\alpha_j - \alpha_{j+1}) \\
&\leq \sum_{k \notin S} b_{k+1} \alpha_k
\end{aligned}$$

where the final inequality is a telescoping sum. We therefore conclude that

$$\sum_{i \in S} b_i \alpha_{\sigma(i)} - \sum_{i \in S} b_i \alpha_i \leq \sum_{k \notin S} b_{k+1} \alpha_k$$

and rearranging yields the desired inequality (6.2). \square

6.1.2 Gradient-Based Pacing Dynamics

With the above formalism, we may now define the class of algorithms our analysis will apply to. We first consider a budget-pacing algorithm motivated by stochastic gradient descent, which was introduced and analyzed by Balseiro and Gur [17] in the context of second-price auctions. See Algorithm 3. Each

bidder k maintains a *pacing multiplier* $\mu_{k,t} \in [0, \bar{\mu}]$ for each round t . Multiplier $\mu_{k,t}$ is determined by the algorithm before the value $v_{k,t}$ is revealed. The bid is set to $v_{k,t}/(1 + \mu_{k,t})$, or the remaining budget $B_{k,t}$ if the latter is smaller. Once the round's outcome is revealed, the multiplier is updated as per Line 5, where $P_{[a,b]}$ denotes the projection onto the interval $[a, b]$. Intuitively, the agent's goal is to keep expenditures near the expenditure rate ρ_k . Hence, if ρ_k is above (resp., below) the current expenditure $z_{k,t}$, the agent decreases (resp., increases) her multiplier, by the amount proportional to the current "deviation" from the target expenditure rate ρ_k .

Algorithm 3: Gradient-based pacing algorithm for agent k

Input: Budget B_k , time horizon T , step-size $\epsilon_k > 0$, pacing upper bound

$\bar{\mu}$

- 1 Initialize $\mu_{k,1} = 0$ (pacing multiplier), $B_{k,1} = B_k$ (remaining budget),
 $\rho_k = B_k/T$ (target spend rate).
 - 2 **for** round $t = 1, \dots, T$ **do**
 - 3 Observe value $v_{k,t}$, submit bid $b_{k,t} = \min \{v_{k,t}/(1 + \mu_{k,t}), B_{k,t}\}$;
 - 4 Observe expenditure $z_{k,t}$;
 - 5 Update $\mu_{k,t+1} \leftarrow P_{[0,\bar{\mu}]}(\mu_{k,t} - \epsilon_k(\rho_k - z_{k,t}))$ and $B_{k,t+1} \leftarrow B_{k,t} - z_{k,t}$.
-

Parameters $\bar{\mu}$ and ϵ_k will be specified later. While the upper bound $\bar{\mu}$ could in general depend on the agent k , as in $\bar{\mu} = \bar{\mu}_k$, we suppress this dependence for the sake of clarity. Our bounds depend on $\max_k \bar{\mu}_k$.

Remark 6.2. *One key feature of Algorithm 3 is that it never overbids, in the sense that the bids $b_{k,t}$ are always upper-bounded by the respective values $v_{k,t}$. This restriction is provably necessary with quasilinear utilities in truthful auctions, like second-price, but this may not hold in nontruthful auctions like first-price. Nonetheless, our results will*

still apply in these settings where pacing is viewed as a natural behavioral assumption.

Motivation for Algorithm 3. While the motivation behind Algorithm 3 will not be needed to understand our arguments, we briefly explain where this algorithm comes from. The below discussion follows Balseiro and Gur [17]. For simplicity, we focus on a single agent so drop the k indexing and let $(v_1, \dots, v_T) \in [0, \bar{v}]^T$ and $(d_1, \dots, d_T) \in [0, \bar{v}]^T$ be any *fixed* set of valuations and maximum competing bids for an agent. The best *any* algorithm could do, with hindsight, is precisely the solution to the following integer program:

$$\pi^H(\mathbf{v}, \mathbf{d}) \triangleq \max_{x_1, \dots, x_T \in \{0, 1\}} \sum_{t=1}^T x_t(v_t - d_t) \quad (6.3)$$

$$\text{such that } \sum_{t=1}^T x_t d_t \leq B. \quad (6.4)$$

Ignoring the integrality constraint⁶, the Lagrangian dual of Equation (6.3) is exactly

$$\mathcal{L}(\mathbf{x}, \mu) \triangleq \left(\sum_{t=1}^T x_t(v_t - d_t) \right) + \mu \left(\rho_k T - \sum_{t=1}^T x_t d_t \right) \quad (6.5)$$

$$= \sum_{t=1}^T [x_t v_t - (1 + \mu)x_t d_t + \rho_k \mu] \quad (6.6)$$

For any fixed value of μ , the optimal value of x is given by $x_t = 1$ if and only if $v_t \geq (1 + \mu) \cdot d_t$, and therefore, we find from duality that

$$\pi^H(\mathbf{v}, \mathbf{d}) \leq \inf_{\mu \geq 0} \sum_{t=1}^T [(v_{k,t} - (1 + \mu)d_t)^+ + \rho_k \mu] = \inf_{\mu \geq 0} \sum_{t=1}^T \psi_t(\mu). \quad (6.7)$$

The insight of prior work is that one can construct algorithms that tries to find the optimal value of μ of the RHS. Note that a subdifferential of ψ_t is

$$\partial \psi_t(\mu) = -\mathbf{1}\{v_t \geq (1 + \mu)d_t\}d_t + \rho_k = \rho_k - z_t.$$

⁶The integrality constraint does not meaningfully affect the analysis so long as valuations are much smaller than the budget.

In particular, if one tried to run a gradient descent-style algorithm to tune μ , one would update by subtracting a small multiple of the above. This exactly leads to Algorithm 3.

Balseiro and Gur show that Algorithm 3 has a number of favorable properties for suitable choices of the learning rate: for one, under adversarial valuations and competing bids, their algorithm attains the optimal approximation ratio in general. Moreover, under various convexity assumptions, they show that their algorithm attains no-regret with respect to sufficiently well-behaved stochastic competition and actually converges quickly to the (unique) pacing equilibrium under strong monotonicity when all agents use such algorithms.

6.1.3 Liquid Welfare

The notion of welfare that we will consider in this chapter is *liquid welfare*, a notion first introduced by Dobzinski and Paes Leme [54] in the context of auctions with budgets. We use the notation introduced in the previous section.

Definition 6.2. Given a sequence of valuation profiles $\mathbf{v}_1, \dots, \mathbf{v}_T \in \mathbb{R}_+^n$ and a sequence of allocations $\mathbf{x}_1, \dots, \mathbf{x}_T \in X$, the **liquid welfare** obtained by agent i is given by

$$\text{LW}_i(\mathbf{x}_1, \dots, \mathbf{x}_T; \mathbf{v}_1, \dots, \mathbf{v}_T) = \min \left\{ \sum_{t=1}^T x_{i,t} v_{i,t}, B_i \right\}. \quad (6.8)$$

The aggregate liquid welfare is defined as

$$\text{LW}(\mathbf{x}_1, \dots, \mathbf{x}_T; \mathbf{v}_1, \dots, \mathbf{v}_T) = \sum_{i=1}^n \text{LW}_i(\mathbf{x}_1, \dots, \mathbf{x}_T; \mathbf{v}_1, \dots, \mathbf{v}_T).$$

To understand this definition, observe that the first term in Equation (6.8) is

precisely the value of agent i for her received allocation. With quasilinear utilities, this term is therefore equivalent to the maximum amount she would be willing to pay to receive this allocation, or equivalently, the maximum amount of revenue an omniscient auctioneer could extract from her. By taking the minimum, the same interpretation holds in the budgeted setting; the liquid welfare of an agent denotes her willingness to pay for her allocation while respecting her budget constraints. The perspective of liquid welfare is that it is incongruous to “give credit” for an agent’s valuation for an allocation should she not be willing to pay for this valuation due to budget constraints.

As a brief aside, we briefly comment on the myriad advantages of the above formulation of welfare. Beyond the interpretation discussed above, notice that when agents are not budget-constrained so that $B_i = \infty$, the above specializes to standard utilitarian welfare. When budgets are nontrivial, these standard welfare notions are easily seen to be impossible to achieve via any market outcome: low valuation bidders with large budgets can dominate the allocation by outbidding high valuation bidders with small budgets, ensuring that goods will often be misallocated to low valuation buyers. Liquid welfare corrects for this by capping the counterfactual welfare that could be attained by such buyers under any allocation rule, as these allocations are not reasonable. In particular, liquid welfare is known to be an achievable benchmark in price of anarchy analyses [54, 13, 101, 61].

6.2 Main Result: Optimal Liquid Welfare for GBP

In this section, we finally turn to the statement and proof of our main result: gradient-based pacing achieves optimal liquid welfare guarantees. We show that when all agents use Algorithm 3, the expected liquid welfare is at least half of the optimal, minus regret term. In fact, our guarantee applies to a wider class of algorithms that do not overbid ($b_{k,t} \leq v_{k,t}$), bid their full value when $\mu_{k,t} = 0$, and are not constrained otherwise.

Definition 6.3. Consider a measurable bidding algorithm which inputs the same parameters as Algorithm 3, internally updates multipliers $\mu_{k,t}$ in the same way, and never overbids: $b_{k,t} \leq v_{k,t}$ for all rounds t . Call it a **generalized pacing algorithm** if $(\mu_{k,t} = 0) \Rightarrow (b_{k,t} = v_{k,t})$ for all rounds t (“no unnecessary pacing”).

We use a stronger benchmark: optimal *ex ante* liquid welfare, which is each agent’s willingness to pay for her *expected* allocation sequence. It upper-bounds the expected (ex post) liquid welfare by Jensen’s inequality.

Definition 6.4. Fix distribution F over valuation profiles and allocation rule $\mathbf{y} : [0, \bar{v}]^n \rightarrow X$. The *ex ante liquid value* of each agent k is $\overline{\text{LW}}_k(\mathbf{y}, F) = T \times \min\{\rho_k, \mathbb{E}_{\mathbf{v} \sim F}[y_k(\mathbf{v}) v_k]\}$ and the ex ante liquid welfare is $\overline{\text{LW}}(\mathbf{y}, F) = \sum_{k=1}^n \overline{\text{LW}}_k(\mathbf{y}, F)$.

In this definition, the restriction that the same allocation rule \mathbf{y} is used in every round is without loss of generality. Indeed, given any allocation sequence rule there is a single-round allocation rule with the same ex ante liquid welfare:

Lemma 6.3. Let $\tilde{\mathbf{y}} : [0, \bar{v}]^{nT} \rightarrow X^T$ be an allocation sequence rule that takes in the entire sequence $\mathbf{v}_1, \dots, \mathbf{v}_T$ and allocates $\tilde{y}_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T)$ units to agent k at time t . Then there

exists a (single-round) allocation rule $y : [0, \bar{v}]^n \rightarrow X$ such that

$$\begin{aligned} \overline{\text{LW}}(\tilde{\mathbf{y}}, F) &\triangleq \sum_{k=1}^n \min \left\{ B_k, \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_T \sim F} \left[\sum_{t=1}^T \tilde{y}_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) v_{k,t} \right] \right\} \\ &= \sum_{k=1}^n T \cdot \min \{ \rho_k, \mathbb{E}_{\mathbf{v} \sim F} [y_k(\mathbf{v}) v_k] \} = \overline{\text{LW}}(\mathbf{y}, F). \end{aligned}$$

Proof. For each t , by slightly abusing notation, we define an allocation rule $\hat{\mathbf{y}}_t : [0, \bar{v}]^n \rightarrow [0, 1]^n$ by

$$\hat{y}_{k,t}(\mathbf{v}_t) \triangleq \mathbb{E}_{\mathbf{v}_{-t} \in F^{T-1}} [y_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) | \mathbf{v}_t],$$

Note that this is a feasible allocation rule as the set of feasible allocations is convex and closed. We have

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_T \sim F} [y_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) \cdot v_{k,t}] &= \mathbb{E}_{\mathbf{v}_t} [\mathbb{E}_{\mathbf{v}_{-t}} [y_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) \cdot v_{k,t} | \mathbf{v}_t]] \\ &= \mathbb{E}_{\mathbf{v}_t} [\hat{y}_{k,t}(\mathbf{v}_t) \cdot v_{k,t}] = \mathbb{E}_{\mathbf{v} \sim F} [\hat{y}_{k,t}(\mathbf{v}) \cdot v_k]. \end{aligned}$$

Now we define the allocation rule $\tilde{\mathbf{y}}$ by setting $\tilde{y}_k = \frac{1}{T} \sum_{t=1}^T \hat{y}_{k,t}$ for each $k \in [n]$, which is again feasible because the set of feasible allocations is convex. By the linearity of the expectations operator, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_T \sim F} \left[\sum_{t=1}^T y_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) \cdot v_{k,t} \right] &= \sum_{t=1}^T \mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_T \sim F} [y_{k,t}(\mathbf{v}_1, \dots, \mathbf{v}_T) \cdot v_{k,t}] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathbf{v} \sim F} [\hat{y}_{k,t}(\mathbf{v}) \cdot v_k] = \mathbb{E}_{\mathbf{v} \sim F} \left[\sum_{t=1}^T \hat{y}_{k,t}(\mathbf{v}) \cdot v_k \right] = T \cdot \mathbb{E}_{\mathbf{v} \sim F} [\tilde{y}_k(\mathbf{v}) \cdot v_k]. \end{aligned}$$

This proves the desired result. \square

Now we are ready to state the main result: liquid welfare guarantees for pacing dynamics.

Theorem 6.4. *Fix any core auction and any distribution F over agent value profiles. Suppose that each agent k employs a generalized pacing algorithm to bid, possibly with a*

different step-size ε_k . Write $\mathbf{x}: [0, \bar{v}]^{nT} \rightarrow X^T$ for the corresponding allocation sequence rule. Then for any allocation rule $\mathbf{y}: [0, \bar{v}]^n \rightarrow X$ we have

$$\mathbb{E}_{\mathbf{v}_1, \dots, \mathbf{v}_T \sim F}[\text{LW}(\mathbf{x}(\mathbf{v}_1, \dots, \mathbf{v}_T))] \geq \frac{\overline{\text{LW}}(\mathbf{y}, F)}{2} - O\left(n\bar{v}\sqrt{T \log(\bar{v}nT)}\right). \quad (6.9)$$

We note that this factor of 1/2 is essentially optimal; it is known (c.f. Theorem 6.5) that this factor is optimal even for the class of offline pacing equilibria (c.f. Definition 6.5 when all valuations and budgets are known). Therefore, Theorem 6.4 achieves essentially optimal performance guarantees for a natural class of learning algorithms that have been extensively employed in theory and practice.

6.3 Overview of Proof and Challenges

In this section, we give an overview of the challenges and ideas that go into the proof of Theorem 6.4. We first discuss why two natural approaches to this problem seem to fail before discussing our approach.

Idea #1: Welfare from Equilibrium Convergence. A natural approach to this problem would be to argue as follows: suppose that the learning dynamics induced by all agents employing Algorithm 3 *converges* to a tuple of pacing multiplier $\boldsymbol{\mu}^* = (\mu_1^*, \dots, \mu_n^*)$ that forms a *pacing equilibrium* in our setting (as defined below). In that case, we could reduce the problem of understanding the aggregate guarantees of learning dynamics to the aggregate guarantees of equilibrium, a completely static notion.

The notion of pacing games and pacing equilibria was initially introduced by Conitzer, et al [43]. For us, the relevant definition is as follows:

Definition 6.5 (Definition 2 of [43]). Given valuations $\mathbf{v}_1, \dots, \mathbf{v}_T \in [0, \bar{v}]^n$ for all agents at each time, a *pacing equilibrium* is an allocation \mathbf{x} and set of pacing multipliers $\alpha \in [0, 1]^n$ satisfying

1. For all t , $\sum_{k=1}^n x_{k,t} \leq 1$ with equality if and only if some agent has $v_{k,t} > 0$ and such that $x_{k,t} > 0$ implies $b_{k,t} = \alpha_k v_{k,t}$ is (tied) for the highest bid.
2. If $x_{k,t} > 0$, the per-unit price p_t of the t 'th item is the second highest bid for the item.
3. For all $k \in [n]$, $\sum_{t=1}^T x_{k,t} p_t \leq B_k$; moreover, if $\sum_{t=1}^T x_{k,t} p_t < B_k$, then $\alpha_k = 1$.

Conitzer, et al [43] show that pacing equilibria always exist for any sequence of valuations; this result has since been generalized to satisfy additional constraints [3, 14]. One can extend this equilibrium notion to hold for a distribution over values that are sampled at each time (see [17] for such an extension). Regardless, in our notation, we might hope that the pacing multipliers converge in some sense to a vector $\boldsymbol{\mu}^* = 1 + 1/\alpha$, where we interpret this expression coordinate-wise, and α is such a pacing equilibrium.

Beyond converting a dynamic learning problem to a static equilibrium problem, the aggregate guarantees of pacing equilibria have been well-studied in the last few years. It has been shown by Aggarwal, Badanidiyuru, and Mehta [3] and Babaioff, et al [14] in somewhat more general settings that any pacing equilibrium achieves at least half the optimal liquid welfare, and that this factor cannot be improved:

Theorem 6.5 (Theorem 4.1 of [14], simplified). *Let \mathbf{x}^* be the (possibly fractional) allocation at any pacing equilibrium α^* and let \mathbf{y} be any other allocation. Then*

$$\text{LW}(\mathbf{x}^*) \geq \frac{\text{LW}(\mathbf{y})}{2}. \quad (6.10)$$

In particular, the liquid welfare at equilibrium is at least twice the optimal liquid welfare over any allocation. Moreover, this factor cannot be improved.

The idea behind Theorem 6.5 is a *charging* argument that crucially relies on the following observation: at pacing equilibrium, agents can be partitioned into *pacers* (who exhaust their budget) and *non-pacers* (that. Any agent that *does not* spend their whole budget must bid their true value at each time, and therefore must set a comparatively high price on any item she would receive in an arbitrary allocation y but not in equilibrium. Therefore, we may charge her welfare loss to the prices paid by the other agents. Meanwhile, agents that exhaust their budget must be attaining maximal liquid welfare, as their expenditure lower bounds their valuation.

To use this approach, we also note that Balseiro and Gur further show that Algorithm 3 *will* actually converge to a pacing equilibrium in an appropriate sense (Theorem 4.3 of [17]). The caveat is that this result requires extremely strong vector generalizations of strong convexity (known as strong monotonicity) to hold for the associated expected Lagrangian functions that essentially enforce that joint gradient based methods will converge to equilibrium (Assumption 4.1 of [17]). To employ this reduction, we would need to argue that these convexity assumptions are not necessary.

However, recent work by Chen, Kroer, and Kumar [38] shows that this is rather unlikely; it turns out that task of computing even approximate pacing equilibria is PPAD-hard for general valuations. We omit a more formal discussion of the complexity class PPAD, but this result implies that we should not expect any efficient algorithm to construct approximate pacing equilibria, let alone decentralized learning dynamics as in Algorithm 3. Therefore, it would

seem that we need a new approach.

Idea #2: Welfare from No-Regret. Since a reduction to an equilibrium analysis appears to be out of the question, an alternative approach would be attempt to leverage no-regret properties of learning dynamics. One can show that Algorithm 3 will satisfy some form of no-regret condition with respect to a certain expected Lagrangian dual function that was introduced by Balseiro and Gur. The reason is simply that this algorithm runs stochastic gradient descent on the expected Lagrangian, implying some form of no-regret (see, for example [88]). However, it is not immediately clear how to relate these performance guarantees of the expected Lagrangians, which themselves will depend on the pacing multipliers of the other agents, to our goals.

In fact, it turns out that even if Algorithm 3 satisfied no-regret with respect to quasilinear utility or value maximization for each agent, it is not possible to conclude any constant factor approximation to liquid welfare. To see this, we present a simple example showing that a dynamic bidding strategy that achieves vanishing (in fact, zero) regret for all agents might still have an unbounded approximation factor with respect to the liquid welfare of the allocation sequence. In particular, it does not seem possible to directly leverage no-regret properties of learning dynamics to the liquid welfare in our setting.

Proposition 6.6. *There exists an instance of a second-price auction for a single good and two agents, described by a distribution over valuations, and a pair of no-regret bidding strategies such that when both agents follow the pair of no-regret bidding strategies, the resulting expected liquid welfare is arbitrarily small compared to the optimal liquid welfare.*

Proof. We begin by showing something simpler: there exists a distribution over

valuations and a bidding strategy for each agent such that, under these strategies, both agents have zero regret. We will then show how to extend this construction to one in which each agent is using a truly no-regret bidding strategy.

The per-round auction in our example is a second-price auction for a single good. There are two agents. The distribution F over value profiles is such that $(v_1, v_2) = (2, 1)$ with probability 1. The target per-round spend rates for the agents are $(\rho_1, \rho_2) = (1/(1 + \bar{\mu}), 1)$, where $\bar{\mu}$ is some arbitrarily large constant independent of T .

Consider the following bidding strategies for the agents. Agent 1 bids value 2 for all periods and agent 2 bids 0 for all periods. Under this strategy profile, agent 1 receives all the items over the T rounds, and both agents pay nothing. Agent 1 actually achieves zero regret as she obtains the maximum possible value. Agent 2 likewise has zero regret, since no choice of bid less than v_2 can cause her to win in any round. Note that the agents would still have zero regret if their objective were changed to maximizing value minus (any scalar multiple $\lambda \in [0, 1]$ times) payments.

The liquid welfare of this equilibrium is $T/(\bar{\mu} + 1)$, the total budget of agent 1. However, allocating all goods to agent 2 achieves a liquid welfare of T . Since $\bar{\mu}$ is an arbitrarily large constant, this approximation factor is unbounded.

This idea can easily be modified so that the bidding strategies of both agents are no-regret with respect to *any* adversary, not just their current opponent. Let σ be any no-regret bidding strategy against any adversary (ignoring any issues of existence). Let σ_1 be the bidding strategy that always bid value v if the historical bids of the opponent is always 0, and follows the bidding strategy σ

otherwise. Let σ_2 be the bidding strategy that always bid 0 if the historical bids of the opponent is always 2, and follows the bidding strategy σ otherwise. Note that given any adversary, strategy σ_1 achieves no-regret for agent 1 given any adversary, and strategy σ_2 achieves no-regret for agent 2. This is because both agents achieve no-regret when the adversary proceeds as anticipated, and the remaining budget of the agent until detecting any deviation does not decrease. Finally, when both agents follow the strategy profile (σ_1, σ_2) , no agent will deviate to strategy σ and the equilibrium liquid welfare is $T/(\bar{\mu} + 1)$, which can be unbounded smaller than the optimal liquid welfare. \square

Our Approach: Epoch Decompositions: Due to the apparent difficulties with either of the previous standard approaches, we instead take a different route. One easy observation is that since B_k is an upper bound on the liquid welfare obtainable by agent k , any agents who are (approximately) exhausting their budgets in our dynamics are achieving optimal liquid welfare. We therefore focus on agents who do not exhaust their budgets. We'd like to argue that such agents are often bidding very high, frequently choosing pacing multipliers equal to 0 (i.e., bidding their values). This would be helpful because our auction is assumed to be a core auction, which implies that either the high-bidding agents are winning (and generating high liquid welfare) or other, budget-exhausting agents are generating high revenue for the seller (which likewise implies high liquid welfare).

Why should agents who are underspending their budget be placing high bids in many rounds? While it's true that the pacing dynamics increases the next-round bid whenever spend is below the per-round target, this is only a local adjustment and does not depend on total spend. So how do we analyze

bidding patterns in aggregate across rounds?

It is here where we use the fact that agents use generalized pacing algorithms. We introduce the notion of an *epoch* in Definition 6.6, which roughly corresponds to a maximal contiguous sequence of rounds in which an agent's pacing multiplier is strictly greater than 0. The key idea behind this definition is that on each distinct epoch, we may directly reason about the value being obtained by the agent by comparing it to the evolution of their pacing multiplier. In Lemma 6.8 we show that an agent k 's total spend over an epoch of length t must be (approximately) $t\rho_k$. In other words, the average spend over an epoch approximately matches the target per-round spend. This is a direct implication of the update rule for the pacing multiplier: since multiplier increases and decreases balance out (approximately) over the course of an epoch, the budget deficits and surpluses must balance out as well. An immediate implication is that an agent whose total spend is much less than $T\rho_k$ must often have her pacing multiplier set equal to 0. To get the optimal constant bound, we will actually require a simple, but somewhat subtle strengthening of this value lower bound.

To summarize, each agent either spends most of her budget by time T or spends many rounds bidding her value. It might seem that by combining these two cases we should obtain a constant approximation factor, and not just in expectation but for every realized value sequence \mathbf{v} . But this is too good to be true. Indeed, this proof sketch misses an important subtlety: whether an agent exhausts her budget or not depends on the realization of the value sequence, which is also correlated with the benchmark allocation \mathbf{y} . For example, what if an agent under-spends her budget precisely on those value sequences where it would have been optimal (according to the benchmark \mathbf{y}) for her liq-

uid welfare to equal her budget? This could result in a situation where, conditional on $\mu_{k,t} = 0$, the value obtained by the benchmark is much higher than expected and cannot be approximated. To compare against the liquid welfare of \mathbf{y} we must control the extent of this correlation. We employ a variation of the Azuma-Hoeffding inequality (Lemma 6.9) to argue that since value realizations are independent across time, the pacing sequence and the benchmark allocation cannot be too heavily correlated.

6.4 Proof of the Main Result

6.4.1 Proof Preliminaries

Before getting into the details of the proof of Theorem 6.4, we first introduce some notation and define an epoch more formally. Write $\mu_{k,t} = \flat$ if at time t , the agent's algorithm has stopped, i.e., if the agent is out of money or if $t = T + 1$. For $t_1 \leq t_2 \in \mathbb{N}$, we slightly abuse notation and write $[t_1, t_2] \triangleq \{t_1, \dots, t_2\}$ to be the set of integers between them, inclusive, when the meaning is clear. Similarly, we write $[t_1, t_2) \triangleq \{t_1, \dots, t_2 - 1\}$ for the half-open set of integers between them and analogously for $(t_1, t_2]$ and (t_1, t_2) .

Definition 6.6. Fix an agent k , time horizon T , and a sequence of pacing multipliers $\mu_{k,1}, \dots, \mu_{k,T}$. A half-open interval $[t_1, t_2)$ is an **epoch** with respect to these multipliers if it holds that $\mu_{t_1} = 0$ and $\mu_t > 0$ for each $t_1 < t < t_2$ and t_2 is maximal with this property.

Remark 6.7. *Because we initialized any generalized pacing algorithm at $\mu_{k,1} = 0$ for all $k \in [n]$, the epochs completely partition the set of times the agent is bidding. Moreover,*

if the pacing multipliers at times t_1 and $t_1 + 1$ are both zero, then $[t_1, t_1 + 1)$ is an epoch; we refer to this as a trivial epoch.

The following lemma shows that an agent's total spend over a maximal epoch can be bounded from below by an amount roughly equal to the target spend rate ρ_k times the epoch length, plus an adjustment for the first round of the epoch.

Lemma 6.8. *Fix an agent k , fix any choice of core auction and any sequence of bids of other agents $(\mathbf{b}_{-k,1}, \dots, \mathbf{b}_{-k,T})$. Fix any realization of values $(v_{k,1}, \dots, v_{k,T})$ for agent k and suppose that $\mu_{k,1}, \dots, \mu_{k,T}$ is the sequence of multipliers generated by the generalized pacing algorithm given the auction format and the other bids. Then for any epoch $[t_1, t_2)$ where $\mu_{k,t_2} \neq 0$, we have*

$$\sum_{t=t_1}^{t_2-1} x_{k,t} v_{k,t} \geq x_{k,t_1} v_{k,t_1} - z_{k,t_1} + \rho_k \cdot (t_2 - t_1 - 1). \quad (6.11)$$

Proof. The assumption that $\mu_{k,t_2} \neq 0$ means that the agent participates in all periods of $[t_1, t_2)$. Moreover, if $t_2 = t_1 + 1$, then Equation (6.11) is trivial as the third term on the right hand of side of inequality (6.11) is zero and $z_{k,t_1} \geq 0$. Therefore, we may assume $t_2 \geq t_1 + 2$.

By the definition of an epoch, there is no negative projection in the dynamics on the epoch until possibly time t_2 . I.e., $\mu_{k,t} > 0$ for all $t_1 \leq t < t_2$. The pacing recurrence condition implies that

$$\begin{aligned} 0 < \mu_{k,t_2-1} &= P_{[0,\bar{\mu}]}\left(\mu_{k,t_2-2} + \epsilon(z_{k,t_2-2} - \rho_k)\right) \leq \mu_{k,t_2-2} + \epsilon(z_{k,t_2-2} - \rho_k) \\ &\leq \epsilon \sum_{t=t_1}^{t_2-2} (z_{k,t} - \rho_k) = \epsilon \left(\sum_{t=t_1}^{t_2-2} z_{k,t} - (t_2 - t_1 - 1)\rho_k \right). \end{aligned}$$

The first and second inequalities follow because the multipliers are positive during an epoch. The third inequality follows from applying the second inequality repeatedly. Note that the inequality holds even if there is a positive projection during the epoch (i.e., if $\mu_{k,t} = \bar{\mu}$ for some $t \in [t_1, t_2)$).

Thus, the expenditure of agent k on this epoch is at least $\sum_{t=t_1}^{t_2-2} z_{k,t} \geq (t_2 - t_1 - 1)\rho_k$. Let us now consider the *value* obtained by the agent on this epoch. Because agents never overbid in the pacing algorithm and because payments are always lower than the value in a core auction, the value obtained by the agent on the epoch is at least the expenditure, which we just lower bounded. To get a slight sharpening of this, note that on the first period of the epoch, the agent actually receives $x_{k,t_1} v_{k,t_1}$ value and pays z_{k,t_1} , which we know is at most $x_{k,t_1} v_{k,t_1}$ from the first property of core auctions and the no overbidding condition. Therefore, we can trade z_{k,t_1} expenditure for $x_{k,t_1} v_{k,t_1}$ value. It follows from the above bound and this observation that

$$\sum_{t=t_1}^{t_2-1} x_{k,t} v_{k,t} \geq x_{k,t_1} v_{k,t_1} - z_{k,t_1} + \sum_{t=t_1}^{t_2-2} z_{k,t} \geq x_{k,t_1} v_{k,t_1} - z_{k,t_1} + (t_2 - t_1 - 1)\rho_k. \quad \square$$

This result will prove vital towards proving Theorem 6.4, as it provides a very convenient lower bound on the values obtained by pacing on an epoch-by-epoch scale.

Next, motivated by the intuition in Section 6.3, we introduce a concentration inequality that will be helpful for our analysis. Roughly speaking, we will use this lemma to show that the sequence of values obtained by agent k in the benchmark on rounds in which $\mu_{k,t} = 0$ are not “far from expectation,” in the sense that the total expected value obtained over such rounds is not much greater than ρ_k per round.

Lemma 6.9. Let Y_1, \dots, Y_T be random variables and $\mathcal{F}_0 \subseteq \dots \subseteq \mathcal{F}_T$ be a filtration such that:

1. $0 \leq Y_t \leq \bar{v}$ with probability 1 for some parameter $\bar{v} \geq 0$ for all t .
2. $\mathbb{E}[Y_t] \leq \rho$ for some parameter $\rho \geq 0$ for all t .
3. For all t , Y_t is \mathcal{F}_t -measurable but is independent of \mathcal{F}_{t-1} .

Suppose that $X_1, \dots, X_n \in [0, 1]$ are random variables such that X_t is \mathcal{F}_{t-1} -measurable.

Then

$$\Pr\left(\sum_{t=1}^T X_t Y_t + (1 - X_t)\rho \geq \rho \cdot T + \theta\right) \leq \exp\left(\frac{-2\theta^2}{T\bar{v}^2}\right). \quad (6.12)$$

The proof of Lemma 6.9 is deferred to the Chapter Notes, as it is a somewhat technical but straightforward application of concentration bounds.

We may finally proceed with the proof:

Proof of Theorem 6.4. We first introduce some notation. Fix the generalized pacing dynamics algorithm used by each agent. We will then write $\mathbf{x} = \{x_{k,t}\}_{k \in [n], t \in [T]}$ for the random variable corresponding to the allocation obtained under these bidding dynamics, given the values $\{v_{k,t}\}$. We also write $\mu_{k,t}$ for the pacing multiplier of agent k in round t , and $z_{k,t}$ for the realized spend of agent k in round t , which are likewise random variables. For notational convenience we will write $\text{WEL}_{\text{GPD}}(\mathbf{v})$ for the liquid welfare (where GPD stands for ‘‘Generalized Pacing Dynamics’’) given valuation sequence \mathbf{v} . That is,

$$\text{WEL}_{\text{GPD}}(\mathbf{v}) \triangleq \sum_{k=1}^n \min\left\{B_k, \sum_{t=1}^T x_{k,t} v_{k,t}\right\}. \quad (6.13)$$

We also write $\text{WEL}_{k,\text{GPD}}(\mathbf{v})$ for the liquid welfare obtained by agent k , and $\text{WEL}_{\text{GPD}}(F)$ for the total expected liquid welfare $\mathbb{E}_{\mathbf{v} \sim F^T}[\text{WEL}_{\text{GPD}}(\mathbf{v})]$.

We claim that to prove theorem 6.4, it is sufficient to show that inequality (6.9) holds for any allocation rule \mathbf{y} such that

$$\mathbb{E}_{\mathbf{v} \sim F}[y_k(\mathbf{v})v_k] \leq \rho_k \quad \text{and} \quad \overline{\text{LW}}_k(\mathbf{y}, F) = T \cdot \mathbb{E}_{\mathbf{v} \sim F}[y_k(\mathbf{v})v_k] \leq \rho_k \cdot T \quad \forall k \in [n]. \quad (6.14)$$

This is sufficient because if one of the above conditions is violated, one can always decrease the allocation for agent k , which maintains the feasibility (since we assume that the set of feasible allocations X is downward closed) without affecting the ex ante liquid welfare $\overline{\text{LW}}(\mathbf{y}, F)$. We will therefore assume without loss that \mathbf{y} satisfies (6.14).

Preliminaries completed, we now prove Theorem 6.4 in three steps. First, we will define a “good” event in which the benchmark allocations are not too heavily correlated with the pacing multipliers of the generalized pacing dynamics, and show that this good event happens with high probability. Second, for all valuation sequences \mathbf{v} that satisfy the good event, we bound the liquid welfare obtained by the pacing dynamics in terms of the benchmark allocation and the payments collected by the auctioneer. In the third and final step we take expectations over all valuation profiles to bound the expected liquid welfare.

Step 1: A Good Event. For each agent $k \in [n]$, we define the following quantity, whose significance will become apparent shortly:

$$R_k(\mathbf{v}) \triangleq \sum_{t=1}^T [\mathbf{1}\{\mu_{k,t} = 0\}y_k(\mathbf{v})v_{k,t} + \mathbf{1}\{\mu_{k,t} \neq 0\}\rho_k].$$

$R_k(\mathbf{v})$ is the total value obtained by agent k under allocation rule \mathbf{y} , except that in any round in which $\mu_{k,t} \neq 0$ this value is replaced by the target spend ρ_k .

We would like to apply Lemma 6.9 to bound $R_k(\mathbf{v})$. Some notation: we will write $Y_t = y_k(\mathbf{v}_t)v_{k,t}$ and $X_t = \mathbf{1}\{\mu_{k,t} = 0\}$ with $\mathcal{F}_t = \sigma(\mathbf{v}_1, \dots, \mathbf{v}_t)$. Then because $\mu_{k,t}$

is \mathcal{F}_{t-1} measurable and the sequence $\{v_j\}$ is a sequence of independent random variables, Lemma 6.9 implies that with probability at least $1 - 1/(\bar{v}nT)^2$, we have

$$\sum_{t=1}^T [\mathbf{1}\{\mu_{k,t} = 0\}y_k(\mathbf{v})v_{k,t} + \mathbf{1}\{\mu_{k,t} \neq 0\}\rho_k] \leq \rho_k \cdot T + \bar{v} \sqrt{T \log(\bar{v}nT)}.$$

Taking a union bound over $k \in [n]$, with probability at least $1 - 1/(\bar{v}nT)^2$ over the randomness in the sequence $\mathbf{v} = (v_1, \dots, v_T)$, we have that

$$R_k(\mathbf{v}) \leq \rho_k \cdot T + \bar{v} \sqrt{T \log(\bar{v}nT)}, \quad \forall k \in [n]. \quad (6.15)$$

We will write E_{GOOD} for the event in which (6.15) holds. Going back to the intuition provided before the statement of Lemma 6.9, E_{GOOD} is the event that the value each agent obtains in the benchmark allocation \mathbf{y} is not “too high” on rounds in which their pacing multipliers are 0.

Step 2: A Bound on Liquid Welfare For “Good” Value Realizations. Fix any realized sequence v_1, \dots, v_T such that Equation (6.15) holds. We will now proceed to derive a lower bound on the liquid welfare of the agents (under allocation \mathbf{x}) by considering the two different possible cases for $\text{WEL}_{k,\text{GPD}}(\mathbf{v})$. Recall that the liquid welfare $\text{WEL}_{k,\text{GPD}}(\mathbf{v})$ of any agent k is either $B_k = \rho_k \cdot T$ or is $\sum_{t=1}^T x_{k,t}v_{k,t}$. For any k such that $\text{WEL}_{k,\text{GPD}}(\mathbf{v}) = B_k$, we obtain via Equation (6.15) that

$$\text{WEL}_{k,\text{GPD}}(\mathbf{v}) = B_k \geq R_k(\mathbf{v}) - \bar{v} \sqrt{T \log(\bar{v}nT)}. \quad (6.16)$$

We now consider all of the remaining agents. Let $A \subseteq [n]$ be the set of agents k such that $\sum_{t=1}^T x_{k,t}v_{k,t} < B_k$ on this realized sequence \mathbf{v} , and hence $\text{WEL}_{k,\text{GPD}}(\mathbf{v}) = \sum_{t=1}^T x_{k,t}v_{k,t}$. That is, each of their contributions to the liquid welfare on this realized sequence is uniquely determined by their true value for winning the items in the auction. Note that this further implies that no agent in

A runs out of budget early because the generalized pacing algorithm does not allow overbidding (i.e., bidding above the value). Thus for all $k \in A$, $\mu_{k,T} \neq b$ and $\mu_{k,t} \geq 0$ surely for all t . We claim that

$$\sum_{k \in A} \text{WEL}_{k, \text{GPD}}(\mathbf{v}) = \sum_{k \in A} \sum_{t=1}^T x_{k,t} v_{k,t} \geq \sum_{k \in A} R_k(\mathbf{v}) - \sum_{k \in [n]} \sum_{t=1}^T z_{k,t}. \quad (6.17)$$

To show that the inequality holds, we partition the interval $[1, T]$ into maximal epochs for each agent k and bound the value obtained by agent k on each maximal epoch separately. Fix any agent $k \in A$ and suppose that $[t_1, t_2)$ is a maximal epoch inside $[1, T]$. Since agent k does not run out of budget on the entire interval $[1, T]$, she does not run out of budget on this epoch in particular. We can therefore apply Lemma 6.8 to obtain

$$\sum_{t=t_1}^{t_2-1} x_{k,t} v_{k,t} \geq x_{k,t_1} v_{k,t_1} - z_{k,t_1} + \rho_k \cdot (t_2 - t_1 - 1).$$

Since $[1, T]$ can be partitioned into maximal epochs for each agent $k \in A$, we can sum over all time periods and apply the definition of a maximal epoch to conclude that

$$\sum_{t=1}^T x_{k,t} v_{k,t} \geq \sum_{t=1}^T [\mathbf{1}\{\mu_{k,t} = 0\}(x_{k,t} v_{k,t} - z_{k,t}) + \mathbf{1}\{\mu_{k,t} \neq 0\} \cdot \rho_k].$$

Summing over all $k \in A$ and changing the order of the summations implies that

$$\sum_{k \in A} \sum_{t=1}^T x_{k,t} v_{k,t} \geq \sum_{t=1}^T \sum_{k \in A} [\mathbf{1}\{\mu_{k,t} = 0\}(x_{k,t} v_{k,t} - z_{k,t})] + \sum_{k \in A} \sum_{t=1}^T \mathbf{1}\{\mu_{k,t} \neq 0\} \cdot \rho_k. \quad (6.18)$$

We will now use the assumption that the auction is a core auction. For any $t \in [1, T]$, let $S \subseteq A$ be the set of agents in A for which $\mu_{k,t} = 0$. We have that

$$\begin{aligned} \sum_{k \in A} [\mathbf{1}\{\mu_{k,t} = 0\}(x_{k,t} v_{k,t} - z_{k,t})] &= \sum_{k \in S} (x_{k,t} v_{k,t} - z_{k,t}) \\ &\geq \sum_{k \in S} y_k(\mathbf{v}_t) v_{k,t} - \sum_{k=1}^n z_{k,t} = \sum_{k \in A} \mathbf{1}\{\mu_{k,t} = 0\} y_k(\mathbf{v}_t) v_{k,t} - \sum_{k=1}^n z_{k,t}. \end{aligned}$$

The inequality follows from the definition of a core auction and the no unnecessary pacing condition (which implies $b_{k,t} = v_{k,t}$ for all $k \in S$), by considering the deviation in which the agents in S jointly switch to allocation $\{y_k(\mathbf{v}_t)\}$. Substituting the above inequality into Equation (6.18) and rearranging yields Equation (6.17).

Summing over Equation (6.16) for each $k \notin A$ and combining it with Equation (6.17), we obtain that as long as inequality Equation (6.15) holds (i.e., event E_{GOOD} occurs), we have

$$\sum_{k \in [n]} \text{WEL}_{k,\text{GPD}}(\mathbf{v}) \geq \sum_{k \in [n]} R_k(\mathbf{v}) - \sum_{k \in [n]} \sum_{t \in [T]} z_{k,t} - n\bar{v} \sqrt{T \log(\bar{v}nT)}. \quad (6.19)$$

Step 3: A Bound on Expected Liquid Welfare. Since the liquid welfare is non-negative, we can take expectations over $\mathbf{v}_1, \dots, \mathbf{v}_T$ to conclude from (6.19) that

$$\begin{aligned} \mathbb{E} \left[\sum_{k=1}^n \text{WEL}_{k,\text{GPD}}(\mathbf{v}) \right] &\geq \mathbb{E} \left[\mathbf{1}\{E_{\text{GOOD}}\} \cdot \sum_{k=1}^n \text{WEL}_{k,\text{GPD}}(\mathbf{v}) \right] \\ &\geq \mathbb{E} \left[\mathbf{1}\{E_{\text{GOOD}}\} \cdot \sum_{k \in [n]} R_k(\mathbf{v}) \right] - \mathbb{E} \left[\sum_{k=1}^n \sum_{t=1}^T z_{k,t} \right] - n\bar{v} \sqrt{T \log(\bar{v}nT)}, \end{aligned} \quad (6.20)$$

where the last inequality holds via Equation (6.19).

It remains to analyze the expectations on the right side of the inequality. First, note that

$$\begin{aligned} \mathbb{E} \left[\mathbf{1}\{E_{\text{GOOD}}\} \sum_{k \in [n]} R_k(\mathbf{v}) \right] &= \mathbb{E} \left[\sum_{k \in [n]} R_k(\mathbf{v}) \right] - \mathbb{E} \left[(1 - \mathbf{1}\{E_{\text{GOOD}}\}) \sum_{k \in [n]} R_k(\mathbf{v}) \right] \\ &\geq \mathbb{E} \left[\sum_{k \in [n]} R_k(\mathbf{v}) \right] - \frac{n\bar{v}T}{n\bar{v}T^2} = \mathbb{E} \left[\sum_{k \in [n]} R_k(\mathbf{v}) \right] - 1/T. \end{aligned}$$

The inequality holds due to the fact that $R_k(\mathbf{v}) \leq \bar{v}T$ as well as our bound on the probability that event Equation (6.15) does not hold. Let $q_{k,t}$ be the unconditional probability that $\mu_{k,t} = 0$. Then

$$\begin{aligned} \mathbb{E} [R_k(\mathbf{v})] &= \sum_{t=1}^T \mathbb{E} \left[\mathbf{1}\{\mu_{k,t} = 0\} \mathbb{E}[y_k(v_{k,t})v_{k,t} | \mathcal{F}_{t-1}] + \mathbf{1}\{\mu_{k,t} \neq 0\} \rho_k \right] \\ &= \sum_{t=1}^T [q_{k,t} \mathbb{E}[y_k(v_k)v_k] + (1 - q_{k,t})\rho_k] \geq \sum_{t=1}^T \mathbb{E}[y_k(v_k)v_k] = \overline{\text{LW}}_k(\mathbf{y}, F). \end{aligned}$$

The first inequality uses the conditional independence of $y_k(v_{k,t})v_{k,t}$ on $\mu_{k,t}$, as this is already determined by time $t - 1$. The inequality holds since $\mathbb{E}[y_k(v_k)v_k] \leq \rho_k$ according to our assumption on \mathbf{y} . Substituting the inequalities into Equation (6.20), we obtain that

$$\text{WEL}_{\text{GPD}}(F) \geq \sum_{k=1}^n \overline{\text{LW}}_k(\mathbf{y}, F) - \mathbb{E} \left[\sum_{k=1}^n \sum_{t=1}^T z_{k,t} \right] - n\bar{v} \sqrt{T \log(\bar{v}nT)} - 1/T. \quad (6.21)$$

Recall that

$$\mathbb{E} \left[\sum_{k=1}^n \sum_{t=1}^T z_{k,t} \right] = \sum_{k=1}^n \sum_{t=1}^T \mathbb{E} [z_{k,t}] = \sum_{k=1}^n \mathbb{E}[P_k] \leq \text{WEL}_{\text{GPD}}(F), \quad (6.22)$$

where P_k is the total expenditure of agent k , and this is upper bounded by the liquid value they obtain. Substituting into (6.21) and rearranging the terms, then noting that $\text{WEL}_{\text{GPD}}(F)$ is precisely the left-hand side of (6.9), we conclude that inequality (6.9) holds. \square

6.5 Chapter Notes

The results of this chapter originally appeared as an extended abstract in [70](manuscript version at [69]), joint work with Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Compared to the original version, we have primarily focused on the aggregate guarantees of dynamic pacing algorithms. The original paper also proves new *individual* guarantees of dynamic pacing in auctions satisfying a so-called “monotone bang-for-buck” property, further justifying their adoption in practice. This is achieved by generalizing the viewpoint of pacing as running stochastic gradient descent on a suitable objective function as done by Balseiro and Gur [17] and relating the performance with this artificial objective to performance with respect to value maximization.

We have chosen to focus on the aggregate guarantees for cohesiveness with the themes of the remainder of this part of the thesis. As a result, the exposition and organization somewhat differs from the original paper while the proofs are essentially identical.

6.5.1 Deferred Proofs

Proof of Lemma 6.9. Let $\Delta_t = X_t Y_t + (1 - X_t)\rho - (\mathbb{E}[X_t Y_t + (1 - X_t)\rho | \mathcal{F}_{t-1}])$. Clearly, the sequence Δ_t forms a \mathcal{F}_t -martingale difference sequence by construction. Moreover, we observe that

$$\mathbb{E}[X_t Y_t + (1 - X_t)\rho | \mathcal{F}_{t-1}] = X_t \mathbb{E}[Y_t | \mathcal{F}_{t-1}] + (1 - X_t)\rho = X_t \mathbb{E}[Y_t] + (1 - X_t)\rho,$$

where we use the facts that X_t is \mathcal{F}_{t-1} -measurable and t , Y_t is independent of \mathcal{F}_{t-1} . It follows that $\Delta_t \in [-\mathbb{E}[Y_t], \bar{v} - \mathbb{E}[Y_t]]$. As an immediate consequence of the Azuma-Hoeffding inequality, we obtain

$$\Pr\left(\sum_{t=1}^T \Delta_t \geq \theta\right) \leq \exp\left(\frac{-2\theta^2}{T\bar{v}^2}\right). \quad (6.23)$$

But observe that

$$\begin{aligned} \left\{\sum_{t=1}^T \Delta_t \geq \theta\right\} &= \left\{\sum_{t=1}^T [X_t Y_t + (1 - X_t)\rho] \geq \theta + \sum_{t=1}^T [X_t \mathbb{E}[Y_t] + (1 - X_t)\rho]\right\} \\ &\supseteq \left\{\sum_{t=1}^T [X_t Y_t + (1 - X_t)\rho] \geq \theta + T \cdot \rho\right\}, \end{aligned}$$

using the assumption $\mathbb{E}[Y_t] \leq \rho$. This inclusion together with Equation (6.23) yields Equation (6.12). \square

6.5.2 Related Work

Our work is heavily motivated by the results of Balseiro and Gur [17], who attain convergence guarantees in repeated second-price auctions under strong convexity-like assumptions. Algorithm 3 appears to have been introduced in their work. A similar tatonnement algorithm was introduced and analyzed by Borgs, et al [27], who attain a similar result for first-price auctions without convexity assumptions. Our emphasis on welfare guarantees *without requiring convergence* appears novel, and possibly necessary given the aforementioned PPAD-hardness result [38].

A static (single-shot) game between budget-constrained bidders who tune their pacing multipliers, a.k.a. the *pacing game*, along with the appropriate equilibrium concept was studied in for second-price auctions by Conitzer, et al [43], and later by Conitzer, et al [42] in first-price auctions. It has since been extended to more general payment constraints [3] and utility measures [14]. In particular, any pure Nash equilibrium of this game achieves at least half of the optimal liquid welfare when the underlying auction is truthful. In a related contextual auction setting and simultaneously with our work, [18] establishes a similar bound on liquid welfare at any (possibly non-linear) Bayes-Nash equilibrium for i.i.d. bidders, for a class of standard auctions that includes first-price and second-price auctions. In contrast to these results, our efficiency result does not rely on convergence to equilibrium, and applies to all core auctions.

An alternative approach to budget pacing is known as *throttling* (also *probabilistic pacing*). Instead of pacing their bids, agents participate in only a fraction of the auctions [16]. Very recently, [39] proved that throttling converges to a Nash equilibrium in the first-price auctions (albeit without any stated implica-

tions on welfare or liquid welfare). On the other hand, they show that no such convergence is possible for second-price auctions since a Nash equilibrium is also PPAD-hard to compute [38]. The equilibria obtained by throttling and pacing dynamics in the first-price auction can differ in revenue by at most a factor of 2 [39].

As mentioned above, competing in repeated auctions with a budget is a special case of the problem of adversarial bandits with knapsacks [93]. It is known that in general, one cannot attain admit regret bounds: instead, one is doomed to approximation ratios, even against a time-invariant benchmark and even in relatively simple examples [93]. Balseiro and Gur [17] obtain related results specifically for repeated budget-constrained bidding in second-price auctions. The essential reason for this impossibility is the *spend-or-save dilemma* [93], whereby the algorithm does not know whether to spend the budget now or save it for the future.

Subsequent Work. Following the initial appearance of these results, two closely related works have expanded on the themes considered here. Fikioris and Tardos [57] show that in repeated *first-price auctions* with budgets, individual performance guarantees imply aggregate liquid welfare guarantees by leveraging smoothness-style arguments. In particular, if all agents achieve a multiplicative $\gamma \geq 1$ approximation of the performance of the best fixed pacing multiplier, the liquid welfare obtained by the agents is within a multiplicative factor of ratio $R = \gamma + O(1)$ on liquid welfare. They achieve $\gamma = T/B$ and $R \approx T/B + 1/2$ by plugging in a recent result on bandits with knapsacks [34]. Their guarantees apply to an arbitrary sequence of valuation profiles \mathbf{v}_t .

Lucier, et al [103] show that Theorem 6.4 can be extended to settings where

agents must also satisfy ROI (return-on-investment) constraints in addition to budgets via a variation of Algorithm 3. Their result suffers no quantitative loss and the algorithm corresponds with ours when there are no ROI constraints. Their results hold for any auction format such that a single item is sold to the highest bidder and the payment is between first and second price.

Part II

Opinion Dynamics: Beyond Consensus

CHAPTER 7

OPINION DYNAMICS

In the first half of this thesis, we analyzed stateful social systems that were primarily driven by strategic behaviors. By understanding how these game-theoretic agents interacted, we were able to achieve fairly strong results on the long-run outcomes that must arise. The general philosophy of the above approach was that even though selfish agents seek to learn or optimize for their own benefit, these objectives are sufficiently well-aligned with desirable global objectives that take into account all agents. These strategic behaviors were then shown to be provably benign, in an appropriate sense, via our price of anarchy or stability bounds. However, not all social systems are driven by inherently *strategic* behaviors, and relatedly, the kinds of important long-run properties of these systems will not be expressible via utilities. The game-theoretic language of the previous chapters will thus prove inappropriate for their study.

In the remainder of this thesis, we study a different model of social systems, where we must employ a very different kind of analysis: models of opinion formation in networks. In these models, agents do not act strategically, but instead apply simple update rules to adjust their opinions based on their observations in the network. Nonetheless, we will remain interested in providing provable guarantees of the *macroscopic phenomena and emergent properties* of these simple dynamics. Rather than welfare, we will investigate two key emergent properties of these dynamics, namely polarization and discord, to better understand recent empirical phenomena. The localized dynamics of these simpler updating rules lead to significant dependencies in the long-run evolution, imposing new, but similar technical challenges to that of Part I.

In this chapter, we provide a brief introduction to this area and give a simple abstraction of all the models of opinion formation we will consider in Part II. We then specialize this to a classical and highly tractable model of opinion formation, the DeGroot model of opinion dynamics [47]. However, we conclude by demonstrating how these kinds of models are poorly-suited to understanding real-world phenomena that go beyond consensus. These considerations will motivate the rest of our study in Part II.

The study of opinion dynamics has a long history in the economics, social science, and computer science communities. In lieu of a comprehensive discussion of this line of work at this time, we refer to the excellent surveys of Castellano, Fortunato, and Loreto [33] and Acemoglu and Ozdaglar [2] for more information on classical models of opinion formation and social learning.

7.1 Opinion Evolution in Networks

People's beliefs evolve in response to a variety of internal or external factors. Beliefs naturally adapt in relation to the beliefs of other people: after meeting with other members of their social network, like friends, family, neighbors, or online communities, their beliefs might change to reflect these interactions. Of course, opinions may also change due to more explicit modes of persuasion. Political figures or advertising campaigns may also explicitly seek to change the overall sentiments on one or more issues throughout the population. Random events may also perturb beliefs in stochastic ways.

A classical way to model this kind of opinion formation is to view opinion evolution as a discrete-time process that takes place on a weighted graph

$G = (V, E, w)$. Each agent is represented by a node in the graph and the underlying network topology provides the (weighted) graph of connections each agent has to others in the population. Each agent $i \in V$ then holds a numerical or vector-valued opinion $\mathbf{x}_i \in \mathbb{R}^d$, where d is the underlying dimension of the space of opinions. Each coordinate of an opinion vector \mathbf{x} may represent an agent's opinion on a single issue or topic. For instance, if each opinion vector is \mathbf{x} is single-dimensional, a positive value may refer to more liberal beliefs, while a negative or zero value may refer to conservative or mixed beliefs, respectively. In the higher-dimensional case, the separate issues may refine this to conservative or liberal beliefs on separate issues of importance, like foreign policy or health care.

In this case, a general way to model opinion evolution is via the following process, given in Algorithm 4. In words, at each time, an agent observes the

Algorithm 4: Opinion Formation on Networks

Input: Weighted graph $G = (V, E)$, initial opinions $\mathbf{x}_1^1, \dots, \mathbf{x}_1^n \in \mathbb{R}^d$,
 updating functions f^1, \dots, f^n , distribution \mathcal{D} over \mathbb{R}^k

1 **for** $t = 1, \dots$ **do**

2 Sample $\xi_t \sim \mathcal{D}$ i.i.d.

3 **for** $i \in V$ **do**

4 Set

$$\mathbf{x}_{t+1}^i = f^i(\mathbf{x}_t^i, (\mathbf{x}_t^j)_{j \in N(i)}, \xi_t). \quad (7.1)$$

opinions of each of their neighbors in G , as well as some random vector drawn i.i.d. from some underlying distribution. They then use this information by updating their present opinion using some fixed updating function over their current opinion and these observations at the present time.

Note that including the underlying graph G is not strictly necessary. We could just have easily encoded the restrictions to the network connections via the updating function f^1, \dots, f^n . However, it is more convenient to make this dependence explicit. This model could easily be updated to allow the updating functions to themselves be time-dependent and further allow dependencies on all previously observed opinions of neighbors. However, the above formulation is already enough to obtain rich behaviors while capturing every model we consider in the remainder of Part II.

7.2 A Classical Model of Opinion Formation: DeGroot Dynamics

We now present the simplest possible formulation of the framework: the DeGroot model of opinion formation [47]. In this model, agents simply update opinions by taking weighted averages of the opinions of their neighbors. In the model given in Algorithm 4, this is equivalent to setting f^i to be the graph-weighted average of the opinions of the heads of the outgoing edges from the node.

Definition 7.1. Given a directed, weighted graph $G = (V, E, w)$ and initial opinions $\mathbf{x}_1^1, \dots, \mathbf{x}_n^1$, the **DeGroot dynamics** are given by

$$\mathbf{x}_{t+1}^i = \frac{\sum_{j \in N(i)} w(i, j) \cdot \mathbf{x}_t^j}{\sum_{j \in N(i)} w(i, j)}.$$

The interpretation of the DeGroot model is that the underlying graph weights represent the relative trust that agents have in the beliefs of their neighbors. Outgoing edges with relatively high weight reflect high levels of social

trust, while outgoing edges with less relative weight (or pairs with no edge at all) reflect comparatively little social trust. This manifests in the updating rule, where the associated weights are proportional to these levels of trust.

The DeGroot dynamics enjoy a number of favorable properties from a theoretical standpoint. First, the dynamics admit a particularly simple recursion. To see this, define the transition matrix $A \in \mathbb{R}^{n \times n}$ by:

$$A_{i,j} = \frac{w(i,j)}{\sum_{j \in N(i)} w(i,j)}.$$

Notice that A is a row-stochastic matrix, meaning the entries are nonnegative and each row sums to one: indeed, A is simply the transition matrix of the Markov chain on G where transitions from a node are proportional to the weights of the out-edges. With this notation, the DeGroot dynamics simply follow the recursion:

$$\mathbf{X}^{t+1} \triangleq \begin{bmatrix} -\mathbf{x}_{t+1}^1- \\ \vdots \\ -\mathbf{x}_{t+1}^n- \end{bmatrix} = A^t \mathbf{X}^1. \quad (7.2)$$

Under standard conditions from the theory of discrete-time finite Markov chains, namely *aperiodicity and irreducibility* (see, e.g. [55]), it is well-known that

$$A^t \rightarrow \mathbf{1}\pi^T = \begin{bmatrix} -\pi- \\ \vdots \\ -\pi- \end{bmatrix}, \quad (7.3)$$

where π is the *unique* stationary distribution of the Markov chain whose transitions are A ; that is, the unique distribution π satisfying $A\pi = \pi$. By plugging in Equation (7.3) into Equation (7.2), we find that the limiting beliefs of the DeGroot dynamics is *consensus* at a weighted average belief in each dimension,

where the weights depend *only* on the underlying graph G . The weights given by π thus reflect the limiting influences of the beliefs of each agent at this final consensus, uniformly over the initial vectors of opinions. These observations, in the Bayesian setting, lead to the so-called “wisdom of the crowd” so long as no small set of agents has overly large influences asymptotically [79].

7.3 Challenges: Consensus, Discord, and Polarization

While the DeGroot model enjoys a clean mathematical formulation that makes it especially suitable for theoretical analysis, how reasonable are the qualitative predictions it makes? Certainly the prediction that people reach consensus in their opinions is rather wildly optimistic at best, empirically refuted at worst. However, since the prediction of consensus holds asymptotically, one could argue that the underlying dynamics themselves may be reasonably sound and consensus would be achieved after sufficiently long time.

However, a more alarming feature is that the DeGroot model does not just imply that consensus is asymptotically achieved. A significantly more refined prediction of the dynamics is that opinions necessarily get *closer* together in *every round*. This follows simply because the opinions at each given time are multiplied by the row-stochastic matrix matrix A .

Lemma 7.1. *Let $\mathbf{x} \in \mathbb{R}^n$ and A be row-stochastic. Define $\mathbf{y} = \mathbf{Ax}$. Then*

$$\max_{i,j} |y_i - y_j| \leq \max_{i,j} |x_i - x_j|.$$

Proof. It suffices to show that $\max_i y_i \leq \max_i x_i$ and $\min_i y_i \geq \min_i x_i$. But this is immediate because each entry of \mathbf{y} is a convex combination of entries in \mathbf{x} by

row-stochasticity. □

To put it a different way, DeGroot dynamics is an inherently *contractive* model of opinion formation. This kind of oblivious repeated averaging is evidently well-suited to reasoning about phenomena that go beyond at least weak consensus, simply because the dynamics themselves do not themselves provide a promising mechanism to do so.

Our goal in the remainder of this thesis will be to go past these simplistic predictions of asymptotic consensus to understand the theoretical properties of *polarization and discord*. Roughly, we say that a set of opinions polarizes if the limiting opinions *cluster* in opposite directions, in either the single-dimensional or multi-dimensional setting. Similarly, we will measure *discord* in terms of the numerical disagreement in limiting opinion vectors.

As we will discuss in later chapters, both of these properties of real-world opinions have been exhibited and exacerbated in recent years. But in the DeGroot model, these properties are trivial; no polarization nor discord can arise if opinions reach consensus. We thus cannot even begin to reason about how the underlying dynamics of opinion formation contribute to these important phenomenon without going past these kinds of models.

In the remainder of Part II, we will thus analyze different models and environments of opinion formation, instantiated through Algorithm 4, that provably can predict discord and polarization. Our analysis will require us to identify new mechanisms where these features can provably arise; the cost of doing so is that we will no longer be able to enjoy quite the tractability of the DeGroot model. Along the way, we will encounter many of the same probabilistic chal-

lenges as in Part I to deal with these complexities.

CHAPTER 8

ADVERSARIAL PERTURBATIONS OF FRIEDKIN-JOHNSEN DYNAMICS

As described in the previous chapter, a typical feature of real-world networks is that even approximate consensus is rarely, if ever, achieved. While classical models of opinion formation, like the DeGroot dynamics described in Section 7.2, provide a fascinating theory of the interplay between network structure and limiting beliefs, it cannot immediately shed light on these types of long-run properties.

To help address this gap, richer models like the Friedkin-Johnsen (FJ) variation of the standard DeGroot model [63] or the bounded-confidence Hegselmann-Krause model [90] have been proposed to give plausible explanations for how these long-term disagreements can persist while maintaining the strong intuition that one will account for and adjust to the opinions and expressed behaviors of those near them. The FJ model will be the subject of our analysis in this chapter; informally, the underlying updating of opinions follows the DeGroot model, but each node also holds their own immutable internal opinion which always gets incorporated during the repeated averaging. Since these fixed internal opinions do not themselves evolve, their net effect is to provide a persistent pull towards the initial beliefs of the agent that prohibits consensus except in trivial cases.

In this chapter, we study a topical mechanism of polarization: external actors artificially “seed” initial opinions with the aim of inducing discord among the opinions held by the nodes in the network after the dynamics. Note that this objective, of sowing discord, is somewhat more abstract than more well-studied objectives like shifting opinions (see, for instance [75] for an example in FJ dy-

namics). An external agent interested in these objectives may aim to shift the opinions of a small set of nodes with the hope that these shifts propagate in the same direction throughout the network. On the other hand, an adversary that aims to induce discord need not have a preferred direction.

The theoretical formulations we pose and analyze in this chapter are directly motivated by real-world behaviors. For instance, the U.S. Department of Justice Special Counsel’s Office argues in their 2017 indictment of the Russian Internet Research Agency (IRA) that the IRA used social media and targeted advertising with “a strategic goal to sow discord in the U.S. political system, including the 2016 U.S. presidential election” [1]. This behavior has been seen in multiple countries, and is not limited to any one actor. For example, the doctrine of offensive information warfare by the U.S. intelligence community includes provisions for the instigation of discord between opposing parties [108]. More recently, Twitter disclosed in 2019 that external bots made a coordinated effort to “sow political discord in Hong Kong” with the aim of making protesters less effective at organizing during the Hong Kong independence movement [128]. It is to be expected that as social media becomes more and more prevalent, these sorts of external influences and behaviors will become more frequent.

Note that an adversary interested in these objectives of inducing discord may exhibit somewhat more subtle behavior than an actor that simply seeks to shift opinions in a preferred direction, a finding that we will extensively explore theoretically in this chapter. However, these kinds of behaviors have been observed empirically as well. For instance, analysis of Twitter data in the past few years has revealed how bots sponsored in some cases by state actors managed to play both sides of some existing issue to exploit divisions along racial

lines, pro- and anti-vaccine groups, NFL kneeling, and gun reform, among other matters [119, 30, 53].

Our first task in this chapter is to provide a formal model of adversarial optimization in the Friedkin-Johnsen dynamics. In our formulation, we consider instances of the Friedkin-Johnsen model in which all nodes' internal opinions are initially at equilibrium, e.g., they are all equal to 0, and then we will allow an adversary to perturb these internal opinion subject to some radius constraint modeling resource constraint of the attacker. These seeded internal opinions then diffuse across the network to reach an equilibrium under the Friedkin-Johnsen dynamics, at which point some measure of disagreement among adjacent nodes is measured. Our main goals are to characterize the interplay between graph structure and the adversary's power: which networks are most susceptible to or resilient to these kinds of perturbations; how can the adversary most efficiently spend a bounded amount of resource in order to produce the maximum disagreement; and how can we most effectively defend a network against such attacks?

8.1 Overview of Results and Techniques

8.1.1 Spectral Properties of Discord

The first main result of this chapter is formulating and analyzing this model of an adversary seeking to maximize some measure of *discord* among the nodes of a network. In our model, the adversary may perturb the internal opinions of the nodes (for instance, by targeted advertising) by a bounded amount (corresponding to a budget constraint) to maximize this measure of discord, knowing

that these perturbations will diffuse according to the FJ dynamics.

We first show that many such objectives for such an actor can be naturally posed as the maximization of an associated quadratic form of the underlying graph Laplacian. The simple, but crucial observation is that this directly connects the adversary's power in attaining their objective with the eigenvalues of some function of the underlying graph. The graph Laplacian is an extremely well-studied object, whose spectral properties are known to be good approximations of combinatorial features of the underlying network; this connection will allow us to qualitatively understand what structural properties of the graph matter in these different settings.

For different discord objectives, we find that, somewhat surprisingly, different regimes of the spectral structure of the underlying networks characterize the ability of this adversary to succeed in their aim. In perhaps the most natural case, we define the disagreement on an edge of the network to be the squared difference of the opinions at the edge's endpoints; and we define the disagreement of the network to be the sum of the disagreements on all its edges. The challenge for an adversary in maximizing disagreement is that the same network is being used both by the opinion dynamics to average away the disagreement, and by the objective function to measure the disagreement. We show in Theorem 8.1 that the net result of this tension is that, the adversary's power to induce disagreement is bounded uniformly over graph topologies, and *any* eigenspace of the graph Laplacian can govern the adversary's power under some scaling of the weights of the graph interactions. This complete dependence on the eigenstructure stands in stark contrast to typical results in spectral graph theory, where usually the extremal eigenspaces hold the most physical signifi-

cance.

Beyond disagreement, we also consider a variety of other adversarial objectives: namely, we consider variants of disagreement and polarization (as introduced by Musco, Musco, and Tsourakakis [107]), which essentially measures the variance of the resulting opinion vector. In these cases, we again show that the spectral properties of the Laplacian govern the adversary’s power. In fact, we show that spectral analyses can provide useful (and sometimes tight) bounds even beyond the case of ℓ_2 -constrained adversaries. Finally, by characterizing the adversary’s power to induce these different forms of discord in terms of spectral structure, we identify extremal robust and susceptible graph structures by leveraging the well-known connections between spectral and combinatorial properties of graphs.

8.1.2 Network Defense

The results of the previous section only consider the interventions of a single adversarial actor. Towards providing positive results, in Section 8.4, we consider the algorithmic task of a network defender to provide an initial intervention to help *minimize* the ability of such an adversary to promote their objective. Our general approach will be to assume that the graph topology is itself immutable; however, we allow a network defender to *insulate* the individuals in the network heterogeneously. Concretely, the network defender can provide a weighting of the ℓ_2 -constraints of the adversary, corresponding to making certain individuals more expensive to influence. This might be achieved, for instance, through technological literacy initiatives to make it more difficult for certain individu-

als to be susceptible to the normal modes by which an external actor may seed internal opinions, like bots or fake news.

Formally, we will consider a two-player min-max game, where a network defender can first heterogeneously *weigh* different individuals in the network (subject to some explicit normalization). The adversary then solves their objective function as before, but subject to these new weights that govern their possible interventions. The constraints on the defender's choice will include the option to keep the individuals' weights the same as above, so that this optimization problem for the defender is upper-bounded by the previous analyses. On the algorithmic side, we show in Theorem 8.9 that while the adversary's inner maximization problem seemingly cannot be reduced to a spectral analysis of the original graph, the network defender's problem can be globally solved using convex programming for many choices of the adversary's objective and weight normalization.

8.1.3 Mixed-Graph Objectives

Finally, in Section 8.5, we generalize the original adversarial optimization problem to better align. Recall that in our first analysis of adversarial power, the underlying network both averages opinions of neighboring nodes and is the topology with respect to which discord is measured (in particular, along the edges). This feature led to uniform bounds on an adversary's power independent of graph structure. However, in real-world opinion diffusion, there is no explicit need for the graph to play both roles. In fact, there is good reason to believe that the structure of opinion diffusion can be entirely orthogonal from the structure

along which an adversary seeks to induce discord. For instance, opinion formation may take place primarily via online interactions; these connections may be completely different from real-world social networks of everyday interactions, like with coworkers or neighbors.

These considerations suggest that we return to the adversary's disagreement maximization problem, but when the network where disagreement is measured can be different from the network where the Friedkin-Johnsen opinion dynamics take place. The individuals and communities one interacts with on social media can be quite global in nature, as well as relatively anonymous and self-selecting; in contrast, one's real-world interactions are far more dictated by other factors like geography and occupation. It is thus quite natural to think of opinion formation as occurring somewhat independently, or even nearly disjointly, to these real-world interactions. To sow discord, an adversary might hope to induce disagreement measured via these real-world connections, but must do so filtered through the opinion dynamics in the online world. Intuitively, if the opinion graph which does the smoothing of internal opinions and the disagreement graph look very different, an adversary will be much more powerful than what is possible in the basic model when these graphs must be the same.

To understand this setup, we generalize the adversary's problem with *mixed-graph* objective functions in the natural way by considering two distinct graphs at once. One might expect that the adversary is most limited when the opinion graph and disagreement graphs are the same, as then the opinion dynamics smooth out the seeded opinions along the exact same network where disagreements will be measured. Somewhat surprisingly, we show that it is not always

the case, even subject to some natural degree normalization; there are explicit, simple disagreement graphs where a different opinion formation graph provably lowers the adversary’s ability to induce disagreement by a non-negligible amount. However, we also give a general lower bound that suggests that these examples are somewhat pathological in that, typically, having different graphs will help the adversary in this aim compared to the single-graph setting. For other disagreement graphs, it will in fact turn out that they are optimal for themselves.

In general, we connect the newfound power of the adversary in the mixed-graph setting to the *spectral similarity* of the disagreement and diffusion graphs, appropriately defined. In particular, we show in Theorem 8.18 if the two graphs are good spectral approximations for each other, then the objective of the adversary is essentially unchanged. This implies that it is not necessary at all for the two graphs to be *physically* quite similar, edge-by-edge but rather only *spectrally* similar; indeed, it is well-known celebrated results of Batson, Spielman, and Srivastava that every graph has a spectral approximation with linear number of edges [19]. Conversely, it is still quite intuitive that the more dissimilar the two graphs are, the more an adversary will be able to induce disagreement. We show that spectral dissimilarity, again appropriately defined, is sufficient for this purpose in Proposition 8.20. It will nontrivially follow (Corollary 8.21) if the two graphs are misaligned in the sense that their cut structure is sufficiently different anywhere, then the adversary gains significant power to induce discord. This provides some theoretical justification for why external actors have been able to effectively sow discord; our formalism shows that once opinion formation and disagreement become uncoupled, as is possible with social media platforms, large discord is achievable.

8.2 Preliminaries and Friedkin-Johnsen Dynamics

In this chapter, we assume that opinions evolve using the Friedkin-Johnsen dynamics [63] (FJ dynamics), as defined below:

Definition 8.1. The Friedkin-Johnsen dynamics are specified by an undirected simple graph $G = (V, E, w)$ and an initial, internal opinion vector $\mathbf{s} \in \mathbb{R}^n$. Starting with $\mathbf{z}^{(0)} = \mathbf{s}$, each node $i \in V$ updates her opinion by taking the weighted average of her neighbors in G , as well her own internal opinion:

$$\mathbf{z}^{(t+1)}(i) = \frac{\mathbf{s}(i) + \sum_{j:(i,j) \in E} w(i, j) \mathbf{z}^{(t)}(j)}{1 + \sum_{j:(i,j) \in E} w(i, j)}.$$

Notice that these equations implicitly normalize a weight of 1 on one's private opinion. It is well-known that these dynamics converge to a fixed point, and the limiting final opinion vector is given by

$$\mathbf{z} = (I + L)^{-1} \mathbf{s},$$

where L is the Laplacian of G .

We will consider a variety of different adversarial objectives. For this, we require the following notions of disagreement and polarization of opinion vectors, as was introduced by Musco, Musco, and Tsourakakis [107]:

Definition 8.2. Given any opinion vector $\mathbf{x} \in \mathbb{R}^n$ and undirected graph G , the *disagreement* of G with opinions \mathbf{x} as

$$\sum_{(i,j) \in E} w(i, j) (\mathbf{x}(i) - \mathbf{x}(j))^2 = \mathbf{x}^T L \mathbf{x},$$

where L is the graph Laplacian of G .

Definition 8.3. Given an opinion vector \mathbf{x} , let

$$\bar{\mathbf{x}} \triangleq \mathbf{x} - \frac{\mathbf{x}^T \mathbf{1}}{n} \cdot \mathbf{1} = \left(I - \frac{1}{n} \mathbf{1}\mathbf{1}^T\right) \mathbf{x}$$

be the *de-meaned* version of \mathbf{x} obtained by subtracting off the average of \mathbf{x} from each component. Then the *polarization* of \mathbf{x} is

$$\|\bar{\mathbf{x}}\|_2^2 = \sum_{i=1}^n \bar{x}(i)^2.$$

In words, the polarization is a measure of variance for \mathbf{x} .

8.3 Adversarial Optimization in FJ Dynamics

In this section, we consider a variety of different adversarial objectives of an entity that perturbs Friedkin-Johnsen dynamics from an initially all zero state. Concretely, we consider the following problem: what initial opinions should an adversary supply in the FJ dynamics (e.g. via heterogeneous targeting) to maximize these adversarial objects after diffusion, subject to constraints on the overall magnitude of these perturbations? Our goals will be, for each objective we consider, to obtain bounds on the overall effect that an adversary can induce that depend on the underlying structural properties of the network.

As we will show, due to the special closed-form of the FJ dynamics, most of these objectives can be expressed in terms of positive semi-definite quadratic forms in the graph Laplacian L . What we will find is that the relationship between the spectral properties and an adversary's ability to effect their desired objective can be rather subtle: in different settings, different regimes of the eigenstructure of L will govern the adversaries power. To be slightly more formal, we will show how several natural adversarial objectives can be formulated

via the following optimization problem:

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\| \leq R} \mathbf{s}^T (I + L)^{-1} f(L)(I + L)^{-1} \mathbf{s}, \quad (8.1)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $f(y) \geq 0$ for $y \geq 0$ and $\|\cdot\|$ is some norm; the former restriction is made to ensure that the above quadratic form is nonnegative, while the condition imposes some normalization. Because our goal is to understand the relationship between the graph structure and the adversarial objective, and because Equation (8.1) is homogeneous of degree 2 in R , one may safely think of $R = 1$. When $\|\cdot\| = \|\cdot\|_2$, the above optimization is non-convex but is easily related to the eigenstructure of L . However, we will see that this spectral analysis can be highly fruitful even extending beyond this case.

The interpretation of Equation (8.1) is that on some issue that is currently at consensus in the graph, the adversary will first supply initial opinions \mathbf{s} , for instance via fake news or targeted advertisements. This is done subject to a fixed norm constraint, which corresponds to a limited budget to perturb initial opinions in this way. These opinions then diffuse and become “smoothed” through the underlying graph G via the Friedkin-Johnsen dynamics. The goal of the adversary is to choose these initial opinions in order to maximally induce some desired effect, *knowing* that whatever initial opinions it seeds, the opinion dynamics dictated by the underlying graph will inevitably partially equilibrate them. Of course, we will be interested in functions f where the above holds some sort of physical meaning.

8.3.1 Disagreement

First, we consider an adversary that seeds initial opinions subject to a norm constraint with the goal of *maximizing disagreement*. In the framework given by Equation (8.1), this can be realized via the choice $f(y) = y$, and yields the simple objective function

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I + L)^{-1} L (I + L)^{-1} \mathbf{s}. \quad (8.2)$$

The crucial feature of this objective is the fact that L both dictates the measurement of disagreement, and the opinion dynamics themselves. Somewhat surprisingly, for this objective function, the *entire* eigenstructure of the graph can matter; as L is scaled in importance compared to I , the adversary passes through each eigenspace in order as they optimize.

Theorem 8.1. *For any graph G , Equation (8.2) is upper bounded by $\frac{R^2}{4}$. Moreover, if G is connected and we consider the family of problems given by Equation (8.2) with $L(t) = tL$ for $t > 0$, then for any $i > 1$, there exists a $t > 0$ such that $R \cdot V_i$ is an optimizer for this value of t .*

Proof. By the variational characterization of eigenvalues of symmetric matrices, this amounts to understanding the spectrum of the matrix $(I + L)^{-1} L (I + L)^{-1}$. This matrix has eigenvalues

$$\frac{\lambda_i(L)}{(1 + \lambda_i(L))^2}, \quad i = 1, \dots, n,$$

with the same corresponding eigenvectors as L . In particular, the optimal value of Equation (8.2) turns out to be $R^2 \cdot \left(\max_{i \in [n]} \frac{\lambda_i(L)}{(1 + \lambda_i(L))^2} \right)$. Consider now the function $g(x) = \frac{x}{(1+x)^2}$. It is easy to compute $g'(x) = \frac{1-x}{(1+x)^3}$, from which it immediately follows that g is increasing for $0 \leq x \leq 1$ and decreasing for $x > 1$, attaining a

peak of $1/4$ at $x = 1$. This immediately gives the upper bound, where equality holds if and only if $\lambda_i(L) = 1$ for some $i \in [n]$.

For the second statement, as we vary $L(t) = t \cdot L$ with $t > 0$, the eigenvectors of L remain the same, but the eigenvalues scale as $t \cdot \lambda_2, \dots, t \cdot \lambda_n$. As t varies strictly between 0 to infinity, each nonzero eigenvalue of $t \cdot L$ passes through the peak of the function $g(x)$ at $x = 1$ when $t = 1/\lambda_i(L)$. This implies that for this value of t , the set of optimizers is $R \cdot V_i$, the length R vectors in V_i , the λ_i -eigenspace.¹ \square

One way to interpret the previous theorem is by first considering the extreme ranges of tL . For $t \approx 0$, tL is negligible compared to I , and so $(I + tL)^{-1} \approx I$. Physically, this corresponds to each individual not really listening to their neighbors over their own initial opinion, so the opinion dynamics do not substantially change the seeded opinion by the adversary. When this is the case, the above shows that the optimal strategy for the adversary that seeks to induce maximum disagreement is to seed vectors in the direction of V_n . In this case, by the quadratic form Equation (8.2), the adversary's strategy is to simply feed in opinions that directly maximize disagreement in G , as the opinion dynamics are themselves negligible. The actual vectors in the set V_n attempt to place different values on the two sides of each edge; in that sense, the vectors in V_n can be thought of as solving a type of graph coloring problem [4]. The quantity $\lambda_n(L)$ itself is also known to be closely related to the size of maximum cuts in graphs, see for instance [50, 78, 127].

On the other extreme, for large t , $(I + tL)^{-1}(tL)(I + tL)^{-1} \approx \frac{1}{t}L^\dagger$, where L^\dagger is the

¹We remark that for certain values of t , there may exist optimizers that are *not* eigenvectors of L . The reason is that because g is not injective, there may exist distinct eigenvalues of L that are maximal for g , so that both eigenspaces are optimal, and therefore the span of these eigenspaces are optimal as well, even though most such vectors are not eigenvectors of L .

pseudo-inverse of L . In this regime, the largest eigenvalue of L^\dagger is the inverse of the smallest nonzero eigenvalue of L , which is just λ_2 , with corresponding optimizer proportional to \mathbf{v}_2 . In general, it is well-known that λ_2 and \mathbf{v}_2 are intimately connected to the *normalized sparsest cut* of G ; if G is d -regular, then the famous discrete Cheeger inequality asserts that λ_2 is an approximation to the normalized sparsest cut to a factor of $\Theta(\sqrt{d})$ [80]. It is also known that some sweep cut of \mathbf{v}_2 yields a bipartition that attains this bound. Therefore, in this large t regime where graph neighbors are higher weighted than internal opinions, the initial opinion vector inducing maximal disagreement roughly corresponds to a sparse cut. Because the graph interactions in the Friedkin-Johnsen dynamics are so strong compared to the weight of the internal opinions, the optimal strategy of the adversary is roughly to induce disagreement on along some sparse cut of the graph.

As t varies between these two regimes, the above result implies that *every* nontrivial eigenvector becomes relevant. In the intermediate range of t , the relative effects of tL in measuring disagreement and tL in smoothing the initial opinions via the opinion dynamics directly conflict, which causes the adversary to pass through each eigenspace. For different regimes of how the each individual weighs their internal opinion to those of their neighbors, these effects balance in different ways, leading to this more interesting connection between the eigenstructure of L and the adversary's ability to induce large disagreement in a network with these dynamics.

8.3.2 Repeated Disagreement

As a simple extension of one-period disagreement considered in Section 8.3.1, it is natural to consider a similar objective taken over a longer time horizon. An analogous multi-period setting is the following: the adversary supplies the initial opinions \mathbf{s} , which the Friedkin-Johnsen dynamics take to $(I + L)^{-1}\mathbf{s}$. In the first period, the disagreement of this equilibrium opinion vector is measured, as before. In the next period, the last period final opinions $(I + L)^{-1}\mathbf{s}$ become the new initial opinions, which subsequently get updated by the dynamics to $(I + L)^{-2}\mathbf{s}$; the disagreement of these opinions is then measured and added to the first disagreement. This process is repeated for $T + 1$ periods, where $T \in \mathbb{N} \cup \{\infty\}$.

One natural setting for this objective is the following: consider a setting where opinions solidify, then update on a cyclic basis. For instance, during a U.S. election cycle, opinions might form along network structure according to the FJ dynamics; however, once the elections happen, the final opinions can be viewed as priors, or innate opinions, for the next cycle, at which point the dynamics will take hold once again and so on. These sorts of repeated dynamics are considered by Chitra and Musco [40] in their work on understanding filter bubbles in the FJ dynamics.

In this multi-period setting, the adversary's problem is to maximize the total disagreement across all time periods: putting this in our framework, this is

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I + L)^{-1} \left(L + (I + L)^{-1} L (I + L)^{-1} + \dots + (I + L)^{-T} L (I + L)^{-T} \right) (I + L)^{-1} \mathbf{s}.$$

To understand this setting, notice that we may work orthogonal to $\mathbf{1}$, as this gives zero objective value. In this case, the relevant f function that acts on the

eigenvalues becomes

$$f(x) = \sum_{i=0}^T \frac{x}{(1+x)^{2i}} = \frac{-(1+x)^{-2T} + (1+x)^2}{2+x},$$

with the convention that the first term is zero when $T = \infty$. We thus find that

$$\begin{aligned} \max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I+L)^{-1} \left(L + (I+L)^{-1}L(I+L)^{-1} + \dots + (I+L)^{-T}L(I+L)^{-T} \right) (I+L)^{-1} \mathbf{s} \\ = R^2 \cdot \max_{1 \leq i \leq n} \left\{ \frac{-(1 + \lambda_i(L))^{-2(T+1)} + 1}{2 + \lambda_i(L)} \right\} \end{aligned}$$

Note that for $T = 0$, this reduces to the analysis of the first subsection. Rather curiously, for $T = \infty$, this implies that the optimizer of this infinite-period game is precisely $R \cdot V_2$, with corresponding objective value $\frac{R^2}{2 + \lambda_2(L)}$. While in the one-shot game considered above, any one of the eigenvectors could be the relevant optimizer, the optimizer of this problem necessarily lies in the direction of V_2 .

8.3.3 Polarization and Disagreement

For a different objective, Musco, Musco, and Tsourakakis [107] considered the following cost metric on opinions, which they term the *polarization-disagreement index*, obtained by taking the sum of the disagreement of \mathbf{x} in G and the polarization of \mathbf{x} ; the authors show that when this measure is done with final opinions $\mathbf{x} = (I+L)^{-1}\mathbf{s}$, the polarization-disagreement index of the Friedkin-Johnsen dynamics with graph G and initial opinions \mathbf{s} can be simplified to

$$\bar{\mathbf{s}}^T (I+L)^{-1} \bar{\mathbf{s}}.$$

Note that this is not quite of the form Equation (8.1), though it is quite similar.

Suppose that an external adversary now chooses the internal opinions \mathbf{s} to maximize the polarization-disagreement index of the final opinions after under-

going the Friedkin-Johnsen dynamics:

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \bar{\mathbf{s}}^T (I + L)^{-1} \bar{\mathbf{s}} = \max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I - \frac{1}{n} \mathbf{1}\mathbf{1}^T) (I + L)^{-1} (I - \frac{1}{n} \mathbf{1}\mathbf{1}^T) \mathbf{s}.$$

Theorem 8.2. *For any graph G , the maximizer of the above maximization problem is $\pm R \cdot V_2$, and the resulting objective value is $\frac{R^2}{1+\lambda_2}$.*

Proof. First, notice that for any vector \mathbf{s} , we have $\|\mathbf{s}\|_2 \geq \|\bar{\mathbf{s}}\|_2$; this follows from the Pythagorean theorem and decomposing \mathbf{s} into the parts orthogonal to $\mathbf{1}$ (namely, $\bar{\mathbf{s}}$) and the projection onto $\mathbf{1}$. Because the above maximization function only depends on the de-meanned version of \mathbf{s} and is homogenous in \mathbf{s} , we may assume the adversary restricts to the subspace orthogonal to $\mathbf{1}$. As such, the problem becomes

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R, \mathbf{s} \perp \mathbf{1}} \mathbf{s}^T (I + L)^{-1} \mathbf{s}.$$

The eigenvalues of $(I + L)^{-1}$ are $\frac{1}{1+\lambda_i(L)}$ for $\lambda_i(L)$ the eigenvalues of L ; as $\mathbf{1}$ is the eigenvector for the largest eigenvalue 1 of this matrix, the variational characterization of eigenvalues implies that the above is exactly

$$R^2 \cdot \lambda_{n-1}((I + L)^{-1}) = \frac{R^2}{1 + \lambda_2(L)},$$

and this is attained by the set of vectors $R \cdot V_2$, as desired. \square

In particular, it follows that when an adversary chooses \mathbf{s} to maximize the polarization-disagreement index of G under these dynamics, the only relevant structure of the network that determines its robustness to these adversarial perturbations is precisely determined by the second smallest eigenvalue, and the initial opinion vector inducing this are the corresponding eigenvectors. Under the connection between the second smallest eigenvalue and eigenvectors of L and sparse cuts discussed above, Theorem 8.2 essentially asserts that the ability

of an adversary to induce polarization-disagreement in a graph G is essentially determined by the existence of small normalized cuts. Moreover, the actual optimizer for initial opinion vector roughly places large values on one side of a small cut and smaller values on the other side. As an immediate corollary of the above result, we can obtain the following intuitive, but nontrivial fact:

Corollary 8.3. *Let \mathcal{L} be the set of Laplacians of weighted n -node graphs subject to the total edge weight normalization $\text{Tr}(L) = 2m$. Then*

$$\arg \min_{L \in \mathcal{L}} \left\{ \max_{\mathbf{s} \in \mathbb{R}^n, \|\mathbf{s}\|_2 \leq R} \bar{\mathbf{s}}^T (I + L)^{-1} \bar{\mathbf{s}} \right\} = \frac{2m}{n(n-1)} \cdot L_{K_n},$$

where L_{K_n} is the Laplacian of the unweighted simple complete graph K_n .²

Proof. From the proof of Theorem 8.2, for any fixed Laplacian L , the inner maximization yields the objective value $\frac{R^2}{1+\lambda_2(L)}$; therefore, the claim is equivalent to showing that the Laplacian $\frac{2m}{n(n-1)} \cdot L_{K_n}$ has the maximum second-smallest eigenvalue among all $L \in \mathcal{L}$.

To see this, observe that for any such L , $\lambda_1(L) = 0$, and therefore

$$2m = \text{Tr}(L) = \sum_{i=2}^n \lambda_i(L) \geq (n-1) \cdot \lambda_2(L). \quad (8.3)$$

It immediately follows that for any such Laplacian L , $\lambda_2(L) \leq \frac{2m}{n-1}$; we will be done if we can show that this is attained for the claimed weighted complete graph. But it is not difficult to check that

$$\lambda_2(L_{K_n}) = n; \quad (8.4)$$

after scaling by $\frac{2m}{n(n-1)}$ so that the trace condition is satisfied, we see that this upper bound on λ_2 is exactly attained, proving optimality. Moreover, this occurs

²The factor $\frac{2m}{n(n-1)}$ is just to satisfy the total edge weight condition.

if and only if all of the nonzero eigenvalues are $\frac{2m}{n-1}$ by virtue of Equation (8.3), which occurs if and only if the graph is the scaled complete graph. \square

This corollary thus states that the complete graph, appropriately weighted, is min-max optimal given the adversary's objective to induce maximal polarization-disagreement index when running the Friedkin-Johnsen dynamics. More generally, any *spectral expander* will be robust to these adversarial perturbations, see for instance [92].

8.3.4 Absolute Displacement

Along these lines, suppose now that the attacker simply seeks to displace opinions maximally from the consensus at 0, measured in Euclidean norm. This too can be realized in the above setting: suppose that $f(y) \equiv 1$, so that $f(L) = I$. In this case, the adversary solves the following problem:

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I + L)^{-2} \mathbf{s} = \max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \|(I + L)^{-1} \mathbf{s}\|_2^2.$$

This latter identity shows that for this choice of f , the adversary's goal is indeed to maximize the ℓ_2 -norm of the final opinion vector $(I + L)^{-1} \mathbf{s}$, or equivalently, to *displace* the final opinions from the initial consensus at $\mathbf{0}$ as much as possible in Euclidean distance. However,

$$(I + L)^{-2} = \sum_{i=1}^n \frac{1}{(1 + \lambda_i)^2} \mathbf{v}_i \mathbf{v}_i^T.$$

As before, we thus obtain

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq R} \mathbf{s}^T (I + L)^{-2} \mathbf{s} = R^2 \cdot \lambda_n((I + L)^{-2}).$$

But recall that for any graph, the smallest eigenvalue of L is 0, with corresponding eigenvector $\pm \frac{1}{\sqrt{n}} \cdot \mathbf{1}$ (and this is unique if G is connected); as a result,

$$\lambda_{\max}((I + L)^{-2}) = 1,$$

and the unique maximizer (up to sign) is $\mathbf{s} = \frac{R}{\sqrt{n}} \mathbf{1}$. In particular, for this optimization problem, the network topology plays no role at all. This observation is quite similar to one made in the context of opinion maximization in [75].

8.3.5 Other Adversary Constraints

In the previous sections, we focused on the ℓ_2 -budget constraint, as this most cleanly elucidates the nontrivial interplay between spectral structure and the optimal choices of the adversary for given objectives of discord. In this section, we provide positive results in other settings; the first comes from the independent work of Chen and Rácz [37], who consider the same disagreement objective as we do, but instead makes the adversary sparsity-constrained, as well as restricted to a cube. We first show how a combination of the spectral ideas above with their analysis can lead to a more refined bound on the adversary's power with these constraints.

We then consider the problem of the adversary choosing weights subject to an ℓ_∞ constraint, say $\|\mathbf{s}\|_\infty = 1$; it is not difficult to show that the adversary will choose a seed such that each component has absolute value 1. One can think of this setting as allowing the adversary to partition the agents, and then seed them separately a bounded amount. In this case, one can of course give a trivial spectral bound for each objective. However, using known techniques from combinatorial optimization, we show that it is possible to efficiently obtain

an $O(1)$ approximation to the adversary’s power in this setting, and moreover, using known algorithms in this setting, obtain an explicit near-optimal vector for the adversary.

Finally, we consider the problem of an ℓ_1 constrained adversary. However, in this version of the problem, the problem becomes somewhat trivial, in that the adversary must pick from a fixed set of vectors that themselves are graph-independent; still, we give a spectral argument that nonetheless gives some bound on the adversary’s power, that is provably tight for the class of vertex-transitive graphs.

Sparsity-Constrained Adversary

We now explain a *sparsity-constrained* setting for adversarial optimization that was introduced in independent and concurrent work of Chen and Rácz [37]. We further show how combining the techniques considered here with their methods can provide a slight strengthening of their theoretical results.

The setup of Chen and Rácz is as follows: the network may have nonzero initial starting opinions $\mathbf{s}_0 \in [0, 1]^n$, and the adversary may then perturb the initial opinions in $[0, 1]^n$ arbitrarily subject to a k -sparsity constraint. That is, in their model, they consider the following optimization problem for the adversary:

$$\max_{\mathbf{s} \in [0, 1]^n: \|\mathbf{s} - \mathbf{s}_0\|_0 \leq k} \mathbf{s}^T (I + L)^{-1} L (I + L)^{-1} \mathbf{s}.$$

We now show how to combine their analysis with our spectral techniques to obtain slightly sharper bounds than the results in their paper on the amount an adversary can induce discord.

Lemma 8.4 (Lemma 4.1 of [37]). *For any \mathbf{s} satisfying the above constraints,*

$$\|(I + L)^{-1}(\mathbf{s} - \mathbf{s}_0)\|_1 \leq k.$$

Proof. It is well-known that matrix norm induced by the ℓ_1 vector norm is exactly the maximum largest column sum. As $\mathbf{1}$ is a right and left eigenvector of $(I + L)^{-1}$ with eigenvalue 1, and $(I + L)^{-1}_{ij} = \mathbf{e}_i^T (I + L)^{-1} \mathbf{e}_j \geq 0$ by the repeated-averaging interpretation of the dynamics, this implies that $I + L$ is doubly stochastic, so every absolute column sum is 1. Thus $\|(I + L)^{-1}(\mathbf{s} - \mathbf{s}_0)\|_1 \leq \|\mathbf{s} - \mathbf{s}_0\|_1 \leq \|\mathbf{s} - \mathbf{s}_0\|_0 \leq k$, where the last fact uses the fact the difference componentwise is at most 1 and the sparsity constraint. \square

Lemma 8.5. *For $\mathbf{s}_0 \in [0, 1]^n$,*

$$\|L(I + L)^{-1}\mathbf{s}_0\|_\infty \leq d_{\max},$$

where d_{\max} is the largest degree in G .

Proof. Note that $\mathbf{z}_0 := (I + L)^{-1}\mathbf{s}_0 \in [0, 1]^n$ as we have seen that $(I + L)^{-1}$ is doubly stochastic, so is a positive contraction in ℓ_∞ (this also follows from the repeated-averaging interpretation). Then

$$|(L\mathbf{z}_0)_i| = \left| \sum_{j=1}^n L_{i,j}(\mathbf{z}_0)_j \right| = \left| d_i(\mathbf{z}_0)_i - \sum_{j \neq i} A_{ij}(\mathbf{z}_0)_j \right| \leq \max\{d_i(\mathbf{z}_0)_i, d_i \max_j(\mathbf{z}_0)_j\} \leq d_{\max}.$$

\square

From these lemmas, we can now combine the ideas in [37] with the spectral approach of this work to obtain a sharpened bound in this setting when $k = \Omega(\sqrt{n})$. For convenience, we will write $(I + L)^{-1}L(I + L)^{-1}$ as Σ .

Theorem 8.6 (Theorem 1.3 of [37]). *For any graph G , the amount of disagreement an adversary can add is bounded by*

$$\lambda_{\max}(\Sigma)k + \sqrt{k} \min\{2d_{\max} \sqrt{k}, 2 \cdot \lambda_{\max}(\Sigma) \sqrt{n}\}.$$

Proof. The difference between the disagreement of \mathbf{s} and \mathbf{s}_0 after running the dynamics is easily seen to be

$$2(\Delta\mathbf{s})^T(I+L)^{-1}L(I+L)^{-1}\mathbf{s}_0 + (\Delta\mathbf{s})^T(I+L)^{-1}L(I+L)^{-1}(\Delta\mathbf{s}),$$

where $\Delta\mathbf{s} = \mathbf{s} - \mathbf{s}_0$. Applying Cauchy-Schwarz to the last term with the result in Theorem 8.1 gives an upper bound of

$$\|(I+L)^{-1}L(I+L)^{-1}\|\|\Delta\mathbf{s}\|_2^2 \leq \lambda_{\max}(\Sigma)k.$$

The first term can be bounded in a couple of ways: first, using Hölder's inequality and Lemma 8.4 and Lemma 8.5, one obtains

$$2(\Delta\mathbf{s})^T(I+L)^{-1}L(I+L)^{-1}\mathbf{s}_0 \leq 2\|(I+L)^{-1}(\Delta\mathbf{s})\|_1\|L(I+L)^{-1}\mathbf{s}_0\|_\infty \leq 2kd_{\max}. \quad (8.5)$$

Another way to bound this first term is simply by Cauchy-Schwarz and Theorem 8.1 to get

$$\begin{aligned} 2(\Delta\mathbf{s})^T(I+L)^{-1}L(I+L)^{-1}\mathbf{s}_0 &\leq 2\|(\Delta\mathbf{s})\|_2\|(I+L)^{-1}L(I+L)^{-1}\mathbf{s}_0\|_2 \\ &\leq 2 \cdot \lambda_{\max}(\Sigma) \sqrt{k}\|\mathbf{s}_0\|_2 \\ &\leq 2 \cdot \lambda_{\max}(\Sigma) \sqrt{nk}. \end{aligned}$$

We conclude that the increase in disagreement is bounded by

$$\lambda_{\max}(\Sigma)k + \sqrt{k} \min\{2d_{\max} \sqrt{k}, 2 \cdot \lambda_{\max}(\Sigma) \sqrt{n}\}. \quad (8.6)$$

□

In [37], the authors achieve a bound of $8d_{max}k$. From this combined result here, coupled with the previous sections, one obtains an improvement of the constants and dependence on d_{max} . Even stronger, the analysis here shows that the dependence on d_{max} is unnecessary for $d_{max} = \Omega(\sqrt{n/k})$. In fact, though the bound attained here may seem to be worse in the setting where $d_{max} \rightarrow 0$, this is not so; it is trivial to see that $\frac{x}{(1+x)^2} \leq x$ for $x \geq 0$, so that $\lambda_{max}(\Sigma) \leq \lambda_{max}(L)$. By the Gershgorin circle theorem, one also has $\lambda_{max}(L) \leq 2d_{max}$, thus giving an improvement in the small d_{max} setting as well.

We note that the linear dependence on k is necessary in general.³ For instance, consider the family of graphs $L_n := \frac{1}{n}L_{K_n}$. All nontrivial eigenvalues are located at $\lambda = 1$, hence are a worst-case example for disagreement. Consider the vectors \mathbf{s}^k obtained by setting the first k entries to 1, with the rest 0. It is easy to see that \mathbf{s}^k can be decomposed as

$$\mathbf{s}^k = \frac{k}{\sqrt{n}} \frac{\mathbf{1}}{\sqrt{n}} + \left(\sqrt{\frac{k(n-k)}{n}} \right) \mathbf{r}^k, \quad (8.7)$$

where \mathbf{r}^k is a unit vector orthogonal to $\mathbf{1}$, hence an eigenvector of Σ with eigenvalue $1/4$. Thus, if we instantiate this sparsity-constrained problem with $\mathbf{s}_0 = \mathbf{0}$, we have

$$\mathbf{s}^k \Sigma \mathbf{s}^k = \frac{k(n-k)}{4n} = (1 - k/n) \cdot \frac{k}{4}. \quad (8.8)$$

In particular, for $k < n/2$ for instance, this is $\Omega(k)$.

³By this, we just mean that there *exists* instances of L and \mathbf{s}_0 where the dependence in k is provably linear. It is easy to see that for any L , there exists some \mathbf{s}_0 where the dependence on k is trivial. For any L , let \mathbf{s}_0 be a vector that maximizes disagreement in the cube $[0, 1]^n$ (as the maximum of a convex function, it will lie in $\{0, 1\}^n$, but that is not used in the argument). Then no perturbation inside the cube can improve it, hence the dependence on k is trivial. It would be interesting, as noted in [37], if one could show that the dependence on k is linear for “typical” instances, as suggested by their empirical results.

ℓ_∞ -Constrained Adversary

Now consider the problem of the adversary choosing weights subject to an ℓ_∞ constraint; for convenience, we will write the disagreement problem⁴ as

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_\infty=1} \mathbf{s}^T (I + L)^{-1} L (I + L)^{-1} \mathbf{s}. \quad (8.9)$$

For simplicity, again write the matrix as Σ . The first simple observation is that the optimal adversary strategy \mathbf{s}^* will satisfy $|\mathbf{s}_i^*| = 1$ for all i ; indeed, if one fixes \mathbf{s}_{-i}^* , then the objective is easily seen to be linear in \mathbf{s}_i , and therefore attains a maximum for $\mathbf{s}_i^* \in \{-1, 1\}$. In particular, the problem can be rewritten as

$$\max_{\mathbf{s} \in \mathbb{R}^n: \forall i \mathbf{s}_i \in \{\pm 1\}} \mathbf{s}^T \Sigma \mathbf{s}. \quad (8.10)$$

In general, there does not appear to be any analytic solution to this problem. One always has the trivial spectral bound $\lambda_n(\Sigma) \cdot n$, just using the fact that the ℓ_2 -norm of the candidate solutions are all \sqrt{n} . However, by leveraging deep results from functional analysis that have long been fruitfully applied to combinatorial optimization, it is possible to give an efficient constant factor approximation:

Theorem 8.7. *There exists a polynomial-time algorithm that computes a $\pi/2$ -approximation to Equation (8.10); moreover, a vector $\mathbf{s}' \in \{-1, 1\}^n$ attaining this bound can be obtained efficiently.⁵*

Proof. To prove the theorem, we apply *Grothendieck's inequality* [5], which asserts for any positive semidefinite matrix Z ,

$$\max_{\mathbf{x} \in \mathbb{R}^n: \forall i \mathbf{x}_i \in \{\pm 1\}} \mathbf{x}^T Z \mathbf{x} \leq \max_{\mathbf{x}_1, \dots, \mathbf{x}_n \in S^{n-1}} \sum_{i,j=1}^n Z_{ij} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \leq \frac{\pi}{2} \max_{\mathbf{x} \in \mathbb{R}^n: \forall i \mathbf{x}_i \in \{\pm 1\}} \mathbf{x}^T Z \mathbf{x}, \quad (8.11)$$

⁴The corresponding results for any other positive semidefinite objective, like the others we considered before, follow in exactly the same way.

⁵We ignore issues of bit complexity; if desired, one may assume that L is given as a rational matrix and that we will be content with solutions up to ϵ accuracy.

where $S^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$. It is well-known that the middle expression is easily attained as the solution to the semidefinite program

$$\begin{aligned} & \max_{X \in \mathbb{R}^{n \times n}} \text{Tr}(ZX) \\ & \text{subject to } X_{ii} = 1 \\ & X \succeq 0, \end{aligned}$$

and therefore can be solved to arbitrary accuracy in polynomial time. By setting Z to be our matrix Σ and applying this result, it follows that the solution to this problem thus gives a $\frac{\pi}{2}$ -approximation to the adversary's problem. We remark that, following the approach for the MAXCUT problem in [50], one can appeal to strong SDP duality to obtain an equivalent, and slightly more analytic, spectral upper bound of $n \cdot \min_{\mathbf{x}: \mathbf{x}^T \mathbf{1} = 0} \lambda_{\max}(\Sigma - \text{diag}(\mathbf{x}))$. This clearly improves on the trivial spectral bound given before, though it is not clear in general what good test vectors \mathbf{x} are to give an explicit better bound.

To actually obtain a ± 1 vector attaining this bound compared to the optimum of the SDP, Alon and Naor [5] show that randomized hyperplane rounding applied to the vectors attaining the SDP optimum can be used to obtain a vector \mathbf{s} that is at least $\frac{2}{\pi}$ of the SDP optimum in expectation, hence of the original problem. □

ℓ_1 -Constrained Adversary

One may also consider an adversary that is bounded in ℓ_1 , which may be viewed as a more natural restriction. However, in this case, the problem becomes trivial; again, we consider the disagreement problem with matrix Σ , but all these results carry over for any other positive semidefinite objective.

Theorem 8.8. *Suppose the adversary is now ℓ_1 -constrained, so that the optimization problem is*

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_1 = 1} \mathbf{s}^T \Sigma \mathbf{s}. \quad (8.12)$$

Then, an optimal adversary strategy is simply to set $\mathbf{s} = \mathbf{e}_i$ where $i \in \arg \max_{j \in [n]} \Sigma_{jj}$ (that is, to put all their budget on an index with largest diagonal term in Σ).

In this case, denoting the adversary's power as D , we have

$$\frac{1}{n} \sum_{i=1}^n \frac{\lambda_i(L)}{(1 + \lambda_i(L))^2} \leq D \leq \max_{i \in [n]} \frac{\lambda_i(L)}{(1 + \lambda_i(L))^2} \leq \frac{1}{4}, \quad (8.13)$$

where the lower bound is sharp in the case of vertex-transitive graphs.

Proof. For an optimal adversary strategy, simply note that the ℓ_1 -ball is the convex hull of the set $\{\pm \mathbf{e}_1, \dots, \pm \mathbf{e}_n\}$. As the maximum of a convex function over a convex set is attained at an extreme point, it suffices to consider this set of strategies; by homogeneity, it suffices to just consider $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$. But when substituting these terms in the optimization problem, one then recovers the diagonal elements of Σ , so the adversary may simply choose the largest such term.

The upper bound follows from the ℓ_2 bound we gave before, using the fact that the ℓ_1 unit ball is contained in the ℓ_2 unit ball. For the lower bound, observe that

$$\sum_{i=1}^n \mathbf{e}_i^T \Sigma \mathbf{e}_i = \sum_{i=1}^n \Sigma_{ii} = \text{Tr}(\Sigma). \quad (8.14)$$

As the trace of any matrix is equal to the sum of the eigenvalues, it follows there must exist a diagonal element that is at least the average of the eigenvalues.

To see the tightness of the lower bound, consider any vertex-transitive graph. By definition, the adjacency matrix is unchanged under the action of a transitive subgroup of the symmetric group. Any vertex-transitive graph is regular, so

the degree matrix is a multiple of the identity, hence also unchanged under the action of any permutation; together, these imply that the Laplacian of the graph must be invariant under the action of a transitive subgroup of the symmetric group. As a result, any rational expression of the Laplacian is also invariant. By applying suitable automorphisms of the graph, this implies that any two diagonal elements of any rational expression of the Laplacian must be equal, in which case they are all equal to the average of the eigenvalues by the trace identity. \square

8.4 Defending the Network

In this section, we consider the problem of *defending* a given network from adversarial perturbations like those considered above. We will view this as a two-player min-max game; first, a network defender will choose how to set some qualitative feature of the network subject to normalization constraints modeling the resource limitation of the defender. Then, the adversary performs the above maximization problem with this choice of settings. The goal of the defender is to choose a setting to minimize the cost of the resulting system (e.g., the measure of disagreement), knowing that the adversary will optimize for this choice. In this section, we show that, in one such formulation, the defender can efficiently do this via solving an appropriate convex optimization problem.

We have generally adopted the convention that the network topology is basically fixed; it is unrealistic to substantively change a real-world network structure. Therefore, in this formulation, the network defender chooses how to vary the cost of the adversary in changing initial opinions of different nodes. That is,

the network defender can choose to weigh each node differently for example, by lessening their exposure to misinformation, so that the adversary pays different costs for perturbing different nodes.

Formally, we consider the following problem: suppose the network defender is resource limited according to a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, such as the ℓ_1 -norm, and is permitted to change node-weights with the restriction of $h(\mathbf{w}) = h(\mathbf{1})$. We will consider, under this resource constraint, what the defender's optimal choice of \mathbf{w} is. More generally, we will assume that h is nonnegative, convex, and radially increasing and homogeneous (i.e. for $\alpha \geq 0$, $h(\alpha \mathbf{x}) = g(\alpha)h(\mathbf{x})$, with $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ an increasing function), as well as a function $f : \mathbb{R} \rightarrow \mathbb{R}$ as in Equation (8.1) that induces a positive semidefinite quadratic form. Then the network defender must solve the following optimization problem:

$$\min_{\mathbf{w} \in \mathbb{R}_{>0}^n : h(\mathbf{w}) = h(\mathbf{1})} \left\{ \max_{\mathbf{s} \in \mathbb{R}^n : \|\mathbf{s}\|_{\mathbf{w},2} \leq R} \mathbf{s}^T (I + L)^{-1} f(L) (I + L)^{-1} \mathbf{s} \right\}. \quad (8.15)$$

In words, the network defender chooses a *weighted* ℓ_2 -norm on the nodes of the network under the resource constraint modeled by h that specifies costs of influencing each individual in the network heterogeneously; with these weights and the same fixed budget R as before, the adversary then optimizes their objective. For instance, if h is the ℓ_1 -norm, then this normalization imposes that $\sum_{i=1}^n w_i = \sum_{i=1}^n 1 = n$, so that the sum of weights on the nodes for the adversary is the same as for the regular ℓ_2 norm. Other natural choices for h include any norm on \mathbb{R}^n or any sum of squares of linear expressions. As a result, $\mathbf{w} = \mathbf{1}$ is a valid choice of the network defender, in which case the inner maximization corresponds to the largest eigenvalue of the relevant quadratic form as we have seen above. However, for other choices of \mathbf{w} , the inner maximization does not have the same interpretation and moreover, will not usually admit a clean

analytical expression as the maximization of a convex objective.

Here, we show that despite this difficulty, this can be reduced to convex optimization via a geometric argument. For convenience, set $R = 1$; this is without loss of generality as the inner maximization is homogeneous. Consider the following procedure:

1. Solve the following convex program with positive semidefinite constraints:

$$\begin{aligned} & \min_W h(\text{diag}(W)) \\ & \text{subject to } 0 \leq (I + L)^{-1} f(L)(I + L)^{-1} \leq W \\ & W_{ij} = 0, \quad \forall i \neq j. \end{aligned}$$

2. Set $\mathbf{w}' = \text{diag}(W^*)$, where W^* is a solution to the above convex program.
3. Let $t \geq 0$ be such that $h(t\mathbf{w}') = h(\mathbf{1})$.
4. Set $\mathbf{w}^* = t\mathbf{w}'$.

We now show that this procedure gives the optimal setting of \mathbf{w} .

Theorem 8.9. *Under the restrictions on h and f , the above algorithm yields the optimal value of the problem given by Equation (8.15).*

Proof. Let h and f be as required and write $\Sigma = (I + L)^{-1} f(L)(I + L)^{-1}$; by our assumption on f , $\Sigma \geq 0$. For any fixed choice of \mathbf{w} , the adversary's optimal choice of $\mathbf{s} \in \mathbb{R}^n$ is obtained by finding the largest level set of the function $\mathbf{x}^T \Sigma \mathbf{x}$ that nontrivially intersects the ellipsoid $\|\mathbf{x}\|_{\mathbf{w},2} \leq 1$. Equivalently, the optimal value of the adversary is the smallest level set of $\mathbf{x}^T \Sigma \mathbf{x}$ that contains the unit

ball of the norm induced by $\|\cdot\|_{\mathbf{w}}$. In particular, the optimal value of the inner maximization for fixed \mathbf{w} is the smallest value $K \geq 0$ such that

$$\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathbf{w},2} \leq 1\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq K\} \quad (8.16)$$

By the restriction on W to being diagonal in the above convex program (with necessarily nonnegative diagonal entries by the PSD constraint), recall that $W \geq \Sigma$ if and only if for all $\mathbf{x} \in \mathbb{R}^n$,

$$\|\mathbf{x}\|_{\mathbf{w}',2}^2 = \mathbf{x}^T W \mathbf{x} \geq \mathbf{x}^T \Sigma \mathbf{x};$$

where $\mathbf{w}' = \text{diag}(W)$; geometrically, this is equivalent to the containment

$$\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathbf{w}',2} \leq 1\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq 1\}. \quad (8.17)$$

In particular, this means that $\Sigma \leq W$ if and only if the unit ball of $\|\cdot\|_{\mathbf{w}'}$ is contained in the unit ball of the (semi)-norm induced by Σ . Let W^* and $\mathbf{w}' = \text{diag}(W^*)$ be as stated, and let $t \geq 0$ be such that $h(t\mathbf{w}') = h(\mathbf{1})$. By the minimality of \mathbf{w}' as well as the positive homogeneity, this implies that if the optimal value of the inner maximization for Equation (8.15) using $t\mathbf{w}'$ is K , then using homogeneity of the containments:

$$\begin{aligned} & \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{t\mathbf{w}',2} \leq \frac{1}{\sqrt{K}}\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq 1\} \\ \iff & \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathbf{w}',2} \leq \frac{1}{\sqrt{t \cdot K}}\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq 1\}; \end{aligned}$$

as such, $K = 1/t$.

Suppose now for a contradiction that an optimizer \mathbf{w}^* of Equation (8.1) has strictly smaller objective value $K^* < 1/t$ than that of $t\mathbf{w}'$. By Equation (8.16), this is equivalent to

$$\begin{aligned} & \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathbf{w}^*,2} \leq \frac{1}{\sqrt{K^*}}\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq 1\} \\ \iff & \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{K^*\mathbf{w}^*,2} \leq 1\} \subseteq \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T \Sigma \mathbf{x} \leq 1\}. \end{aligned}$$

Evidently, $K^* \mathbf{w}^*$ satisfies Equation (8.17) yet

$$h(K^* \mathbf{w}^*) = g(K^*)h(\mathbf{w}^*) < g(1/t)h(t\mathbf{w}') = h(\mathbf{w}').$$

This violates the optimality of \mathbf{w}' , yielding the desired contradiction. \square

Note that for certain choices of $h(\cdot)$, the above can be written as a standard semidefinite program. For instance, if $h(\mathbf{x}) = \|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$, then the above is indeed a regular semidefinite program. If h is instead the squared ℓ_2 -norm, one can similarly write it as a semidefinite program by adding extra positive semidefinite constraints and exploiting Schur complements; we omit the details here.

8.5 Mixed-Graph Objectives

In the previous sections, we have connected an adversary's ability to induce discord in a network with the spectral theory of the underlying graph, as well as considered ways to defend against these attacks. In each of these settings, these results suggest that the opinion dynamics of the network will necessarily "soften" the effect of these attacks, as the disagreement is measured on the edges of the same network that dictate the opinion dynamics. However, one potential explanation of the success of recent adversarial attacks described in the introduction is that the opinion formation graph and disagreement measurement graph need not be the same, and may not even look similar. For instance, opinion formation may take place on the "online" network, via social media, while the disagreement the adversary cares about maximizing may be measured with respect to "real-world" connections. When this occurs, it need

not be the case that the opinion dynamics implicitly equilibrate disagreements measured along the latter graph. When opinion formation and the disagreement graph look quite different, one expects that an adversary will be able to induce significantly more disagreement.

In this section, we explore the degree to which the adversary’s power can increase when the opinion formation graph and measurement graph become independent. First, we provide nontrivial examples that show that, in some cases, having a different graph for the opinion dynamics and for the disagreement measurement can actually *reduce the adversary’s power to induce disagreement*; however, we provide a general lower bound that indicates that typically, the adversary will not be much worse off, if at all. We then show that the relevant relationship that will determine when an adversary gains extra power is an appropriate notion of *spectral similarity* between the two graphs, not necessarily physical similarity. We conclude this section by showing concretely how a large cut misalignment in the graphs will enable an adversary to induce disagreement far beyond what is possible in the spectral theory in the single-graph setting.

Formally, we generalize the previous sections as follows: suppose that there are now *two* relevant graph structures on $[n]$, G_1 and G_2 , with associated Laplacians L and M (which will be the “measurement” Laplacian). The first graph, G_1 , is the graph structure on which the Friedkin-Johnsen dynamics take place while the second graph, G_2 , is the graph where disagreement is measured. In this setting, the adversary chooses initial opinions to maximize the following objective (setting $R = 1$ for notational ease):

$$\max_{\mathbf{s} \in \mathbb{R}^n: \|\mathbf{s}\|_2 \leq 1} \mathbf{s}^T (I + L)^{-1} M (I + L)^{-1} \mathbf{s} = \lambda_{\max}((I + L)^{-1} M (I + L)^{-1}).$$

Typically, we will be interested in settings where L and M are of comparable size, meaning similar total edge weight. If not, say if L has much larger edge weight than M , then the effects of the opinion dynamics will cause all opinions to smooth to a much larger degree compared to the measurement measured via M , so the problem becomes degenerate though not for a theoretically interesting reason.

Remark 8.10. *We note that the ℓ_∞ version of the adversary problem in this setting is easily seen to be NP-hard. The reason is that one can let M be the Laplacian of any graph, and simply set L to be the trivial graph with no edges, at which point the problem can be shown to be precisely MAXCUT (see Section 8.5.1 for the relation to cuts). As such, by appealing to various complexity-theoretic assumptions, one can easily establish hardness-of-approximation results for this version of the problem [98, 87]. In particular, it is NP-hard to attain an approximation within better than a $17/16$ -factor of the optimum in general and UGC-hard to obtain a solution within a factor of better than $\alpha_{GW} \approx 1.14$ of the optimum in general with an ℓ_∞ -adversary.*

In general, it is not obvious how to connect the spectral structure of the above matrix with the spectral properties of the two underlying graphs, unless in the special case where the two graphs commute (and therefore, share an eigenbasis); this is indeed possible in certain special cases.

Example 8.11. *Suppose G_1 is a d -regular, unweighted graph, and let G_2 be the $n-d-1$ -regular, unweighted complementary graph. Then one can check that $M = nI - J - L$, where J is the matrix of all-ones. Every matrix on the right side shares an eigenbasis, hence L and M commute. Similarly, suppose that G_1 or G_2 is a scaled version of the unweighted complete graph K_n . Then it is easy to check that L_{K_n} commutes with every graph Laplacian L as they will share a common eigenbasis.*

One might suspect that, fixing M , $L = M$ is the optimal choice of graph Laplacian (subject to normalization) to minimize the amount of disagreement an adversary can induce. That is, letting \mathcal{L} be the set of graph Laplacians subject to the edge normalization $\text{Tr}(L) = \text{Tr}(M)$, one might guess that

$$M \in \arg \min_{L \in \mathcal{L}} \lambda_{\max}((I + L)^{-1} M (I + L)^{-1}).$$

However, this does not hold in general, via the following simple construction.

Example 8.12. Suppose that G_2 is a complete graph, in the sense that for some $\epsilon > 0$, all off-diagonals of the Laplacian satisfy

$$M(i, j) < -\epsilon.$$

If we write out M in the eigenbasis as

$$M = \sum_{i=2}^n \lambda_i(M) \mathbf{v}_i \mathbf{v}_i^T$$

and further suppose for simplicity that $1 < \lambda_2(M) < \lambda_3(M)$, so that all nonzero eigenvalues lie on the right side of the peak of the function $f(x) = x/(1+x)^2$ at $x = 1$. If we then set $L = M$, we would get an objective value of

$$\frac{\lambda_2(M)}{(1 + \lambda_2(M))^2},$$

as we have seen before. But consider instead the matrix

$$L = (\lambda_2(M) + \eta) \mathbf{v}_2 \mathbf{v}_2^T + (\lambda_3(M) - \eta) \mathbf{v}_3 \mathbf{v}_3^T + \sum_{i=4}^n \lambda_i(M) \mathbf{v}_i \mathbf{v}_i^T$$

for some $\eta > 0$ sufficiently small (depending on ϵ). It is easy to see that $L \in \mathcal{L}$ as the sum of eigenvalues, and therefore the trace, is constant, and moreover, L will still be a Laplacian (with nonpositive off-diagonal entries) of some other graph by continuity. As L and M share an eigenbasis, it is easy to see that for η small enough,

$$\lambda_{\max}((I + L)^{-1} M (I + L)^{-1}) = \frac{\lambda_2(M)}{(1 + \lambda_2(M) + \eta)^2}$$

which is strictly smaller than if $L = M$.

Example 8.13. *In a more interesting example, suppose now that we further require that $(L)_{i,i} = (M)_{i,i}$ for each i . This means that each node has the same weighted degree in both graphs, which in particular implies they have the same trace. First, consider $G_2 = C_4$, the four node unweighted cycle graph. One can numerically check that among all graphs G_1 satisfying this normalization, the mixed objective $(I + L)^{-1}M(I + L)^{-1}$ is minimized when G_1 is a weighted complete graph where each edge in the cycle has weight reduced from 1 to approximately .89, and the remaining two edges are increased from 0 to .22. The mixed objective has value approximately 0.1929, whereas the single-graph objective $(I + M)^{-1}M(I + M)^{-1}$ has largest eigenvalue $2/9 \approx .22$. For comparison, when G_1 is instead set to the appropriately scaled copy of the complete graph, the mixed-objective actually rises to .2975. On the other hand, when $G_2 = P_4$, the unweighted path graph on four nodes, it is numerically optimal for itself under the mixed-graph objective for all graphs satisfying the degree constraint. We are unaware of an analytic reason why this holds.*

However, one expects that these examples are largely pathological. As a first approximation to controlling this quantity, our first result is the following general bound for positive semidefinite matrices. The idea is to apply the Courant-Fischer theorem to subspaces spanned by the eigenvectors of the two matrices to lower bound the spectral norm of the product of matrices.

Lemma 8.14. *Let $B, C \in \mathbb{R}^{n \times n}$ be positive semidefinite matrices with eigenvalues in increasing order. Then*

$$\max_{k \leq n} \left\{ \lambda_{n-k+1}(C)^2 \lambda_k(B) \right\} \leq \lambda_n(CBC) = \|CBC\|_2 \leq \lambda_n(C)^2 \lambda_n(B).$$

Proof. Note that $CBC \geq 0$ by the fact $B \geq 0$. The upper bound follows directly from the submultiplicativity of the operator norm, which for symmetric positive-semidefinite matrices is just the top eigenvalues.

For the lower bound, we use the Courant-Fischer theorem. First, note that if $\lambda_{n-k+1}(C) = 0$, the result is trivial, so suppose it is strictly positive. Let U be the linear subspace spanned by the top k eigenvectors of C , and let V be the subspace spanned by the top $n - k + 1$ eigenvectors of B . These subspaces must intersect non-trivially by a simple dimension argument, so there exists some $\mathbf{z} \in U \cap V$ with unit length. Note that U is an invariant subspace for C , and moreover, C is bijective on U by the nondegeneracy of $\lambda_{n-k+1}(C)$. Now, let $\mathbf{x} = C^{-1}\mathbf{z}$, where we view C^{-1} as restricted to U . By the variational formula of λ_n ,

$$\lambda_n(CBC) \geq \frac{\mathbf{x}^T CBC\mathbf{x}}{\|\mathbf{x}\|_2^2} = \frac{\mathbf{z}^T B\mathbf{z}}{\mathbf{z}^T C^{-2}\mathbf{z}} \geq \lambda_k(B) \frac{\|\mathbf{z}\|_2^2}{\|C^{-1}\mathbf{z}\|_2^2} = \frac{\lambda_k(B)}{\|C^{-1}\mathbf{z}\|_2^2}$$

The second inequality follows from Courant-Fischer, as \mathbf{z} lies in the span of the top $n-k+1$ eigenvectors of B by assumption, so the quadratic form in the numerator gives at least $\lambda_{n-(n-k+1)+1}(B)\|\mathbf{z}\|^2 = \lambda_k(B)\|\mathbf{z}\|^2$. Then $\|C^{-1}\mathbf{z}\|^2 \leq \lambda_{\max}(C^{-1})^2\|\mathbf{z}\|^2 = \frac{\|\mathbf{z}\|^2}{\lambda_{n-k+1}(C)^2}$, as the largest eigenvalue of C^{-1} restricted to U is the inverse of the smallest eigenvalue of C restricted to U . Plugging this in gives the desired inequality. As this holds for all $k \leq n$, it holds for the maximum. \square

From this simple lemma, one can immediately obtain a lower bound in the adversary's optimization problem in the mixed-graph setting.

Corollary 8.15. *Let L, M be as above. Then*

$$\lambda_{\max}((I + L)^{-1}M(I + L)^{-1}) \geq \max_{1 \leq k \leq n} \frac{\lambda_k(M)}{(1 + \lambda_k(L))^2}.$$

Proof. This is immediate from the previous lemma, with $B = M$ and $C = (I + L)^{-1}$, simply noting that

$$\lambda_{n-k+1}((I + L)^{-1}) = \frac{1}{1 + \lambda_k(L)}.$$

\square

Using just this lower bound, in the special case where $M = L_{K_n}$, the Laplacian of the unweighted complete graph on n nodes, it follows that the optimal choice of L subject to having the same trace as L_{K_n} to minimize the mixed-graph objective is just L_{K_n} itself. This holds because for any graph Laplacian L satisfying $\text{Tr}(L) = \text{Tr}(L_{K_n})$, a similar argument to that of Corollary 8.3 implies that some nontrivial eigenvalue of L_1 must be at most n . The lower bound of the previous corollary then asserts that the mixed-graph objective can only increase, with equality if and only if $L = L_{K_n}$.

Moreover, the previous corollary asserts that if the eigenvalues of L and M are only *numerically* similar in the appropriate ordering, then necessarily the objective value will be *approximately at least* the corresponding objective value we considered in Section 8.3.1. Explicitly, this will arise for any graphs with *cospectral* Laplacians; for instance, any isomorphic graphs will have this property (so in particular, if M and L differ by a permutation), as will any strongly regular graphs with same parameters. This suggests that while we have shown explicit examples where having two distinct matrices can even reduce the adversary's power, this case ought be viewed as rather pathological.

8.5.1 Spectral Similarity

The above analysis relied only on a general lower bound involving positive semidefinite matrices. Next, we aim to characterize the relevant structure of L and M that causes the objective function to remain quite close to the value in the single-graph case, and similarly when the objective function will increase. The former case will indicate that an adversary gains little benefit from the misalign-

ment of G_1 and G_2 , while the latter case corresponds to an underlying network that can be sharply exploited to induce large disagreement. Intuitively, if $L \approx M$ component-wise, then

$$\lambda_{\max}((I + L)^{-1}M(I + L)^{-1}) \approx \lambda_{\max}((I + M)^{-1}M(I + M)^{-1}),$$

by the continuity of matrix inverses and eigenvalues. Before proceeding, we need a definition:

Definition 8.4. For any graph $G = (V, E, w)$ and $S, T \subseteq V$ such that $S \cap T = \emptyset$, we define

$$\text{cut}_G(S, T) = \sum_{i \in S, j \in T} w_G(i, j).$$

We will write $\text{cut}_G(S) := \text{cut}_G(S, S^c)$.

For any subset $S \subseteq [n]$, we write χ_S for the ± 1 indicator vector of S , i.e. $\chi_S(i) = 1$ if $i \in S$ and -1 if $i \notin S$. Then it is easy to see that $\|\chi_S\|^2 = n$, and that for any graph G with Laplacian L ,

$$\begin{aligned} \chi_S^T L \chi_S &= \sum_{(i,j) \in E_i} w_G(i, j) (\chi_S(i) - \chi_S(j))^2 \\ &= 4 \sum_{i \in S, j \notin S} w_G(i, j) = 4 \text{cut}_G(S). \end{aligned}$$

We can now provide a quantitative form of this assertion:

Theorem 8.16. *Let G_1, G_2 be n -node graphs with Laplacians M and L , respectively. Suppose that the following holds for some parameters $\eta, \gamma, \epsilon > 0$:*

1. *For each $i \in [n]$, the weighted symmetric difference of their neighborhoods is bounded by η , i.e. for all $i \in [n]$*

$$\sum_{j \neq i} |w_1(i, j) - w_2(i, j)| \leq \eta. \tag{8.18}$$

2. For all $i \in [n]$, the absolute difference in weighted degrees of i in G_1 and G_2 is at most γ .

3. For any disjoint subsets $S, T \subseteq [n]$, we have

$$|\text{cut}_{G_1}(S, T) - \text{cut}_{G_2}(S, T)| \leq \epsilon \sqrt{|S||T|}. \quad (8.19)$$

Then, we have

$$\max_{i \in [n]} \frac{\lambda_i(M) - 2\Delta}{(1 + \lambda_i(M) + \Delta)^2} \leq \lambda_n((I + L)^{-1}M(I + L)^{-1}) \leq \max_{i \in [n]} \frac{\lambda_i(M) + 2\Delta}{(1 + \lambda_i(M) - \Delta)^2}, \quad (8.20)$$

where $\Delta = O(\epsilon \ln(\eta/\epsilon) + \gamma)$. In particular, if all nodes have the same weighted degree in both G_1 and G_2 , $\Delta = O(\epsilon \ln(\eta/\epsilon))$.

Remark 8.17. Before proceeding with the proof, note that if one only assumes the first condition above, and even if every node has the same degree in both graphs (therefore satisfying the second condition with $\gamma = 0$), the best bound one can generically get on the spectral radius of $M - L$ is $O(\eta)$ using the fact that the largest eigenvalue of a matrix is at most the largest ℓ_1 norm of a row. When the combinatorial structures are assumed to be very similar along every subset, as is done here, the dependence on η becomes logarithmic, and gains from the closeness in the ϵ term as well.

Proof. First, by our assumptions, we may apply Lemma 3.3 of Bilu and Linial [23], where we just note that if u is the $\{0, 1\}$ indicator of S and v is the $\{0, 1\}$ indicator for T for some disjoint subsets $S, T \subseteq [n]$, then

$$|u^T(M - L)v| = |\text{cut}_{G_1}(S, T) - \text{cut}_{G_2}(S, T)|. \quad (8.21)$$

We also note that by inspecting the proof of that lemma, one can apply our condition (2) with parameter γ instead of $O(\epsilon \ln(\eta/\epsilon))$ by just paying it in the bound, from which it follows that $\Delta := \|M - L\| \leq O(\epsilon \ln(\eta/\epsilon) + \gamma)$.

From this, we have

$$L - \Delta \cdot I \leq M \leq L + \Delta \cdot I, \quad (8.22)$$

which in turn implies

$$(I + L)^{-1}(L - \Delta \cdot I)(I + L)^{-1} \leq (I + L)^{-1}M(I + L)^{-1} \leq (I + L)^{-1}(L + \Delta \cdot I)(I + L)^{-1}. \quad (8.23)$$

Finally, note that by Weyl's monotonicity theorem, we have $\max\{\lambda_i(M) - \Delta, 0\} \leq \lambda_i(L) \leq \lambda_i(M) + \Delta$. Combining all these bounds with another application of Weyl's monotonicity theorem, we have

$$\max_{i \in [n]} \frac{\lambda_i(M) - 2\Delta}{(1 + \lambda_i(M) + \Delta)^2} \leq \lambda_n((I + L)^{-1}M(I + L)^{-1}) \leq \max_{i \in [n]} \frac{\lambda_i(M) + 2\Delta}{(1 + \max\{\lambda_i(M) - \Delta, 0\})^2}. \quad (8.24)$$

□

That high physical similarity of the graphs implies the problem is not changed significantly is not particularly surprising, though the previous result gives exponentially better dependence on the physical similarity than what can be attained by naive applications of matrix perturbation bounds. However, we now show that high physical similarity edge-by-edge is merely sufficient, but *not necessary*; another relevant property that will ensure that this holds is *spectral similarity*, as defined by Spielman and Teng [125].

Definition 8.5. L and M are ϵ -spectral approximations for each other for some $\epsilon > 0$ if

$$\frac{1}{1 + \epsilon}L \leq M \leq (1 + \epsilon)L.$$

Note that this definition is symmetric in L and M . It is easy to show from this definition that if L and M are ϵ -spectral approximations of each other, then the adversary's objective value cannot differ too much from the single-graph setting with just M .

Theorem 8.18. *Suppose that L and M are ϵ -spectral approximations of each other.*

Then

$$\begin{aligned} \frac{1}{1+\epsilon} \max_{i \in [n]} \min_{c \in [\frac{1}{1+\epsilon}, 1+\epsilon]} \frac{c\lambda_i(M)}{(1+c\lambda_i(M))^2} &\leq \lambda_{\max}((I+L)^{-1}M(I+L)^{-1}) \\ &\leq (1+\epsilon) \max_{i \in [n]} \max_{c \in [\frac{1}{1+\epsilon}, 1+\epsilon]} \frac{c\lambda_i(M)}{(1+c\lambda_i(M))^2}. \end{aligned}$$

Proof. The proof is essentially immediate from the definition: pre- and post-multiplying by $(I+L)^{-1}$, we immediately get from the definition that

$$\frac{1}{1+\epsilon}(I+L)^{-1}L(I+L)^{-1} \leq (I+L)^{-1}M(I+L)^{-1} \leq (1+\epsilon)(I+L)^{-1}L(I+L)^{-1}.$$

By Weyl's monotonicity theorem, this implies the corresponding inequality on each of the eigenvalues. We deduce that

$$\frac{1}{1+\epsilon}\lambda_n((I+L)^{-1}L(I+L)^{-1}) \leq \lambda_n((I+L)^{-1}M(I+L)^{-1}) \leq (1+\epsilon)\lambda_n((I+L)^{-1}L(I+L)^{-1}).$$

To relate this back to the matrix $(I+M)^{-1}M(I+M)^{-1}$, we again use Weyl's monotonicity theorem, as then for each $i \in [n]$,

$$\frac{1}{1+\epsilon}\lambda_i(L) \leq \lambda_i(M) \leq (1+\epsilon)\lambda_i(L).$$

In particular, $\lambda_i(L)$ lies in the $(1+\epsilon)$ -neighborhood of $\lambda_i(M)$. We showed above that

$$\lambda_{\max}((I+M)^{-1}M(I+M)^{-1}) = \max_{i \in [n]} \frac{\lambda_i(M)}{(1+\lambda_i(M))^2};$$

plugging in these ‘‘fuzzy’’ versions of the eigenvalues gives the desired inequalities. □

Remark 8.19. *Note that while the definition of spectral similarity is symmetric, it need not commute nicely with positive rational expressions of the Laplacians. The reason is that in general, positive rational expressions need not be operator monotone, i.e. may not respect the Loewner order. For instance, $0 \leq A \leq B$ does not imply $A^2 \leq B^2$, requiring us to appeal to Weyl's monotonicity theorem to translate between M and L .*

As a corollary, this result shows that it is not necessary for L and M to be extremely close in, say, Frobenius or ℓ_1 norm on each row for the eigenvalues for the adversary's objective value to remain close to the single-graph setting. This is because by seminal results of Batson, Spielman, and Srivastava, *every* graph Laplacian has a weighted ϵ -spectral approximation that corresponds to a graph with $O(n/\epsilon^2)$ edges [19]. Necessarily, these graphs are physically quite different, as they can differ in $\Theta(n^2)$ entries. The previous result shows that this is irrelevant; in the mixed-graph objective function, replacing one of these graphs by the other does not meaningfully change the adversary's power to induce disagreement under the Friedkin-Johnsen dynamics.

8.5.2 Spectral Dissimilarity

In this section, we provide a partial converse to the previous section; we provide a simple condition that will imply that the relevant largest eigenvalue is large that relates to the spectral dissimilarity of L and M . We then show how this can be realized in the special case of cuts in G_1 and G_2 ; it will turn out that if G_1 and G_2 are highly misaligned in the sense of having even one drastically different vertex cut, then the largest eigenvalue is necessarily large.

Definition 8.6. We say L and M are (ϵ, η) -bad spectral approximations if there exists $\mathbf{x} \in \mathbb{R}^n$ with $\|\mathbf{x}\|^2 = n$ such that $\mathbf{x}^T L \mathbf{x} \leq \epsilon$ and $\mathbf{x}^T M \mathbf{x} \geq \eta$.

This definition implies that M and L are not η/ϵ -spectral approximations for each other, but we will crucially be interested in the actual values, not just the ratio. Moreover, notice that this is not symmetric in the directions of the inequalities, and we will actually care about the numerical values, not just the

ratio. For these reasons, the following does not constitute an exact converse, which is essentially immediate from this definition:

Proposition 8.20. *Suppose that L and M are (ϵ, η) -bad spectral approximations. Then*

$$\lambda_n((I + L)^{-1}M(I + L)^{-1}) \geq \frac{\eta}{n + (\|L\| + 2)\epsilon}.$$

Proof. This follows from the variational characterization of eigenvalues:

$$\begin{aligned} \lambda_n((I + L)^{-1}M(I + L)^{-1}) &= \max_{\mathbf{z} \in \mathbb{R}^n} \frac{\mathbf{z}^T (I + L)^{-1} M (I + L)^{-1} \mathbf{z}}{\mathbf{z}^T \mathbf{z}} \\ &= \max_{\mathbf{z} \in \mathbb{R}^n} \frac{\mathbf{z}^T M \mathbf{z}}{\mathbf{z}^T (I + L)^2 \mathbf{z}}. \end{aligned}$$

The proof follows from plugging in the guaranteed vector \mathbf{x} , simply noting that

$$\mathbf{x}^T (I + L)^2 \mathbf{x} \leq \mathbf{x}^T (I + 2L + L^2) \mathbf{x} \leq n + (\|L\| + 2)\epsilon. \quad \square$$

This abstract result shows that if L and M are spectrally misaligned in the above sense, then the largest eigenvalue of the mixed-graph objective is large. Tangibly, one specific way that this can occur is the if L and M have very different cut structure. Plugging in characteristic vectors into Proposition 8.20 and taking the maximum yields

Corollary 8.21. *For any L, M ,*

$$\begin{aligned} \lambda_n((I + L)^{-1}M(I + L)^{-1}) &\geq \max_{S \subseteq V} \frac{4\text{cut}_{G_2}(S)}{n + 4(\|L\| + 2)\text{cut}_{G_1}(S)} \\ &\geq \max_{S \subseteq V} \frac{4\text{cut}_{G_2}(S)}{n + 8(\Delta_{G_1} + 1)\text{cut}_{G_1}(S)}, \end{aligned}$$

where Δ_{G_1} is defined to be the largest degree in G_1 .

Proof. The only new statement comes from noticing $\|L\| \leq 2\Delta_{G_1}$; this follows from the Gershgorin circle theorem, as the maximum absolute row sum of L is at most $2\Delta_{G_1}$. □

By directly analyzing the Rayleigh quotients, we can also obtain a slightly different bound for such vectors:

Proposition 8.22. *For any L, M ,*

$$\lambda_n((I + L)^{-1}M(I + L)^{-1}) \geq \max_{S \subseteq V} \frac{4\text{cut}_{G_2}(S)/n}{(1 + 2\sqrt{2}\text{cut}_{G_1}(S)/\sqrt{n})^2}$$

Proof. The proof proceeds analogously by plugging in χ_S into the Rayleigh quotient for λ_n . Indeed,

$$\lambda_n((I + L)^{-1}M(I + L)^{-1}) \geq \frac{\chi_S^T M \chi_S}{\|(I + L)\chi_S\|_2^2} = \frac{4\text{cut}_{G_2}(S)}{\|(I + L)\chi_S\|_2^2}.$$

It suffices to upper bound the denominator. By the Triangle Inequality,

$$\|(I + L)\chi_S\|_2 \leq \|\chi_S\|_2 + \|L\chi_S\|_2 = \sqrt{n} + \|L\chi_S\|_2.$$

Moreover, again by the Triangle Inequality

$$\begin{aligned} \|L\chi_S\|_2 &= \left\| \sum_{i \in S, j \notin S} 2w_{G_1}(i, j)(\mathbf{e}_i - \mathbf{e}_j) \right\|_2 \\ &\leq 2\sqrt{2}\text{cut}_{G_1}(S). \end{aligned}$$

Plugging in this estimate and factoring out n gives the claim. \square

These results show that if the opinion and disagreement graphs are misaligned on even one large cut of G_2 , then the adversary will be able to induce disagreement far beyond what is possible in the single-graph objective. As an example, consider an extreme case, where G_2 is a complete unweighted bipartite graph on $2n$ nodes, while G_1 is two n -node cliques on both sides of the bipartition with $o(\sqrt{n})$ edges between them. Then if S is one side of the bipartition, $\text{cut}_{G_2}(S) = n^2$, while $\text{cut}_{G_1}(S) = o(\sqrt{n})$. The estimate given by the Proposition 8.22 yields that the adversary can induce disagreement $\approx 2n$, which is tight even up

to constants in light of the upper bound in Lemma 8.14. This is sharper than the generic bound obtained in Corollary 8.21, which is off asymptotically by a factor of $o(\sqrt{n})$.

8.6 Chapter Notes

The results in this chapter originally appear in [67], joint work with Jon Kleinberg and Éva Tardos. Compared to the original version, the proofs are essentially identical while the exposition has been somewhat altered.

8.6.1 Related Work

Several works in the computer science and economics literatures have endeavored to theoretically understand the interplay between network structure and various emergent phenomena from the dynamics; while a complete survey is beyond the scope of this chapter, we highlight some relevant directions here.

In the computer science literature, our results are also similar in spirit to several recent works that consider interventions in opinion dynamics. Closest to the motivation for our work is the aforementioned independent and concurrent work of Chen and Rácz [37] who consider the power of an adversary seeking to induce discord within the Friedkin-Johnsen dynamics, but subject to a sparsity constraint. Their approach relies primarily on carefully bounding the quantities that arise in the associated Laplacian matrices, but without the use of spectral techniques. Gionis, et al [75] establish the NP-hardness of a natural opinion *maximization* problem in the Friedkin-Johnsen dynamics. They also derive a random

walk interpretation of these dynamics to establish submodularity of their problem, enabling a polynomial-time $(1 - 1/e)$ -approximation algorithm using the natural greedy algorithm. Musco, Musco, and Tsourakakis [107] considered the problem, given either a fixed graph topology or fixed initial opinions, of determining the *best* choice of the other to minimize their polarization-disagreement index metric; by contrast, our motivations lead us to a *worst-case*, adversarial analysis where we wish to understand the relation between the fixed graph topology and the adversary's power. Their main results shows that both of these problems can be efficiently solved via convex optimization. In a different direction, prior work by Bindel, Kleinberg, and Oren interprets the Friedkin-Johnsen dynamics as the Nash equilibrium of a natural quadratic cost function [24]; they then characterize the price of anarchy of this equilibrium, namely the ratio between the cost at equilibrium with that of the global optimizer. By using spectral techniques, they establish an upper bound of $9/8$ on this ratio, thereby showing that the equilibrium solution is a good minimizer of the global cost function.

Our techniques in Section 8.3 are similar to those of Galeotti, Golub, and Goyal [73], though with differing motivations. Their work considers optimal interventions in the context of *network games* to induce favorable or unfavorable Nash equilibria from the perspective of social welfare, whereas our motivation comes from opinion dynamics; where these works intersect is via the connection in [24] that the outcome of the Friedkin-Johnsen dynamics can be viewed as the equilibrium behavior of agents in a certain game with quadratic costs. Our results in Section 8.3 considers more general adversaries that attempt to optimize with respect to *different* objectives from the opinion dynamics literature, leading to significantly different conclusions on adversary behavior and graph struc-

ture. We remark that Galeotti, Golub, and Goyal interpret the KKT conditions for optimal interventions for more general initial opinions; however, the connections to graph structure substantially weaken, as the optimal interventions may be determined by the initial opinion rather than the network.

Earlier work by Demarzo, Vayanos, and Zweibel shows that in a variant of the DeGroot dynamics, while beliefs converge eventually, disagreement asymptotically occurs on a one-dimensional axis essentially determined by the second largest eigenpair [52]. Dandekar, Goel, and Lee show that *biased assimilation* in the DeGroot dynamics can provably lead to polarization [44]. Similarly, Golub and Jackson show how *homophily* can affect the rate of convergence of these dynamics, thus offering a behavioral explanation for the failure of consensus. Outside of these related models, recent work by Eliaz and Spiegler studies political polarization and disagreement as a byproduct of the equilibrium behavior of agents in adopting *narratives* that purport to explain long-run correlations in variables generated from some underlying causal model [56].

CHAPTER 9

GEOMETRIC OPINION MODELS: DYNAMICS OF POLARIZATION

In this chapter, we continue our study of opinion dynamics by focusing on *polarization* of opinions, a phenomenon of particular interest in recent years. Informally, we consider opinions to be polarized when agents nontrivially partition into groups holding diametric views. That is, rather than agents holding a rich spectrum of beliefs, individuals instead typically belong to opposite clusters even in cases where beliefs on separate topics ostensibly ought not be correlated. This phenomenon is somewhat widely observed to have accelerated in the last couple of decades; for instance, Gentzkow, Shapiro, and Taddy show that political partisanship is more easily inferrable in recent years than previously from textual analysis using machine learning methods [74]. While polarization and issue realignment is perhaps most familiarly observed along political dimensions (for instance, the correlation between beliefs on climate change and gun rights), these effects are not limited to just these facets [49].

While we analyzed polarization and discord in the previous chapter, an inherent obstruction to a richer understanding of these phenomena in models based on forms of repeated averaging is that such dynamics inherently contract. That is, given an arbitrary set of initial opinions, opinion updating in such models will provably lead to closer limiting opinions. To shed further light on polarization from a theoretical perspective, we will need to move to other kinds of opinion models.

Our analysis of polarization in this chapter builds heavily on a recently introduced model of opinion formation by Hazła, Jin, Mossel, and Ramnarayan (HJMR) [89]. In their model, each agent's opinion is a unit length vector on

a high-dimensional sphere where each dimension represents a particular relevant political axis. At each time, a new random direction representing a new issue, political figure, or influencer is drawn from some fixed distribution. In response, each agent evaluates the correlation between their current opinions and this new issue and moves either toward or away from this new direction depending on the strength and sign of this correlation. The resulting vector is then renormalized to get back to the unit sphere (see Section 9.2 for the exact model). These random directions model the natural intuition that opinion evolution is often driven by concrete, possibly random events. For instance, opinions potentially change dramatically in response to campaign messages or candidates during election seasons, when new political coalitions form and ideologies can themselves prove malleable to accommodate the new compositions of these groups.

One of the key findings of HJMR is that these simple dynamics *strongly polarize* in special cases. For instance, they show that when opinions lie on \mathbb{S}^{d-1} and the new issues are drawn uniformly at random from this same sphere, or when new issues are drawn from one of two “duelling influencers” that are sufficiently close with equal probability, then almost surely the distance between any two starting opinions will converge to 0 or 2 via these random dynamics. Because opinions lie on the unit sphere, this is equivalent to each agent lying on one of two antipodes as required of polarization.

In this chapter, we investigate how robust this phenomenon is in related models of geometric dynamics. For instance, while their analysis can be extended to arbitrary dimensions, the assumption of spherical symmetry of the random issue vectors is paramount in their proof, which relies on a martin-

gale argument. Each agent also updates their opinion independently of every other agent. If these kinds of analyses could be extended to incorporate network structure, as in the previous chapter, geometric dynamics could better connect with more classical models. Our task in this chapter is therefore to further determine whether, or to what extent, polarization can occur more generally in these kinds of geometric models.

To define geometric opinion dynamics more generally, we note the following important features of the HJMR model: (i) the dynamics are random to model the impact of new political issues that arise, and in fact form a Markov chain on the hypersphere, and (ii) the set of polarized vectors is invariant under the dynamics. This latter property is of course necessary to prove that polarization occurs, while the former provides a convenient mathematical framework to analyze these dynamics. This motivates the following generalized definition:

Definition 9.1. Let $n \in \mathbb{N}$ be an arbitrary, but otherwise fixed parameter denoting the number of agents, and let d denote the dimension of opinions. In a general model of **geometric opinion dynamics**, the evolution of each agent's opinion follows a discrete-time Markov chain in \mathbb{S}^{d-1} given by a recurrence of the following form:

$$\mathbf{X}_{t+1}^{(i)} \propto \mathbf{X}_t^{(i)} + f_i(\mathbf{X}_t^{(1)}, \dots, \mathbf{X}_t^{(n)}, \xi_t), \quad (9.1)$$

where $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \dots, \mathbf{X}_0^{(n)}) \in (\mathbb{S}^{d-1})^n$ is a given starting configuration of opinions, $\xi_t \in \mathbb{S}^{d-1}$ is drawn i.i.d. over time (and independent of all other random variables) from some distribution \mathcal{D} , and each $f_i : (\mathbb{S}^{d-1})^{n+1} \rightarrow \mathbb{R}^d$ is an explicit, fixed updating function.

The normalization in this formulation ensures that each opinion returns to the unit sphere. In particular, we assume that the unnormalized quantity above

is nonzero so that the projection onto the sphere is well-defined.

The joint dynamics of this system take place as a Markov chain in $\prod_{i=1}^n \mathbb{S}^{d-1}$. In such models, the random vector ξ_t again denotes a common, random stimulus that each agent in the system updates opinions with respect to, but now possibly also as a function of the opinions of the other agents in the system. For example, ξ_t may model a particular issue or figure during election season that splits society into two camps, “for” and “against,” and opinions update according to these coalitions. Crucially, we will also assume that the distribution \mathcal{D} that these vectors are drawn from remains fixed throughout time so that these dynamics form a time-homogeneous Markov chain, but we do not require uniformity.

With this definition, we extend the results of HJMR in this paper by pursuing a more systematic study of polarization in these general geometric opinion dynamics models that take the form of Equation (9.1). The motivating questions we consider are: *does strong polarization hold more generally, or is it specific to the HJMR model formulation? Can strong polarization be shown with nontrivial network interactions, thereby incorporating a key feature of models like the DeGroot and Friedkin-Johnsen dynamics? Are there other interesting notions of polarization that hold in these settings, even if the strong form does not hold?*

We show that whether polarization holds in such models is rather surprisingly nuanced and requires extending the notion of polarization beyond what is proven in HJMR. In particular, we show that there exist models that *weakly* polarize in a formal sense, but provably do not satisfy the stronger form in HJMR. We show that these weaker forms of polarization are connected to the existence of nontrivial invariant distributions on the induced Markov chain. Nonetheless,

we also prove that strong polarization holds for nontrivial variants of the HJMR model which are also robust to more general distributions of update vectors—in fact, we show that this holds in a model that has network interactions in addition to the random updates. We hope that some of our techniques will prove useful in studying polarization in such models more generally.

9.1 Overview of Results and Techniques

We contribute to the theoretical understanding of polarization in geometric opinion models initiated in [89]. In Section 9.2, we begin by identifying, though not necessarily requiring, several natural natural properties that such dynamics might satisfy and their relation to the problem of polarization. These will be convenient in our later arguments. We also formally define two new geometric opinion models, the **signed HJMR** model and the **party** model that will be the subject of our later analysis.

Robustness of Strong Polarization: Our main results in Section 9.3 show that *strong polarization holds more universally than the HJMR model*. In particular, we show that the strong polarization holds in both of the models we introduce. The *signed HJMR model* is a variant of the HJMR dynamics that replaces the inner product with a sign. With these altered dynamics, we show that strong polarization holds in any dimension and is robust to the choice of distribution on update vectors so long as it is not too far from uniform. We also prove that the same holds in the *party model*, which notably incorporates *network effects* where agents exert influence over each other.

Our analysis of these models substantially differs from that of HJMR, which

primarily relies on either martingale convergence (in the case of the uniform distribution) or deterministic contraction (in two-point models representing “duelling influences”, where the challenge of proving contraction is geometric rather than probabilistic). Our results will rely on general zero-one laws. Our approach takes two modular steps: first, we substantially simplify the problem of proving strong polarization by showing it is enough to simply prove *some* uniform lower bound on the probability of asymptotically polarizing. We then obtain our main results by using clean probabilistic and geometric estimates to give the required bounds, thereby proving strong polarization.

Modes of Polarization: We then consider more general notions of polarization beyond the version considered by HJMR. To do so, we elucidate the connection between various forms of polarization with the general theory of Markov chains in uncountable state spaces in Section 9.4. Such Markov chains can have much subtler behavior than their discrete counterparts. Our new notions of polarization directly arise from standard notions of convergence from probability theory. We first show that in most cases, whether these weaker forms of polarization hold in geometric opinion dynamics is in fact equivalent to the existence of nontrivial invariant distributions of the Markov chain defined by the dynamics. To demonstrate the utility of these new definitions, we then analyze a concrete example: we show that the original HJMR model with update vectors drawn uniformly from an orthonormal basis satisfies our notion of weak polarization, but provably *does not* satisfy strong polarization. In fact, we show that almost surely, almost every starting configuration will not strongly polarize. Our analysis connects the dynamics to pathwise properties of infinite balls-in-bins processes, which we study by applying yet another zero-one law, that of Hewitt-Savage.

9.2 Preliminaries

As stated above, we are primarily interested in the *polarization* properties of these random processes given by Equation (9.1). Throughout this paper, we reserve n to denote the number of agents and d to denote the dimensionality of opinions. We write $\|\cdot\|_p$ for the standard ℓ_p -norm and $\mathbb{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 = 1\}$. We will write $P_{\mathbb{S}^{d-1}} : \mathbb{R}^d \setminus \{\mathbf{0}\} \rightarrow \mathbb{S}^{d-1}$ for the projection onto the unit sphere, i.e. $P_{\mathbb{S}^{d-1}}(\mathbf{x}) = \mathbf{x}/\|\mathbf{x}\|_2$. We also write \mathbf{e}_i for the i th standard basis vector in \mathbb{R}^d , where d will be clear from context. For a set of vectors $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$, we define $\text{cone}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \{\mathbf{z} \in \mathbb{R}^d : \mathbf{z} = \sum_{i=1}^n \alpha_i \mathbf{x}_i, \alpha_i \geq 0\}$. We write $\angle(\mathbf{x}, \mathbf{y})$ for the angle between \mathbf{x} and \mathbf{y} .

For given n, d , we then define $D \subseteq \prod_{i=1}^n \mathbb{S}^{d-1}$ to be the set of diagonal opinion vectors, i.e. the set of elements of the form $(\mathbf{x}, \dots, \mathbf{x})$ for a single vector $\mathbf{x} \in \mathbb{S}^{d-1}$. For a sign vector $\sigma \in \{-1, 1\}^n$, define $\sigma(D) := \{\mathbf{X} \in \prod_{i=1}^n \mathbb{S}^{d-1} : \mathbf{X} = (\sigma_1 \mathbf{x}, \dots, \sigma_n \mathbf{x}), \mathbf{x} \in \mathbb{S}^{d-1}\}$. Finally, we define the set P of polarized vectors by $P = \bigcup_{\sigma \in \{-1, 1\}^n} \sigma(D)$. In other words, P is the set of tuples of vectors such that each vector is equal to the rest up to sign. This definition allows for consensus as a special case, but in many settings, consensus or near-consensus is exponentially unlikely (see Remark 9.5).

For a subset A in some Euclidean space and a point \mathbf{x} , we define $\rho(\mathbf{x}, A) := \inf_{\mathbf{y} \in A} \|\mathbf{x} - \mathbf{y}\|_2$ to denote the distance from a point to a set with respect to the Euclidean metric. With this in mind, we define strong polarization of a geometric opinion dynamics model as follows:

Definition 9.2. Let $\mathbf{X}_t := (\mathbf{X}_t^{(1)}, \dots, \mathbf{X}_t^{(n)})$ be a discrete-time Markov chain as given by eq. (9.1). Then \mathbf{X}_t **strongly polarizes** (from \mathbf{X}_0) if almost surely, $\rho(\mathbf{X}_t, P) \rightarrow 0$;

that is, the distance between \mathbf{X}_t and the set of polarized vectors converges to zero almost surely.

There are several natural properties that one might desire in these dynamics. Below, we consider, though do not require, the following properties which abstracts those of the original HJMR model:

Definition 9.3. A model of geometric opinion dynamics is **continuous** if for each fixed ξ , the functions $f_i(\cdot, \xi)$ are continuous from $\prod_{i=1}^n \mathbb{S}^{d-1}$ to \mathbb{R}^d for all $i \in [n]$.

We will show in Section 9.4 that continuity implies various desirable properties for the dynamics. Note that if this is the case, combined with the fact that we assume the unnormalized update rule is always nonzero, it follows that the dynamics are jointly continuous as transitions from $\prod_{i=1}^n \mathbb{S}^{d-1}$ to itself. This follows because the map taking the joint opinion vector to the unnormalized opinion vectors is continuous and nonzero in every coordinate, which is then composed coordinatewise with a continuous map on $\mathbb{R}^d \setminus \{\mathbf{0}\}$, and so is continuous.

Definition 9.4. The dynamics are **sign-invariant** if, for all $i \in [n]$, the function f_i is *odd* with respect to \mathbf{x}_i , but *even* with respect to the other arguments (i.e. with respect to \mathbf{x}_{-i} and ξ).

Sign-invariance implies each agent reacts to the random update vector and the others the same regardless if any are negated. This feature will be present in all of the models we consider below, including the original HJMR model. The intuition behind sign-invariance is that from the perspective of each agent, if he or she were to react “positively” to the new issue or a different opinion,

she would react “negatively” to the negative of that issue or different opinion—sign-invariance thus asserts that these reactions are balanced.

Definition 9.5. If $f_i(\mathbf{X}_t^{(i)}, \mathbf{X}_t^{(-i)}, \xi) = f_j(\mathbf{X}_t^{(j)}, \mathbf{X}_t^{(-j)}, \xi)$ when $\mathbf{X}_t^{(i)} = \mathbf{X}_t^{(j)}$ and $\mathbf{X}_t^{(-i)} = \mathbf{X}_t^{(-j)}$, we say that the dynamics are **symmetric**, as the updates do not depend on the identities of the agents.

Definition 9.6. If each function f_i does not depend on $\mathbf{X}_t^{(-i)}$ for each i (so that f_i depends only on $\mathbf{X}_t^{(i)}$ and ξ_t), then we say the dynamics are **oblivious**.

In this case, each component of the above process follows a Markov chain, and the joint dynamics form a particular coupling where each component responds to the same update vector. However, any polarization that arises happens indirectly because agents do not influence each other.

In this work, we treat such models in relatively full generality and also specialize to particular models where more specific techniques can establish various desirable properties. The concrete examples we will consider are listed below:

Definition 9.7 (HJMR Model). In the **HJMR Model** [89], the update for each agent i takes the following form for some fixed scalar $\eta > 0$:

$$\mathbf{X}_{t+1}^{(i)} \propto \mathbf{X}_t^{(i)} + \eta \cdot \langle \mathbf{X}_t^{(i)}, \xi_t \rangle \xi_t, \quad (9.2)$$

where $\xi_t \sim \mathcal{D}$ is drawn i.i.d. over time from some distribution \mathcal{D} on \mathbb{S}^{d-1} . In words, each agent moves in the (signed) direction of the random update vector proportionally to the correlation with their current opinion, and then renormalizes.

Note that this model satisfies continuity, sign-invariance, obliviousness, and symmetry (assuming η is a constant over all agents).

Definition 9.8 (Signed HJMR Model). In the **signed HJMR model**, the update rule in Equation (9.2) is amended to

$$\mathbf{X}_{t+1}^{(i)} \propto \mathbf{X}_t^{(i)} + \eta \cdot \text{sgn}(\langle \mathbf{X}_t^{(i)}, \xi_t \rangle) \xi_t. \quad (9.3)$$

Here, we define $\text{sgn}(0) = 0$, but in our applications below, we will assume ξ_t is drawn from a continuous distribution so that almost surely $\langle \mathbf{X}_t^{(i)}, \xi_t \rangle \neq 0$. In this case, the amount the vector updates does not depend on the correlation. This choice is intended to model elections, where one is in favor either towards or against a particular candidate and is drawn “all-or-nothing” towards or against the views of this candidate. This model is *not* continuous due to the sign function, but is still sign-invariant, oblivious, and symmetric (assuming η is a constant over all agents).

Definition 9.9 (Party Model). Suppose each agent i has multipliers $(\eta_1^{(i)}, \dots, \eta_n^{(i)})$ where $\eta_j^{(i)} \geq 0$ measures the influence of agent j on agent i . The **party model** is defined by:

$$\mathbf{X}_{t+1}^{(i)} \propto \mathbf{X}_t^{(i)} + \left(\sum_{j \in [n]: \text{sgn}(\langle \mathbf{X}_t^{(j)}, \xi_t \rangle) = \text{sgn}(\langle \mathbf{X}_t^{(i)}, \xi_t \rangle)} \eta_j^{(i)} \mathbf{X}_t^{(j)} - \sum_{j \in [n]: \text{sgn}(\langle \mathbf{X}_t^{(j)}, \xi_t \rangle) \neq \text{sgn}(\langle \mathbf{X}_t^{(i)}, \xi_t \rangle)} \eta_j^{(i)} \mathbf{X}_t^{(j)} \right). \quad (9.4)$$

That is, each agent moves towards the vectors that came on the same “side” of the random issue ξ_t , and away from those on the opposite “side” of the issue. While this latter assumption may appear non-obvious, there is considerable empirical evidence for the sociological principle that “out-group conflict builds in-group solidarity” [106]: when a binary issue creates disagreement within a collection of people, the people on each side of the disagreement move toward those they agree with and away from those they disagree with [60, 122]. Once again, this model is not continuous due to the sign function, and also

is *not* oblivious as clearly the update rule depends on the values of the other agents. However, it remains sign-invariant, as negating one's own opinions interchanges the sums, and flipping either ξ_t or any other agents opinions only permutes summands.

For any sort of polarization to arise, it is natural to ensure that the dynamics are such that *if* the vector are completely polarized, then they will surely remain so. We provide one simple condition that is easily verified¹:

Lemma 9.1. *Suppose that some geometric opinion dynamics satisfy symmetry and sign-invariance. Then $\sigma(D)$ is invariant under the transitions for every $\sigma \in \{-1, 1\}^n$.*

Finally, note that if the dynamics are oblivious and strongly polarizes for $n = 2$ agents, then it must do so for any finite n by a simple union bound. Note that oblivious and symmetric dynamics ensure that the process is well-defined for any number of agents.

Lemma 9.2. *Suppose that the opinion dynamics are symmetric and oblivious. Then if any form of convergence holds for $n = 2$ agents, the same holds for any finite number of agents.*

9.3 Models with Strong Polarization

In this section, we prove the strong polarization of the signed HJMR model and the party model. To do this, we first establish a simple, but powerful general principle that will significantly simplify the analysis that has the following simple intuition: suppose momentarily that geometric opinion dynamics

¹Note that the converse of Lemma 9.1 trivially fails: one can simply ensure the dynamics are invariant on each such set and otherwise define them arbitrarily.

were a *finite-state* Markov chain and that “polarization” is an absorbing state of the Markov chain. Then from standard and simple Markov estimates, so long as it is possible to reach the state “polarized” from any starting point in some fixed finite number of steps, an easy calculation shows that almost surely the Markov chain will become “polarized.” In that case, it would suffice to show that from any starting state, there is *some* nonzero probability of reaching “polarized” from any starting state in some finite number of steps.

In general, this idea is not so simple to formalize because the geometric opinion dynamics lie in a non-discrete state space, and moreover, it is often only possible to reach P asymptotically, not in any finite number of steps. However, by appealing to more general zero-one laws, we show that this intuition nonetheless holds:

Theorem 9.3. *For any geometric model of opinion dynamics that satisfies the Markov property, the following are equivalent:*

1. *For every choice of starting vector \mathbf{X}_0 , $\Pr_{\mathbf{X}_0}(\rho(\mathbf{X}_t, P) \rightarrow 0) = 1$.*
2. *For every choice of starting vector \mathbf{X}_0 , $\Pr_{\mathbf{X}_0}(\rho(\mathbf{X}_t, P) \rightarrow 0) \geq c$ for some constant $c > 0$.*

Proof. One direction is trivial, so we assume the second condition. Consider the dynamics started at any choice of starting vector \mathbf{X}_0 , and consider the event $A = \{\rho(\mathbf{X}_t, P) \rightarrow 0\}$. Let $\mathcal{F}_t = \sigma(\xi_0, \dots, \xi_{t-1})$ be the filtration generated by the random updates up to time t and let $\mathcal{F}_\infty = \sigma(\xi_0, \dots)$ be the filtration generated by all of them. By standard arguments, $A \in \mathcal{F}_\infty$.

For each $T \geq 0$, define $A_T := \mathbb{E}[\mathbf{1}(A)|\mathcal{F}_T]$. As $\mathbf{1}(A)$ is an indicator random variable, Lévy’s upward theorem (Theorem 4.2.11 of [55]) implies that $A_T =$

$\mathbb{E}[\mathbf{1}(A)|\mathcal{F}_T] \rightarrow \mathbb{E}[\mathbf{1}(A)|\mathcal{F}_\infty] = \mathbf{1}(A)$ almost surely. But observe that by the Markov property, $A_T = \Pr_{\mathbf{X}_T}(\rho(\mathbf{Z}_t, P) \rightarrow 0)$, where \mathbf{Z}_t gives the dynamics started at $\mathbf{Z}_0 = \mathbf{X}_T$. In particular, $A_T \geq c > 0$ pointwise. Because $A_T \rightarrow \mathbf{1}(A)$ almost surely and is bounded from below almost surely by a strictly positive quantity, the only way this can happen is if $\mathbf{1}(A) \equiv 1$ almost surely (over the realizations of the ξ_t), so that A holds almost surely. As \mathbf{X}_0 was arbitrary, this completes the harder direction. \square

9.3.1 Strong Polarization in Signed HJMR Model

Our first main result is that strong polarization holds in the signed HJMR model with a common value of $\eta > 0$, and that this holds for a general class of distributions:

Theorem 9.4. *Suppose there are $n \geq 1$ agents in the signed HJMR model given by Equation (9.3) where each ξ_t is drawn i.i.d. from a distribution \mathcal{D} that is equivalent to the uniform (Haar) measure μ , i.e. there exists $C, C' > 0$ such that for every measurable set $A \subseteq \mathbb{S}^{d-1}$, $C\mu(A) \leq \Pr(\xi_t \in A) \leq C'\mu(A)$. Then this system strongly polarizes from any choice of starting vector \mathbf{X}_0 .*

Remark 9.5. *It can be shown that for any sign-invariant dynamics that strongly polarizes, if \mathbf{X}_0 is drawn uniformly from $(\mathbb{S}^{d-1})^n$, then each possible clustering is equally likely even conditioned on the sequence $\{\xi_t\}_{t=0}^\infty$ almost surely. In particular, the probability over starting configurations and the random updates of consensus is exponentially small in n .*

To set up this result, we establish a sequence of lemmas that will prove useful. We begin with a simple geometric fact:

Lemma 9.6. *Let $\mathbf{x}, \mathbf{y} \in \mathbb{S}^{d-1}$ and suppose $\mathbf{z} \in \mathbb{R}^d$ is such that $\|\mathbf{x} + \mathbf{z}\|_2, \|\mathbf{y} + \mathbf{z}\|_2 \geq 1 + \epsilon$ for some $\epsilon \geq 0$. Then $\|P_{\mathbb{S}^{d-1}}(\mathbf{x} + \mathbf{z}) - P_{\mathbb{S}^{d-1}}(\mathbf{y} + \mathbf{z})\|_2 \leq \frac{\|\mathbf{x} - \mathbf{y}\|_2}{1 + \epsilon}$.*

Proof. We may assume $\mathbf{x} \neq \mathbf{y}$, as otherwise the claim is trivial. Consider the arrangement of vectors $\mathbf{x} + \mathbf{z}, \mathbf{y} + \mathbf{z}, \mathbf{x} - \mathbf{y}$ in the plane spanned by $\{\mathbf{x} + \mathbf{z}, \mathbf{y} + \mathbf{z}\}$ forming a triangle with a vertex at the origin and adjacent sides $\mathbf{x} + \mathbf{z}, \mathbf{y} + \mathbf{z}$. By the assumption that these vectors have length at least $1 + \epsilon$, scaling this triangle by a factor of $r = (1 + \epsilon)^{-1}$ ensures that $r(\mathbf{x} + \mathbf{z})$ and $r(\mathbf{y} + \mathbf{z})$ continue to have at least unit norm, and the distance between them is at most $r(\|\mathbf{x} - \mathbf{y}\|_2)$. As projection onto \mathbb{S}^{d-1} is a contraction in Euclidean distance for vectors of length at least 1, it follows that

$$\|P_{\mathbb{S}^{d-1}}(\mathbf{x} + \mathbf{z}) - P_{\mathbb{S}^{d-1}}(\mathbf{y} + \mathbf{z})\|_2 = \|P_{\mathbb{S}^{d-1}}(r(\mathbf{x} + \mathbf{z})) - P_{\mathbb{S}^{d-1}}(r(\mathbf{y} + \mathbf{z}))\|_2 \leq r\|\mathbf{x} - \mathbf{y}\|_2. \quad \square$$

Next, we show that with some constant probability, a random vector drawn from \mathcal{D} will not split two vectors that form an acute angle, and that the probability of splitting at all tends to zero as the distance tends to zero.

Lemma 9.7. *There exists constants $\beta, \gamma > 0$ depending on d and the measure of equivalence such that the following holds: suppose that $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)})$ satisfies $\langle \mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)} \rangle \geq 0$. Then with probability at least β , it holds that ξ_0 satisfies:*

1. $\text{sgn}(\langle \mathbf{X}_0^{(1)}, \xi_0 \rangle) = \text{sgn}(\langle \mathbf{X}_0^{(2)}, \xi_0 \rangle)$, and
2. $|\langle \mathbf{X}_0^{(i)}, \xi_0 \rangle| \geq \gamma$ for $i = 1, 2$.

Proof. We first show this when ξ_0 is drawn uniformly from the sphere and then simply change constants when moving to any measure that is equivalent to Haar measure. But this is clear: the probability of ξ_0 satisfying the first property is at least $1/2$ under these assumptions, as the direction of ξ_0 in the plane

spanned by $\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}$ is itself uniform and using the acuteness of the two vectors. Moreover, the distribution of $|\langle \xi_0, \mathbf{z} \rangle|$ does not depend on \mathbf{z} by uniformity, so we may choose $\gamma > 0$ small enough so that the probability of the second property is at least $7/8$ for any fixed \mathbf{z} . By a union bound, it follows that the probability ξ_0 has the desired properties is at least $1/4$ under Haar measure. Under the true, equivalent distribution, the probability is thus at least $\beta := C/4 > 0$. \square

Lemma 9.8. *Under the assumptions and notation of Lemma 9.7, the probability that $\text{sgn}(\langle \mathbf{X}_0^{(1)}, \xi_0 \rangle) \neq \text{sgn}(\langle \mathbf{X}_0^{(2)}, \xi_0 \rangle)$ is at most $O(\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2)$ and $O(\angle(\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}))$, where the implicit constant depends only on d and the measure of equivalence.*

Proof. We first show this for Haar measure. By an analogous argument, the set of vectors with the desired property have directions lie in a band of width $O(\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2)$ in the plane spanned by $\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}$, and therefore has probability at most $O(\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2)$ and $O(\angle(\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}))$ by uniformity. For the true distribution, this can increase by a factor of at most C' , which we may absorb into the implicit constant. \square

With this result in hand, we turn to the proof of the main theorem:

Proof of Theorem 9.4. We begin with a series of reductions that simplifies the problem. First, because these dynamics are oblivious, we observe that by Lemma 9.2 it suffices to consider the case $n = 2$ with an arbitrary starting vector $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)})$. Next, by sign-invariance of these dynamics, we may assume that $\langle \mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)} \rangle \geq 0$ by possibly flipping the sign of one of the vectors and noting that both the dynamics and the set P are invariant under these sign changes. Note that this implies that the two starting vectors form an acute angle. Finally, by Theorem 9.3, it suffices to show that the probability that $\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2 \rightarrow 0$ is

bounded below by some nonzero constant $c > 0$, uniformly over the choice of starting vector (though assuming nonnegative inner product).

We use the notation of Lemma 9.7. Note that the good event of Lemma 9.7 and the bad event of Lemma 9.8 are disjoint, though not mutually exhaustive. We claim that by the craps principle, the probability of encountering a random update ξ_t satisfying the good event before the bad event is at least $\beta/(\beta + O(\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2))$. Indeed, while an update need not satisfy either event, on the complement of the bad event, $\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2$ is nonincreasing so long as the bad event does not occur as the unnormalized lengths increase by sign-invariance with respect to ξ_t (so that we may assume both signs are positive) and contractions decrease distances. The claim then follows from the craps's principle, Lemma 9.7, and Lemma 9.8.

Moreover, on this event, the distance between $\mathbf{X}_{t+1}^{(1)}$ and $\mathbf{X}_{t+1}^{(2)}$ decreases by a factor of $(1 + \epsilon)^{-1}$ where $\epsilon = \Omega(\eta\gamma)$ by Lemma 9.6 (using $\mathbf{z} = \pm\eta \cdot \xi_t$) and the fact that the inner products are bounded below on the good event. It follows that on this event, the new distance between vectors is at most $\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2/(1 + \epsilon)$. By the strong Markov property, we may iterate this argument to show that the probability of the good event occurring before the bad event is now at least $\beta/(\beta + O(\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2/(1 + \epsilon))) \geq 1 - O\left(\frac{\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2}{1 + \epsilon}\right)$ where we absorb the constant β . It follows that the probability that the good event occurs infinitely often without the bad event is at least

$$\prod_{k=0}^{\infty} \left(1 - O\left(\frac{\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2}{(1 + \epsilon)^k}\right)\right) \geq \prod_{k=0}^{\infty} \left(1 - O\left(\frac{1}{(1 + \epsilon)^k}\right)\right),$$

where we simply upper bound $\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2 \leq 2$ in the inequality. Note that if this occurs, then Lemma 9.6 show that this implies that $\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2 \rightarrow 0$ as the distance is nonincreasing and geometrically decreases infinitely often. From

standard analysis,

$$\prod_{k=0}^{\infty} \left(1 - \mathcal{O}\left(\frac{1}{(1+\epsilon)^k}\right)\right) > 0 \iff \sum_{k=0}^{\infty} \mathcal{O}\left(\frac{1}{(1+\epsilon)^k}\right) < \infty,$$

and the latter is clearly true as a geometric series. Moreover, these lower bounds are uniform over the value of $\|\mathbf{X}_0^{(1)} - \mathbf{X}_0^{(2)}\|_2$. By the reductions above, this completes the proof. \square

9.3.2 Strong Polarization in the Party Model

We now turn to proving strong polarization in the party model. One complicating factor is that these dynamics are *not* oblivious, unlike the other models where strong polarization is known. Therefore, we have to reason about multiple vectors acting on each other at the same time.

To set up the formal statement of the result, we need the following definition: for any given set of nonnegative coefficients $\boldsymbol{\eta}$, let $A = A(\boldsymbol{\eta})$ be the (directed) $n \times n$ adjacency matrix defined by

$$A_{ij} = \begin{cases} 1 & \eta_j^{(i)} > 0 \\ 0 & \text{else.} \end{cases}$$

In other words, we consider the directed graph with a directed edge from i to j if j influences i . We say that A is irreducible if $A^n > 0$ componentwise. This is equivalent to the existence of a directed path between any two agents i and j . Our main result of this section is as follows:

Theorem 9.9. *Suppose there are $n \geq 1$ agents in the party model where each ξ_t is drawn i.i.d. from a distribution \mathcal{D} that is equivalent to the uniform (Haar) measure μ , i.e. there exists $C, C' > 0$ such that for every measurable set $B \subseteq \mathbb{S}^{d-1}$, $C\mu(B) \leq$*

$\Pr(\xi_t \in B) \leq C' \mu(B)$. Moreover, suppose that $A = A(\boldsymbol{\eta})$ is irreducible. Then this system strongly polarizes from any choice of starting vector \mathbf{X}_0 .

To prove this, we follow a similar high-level plan as that of Theorem 9.4. By sign-invariance, we will assume via Lemma 9.10 that each component of \mathbf{X}_0 lies on one side of a hyperplane with margin strictly bounded below by zero regardless of the individual configuration. We then construct a potential function that is equivalent to the maximum angle between any two agents with the property that, assuming the dynamics do not split up the vectors on the next n iterations, it is guaranteed to decay by some factor strictly bounded above by 1. Since, as we will again see, the probability that a random vector splits up any two components is essentially bounded by the maximum angle between any two components and this quantity is decreasing geometrically, it will follow that the probability of strong polarization is bounded below uniformly. We may then again conclude via Theorem 9.3 that the system strongly polarizes from any starting configuration.

We now carry out this plan. First, we show that there always exists a signing of the starting configuration such that all vectors lie on one side of some hyperplane with nontrivial margin:

Lemma 9.10. *For all $d, n \geq 1$, there exists a constant $\lambda = \lambda(n, d) > 0$ such that for any $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{S}^{d-1}$, there exists $\mathbf{z} \in \mathbb{S}^{d-1}$ such that $|\langle \mathbf{z}, \mathbf{x}_i \rangle| \geq \lambda$ for all $i \in [n]$.*

Proof. Simply choose $\lambda > 0$ such that the probability of a random unit vector drawn from Haar measure does not satisfy the condition for a given i is at most $1/(n+1)$ and then apply a union bound to conclude there exists such a vector. Note that this choice of λ indeed depends on d, n , but not on the choice of vectors

as the distribution of the inner product of a uniformly random vector on the sphere with a fixed vector does not depend on the identity of this fixed vector.

□

Next, we proceed with several purely geometric results that will enable us to show contraction of a suitable potential function.

Lemma 9.11. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{S}^{d-1}$ all lie strictly on one side of a hyperplane in \mathbb{R}^d .*

Define \mathbf{y} by

$$\mathbf{y} \in \arg \min_{\mathbf{v} \in \mathbb{S}^{d-1}} \max_{j \in [n]} \angle(\mathbf{v}, \mathbf{x}_j).$$

Note \mathbf{y} exists as the objective function is continuous and the constraint set is compact.

Let $H = \mathbf{y}^\perp$ be the orthogonal subspace to \mathbf{y} and let P_H be the orthogonal projection onto H . Then $\mathbf{0} \in \text{conv}(P_H(\mathbf{x}_1), \dots, P_H(\mathbf{x}_n))$.

Proof. Observe that the desired claim is implied by $\mathbf{y} \in \text{cone}(\mathbf{x}_1, \dots, \mathbf{x}_n)$. To see this, suppose that this holds: suppose there exists $\alpha_i \geq 0$ such that $\mathbf{y} = \sum_{i=1}^n \alpha_i \mathbf{x}_i$. Note that not all α_i are zero as \mathbf{y} has unit norm by construction. Then applying P_H to both sides, we deduce that

$$\mathbf{0} = P_H(\mathbf{y}) = \sum_{i=1}^n \alpha_i P_H(\mathbf{x}_i).$$

As $\sum_{i=1}^n \alpha_i > 0$, we may renormalize these coefficients so that their sum is one to deduce that $\mathbf{0} \in \text{conv}(P_H(\mathbf{x}_1), \dots, P_H(\mathbf{x}_n))$. Moreover, note that \mathbf{y} can equivalently be defined via $\mathbf{y} \in \arg \max_{\mathbf{v} \in \mathbb{S}^{d-1}} \min_{j \in [n]} \langle \mathbf{y}, \mathbf{x}_j \rangle$. This holds as the inner product is a monotonically decreasing function of angle on $[0, \pi]$ and so the optimization problems are equivalent.

Therefore, suppose that $\mathbf{y} \notin \text{cone}(\mathbf{x}_1, \dots, \mathbf{x}_n)$. By the separating hyperplane theorem, there exists a unit vector \mathbf{u} such that $\langle \mathbf{u}, \mathbf{x}_i \rangle > 0$ for all $i \in [n]$ but $\langle \mathbf{u}, \mathbf{y} \rangle <$

0. Write \mathbb{R}^d in an orthogonal basis that includes \mathbf{u} . Then the coordinate with respect to \mathbf{u} for each of the \mathbf{x}_i is strictly positive, while the coordinate for \mathbf{y} is strictly negative. Therefore, by reflecting \mathbf{y} about \mathbf{u} , we obtain unit $\tilde{\mathbf{y}}$ such that $\langle \tilde{\mathbf{y}}, \mathbf{x}_i \rangle > \langle \mathbf{y}, \mathbf{x}_i \rangle$ for all $i \in [n]$, contradicting the optimality of \mathbf{y} . \square

Lemma 9.12. *Let $\mathbf{z}_1, \dots, \mathbf{z}_n$ satisfy $\|\mathbf{z}_i\| \leq 1$ for all $i \in [n]$ and $\mathbf{0} \in \text{conv}(\mathbf{z}_1, \dots, \mathbf{z}_n)$.*

Then

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \right\|_2 \leq 1 - \frac{1}{n}.$$

Proof. By assumption, $\mathbf{0} = \sum_{i=1}^n a_i \mathbf{z}_i$ for some coefficients a_1, \dots, a_n satisfying $a_i \geq 0$ for all $1 \leq i \leq n$ and $\sum_{i=1}^n a_i = 1$. We may scale this equation so that $\max_i a_i = 1/n$.

Then

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \right\|_2 = \left\| \sum_{i=1}^n \left(\frac{1}{n} - a_i \right) \mathbf{z}_i \right\|_2 \leq \sum_{i=1}^n \left(\frac{1}{n} - a_i \right) \leq 1 - 1/n,$$

where we use the triangle inequality, the assumption that $\|\mathbf{z}_i\|_2 \leq 1$, and that $0 \leq a_i \leq 1/n$ for all i with equality for some index. \square

Next, we show that the irreducibility of A implies that there is geometric decay in the minimal angle in each n steps that the dynamics do not split the vectors. We start with the following crude, but intuitive lemma:

Lemma 9.13. *Let $\mathbf{x}_0 = (\mathbf{x}_0^{(1)}, \dots, \mathbf{x}_0^{(n)})$ be any set of vectors all lying strictly on one side of a hyperplane with $\lambda = \lambda(n, d) > 0$ margin (i.e. there exists unit \mathbf{v} satisfying $\langle \mathbf{v}, \mathbf{x}_0^{(i)} \rangle > \lambda$ for all $i \in [n]$), and consider the update rule given by Equation (9.4) where the second summand is empty.*

Suppose that we apply this update rule n times to obtain \mathbf{x}_n . Then, there exists $\epsilon = \epsilon(n, d, \eta) > 0$ such that $\mathbf{x}_n \propto M \mathbf{x}_0$ for some row-stochastic matrix M depending on \mathbf{x}_0 satisfying

$$M = \epsilon \Pi + (1 - \epsilon) Q, \tag{9.5}$$

where every entry of Π is equal to $1/n$ and Q is some arbitrary stochastic matrix.

Remark 9.14. In the above matrix equation, we interpret \mathbf{x}_t as a $n \times d$ matrix where the i th row is $\mathbf{x}_t^{(i)}$. Moreover, the point of the lemma is that ϵ depends on $n, d, \boldsymbol{\eta}$, but not on the starting configuration.

Proof. We show the following claim by induction: for each $1 \leq k \leq n$, there exists $\epsilon_k = \epsilon(n, d, \boldsymbol{\eta}, k)$ such that $\mathbf{x}_k \propto M_k \mathbf{x}_0$ where M_k is a row-stochastic matrix such that $(M_k)_{ij} > \epsilon_k$ if $j \in \Gamma^k(i)$, where $\Gamma^k(i)$ is the set of nodes reachable from i in k steps in the directed graph induced by $A = A(\boldsymbol{\eta})$ above. This clearly implies the claim by setting $\epsilon = \epsilon_n$ noting that $\Gamma^k(i) \subseteq \Gamma^{k+1}(i)$ by the fact that this matrix has ones on the diagonal.

For the base case $k = 1$, by definition (and absorbing the term $\mathbf{x}_0^{(i)}$ into the $\eta_i^{(i)}$ multiplier): $\mathbf{x}_1^{(i)} \propto \sum_{j=1}^n \eta_j^{(i)} \mathbf{x}_0^{(j)}$. By dividing by $\eta^{(i)} \triangleq \sum_{j=1}^n \eta_j^{(i)}$, we clearly obtain the claim with $\epsilon_1 \triangleq \min_{i \in [n]} \min_{j \in [n]: \eta_j^{(i)} > 0} \eta_j^{(i)} / \eta^{(i)}$. Note that this is independent of \mathbf{x}_0 .

Now suppose it holds for some $k \geq 1$ so that $\mathbf{x}_k \propto M_k \mathbf{x}_0$. By applying the base case to \mathbf{x}_{k+1} (noting that \mathbf{x}_k still lies on the same side of the hyperplane by convexity with same margin),

$$\mathbf{x}_{k+1} \propto M'_1 \mathbf{x}_k \propto (M'_1 M_k) \mathbf{x}_0.$$

By the induction hypothesis, $(M'_1)_{ij} \geq \epsilon(n, d, \boldsymbol{\eta}, 1)$ for all $j \in \Gamma(i)$, while $(M_k)_{ij} \geq \epsilon(n, d, \boldsymbol{\eta}, k)$ for all $j \in \Gamma^k(i)$. For any $j \in \Gamma^{k+1}(i)$, there exists some $j' \in [n]$ such that $j' \in \Gamma^k(i)$ and $j \in \Gamma(j')$. From the definition of matrix multiplication, if $j \in \Gamma^{k+1}(i)$, it follows that $(M'_1 M_k)_{i,j} \geq \epsilon_1 \cdot \epsilon_k$. Letting $\epsilon_{k+1} \triangleq \epsilon_1 \cdot \epsilon_k$ and noting the product of stochastic matrices is stochastic, the claim follows. \square

Finally, we show that one can get the geometric rate of convergence with

respect to a natural potential function.

Lemma 9.15. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{S}^{d-1}$ all lie strictly on one side of a hyperplane in \mathbb{R}^d with margin at least $\lambda = \lambda(n, d) > 0$. Define $\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)$ by*

$$\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n) = \min_{\mathbf{v} \in \mathbb{S}^{d-1}} \max_{i \in [n]} \angle(\mathbf{v}, \mathbf{x}_i). \quad (9.6)$$

Let $\mathbf{x}'_1, \dots, \mathbf{x}'_n$ be the updated vectors after n iterations of the update rule in Equation (9.4) when the second summand is empty. Then there exists $c = c(n, d, \eta) > 0$ such that

$$\Phi(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \leq (1 - c)\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n).$$

Proof. Let \mathbf{y} be an optimizer in Equation (9.6) and let H be the orthogonal complement. Let P_H be the corresponding orthogonal projection onto H and P_y be the orthogonal projection onto \mathbf{y} . Note that for any vector \mathbf{z} , $\mathbf{z} = P_H\mathbf{z} + P_y\mathbf{z}$. By Lemma 9.13, we have $\mathbf{x}' \propto M\mathbf{x}$ for some stochastic matrix M such that $M = \epsilon\Pi + (1 - \epsilon)Q$.

Now, observe that by the definition of Φ and from elementary geometry, $\|P_H(\mathbf{x}_i)\|_2 \leq \sin(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))$ while $\langle \mathbf{y}, \mathbf{x}_i \rangle = \langle \mathbf{y}, P_y\mathbf{x}_i \rangle \geq \cos(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))$ for all $i \in [n]$. Now by definition,

$$\Phi(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \leq \max_{i \in [n]} \angle(\mathbf{y}, \mathbf{x}'_i) = \max_{i \in [n]} \angle(\mathbf{y}, (M\mathbf{x})_i),$$

as angles do not change under positive scaling.

Therefore, we consider the (unnormalized) set of vectors $(M\mathbf{x})_i$. By linearity and the fact M is row-stochastic, we still have $\|P_y(M\mathbf{x})_i\|_2 = \langle \mathbf{y}, (M\mathbf{x})_i \rangle \geq \cos(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))$. On the other hand,

$$\|P_H(M\mathbf{x})_i\| = \|\epsilon P_H(\bar{\mathbf{x}}) + (1 - \epsilon)P_H\tilde{\mathbf{x}}_i\|, \quad (9.7)$$

where $\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$ and $\tilde{\mathbf{x}}_i$ is some arbitrary convex combination. We thus have by linearity

$$\|P_H \tilde{\mathbf{x}}_i\| \leq \sin(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)). \quad (9.8)$$

Moreover, again by linearity, $P_H \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n P_H \mathbf{x}_i$. Recall that by Lemma 9.11, $\mathbf{0} \in \text{conv}(P_H \mathbf{x}_1, \dots, P_H \mathbf{x}_n)$ and therefore by scaling and applying Lemma 9.12 with $\mathbf{z}_i = P_H \mathbf{x}_i$,

$$\|P_H \bar{\mathbf{x}}\|_2 \leq \left(1 - \frac{1}{n}\right) \sin(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)). \quad (9.9)$$

Putting it all together, monotonicity of the function $\tan(\alpha)$ on $[0, \pi/2)$ along with the triangle inequality and Equation (9.7), Equation (9.8), and Equation (9.9) shows that

$$\tan(\Phi(\mathbf{x}'_1, \dots, \mathbf{x}'_n)) \leq \max_{i \in [n]} \frac{\|P_H(M\mathbf{x})_i\|_2}{\|P_Y(M\mathbf{x})_i\|}$$

which in turn is at most

$$\frac{\epsilon \left(1 - \frac{1}{n}\right) \sin(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)) + (1 - \epsilon) \sin(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))}{\cos(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))} = \left(1 - \frac{\epsilon}{n}\right) \tan(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)).$$

We may assume that $\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n) < \pi/2 - \delta$ for some small enough $\delta = \delta(n, d) > 0$ by the assumption that $\mathbf{x}_1, \dots, \mathbf{x}_n$ lie on one side of a hyperplane with margin $\lambda = \lambda(n, d) > 0$. The function $\tan(\cdot)$ has derivative bounded between 1 and some constant depending on n, d on the interval $[0, \pi/2 - \delta]$ and thus by Lemma 9.16, the above geometric decay of $\tan \circ \Phi$ implies that $\Phi(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \leq (1 - c) \cdot \Phi(\mathbf{x}_1, \dots, \mathbf{x}_n)$ for some constant $c = c(n, d, \eta) > 0$. \square

Lemma 9.16. *Let $f : [0, a] \rightarrow [0, b]$ be a differentiable function such that $f(0) = 0$ and $1 \leq f' \leq K$ for some constant K . If $f(x) \leq (1 - c)f(y)$, then $x \leq (1 - c')y$ for $c' = c/2K$.*

Proof. Observe from the Mean Value Theorem, the assumption, and the fact

$f(y) \geq f(x) \geq x$ from the derivative condition that

$$y - x \geq \frac{f(y) - f(x)}{K} \geq \frac{cf(y)}{K} \geq \frac{cx}{K}.$$

It immediately follows that $y \geq \left(1 + \frac{c}{K}\right)x$. As $(1 + z)^{-1} \leq 1 - z/2$ for $z \leq 1$, we obtain the claim. \square

Finally, we can return to the proof of Theorem 9.9:

Proof of Theorem 9.9. By sign-invariance and Lemma 9.10, we may assume that each component of $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \dots, \mathbf{X}_0^{(n)})$ lies strictly on one side of a hyperplane in \mathbb{R}^d with margin at least $\lambda = \lambda(n, d) > 0$; this follows because we may sign the starting configuration arbitrarily, run the dynamics, and then undo the signing by sign-invariance without affecting the polarization properties of the dynamics. We now show that with some constant probability (independent of \mathbf{X}_0), the dynamics monotonically decrease $\Phi(\mathbf{X}_t)$ to zero. Moreover, it is easy to see by the triangle inequality that for any vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$,

$$\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n) \leq \max_{i,j} \angle(\mathbf{x}_i, \mathbf{x}_j) \leq 2\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n). \quad (9.10)$$

Therefore, $\Phi(\mathbf{X}_t) \rightarrow 0$ implies the same of the maximum angle between components and thus implies polarization.

Observe that if the dynamics do not split the vectors on any of the next n iterations, Lemma 9.15 implies that $\Phi(\mathbf{X}_n) \leq (1 - c)\Phi(\mathbf{X}_0)$ for some constant $c > 0$ independent of \mathbf{X}_0 . This occurs with at least some nonzero constant probability δ (again, depending on n, d, η , but not on \mathbf{X}_0) by the margin condition, Lemma 9.8, and the equivalence with the standard Haar measure.

Now, recall from Lemma 9.8 and a union bound that the probability that the dynamics split up any given $\mathbf{x}_1, \dots, \mathbf{x}_n$ that all lie strictly on one side of

a hyperplane with is $O(\max_{i,j} \angle(\mathbf{x}_i, \mathbf{x}_j)) = O(\Phi(\mathbf{x}_1, \dots, \mathbf{x}_n))$ by Equation (9.10), where the implicit constant depends on n, d and the measure of equivalence. The probability that the sequence $\Phi(\mathbf{X}_{t-n}) \rightarrow 0$ is at least the probability that $\Phi(\mathbf{X}_{(t+1)-n}) \leq (1 - c)\Phi(\mathbf{X}_{t-n})$ for each $t \geq 0$. As we have shown that the probability that this fails decays geometrically, it follows that this latter probability is at least

$$\prod_{t=0}^{\infty} (1 - \Theta_{n,d,\eta}((1 - c)')) > c' \quad (9.11)$$

for some constant $c' > 0$ that does not depend on \mathbf{X}_0 , where the inequality follows from the same standard analysis argument as in Theorem 9.4. This immediately implies that strong polarization holds with constant probability on the restriction of the sequence to each n steps. To extend this to the whole sequence, by inspecting the proof of Lemma 9.15, it is easy to see that if the vectors are not split at some time $t + 1$, then $\Phi(\mathbf{X}_{t+1}) \leq \Phi(\mathbf{X}_t)$ for every $t \geq 0$, not just on the subsequence (though we may not have strict contraction). In particular, the event that the dynamics never split up the vectors implies $\Phi(\mathbf{X}_t) \rightarrow 0$, and this holds with constant nonzero probability depending on just n, d, η . By Equation (9.10), Equation (9.11), and Theorem 9.3, we conclude the result. \square

9.4 Markov Chains and Polarization

In the previous section, we showed that the strong polarization observed by HJMR extends to nontrivial variants of geometric opinion dynamics. However, we caution that it is *not* generally true that just any geometric opinion model, even one that is continuous and P -invariant, will strongly polarize from an arbitrary configuration of starting opinions. As a trivial example, suppose that the updates are such that they simply apply a common orthogonal transforma-

tion to each vector at each time. This is clearly continuous and P -invariant, but obviously does not lead polarization except in very pathological examples.

While these trivial examples show that strong polarization does not necessarily hold in such models, it is natural to wonder if it is possible for other dynamics to satisfy other forms of polarization even if they do not satisfy strong polarization. In this section, we consider weaker forms of polarization than the strong form from Definition 9.2 (restated as part (1) in the definition below):

Definition 9.10. Let $\mathbf{X}_t := (\mathbf{X}_t^{(1)}, \dots, \mathbf{X}_t^{(n)})$ be a discrete-time Markov chain as given by eq. (9.1).

1. \mathbf{X}_t **strongly polarizes** (from \mathbf{X}_0) if almost surely, $\rho(\mathbf{X}_t, P) \rightarrow 0$.
2. \mathbf{X}_t **weakly polarizes** (from \mathbf{X}_0) if, for any fixed $\epsilon > 0$, $\Pr(\rho(\mathbf{X}_t, P) \geq \epsilon) \xrightarrow{t \rightarrow \infty} 0$.
3. \mathbf{X}_t **weakly polarizes on average** (from \mathbf{X}_0) if, for any fixed $\epsilon > 0$,

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^T I_{P_\epsilon}(\mathbf{X}_t) \right]}{T} = 1,$$

where I_{P_ϵ} is the indicator of the set $P_\epsilon \triangleq \{\mathbf{z} \in (\mathbb{S}^{d-1})^n : \rho(\mathbf{z}, P) \leq \epsilon\}$.

It is not difficult to see that these forms of convergence are listed in decreasing order of strength:

Proposition 9.17. *Strong polarization implies weak polarization, which in turn implies weak polarization on average.*

Proof. Suppose that \mathbf{X}_t strongly polarizes. Then by definition, for any fixed $\epsilon > 0$, there exists an almost surely finite (but random) T such that for $t \geq T$,

$\rho(\mathbf{X}_t, P) < \epsilon$. For any fixed $\delta > 0$, there exists $K = K(\delta)$ large enough such that $\Pr(T \geq K) < \delta$. It follows that for $t \geq K$,

$$\Pr(\rho(\mathbf{X}_t, P) \geq \epsilon) \leq \Pr(T \geq K) < \delta.$$

As $\delta > 0$ is arbitrary, it follows that the probability of being at least ϵ -far from P tends to zero. As $\epsilon > 0$ is arbitrary, this gives weak polarization.

For the second implication, fix $\delta > 0$ and let K be large enough so that for all $t \geq K$, $\Pr(\rho(\mathbf{X}_t, P) \geq \epsilon/n) \leq \delta/2$. Then for all $T \geq 2K/\delta$,

$$\frac{\mathbb{E} \left[\sum_{t=1}^T I_{P_\epsilon}(\mathbf{X}_t) \right]}{T} \geq \frac{\sum_{t=1}^T \Pr(\rho(\mathbf{X}_t, P) \leq \epsilon/n)}{T} \geq (1 - \delta/2) - K/T \geq 1 - \delta.$$

As this holds for all large enough T and δ was arbitrary, the claim follows. \square

9.4.1 Polarization and Invariant Distributions

With these weaker forms of polarization, we now proceed to show their connection to the theory of invariant distributions in general Markov chains. To begin, we will actually consider the simplest case of $n = 1$ agent, where clearly all forms of polarization are trivial.

Proposition 9.18. *Consider $n = 1$ agent updating opinions via dynamics that are continuous in the current state, i.e. the function $P_{\mathbb{S}^{d-1}}(\mathbf{x} + f_1(\mathbf{x}, \xi))$ is continuous in $\mathbf{x} \in \mathbb{S}^d$ for every fixed realization of ξ . Then there exists an invariant distribution on \mathbb{S}^{d-1} for these dynamics.*

Moreover, assume

1. There exists some point $\mathbf{y} \in \mathbb{S}^{d-1}$ such that for any open neighborhood U of \mathbf{y} ,

and any choice of starting point $x \in \mathbb{S}^{d-1}$ for the dynamics, there exists some $t = t(\mathbf{x}, U)$ such that $\Pr_{\mathbf{x}}(\mathbf{X}_t^{(1)} \in U) > 0$.

2. The dynamics admit a continuous density with respect to the standard Haar measure on \mathbb{S}^{d-1} , i.e. distribution of $\mathbf{X}_1^{(1)}$ given $\mathbf{X}_0^{(1)} = \mathbf{x}$ has a density with respect to standard Haar measure.

Then the above invariant distribution is unique.

Proof. From our continuity assumption, it follows that the dynamics are *Feller*, i.e. for any continuous function $h : \mathbb{S}^{d-1} \rightarrow \mathbb{S}^{d-1}$, the map from \mathbf{x} to $\mathbb{E}[h(\mathbf{x}')]^2$ where \mathbf{x}' takes one step of the dynamics, remains continuous (Definition 2.36 of [81]). Because \mathbb{S}^{d-1} is compact and the dynamics are Feller, it follows that there exists at least one invariant measure (Corollary 4.18 of [81]).

For the second part, the first additional assumption implies there exists an *accessible* point $\mathbf{y} \in \mathbb{S}^{d-1}$, while the second implies that the dynamics satisfy the *strong Feller* property (Definition 2.5 of [82]), i.e. the transition operator maps bounded measurable functions to continuous functions. From Corollary 2.7 of [82], this implies that the dynamics have a unique invariant distribution. \square

We now consider invariant distributions on *the joint opinion vector* and then compare to the above results.

Proposition 9.19. *Consider $n \geq 1$ agents updating opinions, and suppose the dynamics are jointly continuous for every fixed realization of ξ . Then there exists an invariant distribution for the vector $(\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)})$ under these dynamics. Moreover, if the dynamics preserve $\sigma(D)$ for each $\sigma \in \{-1, 1\}^n$, there exists such a distribution supported in $\sigma(D)$ for every $\sigma \in \{-1, 1\}^n$.*

If the dynamics are oblivious and satisfy the stronger assumptions of the previous proposition componentwise, then the i th marginal distribution of every such invariant distribution is unique.

Proof. The first part is identical to the proof of Proposition 9.18 using the Feller property. The second claim holds from the assumed invariance of $\sigma(D)$ and noting that the same argument holds when restricting to this invariant, compact set. The last claim follows from the obliviousness, as then the restriction of any invariant distribution to any component must yield an invariant distribution of the $n = 1$ dynamics. This distribution is unique from Proposition 9.18. \square

Remark 9.20. Note that by Proposition 9.18, it is not true that there exists a unique invariant distribution for $n > 1$ agents for most natural dynamics. The failure of the uniqueness of the invariant distribution arises because the joint dynamics typically will not have either accessible points nor satisfy the strong Feller property, as the set of possible updated values forms a lower-dimensional subset of $\prod_{i=1}^n \mathbb{S}^{d-1}$ and so cannot admit a density.

With these results in hand, we can finally turn to the connection with polarization, which *a priori* is a statement about whether dynamics starting at arbitrary points converge to the polarized set in some suitable sense. The connection is the following:

Theorem 9.21. Suppose that every invariant distribution of the joint dynamics with continuous updates has support contained in P . Then \mathbf{X}_t weakly polarizes on average from any choice of starting vectors.

Proof. The proof of this result is implicit in the proof of the Krylov-Bogolubov theorem (Theorem 4.17 of [81]). Fix any choice \mathbf{X}_0 of starting vectors for the

joint dynamics. Define the family of probability measures for $T = 1, 2, \dots$ by $Q_T(A) := \frac{1}{T} \sum_{t=0}^{T-1} \Pr(\mathbf{X}_t \in A)$. Then because the state space is compact, Prokhorov's Theorem (Theorem 4.15 of [81]) asserts that this sequence of measures contains a weakly convergent subsequence Q_{T_k} for $k = 1, \dots$, with weak limit we denote μ^* ; moreover, from the proof of the Krylov-Bogolubov theorem and the fact that we assumed the dynamics are Feller, μ^* is invariant.

Now, by our assumption, every invariant measure for these dynamics is concentrated on P . For any fixed $\epsilon > 0$, let P_ϵ be as defined above. Because μ^* is supported on P by assumption, $\mu^*(\partial P_\epsilon) = 0$, where ∂A denotes the boundary of a set A . Recall that weak convergence of measures is equivalent to convergence on all continuity sets A that satisfy $\mu^*(\partial A) = 0$. It follows immediately that

$$Q_{T_k}(P_\epsilon) = \frac{\sum_{t=0}^{T_k-1} \Pr(\mathbf{X}_t \in P_\epsilon)}{T_k} \rightarrow \mu^*(P_\epsilon) = 1.$$

Moreover, this same argument holds for *any* subsequence of $\{Q_T\}_{T=0}^\infty$, possibly with a different invariant measure which is nonetheless supported on P by assumption. In particular, for any subsequence of the scalar sequence $\{Q_T(P_\epsilon)\}_{T=1}^\infty$, there exists a further subsequence which converges to 1. From standard analysis, any sequence such that every subsequence contains a further subsequence that converges to a fixed number c must itself converge to c . Therefore, it follows that $Q_T(P_\epsilon) \rightarrow 1$, which by linearity of expectation is the statement of weak polarization on average. \square

The converse always holds even when the dynamics are not continuous (though we do need to assume that there exists *some* invariant distribution).

Theorem 9.22. *Suppose that under the geometric opinion dynamics, there exists an invariant distribution μ^* that puts nontrivial mass on the set of nonpolarized vectors,*

i.e. $\mu^*(P^c) > 0$. Then the opinion dynamics do not weakly converge on average. In particular, if the dynamics are continuous, then the dynamics weakly polarize on average if and only if every invariant distribution is supported on P .

Proof. The assumption and taking limits for $\epsilon \rightarrow 0$ with continuity of measure implies that there exists some $\epsilon > 0$ such that $\mu^*(P_\epsilon^c) > 0$. By invariance, this means that for every $t \geq 0$, $\Pr_{\mathbf{X}_0 \sim \mu^*}(\mathbf{X}_t \in P_\epsilon) = \mu^*(P_\epsilon^c) > 0$, which trivially implies that for every $T \geq 0$,

$$\mathbb{E}_{\mathbf{X}_0 \sim \mu^*} \left[\mathbb{E} \left[\frac{\sum_{t=0}^{T-1} I_{P_\epsilon}}{T} \right] \right] = 1 - \mu^*(P_\epsilon^c),$$

where the inner expectation is over the sequence of random vectors $\{\xi_t\}_{t=1}^T$. By Fatou's lemma,

$$\mathbb{E}_{\mathbf{X}_0 \sim \mu^*} \left[\liminf_{T \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{t=0}^{T-1} I_{P_\epsilon}}{T} \right] \right] \leq \liminf_{T \rightarrow \infty} \mathbb{E}_{\mathbf{X}_0 \sim \mu^*} \left[\mathbb{E} \left[\frac{\sum_{t=0}^{T-1} I_{P_\epsilon}}{T} \right] \right] = 1 - \mu^*(P_\epsilon^c).$$

This implies that there exists some fixed $\mathbf{X}_0 \in \prod_{i=1}^n \mathbb{S}^{d-1}$ such that

$$\liminf_{T \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{t=0}^{T-1} I_{P_\epsilon}}{T} \right] \leq 1 - \mu^*(P_\epsilon^c) < 1,$$

when the process is started at \mathbf{X}_0 . By definition, it follows that the process does not weakly polarize on average with this choice of starting vector. The last claim then follows from combining the above with Theorem 9.21. \square

As an application, HJMR proved that for $d = 2$ and ξ_t is uniform on the unit circle, every finite set of vectors strongly polarizes. By Proposition 9.17, this implies weak polarization on average. Because their dynamics are oblivious, admit a density for each single agent system, and every point is accessible, the unique invariant measure for the dynamics of a single agent is clearly uniform. Theorem 9.22 shows that every invariant measure on the *joint* system must be supported on P .

To summarize, under very natural conditions, these geometric dynamics have invariant distributions in P , essentially by P -invariance of the dynamics. In the oblivious case, the projections of *any* invariant distribution on each component will be unique under natural regularity assumptions and the only ambiguity is in how these distributions are coupled together. If one can show that these P -invariant measures are the *only* such measures, or contrapositively that no invariant measure can put nontrivial mass on the complement of P , then the previous result immediately supplies weak polarization on average.

9.4.2 HJMR Model and Weak Polarization

Given the results in the previous sections, it is perhaps tempting to believe that under some necessary form of irreducibility, strong polarization will arise. In this section, we exhibit a model where *weak polarization* on restricted subsets of initial configurations holds, but strong polarization provably does not. We do so in a very simple version of the original HJMR model by showing how the dynamics in this case relate to basic properties of simple random walks in \mathbb{Z} and more generally, infinite balls-in-bins processes.

In this section, we will consider the HJMR dynamics where the update vector is drawn uniformly from a complete orthonormal set of basis vectors in \mathbb{R}^d ; because the HJMR dynamics commute with orthogonal transformations, we will assume that these are the standard basis vectors. Towards showing that weak, but not strong, polarization holds, we need the following simple calculation showing that the random updates in this process are *commutative*, greatly simplifying updating.

Lemma 9.23. *Let $\mathbf{v} \in \mathbb{S}^{d-1}$ be an arbitrary starting vector. Then for any sequence ξ_1, \dots, ξ_T of update vectors drawn from $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$, the vector \mathbf{z} obtained after T steps of the dynamics given by Equation (9.2) and starting at \mathbf{v} satisfies*

$$\mathbf{z} \propto \left((1 + \eta)^{N_T^{(1)}} v_1, \dots, (1 + \eta)^{N_T^{(d)}} v_d \right),$$

where $N_T^{(i)}$ is the number of times $\xi_t = \mathbf{e}_i$ up to time T .

In particular, with orthonormal update vectors, the vector obtained after T steps of the dynamics depends only on the multiset $\{\xi_1, \dots, \xi_T\}$ and not on the order they arrive.

Proof. Observe that the functional form of the update rule in Equation (9.2) is homogeneous in the starting vector in that scaling the starting vector leads to the same update. As a consequence, one can perform the dynamics as stated by renormalizing after every step, or by applying the update rule without renormalizing until the end.

As a result of this observation, it suffices to show that for any vector $\mathbf{v} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$, the effect of applying the update vector $\xi = \mathbf{e}_i$ is proportional to $(v_1, \dots, (1 + \eta)v_i, \dots, v_d)$. This is sufficient as then we may iterate these updates without renormalizing until only the end of the T unnormalized updates to deduce commutativity. This is now immediate to see from Equation (9.2). \square

Notice that the vector $(N_t^{(1)}, \dots, N_t^{(d)})$ is precisely a balls-in-bins process with d bins and t balls. The identity of Lemma 9.23 shows that to understand the behavior of the process with orthonormal update vectors, we need to understand the sample path properties of an infinite balls-in-bins process. This is done in the following lemma:

Lemma 9.24. Consider an infinite balls-in-bins process with d labeled bins $\{1, \dots, d\}$. Let $N_t^{(i)}$ denote the number of balls in bin i at time t . Then almost surely, the following holds:

1. For each fixed $i \in [d]$, $N_t^{(i)} = \max_{j \in [d]} N_t^{(j)}$ for infinitely many t .
2. For each fixed $i, j \in [d]$, $N_t^{(i)} = N_t^{(j)} = \max_{k \in [d]} N_t^{(k)}$ for infinitely many t .

Proof. Recall from the Hewitt-Savage zero-one law that any event depending on a sequence of i.i.d. random variables that is invariant under finite permutations has probability 0 or 1 (Theorem 2.5.4 of [55]). The first listed event evidently has this property. Clearly by the pigeonhole principle, there surely exists a (random) index $i^* \in [d]$ that is infinitely often the largest, and so by symmetry, the probability that any fixed $i \in [d]$ has this property is at least $1/d$ by a union bound. The Hewitt-Savage zero-one law then implies that the probability is one.

For the second claim, note that the previous claim and the fact that each $N_t^{(i)}$ is nondecreasing in time and changes by at most one in each round implies that infinitely often, there must be at least two indices of bins that are at least as large as the rest (if not, then with positive probability a single index is the unique largest element for all but finitely many times). Again by the pigeonhole principle and symmetry, each fixed pair (i, j) has probability at least $\binom{d}{2}^{-1}$ of having this property, hence by the Hewitt-Savage zero-one law has probability one. □

Lemma 9.25. Consider an infinite balls-in-bins process with d labeled bins $\{1, \dots, d\}$. Let $N_t^{(i)}$ denote the number of balls in bin i at time t . Then for any fixed constant $K \in \mathbb{N}$, the probability that there exists two bins $i \neq j$ such that $|N_t^{(i)} - N_t^{(j)}| \leq K$ is $O(1/\sqrt{t})$, where the implicit constant depends on K, d , but not t .

Proof. Observe that the distribution of $N_t^{(i)} - N_t^{(j)}$ for any fixed $i \neq j$ is equal to that of a lazy, simple random walk on \mathbb{Z} that updates with probability $2/d$. It is well-known (for instance, by the central limit theorem or Littlewood-Offord lemma) that for non-lazy simple random walks on \mathbb{Z} , the probability that the walk has magnitude at most some fixed constant K at time t is $O\left(\frac{1}{\sqrt{t}}\right)$, where the implicit constant depends on K . To account for laziness, note that with all but exponentially small probability, the number of updates to the i or j th bin at time t will exceed, say, t/d . We can then apply the result for non-lazy simple random walks and absorb the error term, just changing the implicit constant to depend on d as well to account for the smaller number of steps. We may then apply a union bound over such pairs. \square

We can now turn to a full analysis of the polarization behavior of this system.

Theorem 9.26. *Consider the HJMR dynamics in \mathbb{S}^{d-1} given by Equation (9.2), where ξ_t is drawn i.i.d. and uniformly from $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$. Let $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}) \in (\mathbb{S}^{d-1})^2$ be such that the two vectors have equal supports but are not equal up to sign. Then:*

1. *The vector \mathbf{X}_t weakly polarizes.*
2. *With probability 1, $\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2 \not\rightarrow 0, 2$. In particular, no nontrivial pair of starting vectors polarizes with nonzero probability.*

Remark 9.27. *If the distribution of random updates is not uniform, but instead has strictly unequal probabilities associated to each basis vector, then it is easy to amend the argument to show that now strong polarization holds between any two starting vectors with the same support, simply because both vectors will get concentrated about the basis vector in their support with the largest probability.*

Moreover, note that there do exist invariant distributions in this system, considered as $\prod_{i=1}^n \mathbb{S}^{d-1}$, that are not supported on P . However, this does not contradict Theorem 9.22 because Theorem 9.26 only applies to vectors with equal supports, a full measure set.

Proof of Theorem 9.26. Let $\mathbf{X}_0 = (\mathbf{X}_0^{(1)}, \mathbf{X}_0^{(2)}) \in (\mathbb{S}^{d-1})^2$ be as stated. Recall that to show weak polarization, we must show that for any fixed $\epsilon > 0$,

$$\Pr\left(\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2 \wedge \|\mathbf{X}_t^{(1)} + \mathbf{X}_t^{(2)}\|_2 > \epsilon\right) = o(1).$$

Recall that from Lemma 9.23, the (unnormalized) i th component of each vector is amplified multiplicatively by a factor of $(1+\eta)^{N_t^{(i)}}$ before renormalization. From this, it is easy to see that for any fixed \mathbf{X}_0 and $\epsilon > 0$, there exists some finite $K = K(\mathbf{X}_0, \epsilon) \in \mathbb{N}$ such that if $N_t^{(i)} \geq N_t^{(j)} + K$ for some $i \in [d]$ and all $j \neq i$, then both components of \mathbf{X}_t are individually $\epsilon/2$ -close to $\pm \mathbf{e}_i$. By the triangle inequality, it would then follow that \mathbf{X}_t satisfies $\|\mathbf{X}_t^{(1)} - \mathbf{X}_t^{(2)}\|_2 \wedge \|\mathbf{X}_t^{(1)} + \mathbf{X}_t^{(2)}\|_2 \leq \epsilon$. Applying Lemma 9.25 with this value of K , we see that the probability of this event is $1 - O\left(\frac{1}{\sqrt{t}}\right)$, where the constant depends only on K, d and hence tends to 1 with t as needed. As $\epsilon > 0$ was arbitrary, this proves the first claim.

For the second, we rely on the other properties we have already shown of the infinite balls-in-bins process. By assumption, as $\mathbf{X}_0^{(1)}$ and $\mathbf{X}_0^{(2)}$ are not equal up to sign, there exists a pair of indices such that the vectors restricted to these indices are not equal up to a constant. Without loss of generality, we assume that these indices are 1, 2. By Lemma 9.24, there almost surely exist infinitely many times t such that $N_t^{(1)} = N_t^{(2)} = \max_{k \in [d]} N_t^{(k)}$. By the unnormalized form of the directions at each such time t given by Lemma 9.23, it follows that at each such time, the first and second components of each vector remain proportional to their value at the beginning of this process but are now scaled by a factor

that is at least 1 after normalization. In particular, at each such time t where this holds, the distance between $\mathbf{X}_t^{(1)}$ and $\mathbf{X}_t^{(2)}$ exceeds some strictly positive constant. Because this occurs infinitely often almost surely, the claim follows. \square

9.5 Chapter Notes

The results of this chapter originally appear in [68], joint with Jon Kleinberg and Éva Tardos. The potential function used in the proof of Theorem 9.9 was graciously suggested by Ron Peretz; the original argument relied on abstract compactness-type results that did not furnish quantitative bounds. Compared to the original version, the proof of Lemma 9.12 has been corrected.

9.5.1 Related Work

As stated above, the particular geometric model of opinion dynamics considered by Hazła, Jin, Mossel, and Ramnarayan [89] motivates many of the considerations in this work. To our knowledge, their work is first to introduce the dynamics in Equation (9.1) in the specific form given by Equation (9.2).

Polarization more generally has been studied in other models of opinion dynamics, typically in the more well-understood DeGroot and Friedkin-Johnsen dynamics as considered in the previous chapter (see Theorem 8.2 for such a result). However, the contractivity of these models makes them poorly suited for understanding polarization. Recent theoretical work as in the previous chapter has attempted to combine the analytically desirable features of these models with questions of polarization by incorporating new elements to the model. For

instance, behavioral biases like *biased assimilation*, which underlies the intuition behind the HJMR model, have also been fruitfully studied in the context of De-Groot dynamics [44]. As in the previous chapter, a related line of work has also attempted to understand polarization by incorporating the influence of external actors or platforms into the dynamics. Polarization then arises either explicitly or implicitly because of the objectives of these external parties [107, 40, 9]. A related model by Hegselmann and Krause [90] circumvents near-consensus by allowing agents to filter out overly dissimilar opinions before averaging at each step, but already this modest extension incurs a high cost: many basic questions about convergence remain open in this model, greatly complicating a more extensive understanding of how polarization arises [21].

As we discuss in Section 9.4, geometric models of opinion dynamics are a special case of the more general theory of Markov chains in non-discrete state spaces. The results in this setting are significantly more involved than corresponding results on discrete state spaces; an accessible reference to some important results in this area can be found in the surveys of Hairer [81, 82]. As we show, polarization is related to the set of invariant distributions on the Markov chain induced by the dynamics. Numerous techniques have been developed to determine the uniqueness of an invariant distribution in these settings, as well as the rate and mode of convergence to it if possible (see for instance, [32]). It would be interesting to find new ways to apply results from this area to provide quantitative bounds on the convergence to polarization when strong polarization holds, particularly in models with network effects.

BIBLIOGRAPHY

- [1] *United States of America v. Internet Research Agency, LLC*. 2018.
- [2] Daron Acemoglu and Asuman E. Ozdaglar. Opinion Dynamics and Learning in Social Networks. *Dyn. Games Appl.*, 1(1):3–49, 2011.
- [3] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with Constraints. In Ioannis Caragiannis, Vahab S. Mirrokni, and Evdokia Nikolova, editors, *Web and Internet Economics - 15th International Conference, WINE 2019, New York, NY, USA, December 10-12, 2019, Proceedings*, volume 11920 of *Lecture Notes in Computer Science*, pages 17–30. Springer, 2019.
- [4] Noga Alon and Nabil Kahale. A Spectral Technique for Coloring Random 3-Colorable Graphs. *SIAM Journal on Computing*, 26(6):1733–1748, 1997.
- [5] Noga Alon and Assaf Naor. Approximating the Cut-Norm via Grothendieck’s Inequality. *SIAM J. Comput.*, 35(4):787–803, 2006.
- [6] Nir Andelman, Michal Feldman, and Yishay Mansour. Strong Price of Anarchy. *Games Econ. Behav.*, 65(2):289–317, 2009.
- [7] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
- [8] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-Optimal Network Design with Selfish Agents. *Theory Comput.*, 4(1):77–109, 2008.
- [9] Jerry Anunrojwang, Ozan Candogan, and Nicole Immorlica. Social Learning Under Platform Influence: Extreme Consensus and Persistent Disagreement. *SSRN*, 2020.
- [10] Sanjeev Arora, Elad Hazan, and Satyen Kale. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory Comput.*, 8(1):121–164, 2012.
- [11] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The Nonstochastic Multiarmed Bandit Problem. *SIAM J. Comput.*, 32(1):48–77, 2002.

- [12] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The Price of Routing Unsplittable Flow. *SIAM J. Comput.*, 42(1):160–177, 2013.
- [13] Yossi Azar, Michal Feldman, Nick Gravin, and Alan Roytman. Liquid Price of Anarchy. In Vittorio Bilò and Michele Flammini, editors, *Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12-14, 2017, Proceedings*, volume 10504 of *Lecture Notes in Computer Science*, pages 3–15. Springer, 2017.
- [14] Moshe Babaioff, Richard Cole, Jason D. Hartline, Nicole Immorlica, and Brendan Lucier. Non-Quasi-Linear Agents in Quasi-Linear Mechanisms (extended abstract). In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 84:1–84:1. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [15] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with Knapsacks. *J. ACM*, 65(3):13:1–13:55, 2018.
- [16] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget-Management Strategies in Repeated Auctions. *Operations Research*, 69(3):859–876, 2021.
- [17] Santiago R. Balseiro and Yonatan Gur. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Manag. Sci.*, 65(9):3952–3968, 2019.
- [18] Santiago R. Balseiro, Christian Kroer, and Rachitesh Kumar. Contextual Standard Auctions with Budgets: Revenue Equivalence and Efficiency Guarantees. In David M. Pennock, Ilya Segal, and Sven Seuken, editors, *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, page 476. ACM, 2022.
- [19] Joshua Batson, Daniel A Spielman, and Nikhil Srivastava. Twice-Ramanujan Sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.
- [20] Lucas Baudin, Marco Scarsini, and Xavier Venel. Strategic Behavior and No-Regret Learning in Queueing Systems. *CoRR*, abs/2302.03614, 2023.
- [21] Arnab Bhattacharyya, Mark Braverman, Bernard Chazelle, and Huy L. Nguyen. On the Convergence of the Hegselmann-Krause System. In Robert D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 61–66. ACM, 2013.

- [22] P. Billingsley. *Probability and Measure, 3rd Edition*. Wiley Series in Probability and Mathematical Statistics. Wiley, 2008.
- [23] Yonatan Bilu and Nathan Linial. Lifts, Discrepancy and Nearly Optimal Spectral Gap. *Combinatorica*, 26(5):495–519, 2006.
- [24] David Bindel, Jon Kleinberg, and Sigal Oren. How Bad is Forming Your Own Opinion? *Games and Economic Behavior*, 92:248–265, 2015.
- [25] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret Minimization and the Price of Total Anarchy. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 373–382. ACM, 2008.
- [26] Shant Boodaghians, Joshua Brakensiek, Samuel B. Hopkins, and Aviad Rubinfeld. Smoothed Complexity of 2-player Nash Equilibria. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 271–282. IEEE, 2020.
- [27] Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of Bid Optimization in Online Advertisement Auctions. In Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider, and Prashant J. Shenoy, editors, *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pages 531–540. ACM, 2007.
- [28] Allan Borodin, Jon M. Kleinberg, Prabhakar Raghavan, Madhu Sudan, and David P. Williamson. Adversarial Queuing Theory. *J. ACM*, 48(1):13–38, 2001.
- [29] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.
- [30] David A Broniatowski, Amelia M Jamison, SiHua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C Quinn, and Mark Dredze. Weaponized Health Communication: Twitter Bots and Russian Trolls Amplify the Vaccine Debate. *American Journal of Public Health*, 108(10):1378–1384, 2018.
- [31] G. W. Brown. Iterative Solutions of Games by Fictitious Play. In *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, 1951.

- [32] Oleg Butkovsky. Subgeometric Rates of Convergence of Markov processes in the Wasserstein Metric. *Annals of Applied Probability*, 24(2):526–552, 2014.
- [33] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical Physics of Social Dynamics. *Reviews of Modern Physics*, 81(2):591, 2009.
- [34] Matteo Castiglioni, Andrea Celli, and Christian Kroer. Online Learning with Knapsacks: The Best of Both Worlds. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 2767–2783. PMLR, 2022.
- [35] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [36] Moses Charikar, Howard J. Karloff, Claire Mathieu, Joseph Naor, and Michael E. Saks. Online Multicast with Egalitarian Cost Sharing. In Friedhelm Meyer auf der Heide and Nir Shavit, editors, *SPAA 2008: Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, Munich, Germany, June 14-16, 2008*, pages 70–76. ACM, 2008.
- [37] Mayee Chen and Miklos Z Racz. Network Disruption: Maximizing Disagreement and Polarization in Social Networks. *arXiv preprint arXiv:2003.08377*, 2020.
- [38] Xi Chen, Christian Kroer, and Rachitesh Kumar. The Complexity of Pacing for Second-Price Auctions. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, page 318. ACM, 2021.
- [39] Xi Chen, Christian Kroer, and Rachitesh Kumar. Throttling Equilibria in Auction Markets. In Michal Feldman, Hu Fu, and Inbal Talgam-Cohen, editors, *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings*, volume 13112 of *Lecture Notes in Computer Science*, page 551. Springer, 2021.
- [40] Uthsav Chitra and Christopher Musco. Analyzing the Impact of Filter Bubbles on Social Network Polarization. In James Caverlee, Xia (Ben) Hu, Mounia Lalmas, and Wei Wang, editors, *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 115–123. ACM, 2020.

- [41] Joel E. Cohen and Paul Horowitz. Paradoxical Behaviour of Mechanical and Electrical Networks. *Nature*, 352(6337):699–701, 1991.
- [42] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolás E. Stier Moses, and Chris Wilkens. Pacing Equilibrium in First-Price Auction Markets. In Anna Karlin, Nicole Immorlica, and Ramesh Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, page 587. ACM, 2019.
- [43] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolás E. Stier Moses. Multiplicative Pacing Equilibria in Auction Markets. *Oper. Res.*, 70(2):963–989, 2022.
- [44] Pranav Dandekar, Ashish Goel, and David T. Lee. Biased Assimilation, Homophily, and the Dynamics of Polarization. *Proc. Natl. Acad. Sci. USA*, 110(15):5791–5796, 2013.
- [45] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- [46] Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. The Complexity of Markov Equilibrium in Stochastic Games. *CoRR*, abs/2204.03991, 2022.
- [47] Morris H. DeGroot. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [48] Ofer Dekel, Ambuj Tewari, and Raman Arora. Online Bandit Learning Against an Adaptive Adversary: from Regret to Policy Regret. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012.
- [49] Daniel DellaPosta, Yongren Shi, and Michael Macy. Why Do Liberals Drink Lattes? *American Journal of Sociology*, 120(5):1473–1511, 2015.
- [50] Charles Delorme and Svatopluk Poljak. Laplacian Eigenvalues and the Maximum Cut Problem. *Mathematical Programming*, 62(1-3):557–574, 1993.
- [51] Erik D. Demaine, Mohammad Taghi Hajiaghayi, Hamid Mahini, and

- Morteza Zadimoghaddam. The Price of Anarchy in Network Creation Games. *ACM Trans. Algorithms*, 8(2):13:1–13:13, 2012.
- [52] Peter M. DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel. Persuasion Bias, Social Influence, and Unidimensional Opinions*. *The Quarterly Journal of Economics*, 118(3):909–968, 08 2003.
- [53] Renee DiResta, Kris Shaffer, Becky Ruppel, David Sullivan, Robert Matney, Ryan Fox, Jonathan Albright, and Ben Johnson. The Tactics & Tropes of the Internet Research Agency. 2019.
- [54] Shahar Dobzinski and Renato Paes Leme. Efficiency Guarantees in Auctions with Budgets. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 392–404. Springer, 2014.
- [55] Rick Durrett. *Probability: Theory and Examples*, volume 49. Cambridge University Press, 2019.
- [56] Kfir Eliaz and Ran Spiegler. A Model of Competing Narratives. *American Economic Review*, 110(12):3786–3816, December 2020.
- [57] Giannis Fikioris and Éva Tardos. Liquid Welfare Guarantees for No-Regret Learning in Sequential Budgeted Auctions. *CoRR*, abs/2210.07502, 2022.
- [58] Giannis Fikioris and Éva Tardos. Approximately Stationary Bandits with Knapsacks. *CoRR*, abs/2302.14686, 2023.
- [59] Jerzy Filar and Koos Vrieze. *Competitive Markov Decision Processes*. Springer Science & Business Media, 2012.
- [60] Ronald J Fisher. Towards a Social-Psychological Model of Intergroup Conflict. In *Ronald J. Fisher: A North American Pioneer in Interactive Conflict Resolution*, pages 73–86. Springer, 2016.
- [61] Dimitris Fotakis, Kyriakos Lotidis, and Chara Podimata. A Bridge between Liquid and Social Welfare in Combinatorial Auctions with Submodular Bidders. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial In-*

telligence Conference, IAAI 2019, The Ninth AAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019, pages 1949–1956. AAAI Press, 2019.

- [62] Daniel Freund, Thodoris Lykouris, and Wentao Weng. Efficient Decentralized Multi-Agent Learning in Asymmetric Queuing Systems. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 4080–4084. PMLR, 2022.
- [63] Noah E Friedkin and Eugene C Johnsen. Social Influence and Opinions. *J. Math. Soc.*, 15(3-4):193–206, 1990.
- [64] Hu Fu, Qun Hu, and Jia’nan Lin. Stability of Decentralized Queueing Networks Beyond Complete Bipartite Cases. In Kristoffer Arnsfelt Hansen, Tracy Xiao Liu, and Azarakhsh Malekian, editors, *Web and Internet Economics - 18th International Conference, WINE 2022, Troy, NY, USA, December 12-15, 2022, Proceedings*, volume 13778 of *Lecture Notes in Computer Science*, pages 96–114. Springer, 2022.
- [65] Drew Fudenberg and David Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [66] Drew Fudenberg and Jean Tirole. *Game Theory*, volume 1 of *MIT Press Books*. The MIT Press, December 1991.
- [67] Jason Gaitonde, Jon M. Kleinberg, and Éva Tardos. Adversarial Perturbations of Opinion Dynamics in Networks. In Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *EC ’20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 471–472. ACM, 2020.
- [68] Jason Gaitonde, Jon M. Kleinberg, and Éva Tardos. Polarization in Geometric Opinion Dynamics. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC ’21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 499–519. ACM, 2021.
- [69] Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget Pacing in Repeated Auctions: Regret and Efficiency Without Convergence. *CoRR*, abs/2205.08674, 2022.

- [70] Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget Pacing in Repeated Auctions: Regret and Efficiency Without Convergence. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPICs*, pages 52:1–52:1. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [71] Jason Gaitonde and Éva Tardos. Stability and Learning in Strategic Queuing Systems. In Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 319–347. ACM, 2020.
- [72] Jason Gaitonde and Éva Tardos. Virtues of Patience in Strategic Queuing Systems. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 520–540. ACM, 2021.
- [73] Andrea Galeotti, Benjamin Golub, and Sanjeev Goyal. Targeting Interventions in Networks. *arXiv preprint arXiv:1710.06026*, 2017.
- [74] Matthew Gentzkow, Jesse M Shapiro, Matt Taddy, et al. *Measuring Polarization in High-Dimensional Data: Method and Application to Congressional Speech*. National Bureau of Economic Research, 2016.
- [75] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. Opinion Maximization in Social Networks. In *Proceedings of the 13th SIAM International Conference on Data Mining, May 2-4, 2013. Austin, Texas, USA*, pages 387–395. SIAM, 2013.
- [76] Chris Godsil and Gordon F Royle. *Algebraic Graph Theory*, volume 207. Springer Science & Business Media, 2013.
- [77] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. Polyhedral Clinching Auctions and the Adwords Polytope. *Journal of the ACM (JACM)*, 62(3):1–27, 2015.
- [78] Michel X Goemans and David P Williamson. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems using Semidefinite Programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [79] Benjamin Golub and Matthew O Jackson. Naive Learning in Social Net-

- works and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2(1):112–49, 2010.
- [80] Fan Chung Graham. *Spectral Graph Theory*. Number 92. American Mathematical Soc., 1997.
- [81] Martin Hairer. *Ergodic Properties of Markov Processes*. 2006.
- [82] Martin Hairer. *Convergence of Markov Processes*. 2010.
- [83] James Hannan. *Approximation to Bayes Risk in Repeated Play*, pages 97–140. Princeton University Press, Princeton, 1958.
- [84] S. Hart and A. M. Colell. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [85] Refael Hassin. *Rational Queueing*. Chapman and Hall/CRC Series in Operations Research Series. CRC Press, 2020.
- [86] Refael Hassin and Moshe Haviv. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. International Operations Resea. Springer US, 2003.
- [87] Johan Hstad. Some Optimal Inapproximability Results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [88] Elad Hazan. Introduction to Online Convex Optimization. *Found. Trends Optim.*, 2(3-4):157–325, 2016.
- [89] Jan Hazła, Yan Jin, Elchanan Mossel, and Govind Ramnarayan. A Geometric Model of Opinion Polarization. *CoRR*, abs/1910.05274, 2019.
- [90] Rainer Hegselmann and Ulrich Krause. Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artificial Societies and Social Simulation (JASSS) vol*, 5(3), 2002.
- [91] Ruben Hoeksma and Marc Uetz. The Price of Anarchy for Utilitarian Scheduling Games on Related Machines. *Discret. Optim.*, 31:29–39, 2019.
- [92] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander Graphs and Their Applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

- [93] Nicole Immorlica, Karthik Abinav Sankararaman, Robert E. Schapire, and Aleksandrs Slivkins. Adversarial Bandits with Knapsacks. *J. ACM*, 69(6):40:1–40:47, 2022.
- [94] Chi Jin, Qinghua Liu, Yuanhao Wang, and Tiancheng Yu. V-Learning - A Simple, Efficient, Decentralized Algorithm for Multiagent RL. *CoRR*, abs/2110.14555, 2021.
- [95] Yaonan Jin and Pinyan Lu. First Price Auction is $1 - 1/e^2$ Efficient. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 179–187. IEEE, 2022.
- [96] Ramesh Johari and John N. Tsitsiklis. Efficiency Loss in a Network Resource Allocation Game. *Math. Oper. Res.*, 29(3):407–435, 2004.
- [97] Adam Tauman Kalai and Santosh S. Vempala. Efficient Algorithms for Online Decision Problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005.
- [98] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [99] Elias Koutsoupias and Christos H. Papadimitriou. Worst-Case Equilibria. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.
- [100] Subhashini Krishnasamy, Rajat Sen, Ramesh Johari, and Sanjay Shakkottai. Regret of Queueing Bandits. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1669–1677, 2016.
- [101] Pinyan Lu and Tao Xiao. Liquid Welfare Maximization in Auctions with Multiple Items. In Vittorio Bilò and Michele Flammini, editors, *Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L’Aquila, Italy, September 12-14, 2017, Proceedings*, volume 10504 of *Lecture Notes in Computer Science*, pages 41–52. Springer, 2017.
- [102] Brendan Lucier and Allan Borodin. Price of Anarchy for Greedy Auctions. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 537–553. SIAM, 2010.

- [103] Brendan Lucier, Sarath Pattathil, Aleksandrs Slivkins, and Mengxiao Zhang. Autobidders with Budget and ROI Constraints: Efficiency, Regret, and Pacing Dynamics. *CoRR*, abs/2301.13306, 2023.
- [104] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016.
- [105] Albert W Marshall, Ingram Olkin, and Barry C Arnold. *Inequalities: Theory of Majorization and Its Applications*, volume 143. Springer, 1979.
- [106] Michael J McCallion. In-Groups and Out-Groups. *The Blackwell Encyclopedia of Sociology*, 2007.
- [107] Cameron Musco, Christopher Musco, and Charalampos E. Tsourakakis. Minimizing Polarization and Disagreement in Social Networks. In Pierre-Antoine Champin, Fabien L. Gandon, Mounia Lalmas, and Panagiotis G. Ipeirotis, editors, *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, pages 369–378. ACM, 2018.
- [108] National Research Council Committee on Offensive Information Warfare. *Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities*. National Academies Press, 2009.
- [109] Denis Nekipelov, Vasilis Syrgkanis, and Éva Tardos. Econometrics for Learning Agents. In Tim Roughgarden, Michal Feldman, and Michael Schwarz, editors, *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 1–18. ACM, 2015.
- [110] Abraham Neyman and Sylvain Sorin. *Stochastic Games and Applications*, volume 570. Springer Science & Business Media, 2003.
- [111] Robin Pemantle and Jeffrey S Rosenthal. Moment Conditions for a Sequence with Negative Drift to be Uniformly Bounded in L . *Stochastic Processes and their Applications*, 82(1):143–155, 1999.
- [112] Julia Robinson. An Iterative Method of Solving a Game. *Annals of Mathematical Statistics*, 54:296–301, 1851.

- [113] Tim Roughgarden. The Price of Anarchy is Independent of the Network Topology. *J. Comput. Syst. Sci.*, 67(2):341–364, 2003.
- [114] Tim Roughgarden. Intrinsic Robustness of the Price of Anarchy. *J. ACM*, 62(5):32:1–32:42, 2015.
- [115] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, USA, 1st edition, 2016.
- [116] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. The Price of Anarchy in Auctions. *J. Artif. Intell. Res.*, 59:59–101, 2017.
- [117] Tim Roughgarden and Éva Tardos. How Bad is Selfish Routing? *J. ACM*, 49(2):236–259, 2002.
- [118] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24. Springer Science & Business Media, 2003.
- [119] United States Senate Select Committee on Intelligence. *Russian Active Measures Campaigns and Interference in the 2016 U.S. Election Volume 2: Russia’s Use of Social Media with Additional Views*.
- [120] Flore Sentenac, Etienne Boursier, and Vianney Perchet. Decentralized Learning in Online Queuing Systems. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 18501–18512, 2021.
- [121] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [122] Muzafer Sherif, O.J. Harvey, Jack B. White, William R. Hood, and Carolyn W. Sherif. *Intergroup Conflict and Cooperation. The Robbers Cave Experiment [1954]*. University of Oklahoma Book Exchange, Norman, 1961.
- [123] John F Shortle, James M Thompson, Donald Gross, and Carl M Harris. *Fundamentals of Queueing Theory*. Wiley, 2018.
- [124] Ziang Song, Song Mei, and Yu Bai. When Can We Learn General-Sum Markov Games with a Large Number of Players Sample-Efficiently? In *The Tenth International Conference on Learning Representations, ICLR, 2022*.

- [125] Daniel A Spielman and Shang-Hua Teng. Spectral Sparsification of Graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- [126] Vasilis Syrgkanis and Éva Tardos. Composable and Efficient Mechanisms. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 211–220. ACM, 2013.
- [127] Luca Trevisan. Max Cut and the Smallest Eigenvalue. *SIAM Journal on Computing*, 41(6):1769–1786, 2012.
- [128] Twitter. Information Operations Directed at Hong Kong, August 2019.
- [129] Neil Walton and Kuang Xu. *Learning and Information in Stochastic Networks and Queues*, chapter 6, pages 161–198.
- [130] Robert Wilson. Computing Equilibria of N-Person Games. *SIAM Journal on Applied Mathematics*, 21(1):80–87, 1971.
- [131] Carsten Witt. Fitness Levels with Tail Bounds for the Analysis of Randomized Search Heuristics. *Inf. Process. Lett.*, 114(1-2):38–41, 2014.